Report Series
No 1

Developments In Data Networks

July 1977

The Butler Cox Foundation

Report Series No I

DEVELOPMENTS IN DATA NETWORKS

JULY 1977

A. Gunton

Butler Cox & Partners Limited
The Press Centre
11th Floor
76 Shoe Lane
London
EC4A 3JB

TABLE OF CONTENTS

I INTRODUCTION

A Readership And Purpose

This report is addressed to data processing managers, corporate planners and other senior staff in large organisations who are involved with or responsible for communication planning and information processing.

Its purpose is to examine the current status and future trends in network system design — to provide an overview from which more detailed structures and implementation plans can develop.

It is the first of a series of reports on the subject of networks to be produced by the Butler Cox Foundation. Further reports will address specific topics such as terminal compatibility and network planning.

B Background

While it has been clear for some years that data communications would play an increasingly important role in computer systems, there are now signs that the rate of change is about to accelerate. On the one hand, a number of factors combine to make more urgent the needs of users for powerful and sophisticated terminal-to-computer communications facilities. On the other, suppliers are developing new products and services which may soon make such facilities more widely available than to the few who can afford them now.

The last statement remains conditional for several reasons. Firstly, network architecture — the new orthodoxy of the computer manufacturers — promises much but has so far delivered little. One reason for this is that its apostles are encumbered with the necessity of coping with the limitations of the past. Secondly, the PTTs are flexing their muscles, in the expectation of exerting a greater influence on events in data communications; it remains to be seen how beneficial this influence will be.

Faced with these realities, user managers must recognise the potential opportunities in networks for furthering their business aims, and must evaluate the options that confront them.

C Report Scope And Structure

The report reviews the following:

- The background to the progress achieved recently in networking techniques, and the important technical issues and concepts involved.
- Some of the more significant networks now in use or under development in three main spheres: computer bureaux, common carriers and private users.
- PTT attitudes, achievements and plans.
- Computer manufacturers' products with particular emphasis on two important but very different products: IBM's SNA and DECNET.

It ends with a short summary and a list of the main implications for management.

II OVERVIEW OF NETWORKS

For many users networks promise to become a more important component of future information systems than the mainframe, and may even displace the latter by becoming the true operational centre.

In this section we summarise the developments which are leading in this direction: the growing needs of users for reliable and adaptable information transmission; the response of the manufacturers; the changing role of the PTTs; and the problems of compatibility and standards which have already arisen and will continue to be a major focus of attention.

A Role Of The Network

The data transmission systems which most users have installed are concerned predominantly, if not solely, with the transport of data. Additional functions, if there are any, are subordinate to that primary purpose. Transport of data will of course continue to be a primary function of the more highly developed networks discussed in this report, but the balance between this and additional services provided by the network is shifting.

The reasons for this are twofold:

1. Exploitation of marginal intelligence
More sophisticated techniques are being introduced to achieve high utilisation of transmission lines and good service levels for network users. This demands greater intelligence on the part of the devices driving the network. A packet switching protocol for example requires more logic at both ends of the link than the polling/select or contention protocols in use today. That same intelligence can then be used at limited marginal cost to improve other aspects of network performance or the service provided to users.

2. The move to heterogeneity
To date networks have typically been homogeneous in nature using techniques specifically tailored for the needs of the application and, often, the operating software of the mainframe concerned. With developments in the range and scope of data transmission-based applications, networks are increasingly becoming heterogeneous. This heterogeneity may take a number of forms:

- of traffic characteristics, for example interactive VDUs sharing transmission lines with remote batch terminals generating bulk traffic

- of terminal type, as one generation supplants another and as new applications are introduced

- of computer supplier, whether resulting from a policy of procurement from more than one vendor, or as formerly stand-alone minicomputers are linked into corporate systems

- of information medium, when voice, image or text traffic is required to share channel capacity or to interact with data systems.

Not all users will have the same needs in these respects, but few large users are likely to be so independent of all of them as to be able to rule out a data network as a possible solution.

Networks are concerned with the transport of information. The new functions that they are tending to undertake can be summarised as management of information, in addition to the pure distribution of information undertaken hitherto. Their role in the corporate environment is illustrated by Exhibit 1.

B The Alternatives

Given a dispersed organisation or a group of users with problems of communicating data internally, a corporate data network is only one of the alternatives open. It is currently at the expensive and risky end of a range of alternatives which stretch to a fleet of vans and a procedures manual at the other end, taking in the relatively passive communications systems which many users now operate somewhere in between. A detailed comparison of these alternatives would be premature at the current stage of development of data networks, and this report does not attempt such a comparison. Its purpose is rather to explain the goals which advanced networks are setting out to attain, and to assess the likelihood of success.

C Trends In On-Line Computing

The characteristics of on-line computer systems are changing. A number of factors are combining to produce this change, the most important being:

- increased intelligence in terminal devices, which enables them to perform more tasks locally and hence calls for higher transmission speeds and generates a less balanced network load
- a demand for higher reliability as dependence on systems grows
- a desire to exploit transmission facilities more effectively, both in terms of the loading of lines and the pattern of traffic, for example by using the overnight capacity of a private network
- an awareness of the limitations of batch systems, and others in which human interaction with the system is conditioned by the equipment capabilities rather than the other way round, in other words a demand for increased workplace-orientation.

Taken together, these forces are creating a demand for more powerful and more sophisticated terminal to computer communications. Existing methods, while not rendered totally obsolete by these changes, are at least revealing their limitations. The kinds of facilities users will be looking for as standard in their on-line systems in the immediate and near future are now only available in purpose-built networks and at unacceptable cost for all but a few.

If such facilities are to be made generally available and at reasonable cost, then network operating systems developed from the bespoke networks operating now will be the medium. Their job will be analogous to that of computer operating systems now, except that the equipment that they control and whose usage they optimise will not be locally attached peripherals, but distributed computers and terminals, and transmission facilities provided by the PTT. It is not unreasonable to expect process-to-process communication to be made as easy via a network as it can be now via an advanced operating system.

Exhibit 1 The role of the network

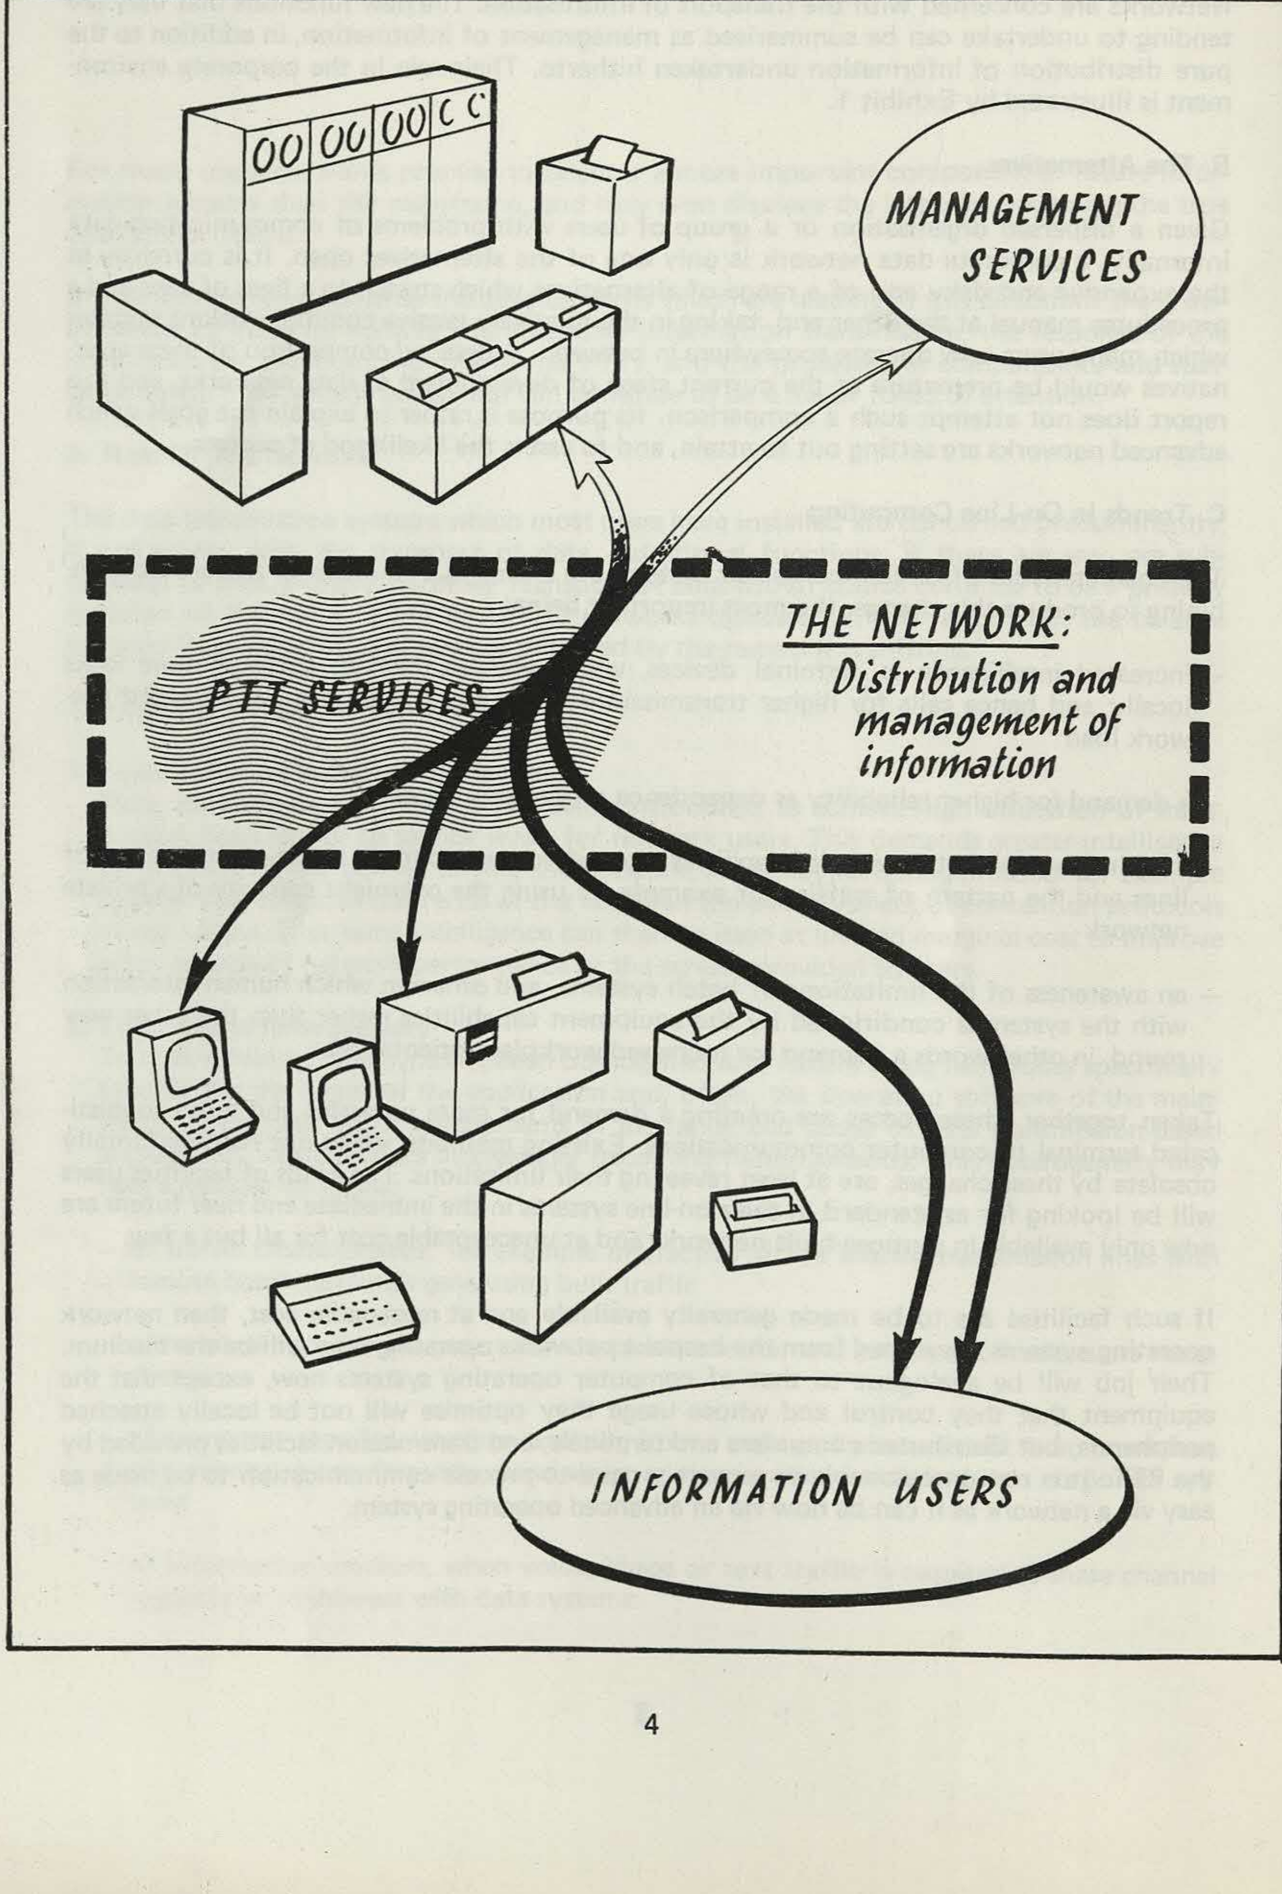

D Network Architecture — The New Orthodoxy

One of the early responses to these changing needs has been the discovery and promulgation of new concepts generally referred to as network architecture. IBM was the first computer manufacturer to preach this new gospel (although as a concept it was no newer than virtual memory which had been espoused a few years earlier), followed by most of the other major manufacturers. The view as expressed by IBM is that the network is a unified structure developed to support user requests for services, rather than merely a mechanism for transferring information from one point to another. By analogy, it represents a transport service rather than just a railway.

In addition, the network can serve to keep the user independent of changes in terminal and mainframe technology, providing a fully transparent access path. It will often incorporate reliability features, such as alternate routing, and may even compensate automatically for incompatibilities between attached devices.

In summary, the network has acquired a range of functions over and above transmission, and has become to some extent a general-purpose tool. For it to perform these functions adequately a formal structure is essential, hence the use of the term 'architecture'.

E The Changing Role Of The PTTs

Previously, PTTs have largely provided only passive components in data networks, concentrating their attentions on the substantially greater voice traffic. Continuing the analogy drawn earlier, they supplied railway track only, together with modems which were a means of getting on and off the track. They too, like the computer manufacturers but for different reasons, have noted that they might usefully provide a complete transport service, in the same way as they already do for voice traffic.

This realisation and the way it is followed up by the PTTs is of profound significance for the evolution cf data networks. Depending on how the complex problems involved are approached and how governments direct their telecommunications agencies to use their monopoly powers, on-line access to computers may either become as commonplace as television, or alternatively become to an even greater extent than now the preserve of those with the strongest financial motivation and muscle — that is large corporations and public bodies.

Networking systems will interact with PTT facilities in two important ways:

1. As data transmission facilities provided by the PTT expand in scope, they will become an increasingly important component of the network. The overall effectiveness of the network will depend heavily on how closely these PTT facilities meet the needs of diverse users, on how the tariffs are structured, and on how their use is regulated.

2. The network can buffer the user from changes in transmission technology, enabling him first of all to preserve his investment in existing on-line applications, and later to make the best use of new PTT facilities as these become available.

F The Standards Battle

Computer users are by now used to such standards as can be agreed for data transmission being rendered ineffective by each computer manufacturer's insistence on introducing his own and unique refinements. And of course, IBM has used its dominance of the world computer market to set de facto standards, which may or may not correspond with what is agreed by standards committees.

As PTTs regard more and more of what goes on within data networks as of legitimate interest, the picture begins to change dramatically. What was previously a conflict of interest between the user community in general, eager to link devices from any source into its networks, and the computer manufacturers, no less eager to preserve their hard-won customers from the predatory attentions of independent terminal suppliers, is now potentially a three-way struggle involving the PTTs as well. The alternatives facing the PTTs are either to carry data according to the rules laid down by the computer suppliers, or to invent the rules themselves.

The CCITT X.25 recommendation on public data networks is an attempt to formulate such rules, and its success in gaining wide acceptance is a measure of the PTTs' greatly extended influence on events in data communications. This is a field where confusion has ruled despite attempts at standardisation. The problem has been both limited acceptance of any one standard, and limitations in the scope of the standards themselves. The latter is of less consequence when the transmission system merely carries bits on behalf of the user – it is then only necessary to know wiring and signalling conventions. As soon as it attempts to regulate traffic at message level, or to supply information management services also, the difficulties increase by an order of magnitude. X.25, which incorporates standards at physical, link and message level, extends control as far as CCITT can hope to, given the need to serve a diversity of users.

However, even if X.25 stands the test of time, it cannot be a complete answer. Without satisfactory standards at user level also, i.e. governing the dialogue between terminal and application program, confusion can still arise as soon as users from different environments wish to communicate. Those environments may differ according to the supplier of the device concerned, according to the design of the application or according to the design of the network procedures. Thus the possibilities for confusion are limitless, despite the existence of standards governing communication at a lower level.

While the PTTs are constrained from carrying standards further by the diversity of their users, the computer manufacturers are similarly constrained by their competitive position. While they actively promote user level standards designed for their own equipment, they have no interest in machine-independent standards. With conspicuous exceptions such as the airlines and the banks, who have merely gone off and made their own arrangements, the user has had limited influence on the definition of standards. According to our analysis, only the user community is in a position to complete the process begun by CCITT in establishing the common standards which data communication requires to be fully effective as a medium for information systems.

III CONCEPTS AND TERMINOLOGY

Networking has brought with it the usual bewildering array of technical terms that seem to accompany every similar development. Some are pure jargon. Some, like those summarised in this section, conceal significant developments in the computing art.

A Functional Layering

The concept of functional layering is fundamental to computer manufacturers' networking products and to network structure in general.

The network procedures are built up layer by layer, each layer having a defined interface with the next and, provided the overall design is sound, each layer functionally independent of the others. The layers generally correspond to the physical transformation which information undergoes in its progress across the network; it passes down through the layers from the source, is transmitted as a bit stream, then is reconstituted for delivery to the recipient — see Exhibit 2.

As can be seen from this example, the total structure is symmetrical, each layer having its counterpart in another part of the network. In theory this arrangement allows changes in any one layer to be absorbed and not transmitted to others.

How the functional layers are mapped onto the hardware will change with time and, if reality matches theory, this change can take place without involving the user in costly re-programming. Examples are shown in Exhibit 3.

B Protocols

A network protocol is a set of conventions used by a particular functional layer to communicate with and co-operate with its counterpart elsewhere in the network. There is a protocol for each functional layer in the network structure, as shown in Exhibit 4, each providing a service for the levels above it or, at the highest level, for the user.

The three lowest levels of protocol are those dealt with in the CCITT X.25 standard. Thus far, widely accepted and practical standards seem likely to emerge. Without some definition of the highest level — that concerning communication between an application program and a terminal — the user is likely to be faced with the same problems that he has today when introducing non-standard equipment into his network, or attempting to inter-work with differing systems.

Whether protocol standards at this level, generally termed high-level protocols and illustrated by Exhibit 5, are either practical or desirable, is a key issue and is a recurring theme throughout the literature on this subject. Our view is that further definition and agreement of high level protocols is a pre-requisite if networks are to realise their potential for the business user.

C Packet Switching

Packet switching was conceived by interactive data processing out of message switching, since it combines the fast response times of the former with the store-and-forward and switching capability of the latter. It is in fact a new form of delivery system using more

Exhibit 2 Functional layering

Computer software	Communications architecture	Typical tasks performed
End user	End user EU	Clerical procedures Terminal operation
Applications environment	Applications environment AE	Presentation of information File access Message processing
Operating system	Access sub-system AS	Allocation of services Access control Format control
Executive program	Transport sub-system TS	Link control Routing Message handling
Processing and I/O hardware	Transmission facility TF	Multiplexing Switching Modulation/demodulation

↕ denotes interfaces between layers

Exhibit 3 Mapping of functional layers onto hardware

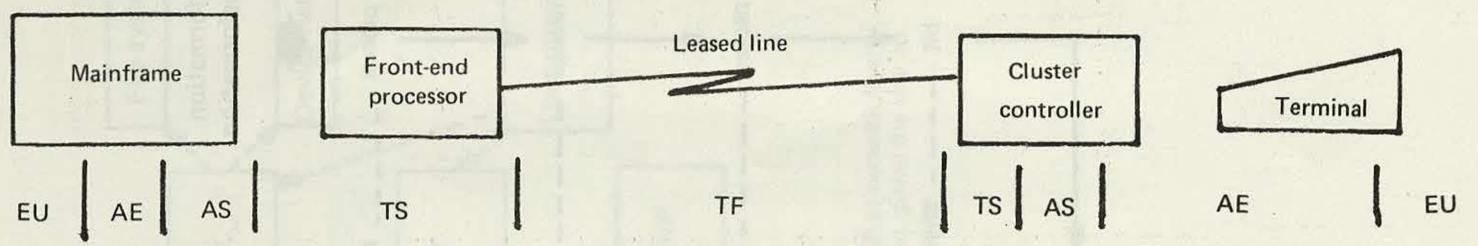

Example 1: an on-line data processing system

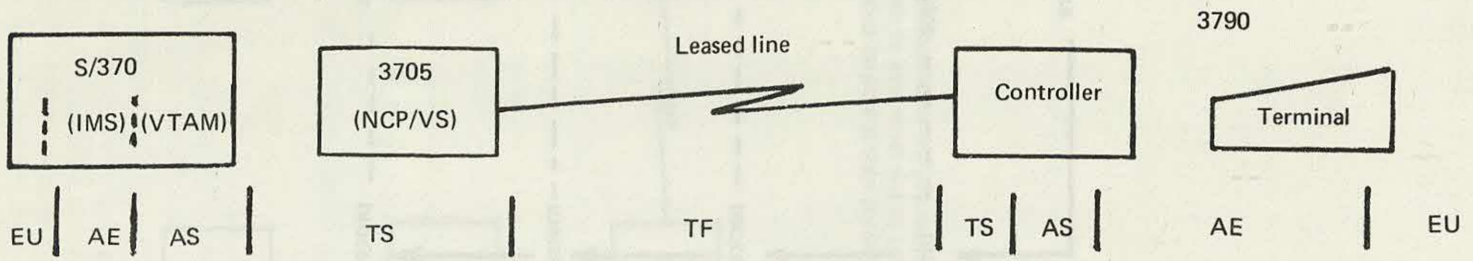

Example 2: an SNA implementation

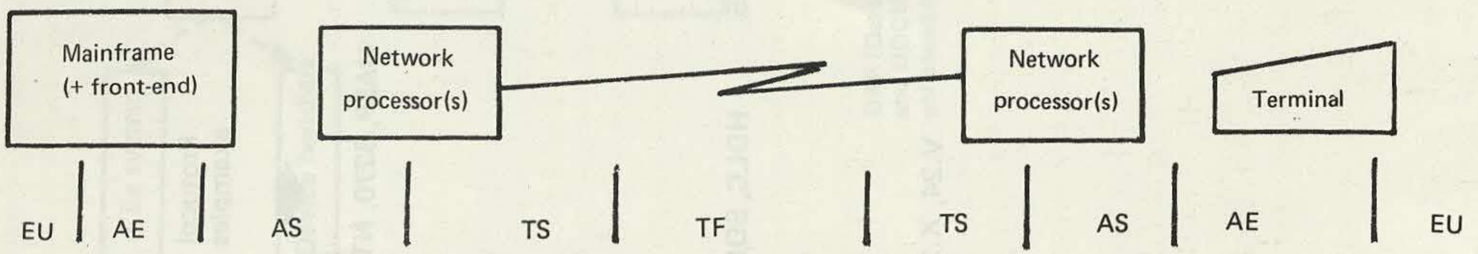

Example 3: an intelligent network system

Exhibit 4 Illustration and examples of protocols

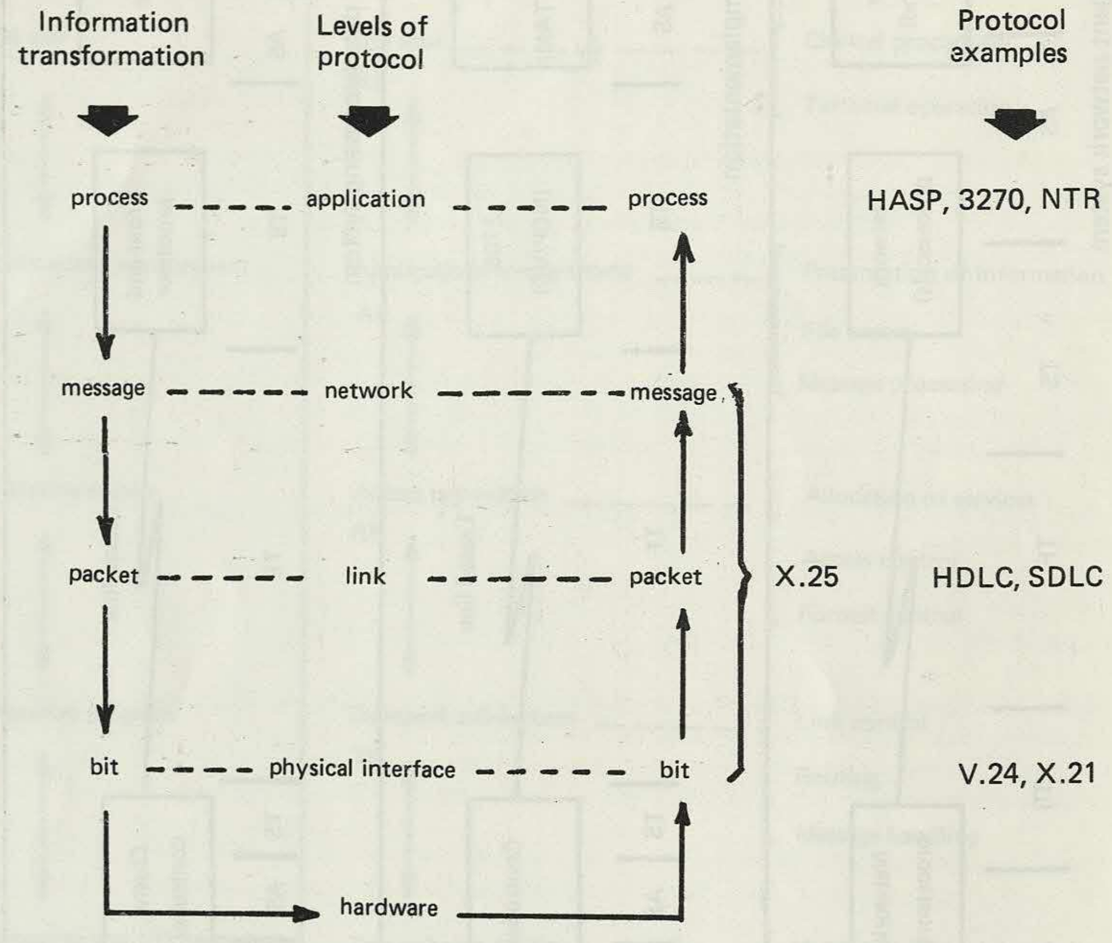

Exhibit 5 Example of a high level protocol – DEC's Data Access Protocol

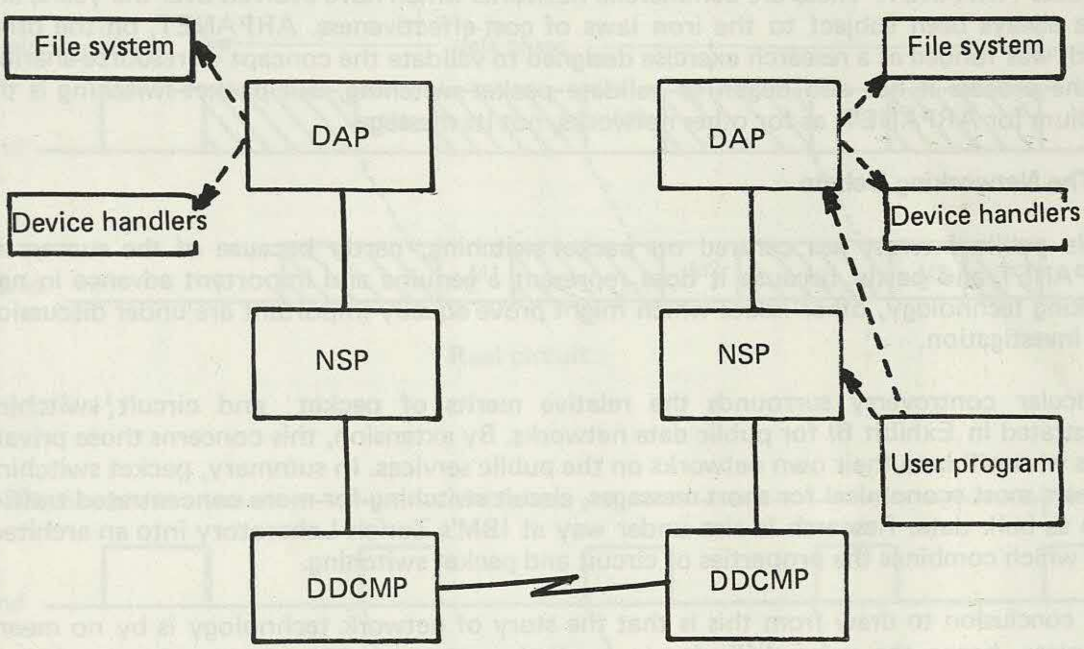

DAP (Data Access Protocols) uses the other protocols (NSP at network level and DDCMP at link level) in the same way as user programs, giving the user an enhanced level of capability for the resources which it handles sophisticated means for multiplexing the information carried over the telecommunication lines. Since the success of ARPANET it has tended to be associated with resource-sharing networks, in which all parties using the network are able to share the facilities of one another's equipment. The ARPANET service is based on packet switching, but this is only one application of the technique.

The next section reviews the development of networking. This review shows that many of the features commonly associated with packet switching can be found in networks which pre-date ARPANET. These are commercial networks which have evolved over the years, and have always been subject to the iron laws of cost-effectiveness. ARPANET, on the other hand, was funded as a research exercise designed to validate the concept of resource-sharing. In the process it has also begun to validate packet switching, but packet switching is the medium for ARPANET, as for other networks, not its message.

D The Networking Debate

While public interest has centred on packet switching, partly because of the success of ARPANET and partly because it does represent a genuine and important advance in networking technology, other issues which might prove equally important are under discussion and investigation.

Particular controversy surrounds the relative merits of packet and circuit switching (illustrated in Exhibit 6) for public data networks. By extension, this concerns those private users who will base their own networks on the public services. In summary, packet switching appears most economical for short messages, circuit switching for more concentrated traffic, such as bulk data. Research is also under way at IBM's Zurich Laboratory into an architecture which combines the properties of circuit and packet switching.

The conclusion to draw from this is that the story of network technology is by no means complete, hence the great difficulty in predicting the cost-effectiveness of new services. There can be little doubt that public data networks will eventually supplant existing services. More doubt surrounds the form of the services, and their presentation to the users. Packet switching is currently a strong contender as the long-term option, with circuit switching a possible interim measure, simplifying users' intial conversion problems.

As far as public packet switched services are concerned, the alternatives being pursued are virtual circuit and datagram. In the former the user establishes a connection, then routes messages along that connection until it is closed, while in the latter addressed messages are posted into the network for delivery without a prior call set-up phase (see Exhibit 7). For the user the relative merits of datagram and virtual circuit revolve around two factors:

1. Datagram leaves the user with the task of performing end-to-end sequence control, while virtual circuit schemes will normally relieve the user from this task; freedom or imposition, depending on your viewpoint.

2. Datagram will be cheaper where short exchanges of information tend to predominate, rather than lengthy dialogues; for the former, the overhead of establishing and closing down connections for virtual circuit working can be heavy.

E Compatibility And Virtual Terminals

Rapid technological advance has meant that substantially increased terminal intelligence has become cost-effective for many applications. Firstly, this has created a gap between the capabilities and mode of operation of terminals widely in use today and new ones available on the market. Secondly, it has made the concept of plug compatibility almost meaningless, since the 'plug' consists not only of hard-wired electrical circuits, but also soft-programmed

Exhibit 6 **Illustration of circuit switching *(real* circuit) and packet switching *(virtual* circuit)**

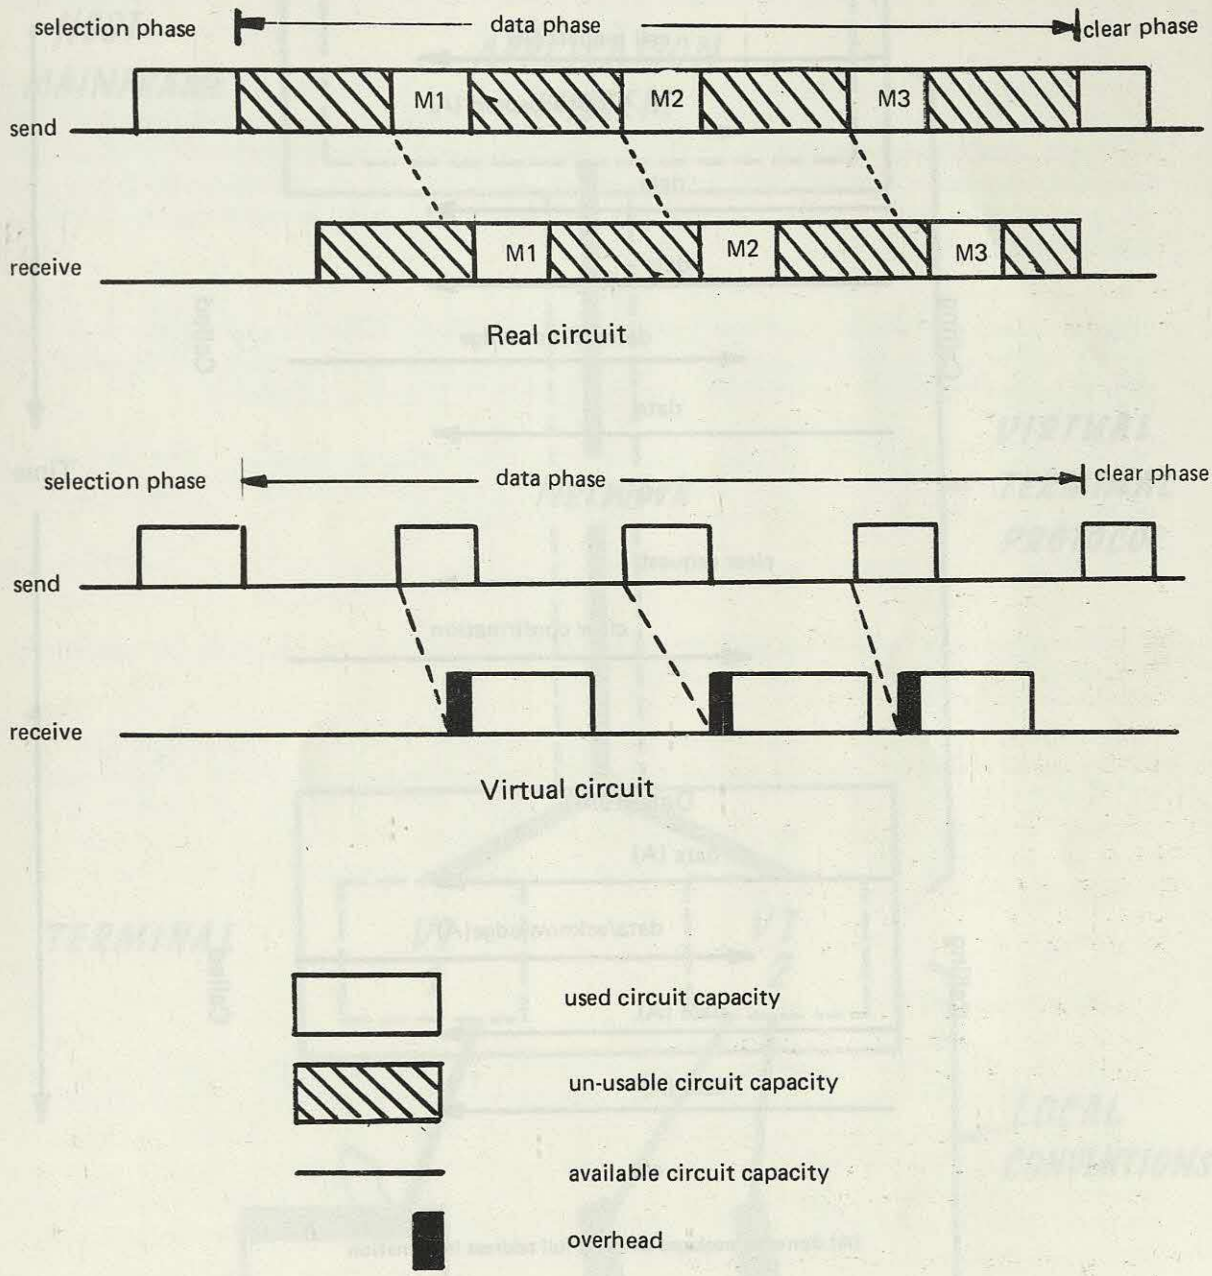

Exhibit 7 Illustration of virtual circuit and datagram

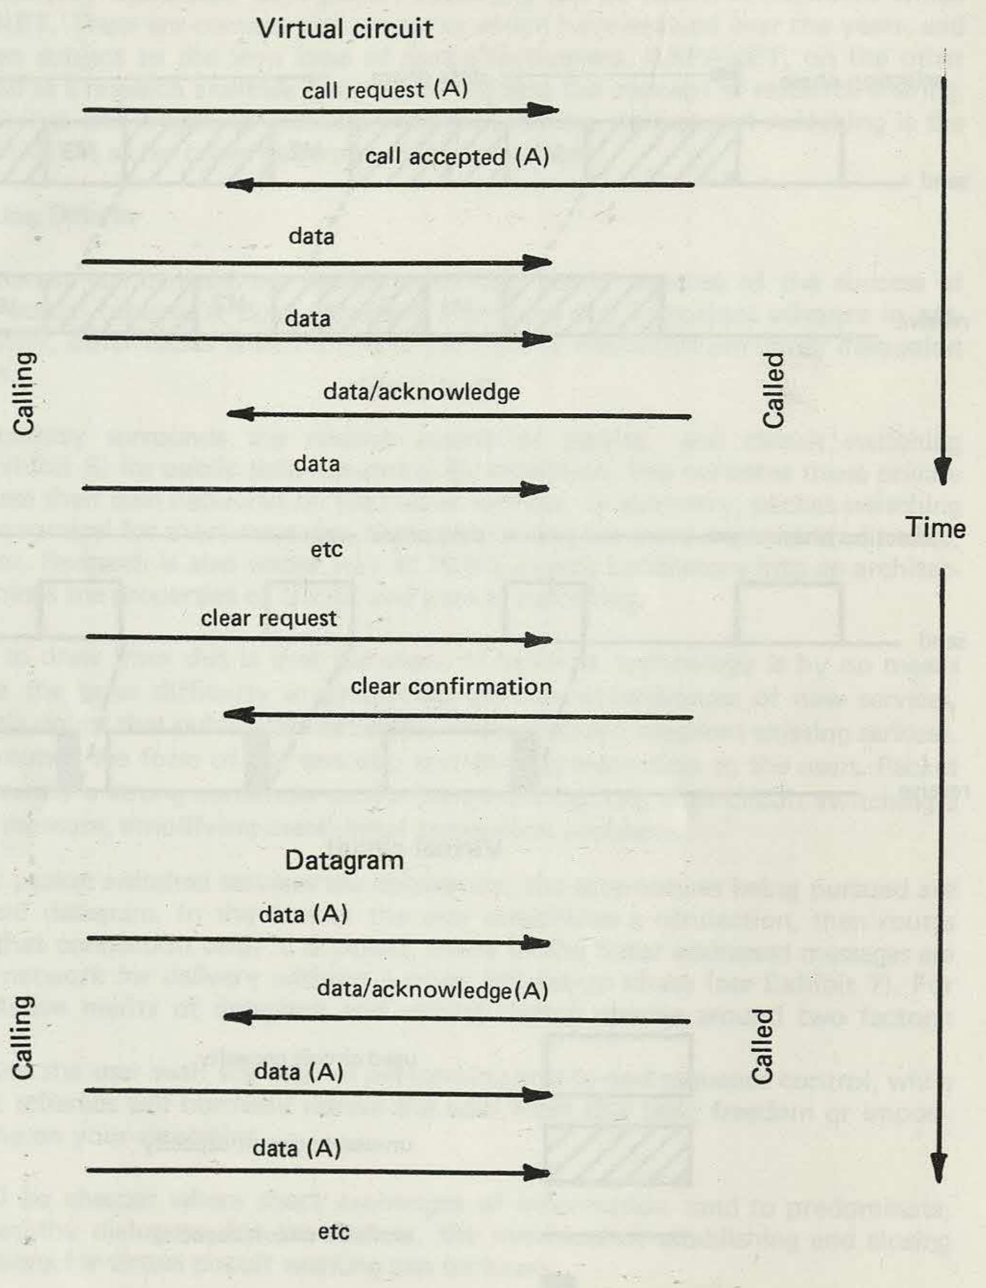

(A) denotes messages carrying full address information

Exhibit 8 The Virtual Terminal concept

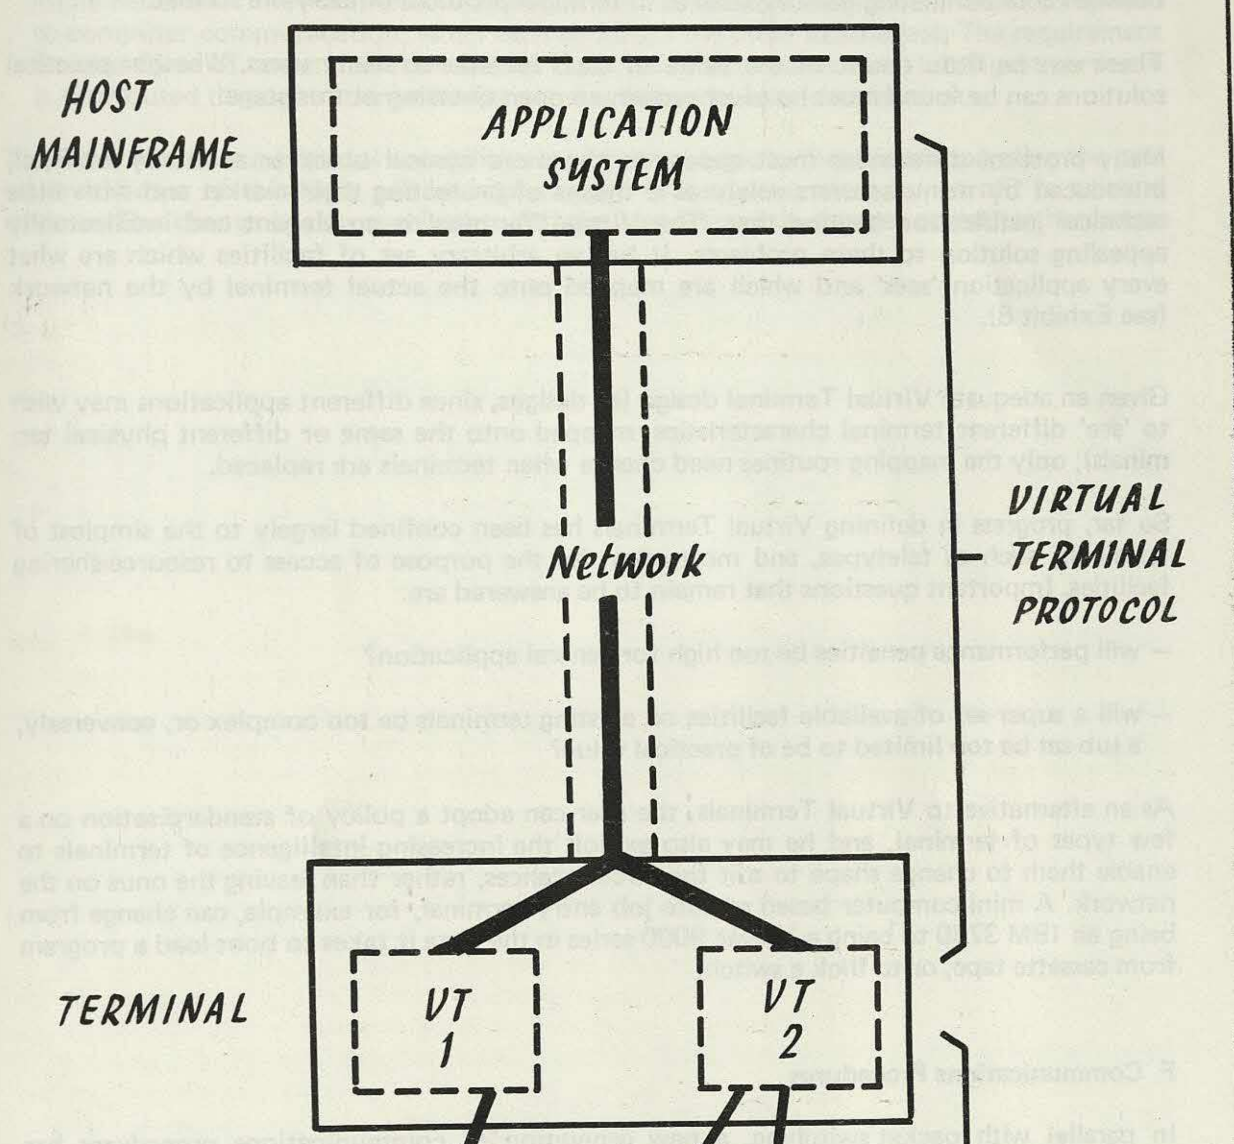

routines. It is difficult to maintain compatibility with something that can change its shape in as little time as is necessary to load a new program.

Interest in the problems of interfacing terminals to resource-sharing networks such as ARPANET has expanded into studies of the compatibility problem in general. Speed conversion and code independence are provided as a matter of course in many data networks, but it might also be the appropriate place to resolve other and more complex differences between communicating devices, such as in terminal protocol or even file format.

There can be little doubt of the value of such features to many users. Whether practical solutions can be found must however remain an open question at this stage.

Many protocol differences must appear to the more cynical observer as totally artificial, introduced by manufacturers solely as a means of protecting their market and with little technical justification beyond that. The Virtual Terminal is an elegant and intellectually appealing solution to these problems. It has an arbitrary set of facilities which are what every application 'sees' and which are mapped onto the actual terminal by the network (see Exhibit 8).

Given an adequate Virtual Terminal design (or designs, since different applications may wish to 'see' different terminal characteristics, mapped onto the same or different physical terminals), only the mapping routines need change when terminals are replaced.

So far, progress in defining Virtual Terminals has been confined largely to the simplest of terminals, such as teletypes, and most often for the purpose of access to resource-sharing facilities. Important questions that remain to be answered are:

– will performance penalties be too high for general application?

– will a super-set of available facilities on existing terminals be too complex or, conversely, a sub-set be too limited to be of practical value?

As an alternative to Virtual Terminals, the user can adopt a policy of standardisation on a few types of terminal, and he may also exploit the increasing intelligence of terminals to enable them to change shape to suit the circumstances, rather than leaving the onus on the network. A mini-computer based remote job entry terminal, for example, can change from being an IBM 3780 to being a Univac 9000 series in the time it takes to boot-load a program from cassette tape, or to flick a switch.

F Communications Procedures

In parallel with packet switching, a new generation of communications procedures has emerged, referred to as Data Link Control procedures. High Level Data Link Control (HDLC) is the ISO standard which has also been adopted within CCITT's X.25 recommendation. Synchronous Data Link Control (SDLC) is IBM's offering, used by new terminals such as the 3790 Communication System and the 3270 VDU. It is also incorporated within Systems Network Architecture (SNA), which is discussed in Section VI B. Other manufacturers have introduced their own procedures, based on or similar to these.

The key features distinguishing these from older procedures such as ISO Basic Mode and IBM's Binary Synchronous Communications (BSC) are as follows:

1. They are fully independent of the data transported. Hence they will support any character code and can also carry pure binary information, such as digitised facsimile or voice.

2. They are capable of full-duplex operation and therefore are able to take full advantage of the higher transmission speeds which are becoming available. Half-duplex procedures such as BSC on the other hand tend to become less efficient as line speeds increase.

3. The older procedures are designed for a situation where a centralised computer is controlling a number of remote terminals; the computer is the master, directing the operation of the slaves. The new procedures still cater for this very common situation, but in some implementations they can also support the more balanced dialogue involved in computer-to-computer communication, when each addresses the other as an equal. The requirement for this form of communication can be expected to increase as more intelligent hardware is distributed throughout organisations and as networks inter-communicate.

The price for these additional capabilities is additional logic in the devices operating the procedures. Some of this logic will be implemented in hardware, often by retrofit to existing devices. The change is also consistent with the trend towards cheaper computer hardware.

IV PRECURSORS AND PIONEERS

Although packet switching has not been with us long (for example ARPANET was conceived in 1963) the precursors of modern networking systems can be found much earlier and outside the realm of packet switching alone.

In this section we trace the history of developments leading to modern networking systems by looking at:

- The airlines' SITA network which has developed over decades and now has most of the features commonly associated with packet switching.
- Computer bureau networks which have been progressively refined over the years.
- Private networks, exemplified by British Steel's corporate network.

Finally, we find that despite the diversity in application and environment of these networks, there are several recurring features. The most striking change is that the network is now seen as a system in its own right, rather than as a component subordinate to the all-powerful mainframe. This frees it to serve the needs of users of its communications services in general, rather than those of a specific piece of equipment.

A Multi-User Message Switching Systems

Message switching has of course existed for decades in telegraphy and, like packet switching, is a form of store-and-forward switching. It differs in original purpose from packet switching in that it was intended primarily for non-real-time people-to-people traffic, whereas packet switching grew from the need for interactive computer-to-computer traffic.

However, if one looks at a message switching system such as the SITA High Level Network, the distinction becomes less clear. SITA started out as a replacement for the torn tape systems used in the early days of remote communications. It has since evolved to carry interactive traffic from visual displays as well, and the facilities it now provides do not differ greatly from those of modern packet switching systems.

The significant difference between SITA and, say, ARPANET, is not therefore speed of transmission, message delivery time or reliability. Whereas SITA serves a specific application, namely airline reservations, ARPANET and its fellows serve a general function, namely access to dispersed computer systems.

Because SITA is specifically directed, cost-effectiveness is more easily achieved than with packet switching. As always, the bespoke solution is more efficient than the general. On the other hand, the techniques pioneered in ARPANET can be applied more widely because of the generality which is implicit in their design. Indeed, the value of packet switching might well be the new possibilities it opens to exploitation, rather than the technique itself.

B Bureau Networks

Time-sharing networks set up by some of the US bureaux have been evolving progressively over the last few years, and form an interesting contrast with the packet switching networks.

CDC's INFONET, for example, by the middle of 1976 served 140 cities in North America and Europe and had 120,000 channel miles of dedicated circuits. It has a virtual line switched architecture (i.e. the route used is invariant for the duration of a given call) combined with data link redundancy, which prevents the loss of a single path from interrupting a call. CDC also claim that the change from leased point-to-point lines to a virtual switched network was achieved in a manner totally transparent to the users.

Tymshare's TYMNET is of interest in other respects. Firstly, it offers another variation on the packet switching theme, using a switching and routing technology which is claimed to be more efficient than ARPANET's. Secondly, it has continued the progression from point-to-point to switched network and become a fully fledged common user data network – the Systems Development Corporation, Lockheed and the National Library of Medicine have attached their computers to TYMNET along with Tymshare's own bureau systems.

C ARPANET

Although networks such as the ones we have just described have been on the scene for some time, it was probably the decision of the Advanced Research Projects Agency of the US Department of Defence to give birth to ARPANET in 1969 which generated the great interest in packet switching and in networking.

ARPANET's significance is that it has validated the packet switching concept on a large scale (now connecting sixty centres), although it is still too early to describe it as a fully proven technology. Much remains to be learnt, for example, about routing and flow control. ARPANET has also contributed to advancements in fields related to packet switching, for example, high level protocols and terminal compatibility.

Although started and funded as a research project, it has led on to commercial developments by revealing the market for value-added networks. These networks use packet switching or a variant as a base for a communications-based service, such as facsimile transmission or access to data banks. This service is of course the value which has been added to the basic data network. Telenet, for example, a direct commercial spin-off from ARPANET, is now expanding at a rate of one node per month.

While there have been some failures it seems probable that these networks will complete the process begun by ARPANET by demonstrating the commercial viability of the technique. They also demonstrate quite clearly the nature of packet switching. As already suggested in this report, principles established over a long period in data communications systems have found a successful general expression in packet switching.

It is this which gives it its great versatility and makes it a suitable base for a range of value-added public networks.

D British Steel's Private Network

Needless to say, what is viable as a value-added network in the US is not necessarily viable elsewhere, nor necessarily for the private user rather than a public service company. Some users in Europe are now beginning to feel the water, and none is probably at a more advanced stage than the British Steel Corporation. Indeed, rather than just feeling the water, they have struck out confidently for the further bank.

Few users will have large enough operations to justify a venture into packet switching networks unaided at this early stage of the game. As it happens, British Steel were aided by generous tariffs offered by the British Post Office for supergroup circuits (240 KHz). Eighty per cent of the bandwidth is being used for voice traffic, so that the 48 KHz data circuits come virtually for nothing. The introduction of the network also ties in well with plans for specialised data centres, providing perhaps a unique combination of favourable circumstances. Nonetheless one cannot but applaud the forthright manner in which the company has made up its mind and set about a task which is not without risk.

Packet switching was selected for the network for three main reasons:

— reliability; apart from the normal alternate routing, facilities have been built in to by-pass failed equipment automatically

— the ease with which equipment can be inter-connected

— better utilisation of transmission lines.

It is also interesting to note how many new and attractive possibilities will present themselves to British Steel once the basic network is operational, enabling them to make use of terminals more easily, to enhance local facilities, and so on. The network is seen not just as a sophisticated transmission system, but as a powerful new weapon in the management services' armoury.

E The Crucial Developments

The single, obvious change which emerges from these diverse developments in the art and practice of networking is that the network can be seen as a system in its own right, rather than as a subsidiary component of a system centering on the all-powerful mainframe computer. The mainframe is now a peripheral along with the terminals and other attached devices all of which revolve around the axis of the network.

This may seem a trivial point, but the change is of profound significance. Among the implications are that:

— it is easier to introduce new applications using the same communications media, without impacting existing applications

— it is easier to use existing terminals to communicate with other devices on the network than the mainframe application which they serve principally

— since the network continues operating regardless of the condition of any one mainframe, more effective fallback and security procedures can be implemented.

There is also a discernible pattern in the objectives motivating those who have constructed the networks reviewed briefly in this section. The two prime factors are:

1. A desire for improved reliability in all its aspects; for example, alternate routing to avoid dependence on any one component, fallback facilities independent of the mainframe, and security checking.

2. Freedom of access to multiple locations, for example to share resources, to enable computer centres to specialise, or to distribute information.

Two other factors only just emerging but nonetheless significant are:

3. Sharing of transmission facilities with other information media such as voice.

4. Support for heterogeneous systems, whether a variety of terminals accessing a common set of applications, or a variety of computer systems with a need to communicate with one another in real time.

V THE PTTs

The future role of the PTTs will not be limited to the provision of passive components in data transmission systems, nor to the enforcement of rules, regulations and standards as to their use. As the scope of the services offered and of the potential market served both widen, so does the ability of the PTTs to influence the course and pace of development.

Depending on how that influence is exercised, data communications may increasingly become the preserve of the powerful few or may become as commonplace and, ultimately, as easy to use as telephone or television. New public services and new attitudes are still largely in the embryo stage, but there are some indications to go by. The purpose of this section is to identify what they are by looking at developments in three key areas:

– In the US, ARPANET has spawned many so-called 'value-added networks', offering a variety of communications-based public services.

– In Canada, a public data network which will provide, in addition to normal features, standard subscriber interfaces for specific terminals and applications such as remote batch, is at a relatively advanced stage of development.

– In Western Europe, most of the PTTs are active, although few have so far committed themselves beyond an experimental first stage. If the situation in the Nordic countries, who are among the most advanced in the field, is anything to judge by, the demand for packet switching or other advanced data transmission services will move development on as quickly as economic considerations allow.

What it will be most difficult for the carriers to decide is the precise facilities needed, and this depends on the user community assessing its own requirements and bringing these forward.

A The Options

The European PTTs and other regulatory bodies such as the FCC are in a position to exercise a profound influence on how terminal and computer networks, and thus communication in general, develop.

Depending on their policies and achievements, the scenario one can visualise a decade or so hence can vary drastically. On the one hand, there might be a proliferation of special-purpose private networks operated predominantly by larger users. Use of data transmission facilities by smaller users and by private individuals could be severely limited by the very high costs involved. On the other hand, universally available and attractively priced public facilities might open up data transmission to all with an interest, with ability to pay not a severely limiting factor.

To some extent the user is compelled to make some estimate of the likely outcome – what facilities will be available, what they will cost, – in order to establish his own policy. Fortunately, this need not be such a guessing game as it has been. International standards like the CCITT X.21 and X.25 recommendations should provide a firm base on which to build one's expectations and plans.

B North American Developments

In the United States, it is the phenomenon of the value-added networks that demands first attention. The type of service offered varies considerably: Telenet provides resource-sharing facilities like its parent ARPANET, but with appropriate enhancements for commercial security and reliability; Graphnet specialises in facsimile transmission; TYMNET is designed to provide terminal access to specific host computers. They will be of interest to a European audience primarily for what they will reveal about the economics of such services.

In Canada, Trans-Canada Telephone Service (TCTS) plans for the Datapac network are well advanced. Over and above the normal features of public networks, such as permanent or switched calls, priority grade of service and so on, TCTS is developing standard subscriber interfaces for specific terminals and applications. The first of these will be for timesharing and remote batch, later ones will provide for transaction processing. The intention now is clearly to woo existing users of leased lines onto the packet switching service. Hence the emphasis on IBM protocols such as HASP and 2741, for which a command language for local editing is also planned. It is interesting to speculate whether TCTS will eventually seek to introduce its own protocols as well.

C European Developments

Because of differences in policies, regulations and tariffs, countries will develop in different ways, and the timescales will vary also. Two contrasting examples can serve to demonstrate how differently events may turn out, depending on local circumstances.

1. Nordic Public Data Network

The Nordic countries have undertaken the joint development of a public data network. The first stage of this will comprise circuit-switching facilities only, in other words facilities equivalent, as far as the user is concerned, to private leased lines.

A detailed schedule of possible future facilities has also been published, and also a broad timescale for development, although no public commitment has been given that any of these additional facilities will be provided until demand for them has been assessed. Current signs are that the demand exists and that implementation will go ahead, possibly at an even faster rate than was envisaged originally.

2. UK's Experimental Packet Switching Service (EPSS)

In the UK, the situation is very different. The Post Office was one of the earliest PTTs to embark on a packet switching service. But, no doubt partly because the service was so clearly marked 'experimental', and partly because it came too early to meet the CCITT X.25 standard and users will therefore be compelled to convert, it has not made a major impact on the potential market. Particularly, it has had little direct support from commercial users — including the major users of data communications such as the banks and the airlines.

The reasons for this are clear and understandable; these organisations have highly-developed data communications systems and can hardly be expected to undertake the risk and the expense of converting them for something as insubstantial as an experimental service.

The danger is, of course, that it will become progressively more difficult for these users to convert at all, as their commitment to their own special-purpose private networks increases. As a consequence, public data communications services may lose their major source of financial support, and latent demand may be stifled by high tariffs. Meanwhile, special-purpose private networks will proliferate, and compound the difficulty.

The problem of the UK's EPSS outlined above may seem pessimistic. It also exists for the Nordic PTTs although on a smaller and more manageable scale. There is no obvious and easy solution.

However, it is vital that all the parties concerned should have complete and detailed information on which to base their decisions and attitudes. This means that users should be aware of the facilities that advanced networking systems can provide, so that they can evaluate the short-term costs of conversion or risk against the long-term benefits. They should also assess their own requirements in terms of those facilities, and bring these requirements to the attention of the PTTs and their other suppliers.

VI THE COMPUTER MANUFACTURERS

The computer manufacturers claim their investment in a total systems capability as justification for their attempts to tie their customers in to their own equipment. Their ability to do so is threatened by the increased activity of the PTTs, outlined in the previous section, and the much more open nature of modern networks. Their own networking products should therefore be seen as ambivalent – part innovation and part defence against inroads into their traditional preserve.

In this section we justify this view by examining two important – but very different – products:

1. IBM's Systems Network Architecture (SNA) led the field and illustrates this ambivalence very well. While it promises the user new freedoms – to distribute computing power, to access facilities regardless of network characteristics – in practice the exercise of these freedoms is rigidly constrained by the supplier. SNA, in its present form, in fact sits firmly within the tradition of mainframe-dominated data communications systems.

2. DEC's product, DECNET, makes an interesting contrast. It is designed for networks of computers rather than networks of terminals. Control is therefore distributed throughout the network rather than remaining centralised as in SNA.

Other manufacturers, both mainframe and mini, offer variations on these basic themes, but all exhibit a preference for the supplier's own equipment and methods, while conceding different degrees of freedom to use alternatives.

A Their Market Stance

Traditionally the computer manufacturers have maintained their grip on their customers by offering 'total systems', so constructed that competitors cannot easily steal the more lucrative sections of the package. This approach is justified by the claim, reasonable on the face of it, that their investment in a total systems capability entitles them to protect themselves from those who have not developed such a capability but seek to share the benefits.

There is now a real prospect of standards originated by the PTTs becoming so widely accepted that the computer manufacturers will find it difficult to ignore or, in the case of IBM, pre-empt them. This threatens to erode their ability to create impregnable total systems.

With telecommunications equipment forming a large and growing component of computer manufacturers' business, they are unlikely to stand idly by while their defences are dismantled by the PTTs, leaving the spoils to be plundered by plug-compatible manufacturers and other specialist suppliers.

The network operating systems offered by the computer manufacturers must therefore be viewed as a necessary defence as well as an extension of their systems capability.

B IBM's SNA

One can only admire the skill with which SNA contrives to acknowledge the trend towards structured, stand-alone networks, while avoiding the dangers discussed above. SNA is presented as 'a unified network structure, developed to support user requests for services'. The terminal user is intended to view the network as an integrated system; interaction with his terminal is to be uncomplicated and defined by the services he requires and not by the network characteristics. Yet SNA sits firmly within the boundaries set and maintained for it by the mainframe system.

For example, intelligence can be distributed to the terminals, but programs are down-line-loaded from the mainframe, or protected by macro generators, which only the bravest of users will tamper with. Again, the protocols and structure hold out the promise of freer exchange of information within the network, but the system at this stage remains resolutely tree-structured, centering on the all-powerful mainframe.

Thus, although it does offer additional capabilities and hints at new freedoms, SNA is solidly within existing traditions and in its present form primarily represents a rationalisation of a confused situation rather than a conceptual breakthrough.

SNA uses a bit-oriented link protocol called SDLC, similar to but not fully compatible with HDLC, which the X.25 standard incorporates. There have been hints and rumours about IBM support for X.25, but nothing substantive has yet resulted.

SNA is in fact a general theory on which the current implementation of the various hardware and software elements which it comprises is based. It incorporates functional layering (see Section III A) which should in theory give the user some freedom to substitute functional layers provided that interfaces are maintained. Exhibit 3 earlier showed the SNA structure mapped on to the equipment in the current implementation. This illustrates how severely constrained the user is from exercising this freedom, because the functional layers at the mainframe are buried deep into the heart of the system, thus making modification an extremely hazardous business.

This view of SNA is perhaps unduly harsh, because it gives little credit to IBM for the undoubted innovation within the product. Nonetheless, it is important to recognise the limitations of SNA in the context of the types of networking system described in Section IV and the benefits which these offer. It is perhaps superfluous to note also that SNA as it exists today is not of course the whole story. We will have to wait on IBM for further instalments.

C DECNET

DECNET provides an interesting contrast with SNA for several reasons:

- Its orientation is towards networks of computers, rather than networks of terminals.
- It provides for distribution of control, not just of computing power. SNA distributes computer power (to some extent) but still centralises control.
- It is heterogeneous in the sense that forms of communication and of processing can be mixed quite freely.

When looking at the diverse forms of communication system that are in use, DEC identified the following major requirements:

- a flexible communications path between programs
- access to peripheral devices and files from programs resident in other systems
- use of a wide variety of link types to satisfy a range of price/performance trade-offs.

The resultant structure has three basic components, each of which will have a component resident in some or all of the computers forming the network. These are:

– DDCMP, the link protocol and broadly equivalent to SDLC

– NSP, which enables any two processes to establish a dialogue

– DAP, a high level protocol which uses the network facilities in the same way as the user may directly; it provides common generic functions for file sharing, which are translated locally to the functions specific to each system.

DECNET provides a much more flexible structure than SNA, which the user can adapt to his particular requirements as he sees fit. By implication, of course, it is more of a framework than a total system. This is of particular interest both because of the wide range of facilities it makes available, albeit sometimes in skeletal form, and because it can be seen as a significant step towards genuinely distributed intelligence.

Although DECNET is still built on existing operating system primitives, it is interesting to speculate about the possibility of making network functions primitives for the operating system, and thus make the operating system less location-dependent internally. This way, the whole of the network would be able to interact in the same way as processes within a single system.

VII THE USERS

Why should a management services manager concern himself with the developments in networking described in this report? If he accepts networks as worthy of attention, why should he do anything now rather than wait a few years when some stability may have been achieved?

In this section we try to answer these questions by emphasising the benefits to be gained by taking positive action:

– Networks can provide independence from many of the petty incompatibilities which bedevil current data communications systems.

– Networks enable users' needs to be met positively and in a controlled manner. This is what primarily distinguishes a networked solution from one using stand-alone units in which control must be exercised by means of a rule book or other imposed restrictions.

– These benefits will not present themselves pre-digested to those who stand and wait – policies must be established to enable the opportunities to be grasped.

Networks are likely to be the crucial component of information systems of the near future. Few users are likely to remain immune from their influence for long.

A Networks: Fad or Revolution?

The mature computer user is used to the frequent 'revolutions' in technique which beset the industry. These usually promise more than they provide and often punish most those whose enthusiasm to exploit the new benefits outweighs their caution. Why should networks not fall into the same category as for example, the integrated management information system which was to revolutionise decision-making, or the database which was to revolutionise data management?

At this stage, it is not possible to say whether networks will disappoint expectations, since this depends both on the expectations which they eventually arouse, and on some of the factors discussed earlier, such as PTT policy. However, it is possible to see what their significance might be, given the right conditions:

1. Freedom of access

Essentially, advanced data networks combine the flexibility of voice communication systems with the control facilities and different characteristics, notably high transmission speeds and short call times, of data systems. Ultimately, they could combine the best features of both, so that the almost limitless freedom of access taken for granted with the telephone would be combined with the more rigorous security, privacy and integrity requirements of most computer systems.

In so doing, they can also free the user from many of the petty and time-consuming technical side-issues with which he is currently concerned – persuading a terminal from one supplier to talk to a mainframe from another, re-capturing information output by one system in an incompatible format for another and so on.

2. Independence from technological change

On another level, an intelligent network can also buffer the applications using it from the characteristics of the transmission services provided by the PTT, whilst also ensuring that the best practical use is made of these facilities as they evolve. This is just one of the advantages which the structure of the network facilitates, and which also make it easier to take advantage of relevant new technology with minimum impact on applications software and on user procedures.

3. Flexibility

By controlling the presentation of services to the end-user, the network enables his needs to be met much more flexibly. Moreover, and this is where the networking solution differs crucially from systems using distributed mini-computers, it can enable the use of those services to be better controlled, whether to meet corporate policy, to optimise efficiency, or to further long-term development plans.

One factor perhaps missing from the cost-justification equation for distributed mini-computer systems is the cost of incoherence, stemming from the problems which may arise when stand-alone computers need to co-operate, because the needs of the organisation change or as applications develop.

B Developments To Watch

There are certain areas of network design in which crucial work is now taking place. Of the two outlined below one is also an area where a contribution from the user community seems vital.

1. High level protocols

Firstly, while the PTT-led moves towards standardisation of communications protocols are of great value, they cannot provide a complete solution to the problem because they only reach so far. Halting standardisation at the highest level dealt with by X.25 for example would be like agreeing the vocabulary for a new language without defining the syntax.

Without high level protocols, such as DEC's Data Access Protocol or the File Transfer Protocol defined for ARPANET, the danger of a proliferation of supplier-dependent and/or terminal-dependent dialects still exists. That is not to say that universal command languages are feasible or even desirable, but rather that users with common interests will benefit by formulating high level protocols to meet their own needs.

The success of common interest networks like SITA, for example, is based not only on commonality of interest and broad similarity of operation, but also on detailed technical and procedural standards. Given the basic foundation, which the CCITT promises to provide in the form of X.25, it seems inherently more likely that workable application interface standards (which is really what high level protocols are) can be formulated by groups of users with common interests, rather than by suppliers whose customer base is wide and diverse.

The group may be the divisions of a large corporation, users from a particular industry sector, or users with the need to exchange large amounts of information. Their common interest is the need to communicate accurately and economically.

It is possible that PTTs in some countries will assist the formation of such interest groups. By encouraging groups of like users to pool their networks and adopt common standards (PTT-approved, of course), they will be able to limit the proliferation of private networks, and thus smooth the path to introduction of public data networks.

2. Terminal compatibility

For many users, the prospect of an intelligent network enabling terminals of different types to address a variety of computer systems with equal ease must be attractive. Many problems remain to be solved before this can be looked on as more than a possibility, and even then there remain doubts as to whether it will prove cost-effective. In essence, cost-effectiveness will depend very heavily on the degree of incompatibility that needs to be resolved.

C Questions Of Policy

1. PTT transmission facilities

During the next several years public data networks will be developing to the point where they become a realistic alternative for many Communications Managers. Meantime most users are likely to be heavily involved in data communications using leased or dialled lines over the analogue voice network. The user who has established an appropriate policy for the use of existing facilities is likely to be best placed to take full advantage of the new ones. The objective must be to create a situation where PTT facilities can be exploited most economically as they develop.

One of the essential functions of a network will be to give the user independence from the precise form of the PTT transmission services, but this freedom will not be unconditional. While relative costs of certain features may change, many of the cost factors will be the same as for today's networks – short response times will cost extra; circuits, whether virtual or physical, can be used inefficiently; unintelligent terminal procedures can introduce heavy penalties, and so on.

So the availability of more powerful and more flexible public data facilities will not change the basis for good design, but it will change some of the rules. These changes must be foreseen and provided for.

2. Standards

Many of the benefits of networks will not be attained effortlessly, nor within a short time-span. They require planning and policy decisions, effectively applied.

For example, the installation and operational cost of a network might be substantially altered if bulk data were transmitted only overnight and were spooled locally, or output at unattended terminals. Such a situation will only come about if remote batch terminals have appropriate facilities, if users are schooled in a particular way of working, if computer loading schedules allow it, and so on.

Again, the overhead of protocol conversion – and hence the cost of the network – could be reduced in two ways:

a By standardising on a limited range of terminals. Such a policy could be presented to users on the basis that new network-based services would only be available via these standard terminals, with continuing, limited support for other devices already in use.

b All terminal suppliers could be presented with a standard interface – a virtual terminal protocol – which they were required to match, thus progressively reducing the degree of incompatibility which the network was required to resolve as terminals were replaced.

D The Vital Ingredients

Our survey of current network activity (see Section IV) indicates that packet switching in itself is not the key discovery. Many of the changes which we traced in the move towards intelligent networks were visible independently of and in parallel to the development of packet switching.

1. The structure of the network

The network can be seen as consisting of a number of functional layers. The major divisions within the network structure, reflected in these functional layers, are:

- between the network proper i.e. the transmission and switching mechanisms, and the routines supporting the devices connected to it

- between the network in total and the PTT facilities which it uses.

This gives rise to a structure in which the network consists of two sub-networks, one of which is a buffer from the PTT facilities, the other a buffer from the connected devices (see Exhibit 9). These can be termed access and transport sub-networks respectively. Within such a structure there could be as many functional layers as necessary to maintain independence from changes occurring at a lower level.

The importance of this structure can be seen if we consider the possibility of the PTTs themselves offering packet switching services at some time in the future. If the right protocols have been adopted within the transport sub-network, it can be discarded entirely, allowing the access sub-network to use the PTT facilities directly. An integrated network structure, on the other hand, would necessitate major surgery.

2. Virtual Terminals

Virtual Terminals promise an enduring solution to the problem of terminal and mainframe incompatibility. Needless to say, such a solution will not be attained without difficulty. It is also possible that less ambitious schemes will prove the most attractive for many users.

E Judgment On Packet Switching

Despite the comments above, packet switching clearly is of major importance, combining as it does the store-and-forward capability of message switching with the fast response times demanded by interactive traffic.

Although not yet totally validated in all its aspects, it has demonstrated its value for public networking services of one kind and another. We have no doubt that it will also prove its worth for private networks.

As things stand at present, the very high front-end cost of packet switching (largely due to the fact that no generalised hardware and software exist, so that every implementation tends to be one-off), effectively limits it to the largest users. Even then special circumstances may be needed to provide a solid cost-justification. We expect this to change over the next two or three years, opening the market to many more users.

Packet switching, as it exists at present, makes relatively hard work of bulk traffic, yet increased terminal intelligence and distribution of computing power are likely to increase the proportion of bulk traffic: an hierarchical database system, for example, might exchange files morning and evening.

Against this background it is interesting to speculate about alternatives to packet switching. For example a recent paper in the Communications of the ACM describes an experiment using very different methods which in many respects seem more appropriate in the corporate environment than packet switching. We also anticipate the telecommunications

Exhibit 9 Network macro-structure

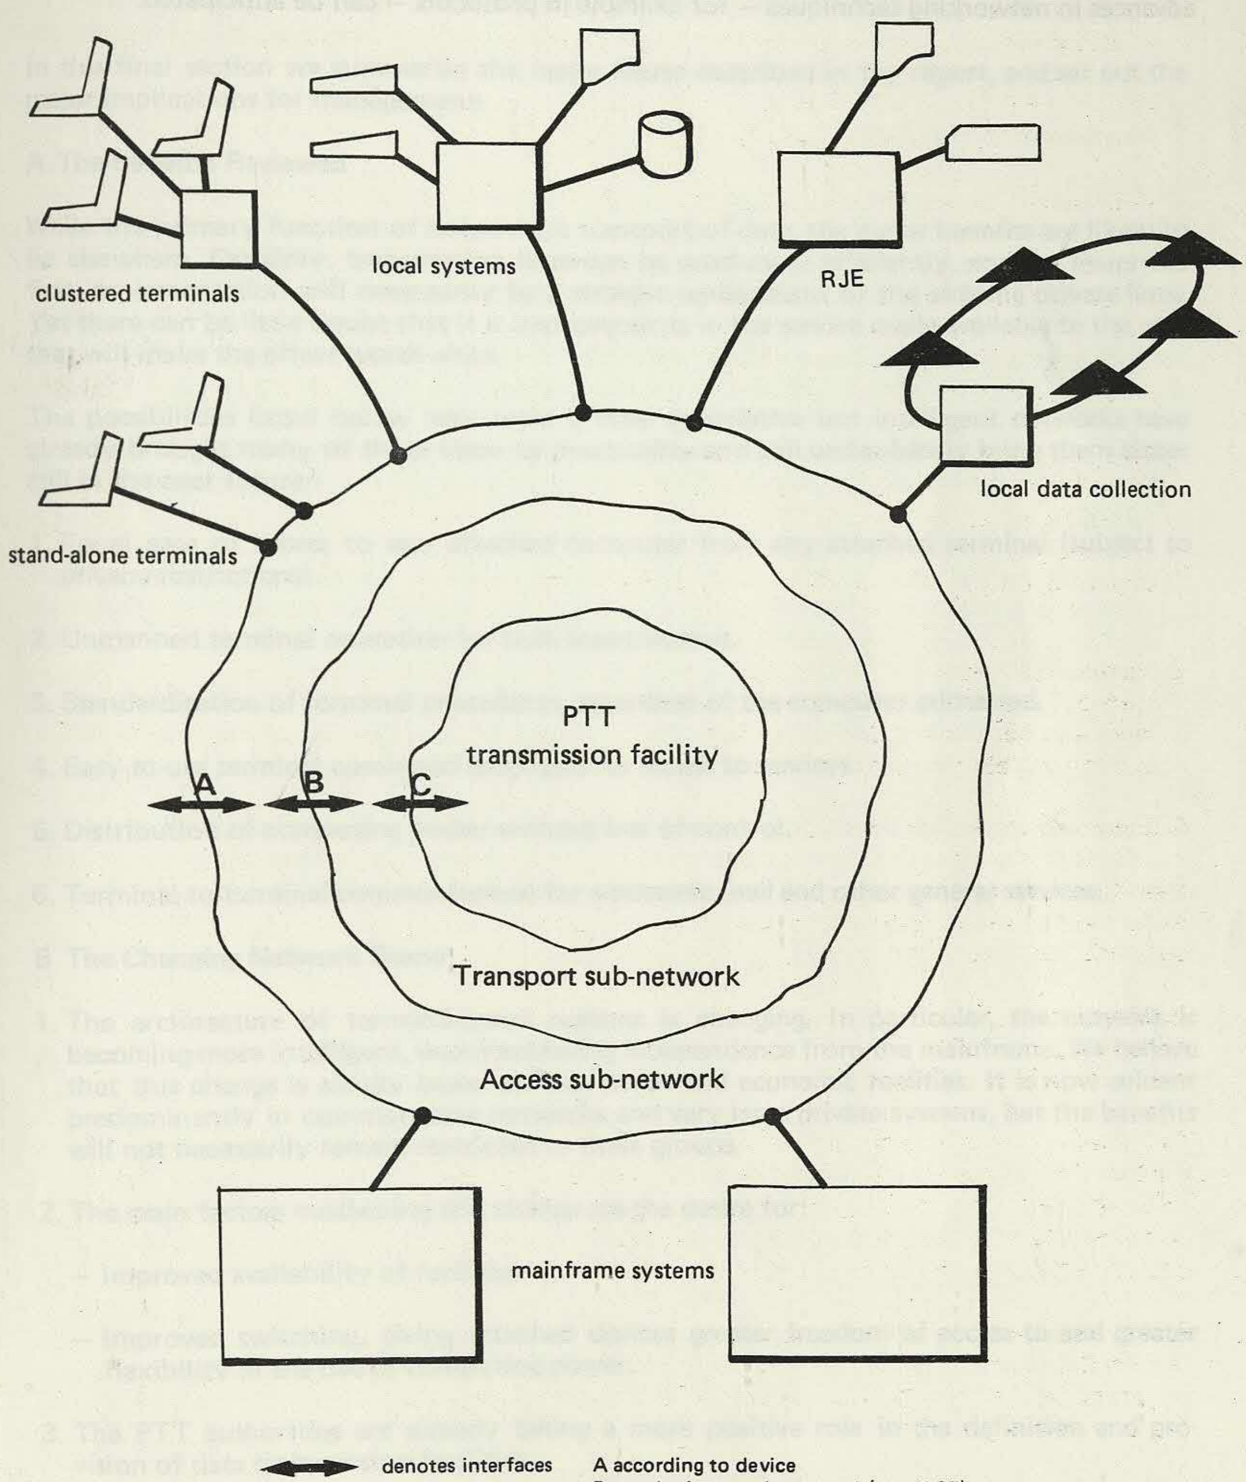

manufacturers influencing events as they face the problems of switching data traffic with their PABXs and other voice-oriented devices.

Nonetheless, packet switching has helped us to see networks in a new light, and is likely to form the basis of public data networks for some years to come. For private users further advances in networking techniques — for example in protocols — can be anticipated.

VIII IMPLICATIONS FOR MANAGEMENT

In this final section we summarise the major issues described in the report, and set out the major implications for management.

A The Benefits Reviewed

While the primary function of networks is transport of data, the major benefits are likely to lie elsewhere. Certainly, transmission lines can be used more efficiently, and for many the first implementation will necessarily be a straight replacement of the existing private lines. Yet there can be little doubt that it is improvements in the service made available to the user that will make the effort worthwhile.

The possibilities listed below may seem a little speculative but intelligent networks have already brought many of them close to practicality and will undoubtedly bring them closer still in the near future:

1. Equal ease of access to any attached computer from any attached terminal (subject to privacy restrictions).
2. Unmanned terminal operation for bulk input/output.
3. Standardisation of terminal procedures, regardless of the computer addressed.
4. Easy-to-use terminal command languages for access to services.
5. Distribution of computing power without loss of control.
6. Terminal-to-terminal communication for electronic mail and other general services.

B The Changing Network Scene

1. The architecture of terminal-based systems is changing. In particular, the network is becoming more intelligent, and is achieving independence from the mainframe. We believe that this change is solidly based on user needs and economic realities. It is now evident predominantly in common user networks and very large private systems, but the benefits will not necessarily remain restricted to these groups.

2. The main factors motivating this change are the desire for:

   – improved availability of facilities

   – improved switching, giving attached devices greater freedom of access to and greater flexibility in the use of computing power.

3. The PTT authorities are already taking a more positive role in the definition and provision of data transmission facilities.

   Their role in the future could crucially influence the development of terminal-based systems, ensuring at one extreme that they become commonplace, at the other that large scale use remains restricted to larger users.

4. A more positive role on the part of the PTTs is likely to erode the ability of computer manufacturers to set and maintain their own standards for data communications. These standards are often defensive in nature and designed to protect the manufacturer's investment in a total systems capability.

5. Computer manufacturers' offerings in the networking field are at present little more than a rationalisation of a confused, even chaotic, past situation. However, they hold out the promise of more far-reaching advances later.

C Management Guidelines

It is too early for the majority of users to begin implementing the type of network outlined in this report. In general, special circumstances will be needed to justify such immediate action and many will gain by waiting the few years during which important issues will be clarified and the cost and risk of implementation will start to diminish.

Nonetheless in the meantime we consider that there are some positive steps which managers can take if they recognise within their own operations the problems which networks attack:

1. They can monitor — or even attempt to influence — developments in high level protocols and virtual terminals, relating to their own organisation's practice and needs for the future.
2. They can refine progressively their own requirements in those areas, and thus move towards a position where they can evaluate fully the major new possibilities which networking can bring. This applies particularly to users with a heavy dependence on data communications, or sharing a mutual interest with other organisations.
3. They can establish a policy on terminal procurement and the use of existing facilities, again to smooth the path towards new opportunities.

4. They can begin an education programme for management services staff and users, to accustom them to the new ideas and new systems possibilities that networks will open up.