THE EUROPEAN COMPUTER NETWORK PROJECT

D.L.A. Barber
National Physical Laboratory
Teddington, Middlesex, England

In November 1971 the Ministers of eight European nations; France, Yugoslavia, Italy, Norway, Portugal, Switzerland, Sweden, and the United Kingdom, together with the Euratom Centre at Ispra, signed an agreement to start a project aimed at building a European computer network. The network, which will use the store-and-forward packet-switching techniques, will initially join five data processing research centres in four countries; but later is expected to be extended to include centres in other Nations.

The paper sets out the arguments that led to the decision to go ahead with a European Computer Network, and discusses the form it will take, the functions it will perform, and the way the project will be conducted.

Introduction

Early in 1969 a working group was established by the European Economic Community to examine research in science and technology. This working group, chaired by a Frenchman, Monsieur Aigrain, made wide ranging proposals for various advanced projects in science and technology to be carried out jointly by members of the community. In April 1970 the EEC working group was widened to include representatives of other European countries and the resulting group of nineteen nations became known as the COST Group. (Co-operation Europeene dens le Domaine de la Recherche Scientifique et Technique).

The COST Group set up a number of expert study groups to consider, in detail, the various projects. One of these, known as Project 11, was concerned with the problems of communication between computers. The Study Group, chaired by the author, put forward a plan to establish a Pilot Computer Communications Network in Europe, to link a number of data processing research centres. This was accepted by several of the member nations, who signed in November 1972 an agreement to undertake a project aimed at building such a network. Initially, the network will link together centres in France, Italy, Switzerland and the United Kingdom, but later is likely to be extended to include centres in other Nations.

This paper sets out the arguments that led to the decision to go ahead with a European Computer network, and discusses the form it will take, the functions it will perform, and the way the projects will be conducted.

The Need for a Network

It is now well accepted that the proper use of computers is one of the most vital factors affecting the economic growth of a nation. In the mid 1960's the association of telecommunications facilities with computers vastly increased their effectiveness and, since then, the use of teleprocessing techniques has grown at a remarkable rate. There is little doubt that use made of communication networks linking users to remote computers will continue to increase sharply during the next few years, and that the transmission of data across Europe will become of vital importance to the economics of European countries in general. The pattern which is beginning to emerge in national data communications will be duplicated across national boundaries, and there are already several instances of private computer communication networks that incorporate international data links.

However, the growing numbers of private data communications networks owned by Industry, Commerce and Government is a cause for concern, because private networks often under utilise telecommunications resources. And, because the number of connections to remote computers is tending to double every year, various national proposals have been made for more efficient networks which may be shared by many users.

When an attempt is made to connect several computer systems by a shared data network fundamental incompatibilities are usually revealed between them. The development of shared networks, therefore, focuses attention on the need for really effective agreement of standards for communication between, and with, computer systems.

The provision of efficient shared data communications facilities can also have a useful effect on advanced computer system development by allowing co-operation between research centres. An excellent example is the Advanced Research Projects Agency (ARPA) network[1] sponsored by the Department of Defence of the United States. The ARPA network is clearly having a most beneficial influence upon the development of computing in America.

Networks of a similar type have been proposed and investigated by Universities and Research Establishments in Europe and there are various plans under consideration for building such networks in some countries. There could be many advantages if these networks were joined by an International network, and this would be

COMMUNICATIONS CHALLENGES

for the 80s

By Dr. Robert R. Fossum
and
Dr. Vinton G. Cerf

THE HISTORY OF electromagnetic communications has, in general, been a history of making things happen faster. It is not surprising, therefore, that computers and electronics are playing an increasingly central role in communications of all types. Indeed, the transistor was invented in the mid-1940s at the Bell Telephone Laboratories and formed the heart of the "computer revolution" of the 1950s and 1960s.

There are two distinct technological threads in the recent historical fabric of electromagnetic communication. The first of these is the use of computers and other special electronics to manage communications circuits. For example, in the telephone system, manual or mechanical switches were replaced by computer-controlled circuit switches, where this was economically feasible. The rapidly dropping cost of electronics is making such replacement more and more cost effective.

Following this thread a little further, we discover that, for all their utility, the automated circuit switches only bridged half of the gap from the manual/mechanical world to the electronic world. The other half of the gap was to move from analog to digital transmission methods. Not only was this mode more compatible with the micro-electronics world, but it also brought substantial opportunities to improve the quality of communication. Physical problems that cause distortion in analog signals could be defended against once the signal was digitized, through the use of pulse-shaping, coding, error correction and the like. So the first thread has led us to the notion of all-digital communications systems.

To trace the second thread, we must go back again to the 1940s and the invention of the electronic computer. By the end of that decade, it was apparent that electronic computing could be used for many commercial as well as military applications. By the late 1950s, the transistor had begun to replace the vacuum

Dr. Robert R. Fossum was appointed to his present position as Director, Defense Advanced Research Projects Agency (DARPA) in 1977. Previously he was the Dean of Science and Engineering at the Naval Postgraduate School, Monterey, CA.
Dr. Vinton G. Cerf has been a Program Manager, Information Processing Techniques Office, Defense Advanced Research Projects Agency (DARPA), since 1976. Prior to his present position he was an Assistant Professor of Computer Science and Electrical Engineering at Stanford University.

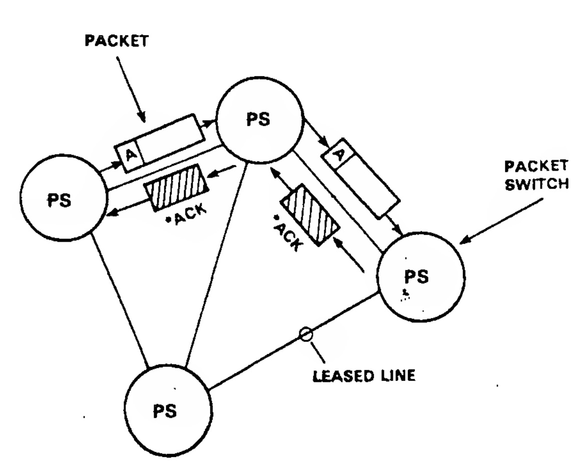

Figure 1. Packet switching

tube in general purpose electronic computers, but the costs of these machines was still very high. In the early 1960s, in the interest of making more efficient use of these expensive machines, the Defense Advanced Research Projects Agency (DARPA) began experimenting with the idea of "interactive computing." The goal of this research was to support rapid, on-line interaction between a user at a terminal and a computer. It soon became apparent that the methods used to provide this interactive service also allowed a computer to be shared by a number of users at once. This notion developed into the concept of "time-sharing" in the 1960s. During this same period, access to distant computing resources, known as "teleprocessing," was also developed to permit remote terminals to access time-shared, central processing facilities.

In the late 1960s, the two threads began to intertwine, each reinforcing and influencing the other. The most relevant development was the concept of "packet switching" instead of circuit switching.[1] The basic concept was not really new, but the implementation was. Rather than using computers to set-up circuits for point-to-point communication through a circuit switching network, a packet switching system forwards short blocks of data, a few hundreds or thousands of bits long, through a network of interconnected computers, called packet switches. Each data block, or "packet," carries addressing information which can be interpreted by each intermediate packet switch. Each packet is stored briefly upon arrival at the next switch, after which it is forwarded to the next switch over a circuit or radio or satellite channel (Figure 1).

The original motivation behind the development of this technology was to find some economical way to interconnect many of the DARPA-funded computing resources so that these resources could be shared on a national basis by all DARPA-sponsored researchers. The result of this effort was the creation of the ARPANET.[2] ARPANET is a nationwide, packet-switched network, based on 50 kilobit/sec leased telephone circuits and mini-computer packet switches called Interface Message Processors (IMPs)[3] by the contractor who built the network, Bolt Beranek and Newman (Figure 2).

The fundamental theme of packet switching is economic. The interconnection of a large number of low duty-cycle computers can be achieved with far fewer, shared, wideband circuits using packet switching technology than with circuit switching.

Packet switching permits on the order of ~3N/2 trunks to be dynamically shared by N hosts, rather than the N(N-1)/2 trunks that would be needed to guarantee full conductivity. Switching times are on the order of hundreds of microseconds for the packet switched technology rather than seconds or tens of seconds for present day circuit switches.

Finally, the cost of the computing power required to share wideband communication circuits through packet switching is dropping much faster than the cost of transmission capacity so packet switching can "pay its way" with relative ease.

Internetting

The packet switching idea extends to many different transmission media, and DARPA has successfully demonstrated packet switching techniques using a single satellite channel shared among several earth stations[4] and using a single, broadband radio channel shared among a number of mobile "packet radios."[5]

Xerox Corporation has demonstrated that this sharing is also effective on a common, broadband co-axial

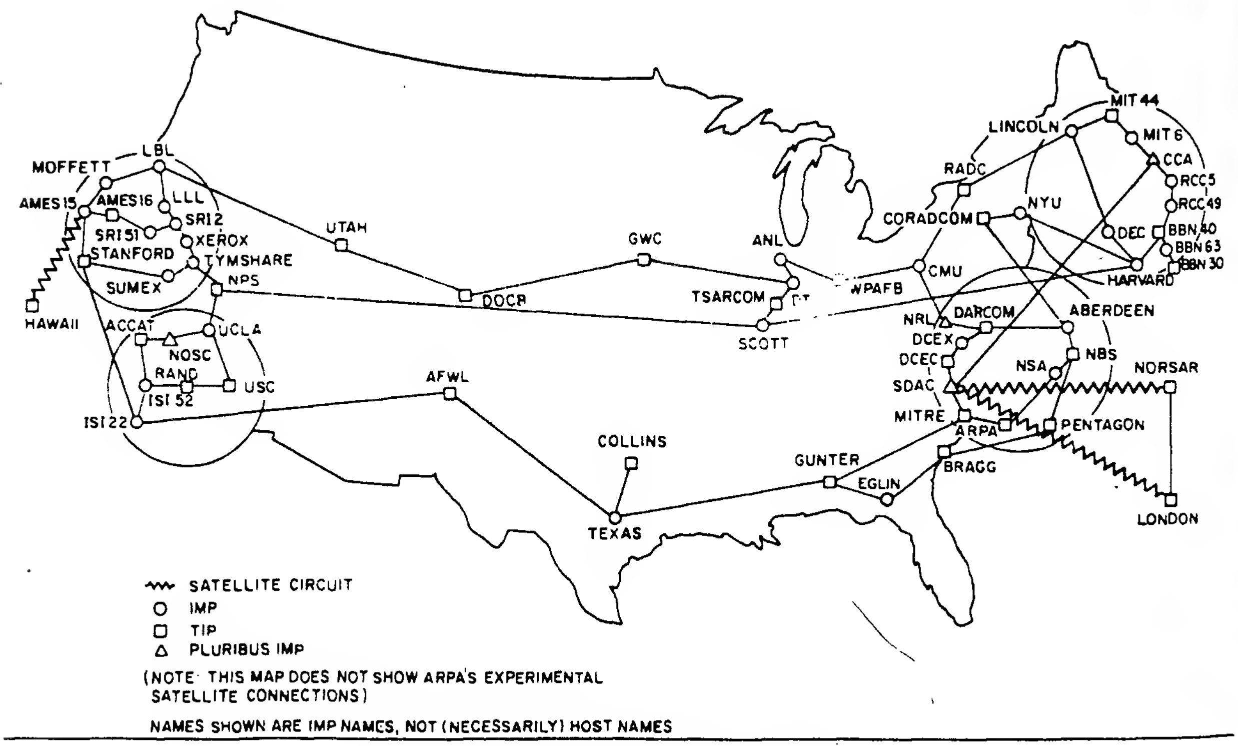

Figure 2. ARPANET geographic map, May 1979

cable.[6] Other organizations have begun to explore these concepts, both in the research, government and commercial world. Packet switched networks are in commercial operation in the United States (Telenet, Tymnet), Canada (Datapac, Infoswitch), France (Transpac, Cyclades/Cigale), Scandinavia (Nordic Data Net), United Kingdom (EPSS, IPSS, PSS), Europe (Euronet, European Informatics Network), and are under development in many other countries (e.g., Germany, Italy, Belgium, Australia, Japan, New Zealand, Austria). The number of private or experimental packet switching networks either in operation or under development is large and growing.

In the mid-1970s, DARPA initiated a new effort to learn how these different types of packet switching networks could be usefully interconnected to form a coherent communication system.[7,8] In military terms, each of the network technologies involved were needed to support the range of military data communications requirements that could be anticipated in the 1980s and 1990s. The packet satellite technology would be needed for long haul, domestic inter-continental and shore to ocean-based data communications, while wire-based systems could be used for intra-continental, inter-city links. Packet radio would be needed for mobile tactical or strategic communications and local co-axial cable nets for intra-building and intra-vehicle data communications.

All of these network types are needed, and they each have different functional characteristics. Satellites and cables are "broadcast" because all participants can hear all transmissions; but the satellite system has higher delay. The packet radio networks, operating with line-of-sight radios, are semi-broadcast, store-and-forward in nature. Wire nets such as ARPANET are non-broadcast.

One major challenge, therefore, has been to develop a rational means of making these networks interoperate. This has been accomplished through the use of "gateways" between networks (Figure 3) and the development of standard protocols which eliminate the need to standardize on network hardware and software while permitting interoperation among computers attached to the different types of networks.

The principal advantage of the protocol standardization strategy is that it permits evolutionary development of new networking techniques (e.g., optical fiber systems) and new protocols, without requiring global changes throughout the system. Such a strategy is essential, considering the unavoidable delays in deployment of new communications technology to ships at sea, widely dispersed land forces, and so on.

The internetting strategy is based on a common Internet Protocol (Figures 4 and 5) which carries internet packets from the source computer, through intermediate networks and gateways to the destination computer. At the source computer, an internet packet is attached. At each gateway, the internet packet is "decapsulated" and its internet destination address examined. The gateway decides where to forward the internet packet and then encapsulates it in a packet of the next network. If necessary, the intermediate gateway can fragment an internet packet into smaller pieces if the internet packet would not fit in the maximum packet of the next network. Such fragments can be reassembled either at the next gateway or by the destination computer.

Gateways exchange information among themselves which permits them to recover gracefully from failures at other gateways or intermediate networks, through alternate routing or retransmission, if necessary.

The computer communications protocol hierarchy now under development within DOD is illustrated in Figure 5. At the lowest level are the so-called "host/network" protocols which are generally unique to each type of network. Next is the Internet Protocol (IP) which is now being standardized by the Office of the Under Secretary of Defense (Command, Control, Communications and Intelligence) through the Defense Communications Agency. This protocol supports basic packet transport services from a source to a destination computer through multiple networks and gateways.

The next level includes a variety of protocols. The Transmission Control Protocol (TCP) supports multiple, reliable, sequenced, end-to-end, point-to-point full duplex virtual circuits between pairs of "processes" or active programs running in any pair of computers in the internet system. This protocol is also now being standardized. The IP and TCP pair form the basis for host computer communication on the AUTODIN II packet network.

DARPA is experimenting with another protocol at this level called a user Datagram Protocol which supports real-time and transaction type communication between pairs of communicating processes. This protocol will be used to support query-response applications, distributed data base management, target tracking and other services which do not require the machinery of the TCP virtual circuits.

The stream (ST) protocol is another experimental protocol for the support of packetized voice communication. This concept will be discussed in more detail in the next section on integrated voice/data networks.

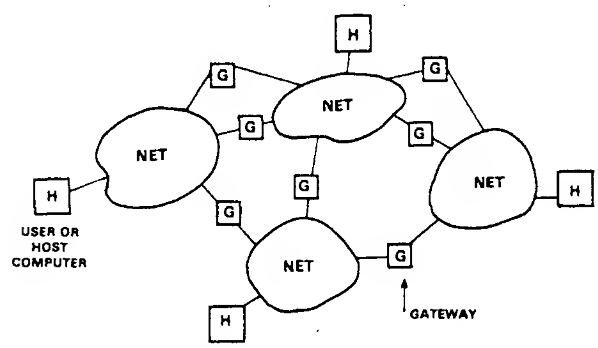

Figure 3. Network "gateways"

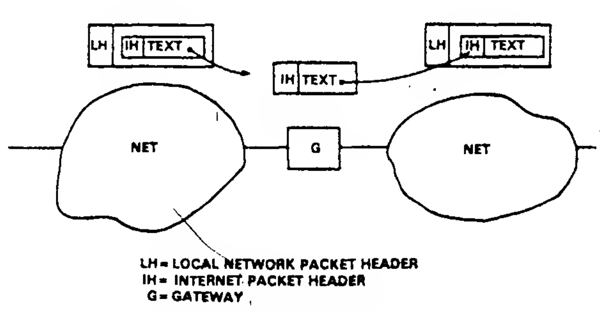

Figure 4. Encapsulation of internet packets

"There is clearly a good deal of research still to be conducted before we understand all the ramifications and opportunities inherent in multimedia, computer-based communication. This is one area of communications research which will receive increasing attention from DARPA in the early 80s."

Above the ST protocol is the Conference Protocol (CP) and the Network Voice Protocol (NVP). Above the TCP level, we find more utility protocols such as TELNET, which permits terminals of any type on one time sharing computer to appear to be directly connected to a distant time sharing computer. This allows users to interact easily with distant computing resources. Terminal characteristics (line width, page length, character set, etc.) are communicated in a standard way using the Network Virtual Terminal (NVT) protocol. The receiving TELENET protocol maps the NVT standard, if necessary, into something specifically understandable to the service host.

Also above TCP are the File Transfer Protocol and Internet Electronic Message Transport Protocol. These use TCP connections to move files and electronic messages among internet computers.

The layered protocol concept, which developed while the ARPANET was first under construction, continues to serve as a good model for preserving flexibility in the development of protocol standards. New protocols can be developed without unnecessarily affecting others already standardized or under development.

Integrated Voice/Data Networks

One of the most challenging ideas in communications is the integration of voice and data networks into a single system capable of serving the requirements of both. The most debated aspect of this notion centers on the choice of switching technique to serve the combined voice/data requirement. Voice has been served through circuit switched techniques since the beginning of telephone communication over a hundred years ago. Data has been transported largely over leased, point-to-point circuits, although improvements in modems have made it possible to use the public switch telephone networks to support computer terminal access to remote service computers.

Only in the last five years has it been considered possible that voice might be carried via packet switching techniques.[9]

Before voice can be carried on digital systems at all, it must be "digitized." The public telephone system now does this routinely on its long-haul satellite and microwave trunk circuits. But there are now a variety of choices for achieving this digitization. The public telephone system uses "pulse code modulation" (PCM) to carry digital voice. The voice waveform is converted from an analog voltage level to a relative digital value at about 8,000 samples per second. Each sample is classified into one of 128 possible values (7 bits) and the resulting 56,000 bits per second are sent in digital form. In fact, it takes 64,000 bits per second to send PCM digitized voice, since one out of eight bits is used for control and signaling.

Digitized voice is of special interest to the Department of Defense, since it is much easier to encrypt digital rather than analog voice, and the quality of the encryption is vastly superior. Reducing the bandwidth requirement for digital voice has been a subject of considerable interest for DOD, as a result.

Ordinary analog speech occupies an analog bandwidth of roughly 3000 Hz. We can typically manage to transmit between 2400 and 9600 bits/second with that amount of bandwidth. But the straightforward PCM approach requires between six and 30 times as much bandwidth as is used with analog systems.

Through the use of special purpose signal processors and microprocessors, a variety of bandwidth reduction methods have been demonstrated. These methods have exotic names such as "continuous, variable slope delta-modulation (CVSD)," "Adaptive Delta Modulation (ADM)," "Linear Predictive Coding (LPC)," "Homomorphic Vocoding," "Sub-band Vocoding," "Channel Vocoding," "Adaptive Predictive Coding," and so on. All of these techniques are realized by filtering and digitizing the analog voice signal, and then analyzing short segments of it, followed by encoding before actual transmission.

Once again, the important role of computer-mediated communication can be seen emerging. The specially programmed microprocessors and other electronics can be designed to detect "silence" and not transmit, to encode at variable rates to adapt to available bandwidth, to switch from one coding scheme to another to adapt either to local conditions or the capability of the decoding electronics at the receiving end, and so on.

In practice, digital speech bandwidths averaging about 1200-2400 bits/second can be reliably achieved and by the early 1980s, it will be possible to do most if not all the digitization, encoding and decoding on one or two large-scale integrated (LSI) circuit "chips."

The prospect of low bandwidth digitized speech led DARPA to begin experimenting in the mid-70s with packetized digital speech transported over packet switched networks. Beginning with ARPANET in 1975, packetized speech has been transmitted/received successfully on all four of the basic packet switching technologies mentioned earlier: ARPANET, packet satellite, packet radio and packet co-axial cable.

These successful demonstrations led to the next important question: "How should digitized voice be intergrated with data from transmission and switching?"

An economic study by Network Analysis Corporation in 1978 led to the conclusion that of the three choices below:[9]

(1) separate circuit and packet switches for voice and data respectively
(2) hybrid circuit/packet switched
(3) all packet switched

the last choice, all packet switched, would be the most economical to service a mixed voice/data requirement even if the relative fraction of data varied from 5 percent to 95 percent of the total traffic requirement!

This is not the expected result, so to determine whether such economies can be realized, DARPA is conducting a joint experiment with DCA to investigate the use of a high speed, 3 mb/s shared domestic packet satellite network to carry large quantities of mixed voice and data. An initial two earth-station network will expand to four earth stations by late 1980, connecting sites on the East and West Coasts of the United States. Mixtures of voice and data terminals will be located at each site and will be concentrated into the high speed earth stations through a high speed local network.

The greatest challenge for all of the packet switching technologies will not, however, come from merely carrying digital voice and data through the same network but rather from the integration of voice and data services.

Computer-Based, Multimedia Communication

Multimedia communication is just that. Multimedia communication combines speech, imagery, text, graphs and facsimile into a single, multifaceted communication service. As for simpler kinds of communication there are two basic cases: immediate, on-line ("live") communications and delayed communications.

For instance, the TWX or TELEX networks are forms of immediate communication. Two (or more) TWX or TELEX terminals are interconnected by a circuit switching system and the operators communicate with each other in real-time.

Electronic message systems, on the other hand, allow messages to be stored in the system if the intended recipient is not on-line when the message is sent. The AUTODIN I system is an example of this, as are the growing number of electronic message systems now planned or offered in the commercial world.

Experiments in some forms of multimedia communication, especially the live, on-line type, are not new. Bell Telephone Laboratory introduced the Picturephone (copyright, AT&T) roughly 20 years ago. Simultaneous use of telephone with other media (e.g., terminals) is common today, especially for testing and debugging complex, distributed systems.

We have only begun to explore the role computers can play in the integration of many of these modalities. One especially interesting possibility is the extension of a conventional electronic message system to include voice, facsimile and other graphics or imagery in addition to text.

To experiment with such an extension, a new work station must be developed which includes voice input and output capability, facsimile input, high resolution CRT display, keyboard, cursor control (for graphics) and access to a high resolution, multifont laser printer. The latter is needed for production of hard copy of the printable portions of messages.

The CRT should be of the "bit-map" variety. That is, each visible portion of the CRT display is generated bit-by-bit from a complete image of the display in the computer memory. This makes it easy to display arbitrary facsimile input and other imagery on the CRT.

A conventional electronic message system provides for the composition, editing, reading, storage and retrieval of text messages. These generally have some basic structure such as destination addresses ("to"), source address ("from"), places or people who should receive copies ("CC"), some indication of content ("subject"), timeliness ("date-time") and actual message ("text"). These structured objects can be manipulated by the electronic message system to aid users in the composition of replies, the search for related, earlier messages, the forwarding of relevant messages to other parties and so on.

The extension of these concepts to new media is a significant challenge, as is the need to transform some of these media into others. For example, it should be possible for the message composer to accept as input a facsimile page. This should be manipulable on the CRT: cropping, rotation, shrinking, expanding, positioning, etc. The image should be combined with text to form a mixed-media page.

Pure text can be manipulated in an infinite variety of ways. The proliferating commercial development of "word processing" equipment makes this obvious.

Editing of mixed voice, text and imagery messages with a common system is a substantially greater challenge. Where we find such mixtures today (e.g., slide show with sound track, books with tape cassettes, etc), we find that the various media are composed and edited separately and then combined, sometimes only loosely.

Transformations from one medium to another will be essential for such mixed mode message systems to be of value. Perhaps not all work stations will have identical capabilities, or a user may be situated at a station which has lost some capability. It will still be important for the message contents to be delivered, even if in an output mode which isn't perfectly matched to the input. For the sake of discussion, we can consider the various possibilities in Figure 6.

By "graphics" we mean structured images produced in hard copy or on a CRT, but based on an internal computer model of the object or scene portrayed. A simple, raster-scanned video image is not graphics, by this definition. The image of a building, based on an

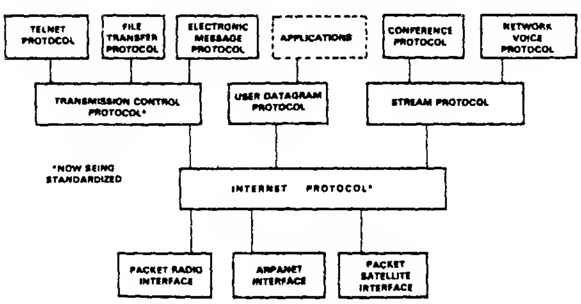

Figure 5. DOD protocol hierarchy

–TO–

–FROM–	VOICE	FAX	TEXT	GRAPHICS
VOICE	–	?	HARD	?
FAX	?	–	?	HARD
TEXT	OK	OK	–	?
GRAPHICS	?	OK	?	–

Figure 6. Multimedia transformations

internal computer model of its blueprints is graphics. So is a scene produced by a flight simulator on a CRT.

The squares in Figure 6 which are marked "?" indicate that either the transformation makes no sense, or we have not thought of any reason to try it. For instance, it just is not clear what would be intended in a transformation from voice to fax. Would it be a "voice print" or would the voice first be transformed to text and then output as a raster-scaled facsimile image? Perhaps it could mean transformation from digital representation to optical sound track?

On the other hand, we have marked some of the squares "hard." Voice-to-text is a good example. DARPA has supported research on speech recognition and understanding systems in the past. The computing tools needed to make such systems operate in real-time are becoming available. With adequate computational capability and memory capacity, a voice stream could be turned into a readable text stream. The problem of sorting our homonyms (such as "so" and "sew" or "to", "two" and "too") and recognizing colloquialisms and fragments of sentences should not be minimized. But it is at least conceivable.

The other technically difficult case, facsimile to graphics requires that the computer examine a raster-scan image and analyze its structure. This is often called scene analysis. Both voice to text and facsimile to graphics require the application of sophisticated artificial intelligence techniques.

The easier cases are text to voice, text to facsimile and graphics to facsimile. Computer-generated speech is not new, although the quality has improved dramatically since the earliest voice responses systems were developed in the 1960s. Simple concatenations of phonemes making up a word will not produce good quality or at times even recognizable speech.

Substantial, real-world knowledge is required to produce well-modulated speech from pure text. This is the other side of the speech understanding coin, but appears to be much more readily achievable.

Text and graphics can easily be transformed into the raster-scanned facsimile format, and these complete the list of "OK" transformtions in Figure 6.

Apart from the simple (!) transformation and editing problems we have raised so far, there is the major problem of interfacing users to such a system. The various input modalities (e.g., speech, keyboard, cursor) might even be put to good use in the control of the message system. Direct substitution of voice and/or cursor for keyboard editing commands is an example of an obvious (but not necessarily simple) extension.

The inclusion of voice in the system provides some interesting challenges. Voice is a temporal phenomenon. We deal with it as a serial process unfolding in time. Text and graphics are spatial phonomena which are linear or two dimensional. Extending our conventional spatial editing tools to include the time dimension will lead to capabilities not unlike those found in film or video type editing stations.

There is clearly a good deal of research still to be conducted before we understand all the ramifications and opportunities inherent in multimedia, computer-based communication. This is one area of communications research which will receive increasing attention from DARPA in the early 80s.

Other Applications

The packet radio concepts mentioned earlier provide a basis for mobile, packet switched communication. These concepts can be realized in other environments, two of which are the subjects of current DARPA research: long-range airborne communication and inter-satellite communication.

The mobile packet radio notion can be applied to long-range communication between aircraft and from ground to aircraft. Vulnerable satellite and fixed asset communication resources on the ground could be backed up using air-to-air and air-to-ground packet radio aboard strategic aircraft which would become airborne prior to a nuclear attack. The ability to rapidly deploy such a network and to interconnect it with other networks using the internetting technology mentioned earlier could make such a system an essential part of a survivable U.S. command, control and communications system.

The second area involves more survivable and affordable satellite systems. Multiple satellites could be placed in low orbit via the Space Shuttle. These non-stationary satellites could operate as a space-borne packet radio system, complete with digital section to effect control. A large number of satellites at low orbit would be harder to jam or knockdown, improving the survivability of the system; the packet switched communication among the satellites and the ground would provide substantial flexibility beyond that which is available today.

Summary

The present trend towards digital transmission and switching is leading towards a much stronger integration of voice and data services. Computers are being used, not only to manage the transmission medium resources, but also to mediate user access to multimedia communication services.

• • • — • —

References

1. L. G. Roberts, "The Evolution of Packet Switching," *IEEE Proceedings*, Special Issue on Packet Communication Networks, Vol. 66, No. 11, Nov. 1978, pp. 1307-1313.
2. L. G. Roberts and B. D. Wessler, "Computer Network Development to Achieve Resource Sharing," *Proc AFIPS SJCC*, V. 36, 1970, pp. 543-549.
3. F. Heart, R. Kahn, S. Ornstein, W. Crowther and D. Walden, "The Interface Message Processor for the ARPA Computer Network," *Proc AFIPS SJCC*, V. 36 1970, pp. 551-567.
4. I. M. Jacobs, R. Binder, E. V. Hoversten, "General Purpose Satellite Networks," *IEEE Proceedings*, Special Issue on Packet Communication Networks, Vol. 66, No. 11, Nov 1978, P. 1448-1467.
5. R. E. Kahn, S. A. Gronemeyer, J. Burchfield and R. C. Kunzelman, "Advances in Packet Radio Technology," *IEEE Proceedings*, Vol. 66, No. 11, Nov. 1978, pp. 1468-1496.
6. R. M. Metcalfe and David R. Boggs, "ETHERNET: Distributed Packet Switching for Local Computer Networks," *CACM*, Vol. 19, No. 7, July 1976, pp. 395-404.
7. V. G. Cerf and R. E. Kahn, "A Protocol for Packet Network Interconnection," *IEEE Trans. Commun. Technol.*, Vol. COM-22 May 1974, pp. 637-641.
8. V. G. Cerf and P. T. Kirstein, "Issues in Packet Network Communications," *IEEE Proceedings*, Vol. 66, No. 11, Nov. 1978, pp. 1386-1408.
9. Network Analysis Corporation, Economic Analysis of Integrated DOD Voice and Data Networks, Final Report Under Contract No. DAHC-15-73-C-0135.