Report Series
No. 28

User Experience with Data Networks

April 1982

The Butler Cox Foundation

MEMO

DATE ________

TO ________

____

FROM ________

____

THE BUTLER COX FOUNDATION

REPORT SERIES NO. 28

USER EXPERIENCE WITH DATA NETWORKS

ISSUED APRIL 1982

Abstract

The purpose of this report is to assist data communication users to select the most appropriate form of network products and services. It documents as case histories the experiences of some typical users of existing data communication products and services, and draws out the general lessons that might help others who are about to establish their own networks.

The report concentrates on data network products and services that are currently available in Western Europe — public data services (circuit-switched and packet-switched), proprietary network architectures provided by the leading computer suppliers, and local area networks.

The report is concerned not only with the technical problems to be faced but also with the management and administrative difficulties encountered in establishing and running a data network. It is intended primarily for those who have responsibility for planning, installing and operating data communication networks.

Research team

The team that researched and authored this report was:

—*John Daly:* a consultant with Butler Cox specialising in telecommunications and network design. He has conducted surveys to establish the market potential for communications products involving both users and manufacturers.

—*Fred Heys:* a senior consultant with Butler Cox specialising in new product and marketing strategy studies. He has led a study of future developments in office communication systems and networks for British Telecom, and he has worked on other projects concerned with data communications and office automation.

The Butler Cox Foundation

© Reproduction by any method is strictly prohibited

Photoset and printed in England by Flexiprint Ltd., Worthing, Sussex

THE BUTLER COX FOUNDATION

BUTLER COX & PARTNERS

Butler Cox is an independent management consultancy and research organisation, specialising in the application of information technology within commerce, government and industry. The company offers a wide range of services both to suppliers and users of this technology. The Butler Cox Foundation is a service operated by Butler Cox on behalf of subscribing members.

OBJECTIVES OF THE FOUNDATION

The Butler Cox Foundation sets out to study on behalf of subscribing members the opportunities and possible threats arising from developments in the field of information systems.

New developments in technology offer exciting opportunities — and also pose certain threats — for all organisations, whether in industry, commerce or government. New types of systems, combining computers, telecommunications and automated office equipment, are becoming not only possible, but also economically feasible.

As a result, any manager who is responsible for introducing new systems is confronted with the crucial question of how best to fit these elements together in ways that are effective, practical and economic.

While the equipment is becoming cheaper, the reverse is true of people — and this applies both to the people who design systems and those who make use of them. At the same time, human considerations become even more important as people's attitudes towards their working environment change.

These developments raise new questions for the manager of the information systems function as he seeks to determine and achieve the best economic mix from this technology.

MEMBERSHIP OF THE FOUNDATION

The majority of organisations participating in the Butler Cox Foundation are large organisations seeking to exploit to the full the most recent developments in information systems technology. An important minority of the membership is formed by suppliers of the technology. The membership is international with participants from the United Kingdom, France, Sweden, Switzerland, Denmark, The Netherlands, Belgium, Italy and the United States.

THE FOUNDATION RESEARCH PROGRAMME

The research programme is planned jointly by Butler Cox and by the member organisations. Each year Butler Cox draws up a short-list of topics that reflects the Foundation's view of the important issues in information systems technology and its application. Member organisations rank the topics according to their own requirements and as a result of this process a mix of topics is determined that the members as a whole wish the research to address.

Before each research project starts there is a further opportunity for members to influence the direction of the research. A detailed description of the project defining its scope and the issues to be addressed is sent to all members for comment.

THE REPORT SERIES

The Foundation publishes six reports each year. The reports are intended to be read primarily by senior and middle managers who are concerned with the planning of information systems. They are, however, written in a style that makes them suitable to be read both by line managers and functional managers. The reports concentrate on defining key management issues and on offering advice and guidance on how and when to address those issues.

FOLLOW-UP TO THIS REPORT

The research team who prepared this report would welcome the opportunity of discussing its findings with small groups of members. If you would like to participate in such a discussion, please contact your local Foundation address shown on the back cover of this report, and let us know the points on which you would like the researchers to expand. We will then contact you to arrange a suitable date and venue.

The Butler Cox Foundation
© Reproduction by any method is strictly prohibited

USER EXPERIENCE WITH DATA NETWORKS

CONTENTS

PREFACE		i
1	DATA COMMUNICATION TRENDS	1
	Data traffic characteristics	1
	The importance of data communications	3
	Management and control of data communications	3
	Data network products and services	3
	Data communication standards	7
2	PUBLIC CIRCUIT-SWITCHING DATA NETWORKS	8
	European circuit-switching data networks	8
	Experience with the Nordic public data network	10
	Other experience with circuit-switching data networks	11
3	PUBLIC PACKET-SWITCHING DATA NETWORKS	12
	European packet-switching data networks	12
	An oil company's experience with Transpac	14
	Other experience with packet-switching data networks	16
4	WIDE-AREA PRIVATE DATA NETWORKS	17
	Computer suppliers' network architectures	17
	Experience in installing an SNA-based system	21
	Experience in installing a Decnet-based system	24
	Other experience of using SNA and Decnet	25
	Experience with other wide-area network products	26
5	LOCAL AREA DATA NETWORKS	28
	The need for local area networks	28
	Local area network products	28
	Experience with a broadband network	32
	Experience with a baseband bus local area network	34
	Experience with a baseband ring local area network	36
	Other experience with local area networks	37
6	DATA NETWORKS — THE KEY ISSUES	38
	Planning and choosing data networks	38
	Operating and controlling data networks	40
	Managing data communication services	42
7	GUIDELINES FOR DATA COMMUNICATION MANAGERS	43
	Long-term plans and strategy	43
	Medium-term plans	44
	Operation of the network	44
CONCLUSION		46
APPENDIX		47

© Reproduction by any method is strictly prohibited

PREFACE

The Foundation's very first report, published in July 1977, was concerned with data networks. At that time there was little real experience of using the proprietary network systems offered by the computer suppliers, and public data networks were at an embryonic stage of development. The report therefore concentrated on the technology of data communication. It did look forward, however, to the opportunities that would be provided by the new data communication services and products.

In the four-and-a-half years since Report No. 1 was published, many of the products and services described there have become commercially available. Experimental public networks operated by the PTTs have moved from trials to become full public services, and proprietary network products, such as IBM's Systems Network Architecture (SNA), are now used by many organisations. The time is now opportune to re-examine data networking in the light of users' experiences.

Foundation Report No. 1 concentrated on what was then future technology. This report concentrates on user experience of current products, and uses the lessons of that experience to provide guidance to organisations about selecting and managing data networks. In addition to those products described in the earlier report, a wide variety of other products is now available in the marketplace. The most prominent of these products are the local area networks that are intended for use within a single site. Despite their recent appearance, there is already some experience of using them.

The purpose of this report

The purpose of this report is to assist data communication users to select the most appropriate form of network products and services. It documents as case histories the experiences of some typical users of existing data communication products and services, and draws out the general lessons that might help others who are about to establish their own networks.

The report is concerned not only with the technical problems to be faced but also with the management and administrative difficulties encountered in establishing and running a data network. It is intended primarily for those who have responsibility for planning, installing and operating data communication networks. The report also provides an insight into the constraints imposed by current data networks on other information processing systems. It should therefore be of value to other data processing managers and analysts, and also to suppliers of data network products and services.

The scope of the report

The report concentrates on data network products and services that are currently available in Western Europe, and in particular on:

- —Public data network services, both circuit-switched and packet-switched.
- —Proprietary network architectures provided by the leading computer suppliers.
- —Local area networks.

The report does not predict the nature of products and services that will become available in the future, although comments on the most likely short-term trends are included where appropriate. Nor does it provide a critique of individual suppliers' products. Any comparisons that are made are between the different approaches to data networking — whether to use public services or proprietary products, for example.

The structure of the report

The framework for the analyses of the users' experiences is set in chapter 1, first by reviewing the need for data communication, and how that need is changing, and then by summarising the different ways in which those needs may be met. Chapter 1 also outlines the current status of data network standardisation. Chapters 2, 3, 4 and 5 then consider particular kinds of network approach in turn:

- —Public circuit-switching data networks.
- —Public packet-switching data networks.
- —Proprietary computer suppliers' network architectures.
- —Local area networks.

Each of these chapters includes a brief description of the salient features of the particular kind of service or product, and also the comments made by users about its problems and benefits.

© Reproduction by any method is strictly prohibited

The key issues faced by data network users, including management and administrative issues, are highlighted in chapter 6. Finally, in chapter 7, our conclusions on the key issues are presented as a set of guidelines for data communication managers.

Much of the material in the report is, by necessity, of a technical nature. Frequent mention is made of CCITT recommendations, and a table of the recommendations that are most relevant to data networks is included as an appendix.

The Butler Cox Foundation
© Reproduction by any method is strictly prohibited

CHAPTER 1

DATA COMMUNICATION TRENDS

The demand for data communication facilities continues to grow at an ever-increasing rate. This growth is illustrated in figure 1 in terms of the number of data network termination points in Western Europe, as predicted in the recent Eurodata forecast. Also, the way in which many organisations are using data communication facilities is changing. Instead of connecting a handful of terminals at one site to a single processor, organisations now need to connect many terminals (which may be dispersed over several sites) to any one of several computers, which may or may not be at the same site as the terminals. Because of this change in user needs, data networks now have to provide many access points and a switching function. They must also be easy to extend as the number of users grows.

Before recounting the results of our researches into users' experiences with the different data network products and services, it is helpful to look at what users are expecting from those products and services. In this chapter, we begin by reviewing the changing characteristics of data transmission traffic. We then highlight the growing importance of data communications to many organisations by showing that data communication costs are accounting for an increasing proportion of the total data processing budget. The next section identifies the management and control issues generated by the increasing use of data communication facilities. The chapter continues with an overview of the main categories of data networks and services (the public telephone network, network control systems, public data networks, private wide-area networks and local area networks). The final section of the chapter outlines the main trends in the development of data communication standards.

Figure 1 Forecast growth in data connections to public networks

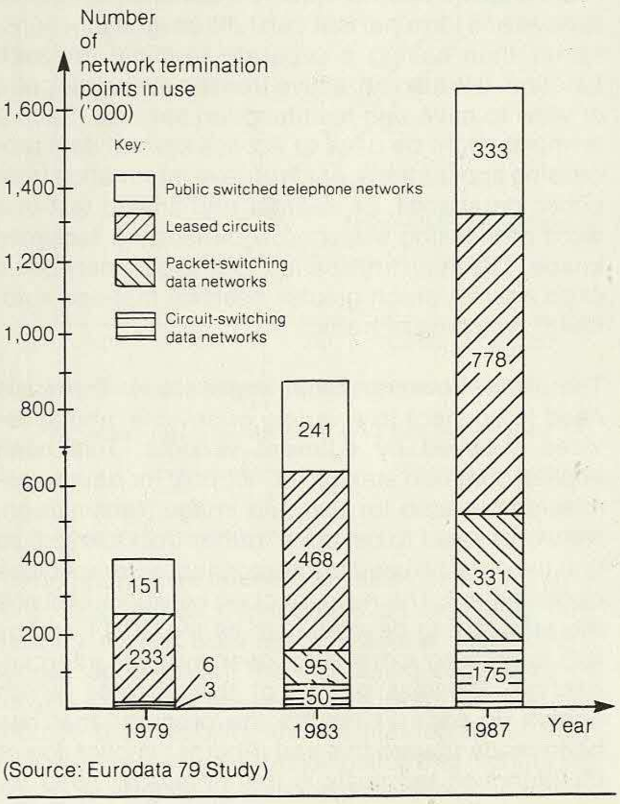

(Source: Eurodata 79 Study)

Much of the material in this chapter will be familiar to those readers who are responsible for data communications planning and operation. The main purpose of the chapter is to provide a basic understanding of data communications for those who are not familiar with the subject. We suggest that those who already are familiar with the subject should omit this chapter and turn to the following chapters, which deal with products and services in detail.

DATA TRAFFIC CHARACTERISTICS

The traffic to be transmitted by data networks is not only growing in volume, it is also changing its nature. This change is a consequence of changes in the applications that need to be catered for.

The first data communication applications were predominantly for remote job entry and for accessing timeshared computers. (Airline reservation systems were a specialised exception.) During the last ten years the price of terminal equipment has fallen, and, as a result, the emphasis has shifted towards transaction applications. More recently, the trend has been to transfer the processing power from a central mainframe computer to several distributed minicomputers and intelligent terminals. Each distributed device may perform separate functions, but it requires access both to common data files and to the processing performed by the other devices.

These trends have several implications for data networks. The main implications are a change in the profile of data calls, and an emerging demand for multifunction terminals.

© Reproduction by any method is strictly prohibited

Data calls

Remote batch traffic consisted largely of lengthy calls during which large volumes of data were transferred to a single destination (or, at the most, to a very few destinations). During a call, the communication traffic was mainly in one direction. High transmission speeds and low error rates were desirable but not crucial, and the speed of call set-up times was not important. Because there was no real switching requirement, lines could be leased from the PTTs to give better quality transmission than could be achieved over the public switched telephone network.

Timesharing application calls had a different profile. Lower volumes of data had to be transmitted during a call, but communication was two-way — a dialogue rather than a monologue. Call durations were shorter than remote batch calls, but tended to be longer than telephone calls. This characteristic could lead to capacity problems if the PABX were to be used for such data traffic. Transmission rates of no more than a few thousand bits per second were adequate for timesharing applications. In addition, call set-up times ideally needed to be faster than for remote batch calls, but they were not significant compared with the normal call duration. Also, as timesharing service bureaux and on-line databases became more prevalent, there was an increasing need for switching in the link to the terminals.

In transaction applications, the calls are shorter than in timesharing. Only a few bits of data may need to be transmitted during a call and, as a consequence, short call set-up times have become more important. In addition, calls from any one terminal may be relatively infrequent, and transaction applications that do not require the frequent use of individual terminals have been encouraged as cheaper terminals have become available. The cost of dedicated leased lines rules out their use for such applications, even though the terminals may need to communicate with only one computer. In an increasing number of transaction applications, however, any one terminal may need to access several computers. For example, a credit card terminal in a shop may need to communicate with several different banks' computers.

Distributed processing places yet further demands on the data network. File transfers, or downloading programs into a local processor from a remote machine, require high-speed and error-free transmission, with fast call set-up times and clear-down times. Transfers such as these are likely to be made from a database shared by many users. The database will be used more efficiently if the communication channels can match the data transfer rates of the storage medium. Distributed processing implies also that there will be many communication paths, and that the paths will form a complex pattern rather than a simple star or a sparse mesh. Fast switching of calls therefore becomes a much more important feature of the data network.

The additional processing power that is now being built into distributed terminals will be used to make several data calls simultaneously. For example, a terminal may call for data from more than one file and then pass the results of its processing to yet another destination. Certain applications will call for very high peak-transfer rates compared with the average rate, but these high rates will be required only for short periods. Peak rates of between 50k bit/s and 500k bit/s (or even higher) may be required.

Communication in distributed processing networks tends also to be between peers rather than between master and slave. Neither of the communicating parties is obviously in control of the dialogue, and calls may be initiated by either of the parties rather than mainly by one of them. Thus, error checking procedures are needed at both ends of the communication path rather than being concentrated in central processors that control the communication. Alternatively, the network itself has to provide its own error checking and correction, and has to guarantee very low error rates in the data transmitted.

Multifunction terminals

From a user's point of view, it is obviously attractive to have one terminal that can fulfil several functions, rather than having a separate terminal for each function. It is also attractive from an economic point of view to have one multifunction terminal. Such a terminal might be used to access several data processing applications, or to retrieve information from public databases, or to enter and amend text in a word processing system, or to display a facsimile image. The main implication for the data network is once again a much greater need for fast and automatic switching of traffic.

There are, however, other implications. Terminals need to connect to a variety of devices, and to devices supplied by different vendors. This need implies common standards, not only for data transmission but also for text and image transmission. Networks need to be 'open' rather than 'closed' so that they can be used to interconnect different types of equipment. The need for open networks explains the effort being devoted both by the CCITT and by ISO to defining a model for open systems interconnection. The last section of this chapter (which begins on page 7) reports the progress that has been made toward this end. Another implication of multifunction terminals is the increasing need for

The Butler Cox Foundation
© Reproduction by any method is strictly prohibited

local area data networks to distribute information within a site as well for wide-area networks to distribute information between sites.

THE IMPORTANCE OF DATA COMMUNICATIONS

The growing use of data communication facilities means that a higher proportion of data processing budgets will be accounted for by communications. We asked several organisations about the extent of their data communication activities and about trends in their expenditure. The results, as shown in figure 2, indicate that most companies are now spending as much as 5 per cent or 10 per cent of their data processing budgets on communications. As a consequence, companies are now more aware of the benefits to be gained from searching both for cheaper networks and for more efficient ways of using their existing networks. Cost is now a most important criterion when companies choose between alternative network products and services.

Figure 2 Relative expenditure on data communications and data processing in a sample of Foundation members

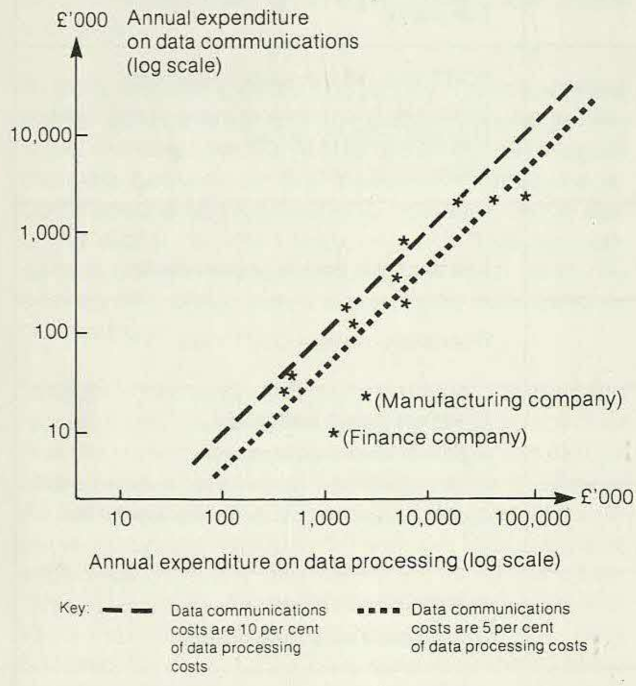

The companies we talked to are also becoming more conscious of their dependence on their data network for routine business operations. This dependence, together with the increasing complexity of networks, puts increasing pressure on suppliers to improve the reliability and maintainability of their equipment. Hence, fault-finding aids and monitoring aids have assumed greater importance.

Data networks are a rapidly developing area of technology, and expertise to control and manage the networks is in short supply. Ideally, therefore, networks should be capable of operation by people who do not possess specialised skills. Networks should be user friendly not only to the end users but also to the operators. This requirement will increase the network complexity and cost — and also the problems faced by the vendors, who are themselves suffering from a lack of expertise to design and develop these more sophisticated products.

MANAGEMENT AND CONTROL OF DATA COMMUNICATIONS

As data networks increase their penetration within an organisation, they create management and control problems of their own. The networks are becoming more of a shared facility rather than a service for specific applications. Consequently, the policy issues to be resolved are similar to those for other company services, such as those for data processing or for the internal telephone network. These policy issues include:

- The way in which data networks should be justified, and how they should be budgeted for and paid for.
- Whether the network and its attached devices should be controlled centrally.
- The way in which a sensible strategy can be evolved so that individual sites can decide on their own systems while retaining the ability to communicate with each other.
- The functional area that should be responsible for data communications. The two choices are the data processing function or a separate communications function that is responsible for all telecommunication facilities.

As part of the research for this report we asked organisations for their views on these policy issues. Before we set out our findings in chapters 6 and 7 as they affect the management and control of networks, we first concentrate in chapters 2, 3, 4 and 5 on user experience with the hardware and the software of different network technologies.

DATA NETWORK PRODUCTS AND SERVICES

In the late 1960s and early 1970s data communication users were constrained by the products that were available but, over the last 15 years or so, equipment suppliers and the PTTs have struggled to keep up with the changing demand from the marketplace described in the previous section of this

© Reproduction by any method is strictly prohibited

chapter. Today, though, there are, at last, some products that are leading the market demand.

Initially, data communication offerings were based on the public telephone network, either by dialling-up or by using dedicated leased lines which could be configured as a private (un-switched) network. The next development in private networks was the provision of network control systems. Later, the PTTs set about broadening the market for data communication services by developing public data switching networks. The technology used by these public networks is now also being used by some of the larger organisations in building their own private data switching networks. All of these data communication offerings concentrate on communication between sites, rather than within a site. More recently, there has been a greater emphasis on the technology of local area networks that can be used to distribute information within a site. We now provide an overview of the development and the technology of each of these five types of data communication offerings.

Use of the public telephone network

Over the years, data transmission speeds available on the public telephone network have increased as modem technology has advanced, and today there are several CCITT recommendations covering a wide range of modem speeds.

In general, the higher speeds have been available for transmissions using leased lines. A leased line dedicated to transmission between two fixed points in the network can be configured to provide higher bandwidth and better transmission characteristics than can be guaranteed over any path that may be selected by the switched network as an individual call is set up. As an example of the data services that can be provided over telephone line connections, figure 3 shows the Datel services available in the United Kingdom.

In most European countries, the PTT has been responsible for providing both the modems and the line to the user's site. Where the PTT has not been

Figure 3 British Telecom's Datel services

Service name	*Modem type*	*Launch date*	*Comments*
Datel 200 Datel 200 Datel 200	2A 13 21	1967 September 1975 September 1980	CCITT V24 (1964) discrete components Limited facilities (4 interchange circuits). Fixed channels call-only modem Improved rack density
Datel 600 Datel 600	1A 20	March 1965 Summer 1975	— Improved rack density (two per shelf)
Datel 1200 Duplex service	27A 27B	October 1980 October 1981	Proprietary modems CCITT V22B
Datel 2400	7A 7B	December 1968 May 1971	Integrated circuit technology Improved speech performance
Datel 2400 Dial up	7C	September 1972	Fixed equaliser to give 2,400 bit/s over PSTN
Datel 2412	12 12MK II	May 1977 September 1979	CCITT V26 and V26 bis Unattended switching to standby
Datel 4832	24A	May 1978	Proprietary modem (Racal Milgo MPS48)
Datel 4800	11	January 1980	—
Datel 48k	8 & 9	July 1970	—
Datel 9600	30A/30B	April 1981	Based on Racal Milgo 9629 modem

(Source: British Telecom)

The Butler Cox Foundation
© Reproduction by any method is strictly prohibited

able to provide its own modems to meet the high transmission speeds required by some users, it has allowed proprietary modems to be used on its leased lines. This has sometimes been a temporary measure until the PTT could develop or acquire its own higher speed modems. The degree to which each PTT has enforced its monopoly to supply modems has varied from country to country.

The user's choice between leased lines or the switched telephone network was dependent on two main factors:

— The expected utilisation of the data link, because the tariff for a leased line is independent of line usage.

— Whether or not the data terminal (or other device) was required to connect to several different computers at different locations.

The first factor can be resolved simply by comparing costs. The higher the utilisation, the more likely that the rental of a dedicated line will be less than the usage costs incurred by switched telephone network connections. Where a terminal accesses only a single computer there is no need for the switching facility provided by the telephone network but, even so, the higher cost of a leased line could prove prohibitive.

If users required a switching facility, there was little choice other than to use the public telephone network. The only alternative was to switch the terminal manually between several leased-line terminations. Some specialised applications, however, made use of the digital switched telex network. This network offered only very low-speed transmission rates (50 bit/s to 100 bit/s), but it did provide international connections.

In order more easily to justify the use of expensive leased lines, multiplexors have been developed so that the transmission capacity of a single line can be shared between several communication channels. As with modems, the transmission speeds of multiplexors have increased and they have become more sophisticated. Early multiplexors shared the capacity of the line by dividing the available capacity into fixed channels. For example, a 4,800 bit/s line could be used for four 1,200 bit/s data channels. More recently, statistical multiplexors have appeared which share the line capacity dynamically according to the actual data traffic on each channel at a given time. In the example just quoted, it is most unlikely that all four channels would be transmitting data at the same time. Other data could be inserted into the transmission gaps, so that more than four 1,200 bit/s channels could share the 4,800 bit/s line.

In addition to the increase in their transmission speeds, modems and multiplexors have become more flexible (and hence more complex). Monitoring and diagnostic facilities have been incorporated within them to assist in fault diagnosis and correction. Remote control facilities have also been incorporated so that networks can be reconfigured from a central point if faults occur or if the requirements change. This facility has become increasingly important as data networks have evolved from simple star formations to more complex mesh networks with multiple routes between the nodes.

Network control systems

Despite the monitoring and remote control facilities that were built into modems and multiplexors, the tasks of locating faults and reconfiguring the network were still manual operations and a high level of technical skill was required by operational staff. The next stage in private network development was to provide centrally located equipment to assist the network operators. These network control systems, as they are called, are now available from several suppliers. Note that this equipment is a tool for the operations staff rather than for the end user. Network control systems do not allow a user to switch his terminal between different processors.

The views of some of the organisations we interviewed on the advantages and disadvantages of using network control systems are included in chapter 4 on page 26.

Public data switching networks

The analogue circuits in the public telephone networks suffer from several disadvantages as far as data traffic is concerned:

— They have limited bandwidth and hence provide limited data transmission speeds.

— They are variable in quality, and subject to noise that increases the incidence of transmission errors.

— The time taken to establish a dialled connection is long compared with the time required to transmit short data transactions.

These disadvantages can be overcome by installing a network that is devoted to data traffic. Most European PTTs have established (or plan to establish) such networks as a public service. The majority of the European public data networks are (or will be) based on packet-switching technology, but a few PTTs have opted for circuit-switching networks. Circuit-switching data networks use techniques similar to those used in telephone networks, but they provide much faster route selection. Chapters 2 and 3 respectively discuss the technology and the user experience of both types of public data network.

© Reproduction by any method is strictly prohibited

Private data switching networks

Instead of using the public data networks for switching traffic between sites, an organisation can now choose from a variety of products that enable it to install its own private data switching system. These products may be grouped into the two broad categories of computer suppliers' network architectures, and data switches. These categories are reviewed below, and are considered in more detail in chapter 4.

Computer suppliers' network architectures
Initially, the major computer suppliers provided software and communications control equipment that facilitated the connection of terminal devices to a particular processor. These initial products were designed to perform the routine functions required by all applications involving data communications. As time has progressed, the products have been enhanced to allow much more flexibility in the interconnection of terminals and processors.

For example, a single terminal may be used to access several different application programs, which could be running on different computers. Equipment now exists to connect the terminal automatically to the appropriate computer. The user does not need to be aware of the computer on which the program is available, and need not be concerned with the links that are used to connect his terminal. Powerful products such as these have normally been specified within an overall network architecture. In effect, the network architecture is a data communications standard for the particular supplier. A main aim of the network architecture is to ensure that all future communication products from that supplier will be compatible with each other. The best-known example of a network architecture is IBM's Systems Network Architecture (SNA).

Data switches
A second option for an organisation that wishes to build a private data network is to buy a smaller version of the switches used in the public data networks. (This option is analogous to the use of private automatic exchanges for telephone traffic.) As in the public data networks, there are two possible switching technologies that can be used — packet switching and circuit switching. Examples of products that use each of these technologies are available in the marketplace. The suppliers of these products tend not to be the major computer suppliers, and their products will support computers and terminals from a variety of manufacturers.

Although we have described public and private data switching networks as two distinct approaches to satisfying the same user needs, in practice there is not such a clear-cut distinction between them. For some applications, terminals may need to be connected to processors which are not included within a private network but which are connected to a public data network. Similarly, terminals connected to public networks may on occasion need to access processors within a private network. Interfaces will therefore be required between public and private networks, and both public and private links may be used within an organisation. A typical example is the use of the public network as a fall-back when a private network communication link fails.

Local area networks

Most data communication applications to date have involved communications between two (or more) sites rather than within a single site. This is reflected in the products referred to earlier in this section, which are intended mainly for wide-area networks (although they may be used also within a large site). In the past, terminals and computers located close to each other would normally have been hard-wired together — an adequate arrangement for terminals that were devoted to particular applications. Within a single site, intermittent users (or the few who did require access to several different systems) could be catered for either by using modems and the internal telephone network, or by using dedicated lines connected manually as required. But the internal telephone system is likely to have disadvantages for data transmission similar to those of the public telephone network, and both types of network have obvious disadvantages once the number of users requiring switched data calls increases beyond a certain level.

Recently, two separate developments have emerged that are aimed at providing a better solution to the problems of intra-site data communication:

—The addition to private telephone exchanges of special facilities to cater for data traffic from extensions that have terminals connected to them.

—The use of an internal digital network devoted to data traffic. Such networks are now commonly referred to as local area networks.

The use of a PABX for data transmission has been described and referred to in previous Foundation reports. Report No. 9 (The Selection of a Computerised PABX), Report No. 21 (Corporate Communication Networks) and Report No. 26 (Trends in Voice Communication Systems) all addressed this issue, and so we have not included any detailed consideration of it in this report. Instead, we have concentrated our attention on local area networks, which are described and reviewed in chapter 5.

The distinction made above between private telephone exchanges and local area networks is reflec-

The Butler Cox Foundation
© Reproduction by any method is strictly prohibited

ted in the products that are currently available, but this distinction may become less clear in the future. It is possible to transmit digitised speech on a local area network, for example, and at least one supplier (Hasler) has a local area network product that combines speech and data transmission. Thus it is conceivable that, in the future, PABX facilities and local area network facilities could both be provided on a single internal network. We do not expect such integrated local area networks to be widely used for several years, and in this report we have confined our discussion of local area networks to their use for data communication.

DATA COMMUNICATION STANDARDS

Earlier in this chapter (on page 2) we referred to the increasing interest in open systems interconnection. The aim of open systems interconnection is to provide a communication environment in which computing equipment, terminals and systems can be freely connected together, irrespective of the type of equipment or the supplier. Foundation Report No. 21 described the International Standards Organisation's seven-level reference model for open systems interconnection. This reference model is intended to provide a structure for the various standards required for open systems interconnection.

Foundation Report No. 21 also contained a discussion of the more important individual standards already recommended by the CCITT. Readers are referred to that report for an explanation of the more important CCITT recommendations that apply to the lowest three layers of the ISO reference model. (See also the appendix to this report.)

Progress in agreeing standards for the higher levels of the reference model has been slow and, as a result, the various suppliers of data network products have implemented their own versions of the standards rather than wait for common standards to be agreed. It is important to remember also that CCITT recommendations do not always define unique standards, and so the adherence of any product to a particular recommendation does not guarantee that it will be completely compatible with other products that adhere to the same recommendation.

The specification of the CCITT X.25 Recommendation has, however, encouraged the adoption of a common standard for the lowest protocol levels of a packet-switching network. Even powerful suppliers, such as IBM, have been forced by market pressure to adapt their existing network products to be able at least to communicate with networks using the X.25 protocol.

Similar standards for local area networks are still in the course of preparation. Unfortunately, there is, as yet, no clear sign of any early agreement on a unique standard.

© Reproduction by any method is strictly prohibited

CHAPTER 2

PUBLIC CIRCUIT-SWITCHING DATA NETWORKS

In chapter 1 we identified the two techniques that can be used to provide the switching function in public data networks as circuit switching and packet switching. This chapter is concerned with public data networks that use the former technique. The first section of the chapter describes the basic facilities provided by the public circuit-switching networks in Europe. It is followed by a case history of a user of the Nordic public data network, and the chapter concludes by reporting some general comments made by other users of public circuit-switching data networks.

EUROPEAN CIRCUIT-SWITCHING DATA NETWORKS

The three most distinctive features of any public circuit-switching data network are:

—One terminal can call automatically any other terminal connected to the network. The time to set up the connection is of the order of one second or less.

—Once a call has been established a circuit is continuously available for the transfer of data between the two terminals until the call is cleared down by either terminal.

—Data calls are charged for on a call duration and distance basis, in much the same way that telephone calls are charged for.

The network is operated by the local PTT, and terminals are connected to the network through data circuit equipment (known as DCEs) that are installed on the user's premises by the PTT. The most common interfaces between the terminal and the DCE are defined by CCITT Recommendations X.21 and X.21 bis. Nodes in the network are known as data exchanges, and DCEs are usually connected to a data exchange via concentrators in the network.

Availability of circuit-switching data networks

The first two operational public circuit-switching data networks in Europe were the Caducée network in France and the Bundespost's Datex-L network in West Germany. Caducée is already a well-established network, with about 2,000 subscribers connected at the end of 1981. Datex-L was opened experimentally in 1978 and became a commercial service in 1979. At the beginning of 1982 there were 5,000 user connections to the Datex-L network.

The Nordic public data network became a commercial service in September 1981, again after more than a year of experimental use. This network is available in Sweden, Norway, Finland, and Denmark and, at the beginning of 1982, it contained nearly 1,000 DCEs. The majority of the initial 31 Swedish organisations to use the network were terminal suppliers, but genuine users now include banks, travel agents, computer bureaux, petrol companies and an airline. Televerket (the Swedish PTT) has ordered 4,500 DCEs from L. M. Ericsson for the first phase of the implementation of the network, but technical problems have apparently led to delays in their delivery. A further 15,000 DCEs are also on order, and the PTTs expect to have installed 25,000 DCEs by 1986.

The network has 25 concentrators and one data exchange (in Stockholm), and further exchanges will be opened in 1982 (in Gothenburg) and in 1983 (in Malmö). By 1985, if the PTTs' expectations for subscribers are fulfilled, there should be 125 concentrators in the network.

Other countries who offer, or plan to offer, circuit-switching data networks are Austria, Italy and Switzerland.

Facilities provided by public circuit-switching data networks

Public circuit-switching data networks provide a wide range of facilities, of which the five most important are:

1. Range of transmission rates
Both synchronous and asynchronous devices may be connected to the DCEs, and both types of devices may transmit at a range of rates. For example, the Bundespost's Datex-L network permits synchronous terminals to transmit at rates of 2,400 bit/s, 4,800 bit/s and 9,600 bit/s, and asynchronous terminals to transmit at rates of 50 to 200 bit/s, and 300 bit/s. Even higher transmission rates may be available in the future. The Swedish PTT is considering offering 48k bit/s DCEs that would be used to connect subscribers' multiplexors to the public network.

The Butler Cox Foundation
© Reproduction by any method is strictly prohibited

2. Closed user groups
To provide added security to the data transmitted, certain subscriber numbers may be defined as belonging to a closed user group. Calls to a number in the group may then be made only from other members of the same group.

3. Call distribution
Several DCEs may be allocated to the same subscriber number so that calls received may be shared between the DCEs. This facility is useful where, for example, a computer is likely to receive several calls simultaneously. The network may also provide a call queuing facility for calls that are directed to a DCE that is already busy.

4. Call back
A transaction received over the public network may, for some applications, initiate a process in a central computer that may take some time to complete. Rather than keep the connection open during this process, the calling terminal terminates the call. The called processor is programmed to call the terminal back when the processing is complete.

5. Test facility
The DCEs may incorporate a test facility that enables network subscribers to check both the operation of the DCE and the link to the network. This facility was being used successfully by the non-technical staff of one of the Nordic public data network users that we interviewed.

Those PTTs who offer public circuit-switching data networks recognise the need both for technical support to new users and for the maintenance of the network so that it provides a high level of availability. For example, large potential users can expect to have a PTT technician assigned to them at no charge to help them design and implement their data network applications. The user may have to pay for this service, however, if it extends beyond an agreed time.

In the future, the Swedish PTT will encourage potential users of the Nordic public data network to approach it for free consultancy advice, but is planning to charge for this service after about eight hours' assistance has been given. It also plans to respond to faults at the user's premises within four to six hours, and to return the DCE connection to service within 24 hours for at least 98 per cent of the faults.

Tariffs for circuit-switching data networks

The charges for using public circuit-switching data networks normally consist of three elements:

—An initial connection charge.

—A monthly rental charge for the line terminating equipment.

—A usage charge based on the connection time. The tariff is a function of the distance between the communicating devices.

Figure 4 Examples of public circuit-switching data network tariffs (in dollars)

Cost item	Network name		
	Caducée	*Datex-L*	*Nordic PDN*
Initial connection	288	92 (200 or 300 bit/s) 182 (2,400 to 9,600 bit/s)	282 (600 bit/s) 940 (9,600 bit/s)
Basic monthly subscription per connection	90 to 300[1]	46 (300 bit/s)[3] 170 (9,600 bit/s)	46 (600 bit/s)[5] 94 (9,600 bit/s)
Call set-up	—	0.02	0.0004 (600 bit/s) 0.008 (9,600 bit/s)
Call charge (per minute)	0.04 (up to 50km)[2] 0.26 (over 300km)	0.18 to 1.10[4]	0.01 to 0.34[4]

Notes: All tariffs have been converted to dollars at the following exchange rates:
— FFr = 5.550
— DM = 2.1875
— SKr = 5.325

1. Tariff depends on distance to nearest Caducée connection point and type of line, but is independent of transmission speed.
2. Tariff depends on distance between points connected, but is independent of transmission speed.
3. Tariffs shown are for X.21 (or X.20) connections. Higher charges are made for other connections.
4. Tariffs depend on distance and transmission speed.
5. Tariffs shown are for X.21 connections. Higher charges are made for X.21 bis connections.

(Source: *Eurodata Foundation Yearbook*, October 1981)

© Reproduction by any method is strictly prohibited

In addition, a call set-up charge may also be made. Each of these elements will have different charging rates for different transmission speeds; some typical charges are shown in figure 4.

EXPERIENCE WITH THE NORDIC PUBLIC DATA NETWORK

In this section we report as a case history the experience of a Swedish bank using the Nordic public data network. The bank uses the network to link cash dispensers with its own computer system. The computer system in turn communicates with the computers in the 13 different banks with which the cash dispenser users have accounts.

Data communication application

Bank-card holders can use any one of between 300 and 400 microprocessor-controlled cash dispensers. To withdraw cash, a user enters his bank card and his own personal identification number into the dispenser. The microprocessor calls the bank's computer via the Nordic public data network, which in turn calls the computer at the bank with which the user has his account. If the transaction is authorised, the cash required is then dispensed from the terminal and the user's account is debited.

Each transaction is completed in less than 30 seconds, and there are between two and three million transactions per month. This number represents an average of 300 transactions per dispenser per day, but there are significant peaks in the traffic. For example, most transactions tend to occur during the lunch period or late afternoon period during the week, and during the late morning period on Saturdays. On Sundays there is virtually no traffic.

Choice of data network

There was no public packet-switching network available in Sweden when the decision was made, and the only possible methods of linking the cash dispensers to the central computer were:

- —To use leased analogue lines.
- —To use the public telephone network.
- —To use the Nordic public data network.

The bank chose to use the public data network because the characteristics of the communication application matched the facilities provided by the network. In particular, the application required:

- —Short transactions, with only a few characters transmitted in either direction.
- —Fast call set-up times.
- —Minimal data transmission errors.
- —Low average utilisation of any one terminal.

The low volume of traffic meant that the cost of leased lines to each terminal could not be justified, and the other characteristics favoured the use of the Nordic public data network.

Implementation of the system

The task of linking the cash dispensers and the central computer to the Nordic public data network was subcontracted to the suppliers of the computer and the terminals. The two suppliers worked with each other and with Televerket to install and commission the data communication system. The bank also employed a firm of technical consultants to resolve any technical problems and to liaise with the PTT and the two equipment suppliers.

The bank started to use the Nordic public data network during the network's experimental phase in 1980. It now has nearly 18 months' experience of using the network, even though the network only began commercial service in September 1981. The network uses the X.21 protocol and the bank uses the High-Level Data Link Control (HDLC) format for all data transfers.

For security and privacy, the bank has chosen to use data encryption techniques rather than the closed user group facility provided within the public network service. Banks, of course, are particularly concerned to prevent unauthorised access to accounts, and to prevent possibly fraudulent transactions being entered.

Problems encountered in using the public network

The data communication system took longer than expected to install due to unforeseen problems in using the Nordic public data network. These problems were caused largely by the pioneering nature of the application. In addition to the early stage of development of the public network itself, the bank's terminals and central computer were also new, and all the suppliers involved (including the PTT) lacked experience of X.21 protocols and equipment. This inexperience led to the familiar situation in mixed supplier systems of each supplier disclaiming responsibility for faults. This situation was compounded because the PTT-supplied modems were of a new design and they also had their share of development problems.

The bank believes that these problems would not be typical of those that future users of the public data network may expect. The PTT, the equipment suppliers and the consultants all gained a great deal of basic experience in implementing their parts of

The Butler Cox Foundation
© Reproduction by any method is strictly prohibited

the system. The bank would recommend that future users subcontract system implementation to the suppliers in the same way that it did.

The bank's experience of using the Nordic public data network has highlighted one general point of system design. All parts of the communication system must be informed of any failures elsewhere in the network, and they must then be able to respond appropriately. For example, if there is a public network fault between a terminal and the computer then the terminal should be isolated until the fault is cleared. The bank also had problems arising from the system not recognising that a terminal had failed to respond to a request.

Another comment on the Nordic public data network by the bank was that the PTT would not yet guarantee the required levels of maintenance service. This could be a problem to users who were depending on a high level of system availability.

Future plans

The bank is satisfied with the basic services offered, or planned, by the PTT on the public data network. For example, the facility to bill a called number rather than the caller will be useful to the bank, because a call from a cash dispenser should be charged to the bank being called rather than to the dispenser.

OTHER EXPERIENCE WITH CIRCUIT-SWITCHING DATA NETWORKS

In addition to the Swedish bank, we interviewed other users both of the Nordic public data network and of the Datex-L network in West Germany. These interviews confirmed the advantages of using circuit-switching public data networks for applications with characteristics similar to those just described. Circuit-switching networks are most suitable for applications that generate short data messages from a large number of terminals used by the general public (for example, point-of-sale terminals). Using a circuit-switching network for such applications avoids the overheads of the extra data that needs to be transmitted when packet-switching networks are used. Our interviews also confirmed that the use of a public circuit-switching data network becomes attractive when the total traffic from any one terminal does not justify the cost of a dedicated leased-line connection. Another important consideration is that the automatic and fast call set-up and clear-down for each call does not discourage the public from using the terminals.

Implementation problems arose mainly from the novelty of the network design and from the lack of experience of both the PTT and suppliers when the networks were first introduced. There has also been a shortage of terminal equipment able to match the X.21 protocols used for connecting to the network. Suitable test equipment also has been in short supply. But these problems are transitory ones. Once the teething troubles are cured, users have found that the networks provide a good service. Also, more equipment designed to X.21 standards will become available as the population of circuit-switching network users in Europe grows. Users should now be able to order DCEs from their PTT for installation and commissioning within one or two months of the order being placed.

As yet, problems do not appear to have occurred as a result of any overloading of the data network exchanges, despite the fact that the number of users is now increasing quite rapidly. The existing exchanges, however, have only a limited capacity to establish call connections. For example, the exchanges used in the Datex-L network currently can handle a maximum of 50 call requests per second.

One disadvantage of public circuit-switching data networks, at least in the short term, is the limited scope they provide for international connections. This limitation is not due to any technical reason — it occurs simply because circuit-switching networks have been implemented in only a few countries.

© Reproduction by any method is strictly prohibited

CHAPTER 3

PUBLIC PACKET-SWITCHING DATA NETWORKS

In this chapter we review the experience with public data networks using packet-switching techniques. The structure of the chapter mirrors that of the previous chapter. Thus it begins with a review of the basic features and facilities provided by a public packet-switching network. It then presents a case history of a user of the Transpac network in France, and it concludes with a review of other users' experiences of public packet-switching networks.

EUROPEAN PACKET-SWITCHING DATA NETWORKS

The techniques of packet switching are now well established. Experimental packet-switching networks were established in the early 1970s, particularly in the United States where their technical viability was clearly demonstrated. Their commercial viability is taking longer to demonstrate, but public networks such as Tymnet and Telenet in the United States and several commercial networks in Europe have been operational now for several years.

The key feature of packet switching is that the data to be transmitted is assembled into a series of packets, which are normally of fixed length — say 128 bytes. Each packet is transmitted independently of others, so that the packets forming a single message or transaction can take different routes through the network. Data defining the destination is added as a header to the information content of each packet. Using packet-switching technology, there is no need for the network to dedicate a circuit to a particular call. Each link in the network can therefore carry data from many different calls by interleaving packets along the link. Furthermore, data in transit may be re-routed should one particular link fail.

The prime benefit from packet-switching technology is an economic one. Each call is charged only for its share of the usage of each link, rather than for all the links and for the total duration of the call. But this economic benefit is dependent on the pattern of calls being such that the transmission paths can be shared effectively between several simultaneous calls. For the network user, the economic benefit of packet switching is that the network provider charges for calls on the basis of the volume of traffic rather than on the basis of the distance over which the data is transmitted.

On the other hand, packet-switching networks have several disadvantages compared with circuit-switching networks:

- —The extra data added to each packet (to denote destination and sequence, for example) imposes an overhead on the volume of data to be transmitted. This is particularly important for short data messages where the length of data is less than the normal packet length.
- —The data has to be assembled and disassembled into and from packets as it enters and leaves the network. These processes impose an added complication for terminal equipment and systems. Terminals not designed for packet operation require an interface known as a packet assembler/disassembler (PAD). A PAD may be provided either by the PTT as part of the network or by the user as a privately supplied item.
- —The use of different routes and the sharing of line capacity means that there will be a variable delay in the transmission of packets across the network. The delay will depend on the number of switching nodes through which the packets pass.

Users may connect to a packet-switching network in a variety of ways. They may, for example, use dedicated analogue leased lines and modems to connect to their nearest network access point. Alternatively they can connect by dialling the nearest access point over the public switched telephone network (again via modems). They can also use the telex or teletex networks which may be interfaced to the packet-switching network via special gateways. Users' terminals may be designed to interface directly to the packet-switching network by forming their own packets in accordance with the network protocols. Alternatively, users may connect non-packet terminals via their own protocol converter (PAD), or via a PAD provided at the packet-switching network nodes.

All new public packet-switching networks are adopting the X.25 protocol (and associated X.28 and X.29 protocols) recommended by the CCITT for such networks. But there is some variation between particular networks and, as a result, they are not yet fully compatible with each other.

The Butler Cox Foundation
© Reproduction by any method is strictly prohibited

Availability of packet-switching data networks

In Europe, Spain pioneered the use of packet-switching for a public data network in the early 1970s. The Spanish RETD network did not originally use the X.25 protocol because, at the time, that protocol had not been standardised. The RETD network has been upgraded recently, however, to include the X.25 protocol. Another early example of a European packet-switching network was Euronet, which was developed by the European Community to link together several of the member countries. Euronet became a commercial service in 1980, and it now services several hundred users. But it is to be superseded by the linking together of the various national packet-switching networks.

France

Transpac in France was the first public data network to use the X.25 protocol. It started operation in 1978 on an experimental basis, and became fully operational in 1979. Connections to the American Telenet and Tymnet networks were established at the end of 1979, and a link to Euronet was provided early in 1980. Users can access Transpac by a dedicated line, by the telephone network or by the telex network. There are 19 switching nodes in the network interconnected by 72k bit/s links, and there are nearly 5,000 user ports serving 900 customers. On current predictions there could be 20,000 ports by 1984. The use of Transpac has been encouraged by its tariffs, which are relatively inexpensive compared with the rental of dedicated lines from the PTT. A high proportion of the computers and terminals linked to Transpac (67 per cent and 75 per cent respectively) support the X.25 protocol, compared with other packet-switching networks.

The Netherlands

The Dutch PTT has three packet-switching exchanges in its recently opened Datanet-1 service. They are supported by 57 concentrator (and multiplexor) access points located in the major cities.

The United Kingdom

An experimental packet-switching service (EPSS), limited to three exchanges, was opened in the United Kingdom in 1975. EPSS, which did not use the X.25 protocol, has been superseded by two new services:

- International Packet Switched Service (IPSS), opened in December 1978.
- Packet Switched Service (PSS), which began full commercial operation in the autumn of 1981 after a year's experimental use.

IPSS was introduced to provide an international service so that users could access the American networks (Tymnet and Telenet, for instance). PSS, on the other hand, is for use within the United Kingdom, although British Telecom is installing a gateway between PSS and IPSS so that all domestic users may make international calls from PSS terminals. Both IPSS and PSS use the X.25 protocol. The inter-network gateway will use the X.75 protocol recommended by CCITT. As with other public packet-switching networks, growth in the usage of PSS and IPSS has exceeded expectations. The number of devices connected to IPSS is growing at about 50 per cent per annum, for example, and there are already orders for 400 packet-mode connections to PSS. (A 1978 forecast had predicted that only 160 packet-mode connections would be required by the end of 1983.) At the end of 1981, PSS had 12 switching nodes, and six more are to be installed in 1982.

Other European countries

The table in figure 5 summarises the status of public data networks in Western Europe at the end of 1981.

Facilities provided by public packet-switching data networks

Like circuit-switching networks, packet-switching networks provide a wide range of facilities for the user. The six most significant facilities are listed below.

Multiple transmission rate

Packet terminals may communicate with the network synchronously and with full duplex working. Transmission rates provided are usually 2,400 bit/s, 4,800 bit/s, 9,600 bit/s and 48,000 bit/s. For such communication the terminal is connected to the network by a dedicated line, with a modem both at the terminal and at the packet network node.

Character terminals may also communicate with the nearest packet network node via dedicated lines and modems. The transmission is asynchronous and is at lower speeds than with packet terminals (for example, 300 bit/s or 1,200 bit/s). Similar speeds are available for connecting character terminals to the packet network via the public switched telephone network.

Conversion of transmission rates and protocols

The network performs any conversions that are needed to allow terminal devices operating at different transmission rates and with different protocols to communicate with each other.

Closed user groups

As with the circuit-switching networks, packet-switching networks permit a closed user group to be defined so that only members of that group may communicate with each other.

Permanent virtual circuits

Packet-switching networks provide a permanent

© Reproduction by any method is strictly prohibited

Figure 5 Public data networks in Europe

Country	Circuit-switching		Packet-switching	
	Network name	Start date	Network name	Start date
Austria	Not known Not known	1978[1] 1982[2]	Not known	1982
Belgium	—	—	Not known	1981
Denmark	Nordic PDN	1981	—	—
Finland	Nordic PDN	1981	—	—
France	Caducée	1972	Transpac	1979
W. Germany	Datex-L	1979	Datex-P	1980
Greece	—	—	—	—
Ireland	—	—	IPSS (international)	1981
Italy	—	—	Not known	1983
Luxembourg	—	—	—	—
Netherlands	—	—	Datanet-1	1981
Norway	Nordic PDN	1981	Norpak	1981
Portugal	—	—	—	—
Spain	—	—	RETD	1971
Sweden	Nordic PDN	1981	Telepak	1981
Switzerland	Datex	1981	EDWP	1982
United Kingdom	—	—	PSS (national) IPSS (international)	1981 1980

Notes: 1. Up to 300 bit/s
2. 2,400 to 9,600 bit/s

(Source: *Eurodata Yearbook 1980/81*)

virtual circuit facility whereby a permanent connection is established between two terminals, instead of each packet being treated in isolation as it is routed across the network. This facility is analogous to a normal leased-line connection, as far as the user is concerned.

Call redirection
Calls may, if required, be diverted automatically to alternative destinations. This facility might be used if a particular terminal were out of operation.

Fast selection
In the normal mode of operation with a packet-switching network, a terminal will initiate a call by transmitting one packet. The network responds by transmitting a packet back to the terminal signifying that it has accepted the call. Data packets are then transmitted across the network, and the terminal transmits a special packet to terminate the call. If the data to be sent is less than the length of one packet, this procedure is very inefficient. Packet-switching networks therefore provide a facility whereby a single packet transmitted from a packet terminal can fulfil all three functions.

Tariffs for packet-switching data networks
Tariffs for packet-switching networks are based on the volume of data transmitted and not on the transmission distance. In addition, the subscriber has to pay an access-time charge which is dependent on the length of his calls, and also the rental and/or usage charges for the telephone or leased line that connects the terminal to the network. Further charges are incurred for the use of any of the special facilities provided by the network, although these charges will normally form only a small proportion of the total cost.

Some typical tariffs for European public packet-switching networks are shown in figure 6.

AN OIL COMPANY'S EXPERIENCE WITH TRANSPAC

This case history reports the experience of a large multinational oil company using the Transpac network for communicating data between its various sites in France.

The Butler Cox Foundation
© Reproduction by any method is strictly prohibited

Figure 6 Examples of public packet-switching data network tariffs (in dollars)

Cost item	*Network name*		
	Transpac	*Datex-P*	*PSS*
Connection charge	—	85 (up to 1,200 bit/s) 169 (over 1,200 bit/s)	833 (2,400 bit/s) 3700 (48k bit/s)
Monthly subscription charge	55 (300 bit/s)[2] 833 (48k bit/s)	42 (300 bit/s) 761 (48k bit/s)	230 (2,400 bit/s) 1541 (48k bit/s)
Usage charges:			
Call set-up	—	0.02	—
Connect-time (per hour)[1]	0.1 (300 bit/s) 2.0 (48k bit/s)	0.25	0.43 to 2.28
Data sent (per kilo segment)[3]	2.0	1.4[4]	0.43

Notes: All charges are for accessing the network via a packet terminal and a dedicated line.
1. Based on tariffs for normal office hours. Discounts are available in France for off-peak usage.
2. Includes modem rental for higher transmission speeds.
3. A segment is 64 8-bit bytes.
4. Discount rates available for off-peak usage.

(Source: *Eurodata Foundation Yearbook, 1980-81*)

Data communication applications

The oil company has a large IBM mainframe installation in Paris, and several regional sales offices and depots throughout France. The sales offices are equipped with terminals that require immediate access to the centrally held files in order to answer customer inquiries. The depots are equipped with IBM 8100 minicomputers that also need to communicate with the mainframe computer in Paris. The initial use of Transpac has been to link four of the sales offices (two at Lyon, one at Toulouse and one at Nantes) and one of the depots (in the South of France) to the IBM mainframe.

Choice of data network

The use of the Caducée network was not considered because the company believed that it was not suitable for its application. The only two practical options for a data network were:

—To use leased analogue lines for the links.

—To use Transpac.

Transpac was chosen because it was the cheaper option. The sales offices are approximately 400 to 500 kilometres from Paris, and the rental for each leased line would have been about 10,000 francs per month. On the other hand, each connection to Transpac costs only 2,000 francs per month. The Transpac data transmission charges incurred by each sales office varies between 600 and 1,000 francs per month. (Each sales office generates an average of only 20 to 30 transactions per hour.)

Implementation of the system

This oil company used IBM Network Interface Adaptors (NIAs) and the corresponding IBM software in the 3705 communications controller (which front-ends the mainframe) to connect the mainframe systems to the X.25 protocol used by Transpac. The rental charge for the software was originally 6,000 francs per month, but was subsequently reduced to 1,800 francs per month. Transpac DCEs were installed at the computer centre and the first sales office was connected to Transpac in March 1980. The other three sales offices were added at monthly intervals.

Problems encountered in using Transpac

The only major problem encountered with the use of Transpac was that the network was subject to frequent failure when the system was first installed. During the worst period there were several failures per day. The failures (caused by the Transpac concentrators in Paris becoming overloaded) were reduced drastically when the original concentrators were replaced. At present, the network fails about once per month, and the down time is usually less than five minutes.

The Transpac connections from the depot also suffered from the early problems with the concentrators, and because of this the company reverted to using a leased line for the connection to Paris. But the company is now reconsidering using Transpac again because the local concentrator has been upgraded and the PTT claims that it now provides a satisfactory performance.

© Reproduction by any method is strictly prohibited

Benefits from using Transpac

The cost of using Transpac has been considerably lower than the cost of using leased lines. One particular benefit in using Transpac is that the company has only one department of the PTT to deal with, and this makes it very much simpler to identify the reasons for transmission problems. The company has found that the Transpac centre in Paris is normally aware of any problems in the network wherever they may be, and the centre can predict accurately the likely duration of the problem. The service provided by the Paris centre is a considerable improvement on the equivalent service when problems are encountered with leased lines. Then, each geographical region of the PTT is responsible for testing the leased lines in its own area, and there is no single reference point to which transmission problems may be referred.

Future plans

In the near future, the customer inquiry system is to be extended to another three sales offices (all of which will be connected to the IBM mainframe via Transpac), and a further five depots will be connected via Transpac during 1982.

OTHER EXPERIENCE WITH PACKET-SWITCHING DATA NETWORKS

In addition to the Transpac user's experience described in the previous section, we interviewed users of Euronet, Datex-P, and the PSS and IPSS networks in the United Kingdom. We also talked to a user of a private packet-switching network, whose experience is reported in the next chapter, which deals with wide-area private networks.

Experience with the Datex-P network in West Germany is very limited because the network has only just changed from an experimental operation to a commercial service. At present, Datex-P is used mainly by equipment suppliers and software suppliers to develop their products. The implementation of Datex-P has progressed remarkably smoothly, largely because its design is based on the Canadian Datapac network. A survey of users showed that 95 per cent of the problems originate in users' systems and equipment, rather than in Datex-P. Particular problems have been experienced in connecting non-packet terminals to the network.

When Datex-P was introduced, the Datex-L circuit-switching network was already providing a switched data network service in West Germany. Datex-P was introduced in addition to Datex-L to cater for those users who wished to make international data calls to other packet-switching network users. Similar user needs have led organisations to make use of IPSS in the United Kingdom and Euronet in other European countries. But the international use of packet-switching networks does give rise to further problems. One organisation that transmits data between its offices in England and the United States has found it difficult to locate faults, because the calls were routed via IPSS through Tymnet to Telenet. If a fault occurred, IPSS would typically display a message "CLR DTE", which should indicate a malfunction in the remote terminal. In fact, the remote terminal (in this case, the host computer in the United States) would be fully operational, and the fault would be located somewhere in the intervening networks. (The company's engineers suspect that Tymnet is the weak link in the connection.)

This problem is of increasing concern in view of plans to link together even more packet networks. The IPSS user that we interviewed was particularly concerned about the plan to use the national service (PSS) as the gateway to IPSS. This change would add yet another link to the three already involved in its international traffic.

Users of IPSS had also suffered in a similar way to the users of Transpac because of the overloading of the network when the service was first introduced. Until the beginning of 1981 the network had insufficient capacity to cater for the number of users connected to it. It was difficult to make a connection with another user, and calls were subject to frequent breakdowns. The network has since been upgraded, and so these problems now occur much less frequently. Users of IPSS and PSS were also very critical of the end-user interface. The need to enter about 50 characters within a very limited time, and the very cryptic coded messages provided by the system, make the services extremely unfriendly to use.

PSS users reported that the tariff advantages of the distance-independent charges (which are not dependent on the time of day) were nullified by the rental charges for a dedicated connection to the network.

Apart from the overloading problems mentioned above, the operation of the packet networks that we investigated has been reliable. Where technical problems have occurred, they have tended to be in the data line between the user's premises and the network. There were familiar complaints from United Kingdom users concerning delays in installing and commissioning links to IPSS and PSS, and concerning the division of responsibilities within British Telecom both for the different parts of the connection to the networks and for the services themselves. Nevertheless, users were, in general, pleased to have a public data switching network that replaced their use of the telephone network.

The Butler Cox Foundation
© Reproduction by any method is strictly prohibited

CHAPTER 4

WIDE-AREA PRIVATE DATA NETWORKS

The previous two chapters have reviewed the use of public data networks for communication between sites. In this chapter, we look at the use of private networking systems for such communication. The physical links for wide-area private data networks are, of course, provided by the PTT as leased lines, but the routeing and control of calls is provided by software and equipment operated by the user.

Although it is conceivable that an organisation could develop its own data network entirely from its own resources, the vast majority of private networks are constructed from software and equipment that has been purchased or leased from suppliers. Many organisations have similar communication requirements, and it is more economic to acquire at least the basic network system as an off-the-shelf product from a data network supplier.

The data network suppliers include:

—The major computer suppliers, who offer network products that support their own range of hardware.

—Software and communications companies, who offer packet-switching network systems capable of linking more than one supplier's terminals and processors. Such network systems are similar to the public packet-switching networks.

—Data communication companies, who offer circuit-switching data exchanges similar to the digital PABXs used for voice traffic.

In this chapter, we concentrate on experience with the data communication products of the major computer suppliers. These products are based on each supplier's proprietary network architecture and, because they appeared on the market earlier than private packet-switching or circuit-switching network products, there is more experience to draw on. After describing the general characteristics of proprietary network architecture products, we examine two architectures in more detail — IBM's SNA and Digital's Decnet. We interviewed several users of SNA and Decnet, and we report the experiences of one representative user of each type of network in the form of case histories.

There then follows a short section that summarises the strengths and weaknesses both of SNA and Decnet as perceived by all the users that we interviewed.

We also interviewed a user of a private packet-switching network, and this user's experience is incorporated in the final section of the chapter, which draws together some general conclusions on wide-area private networks.

Although we have chosen to refer to the above private data networks as wide-area networks, they can, of course, be used within a single large site. Conversely, local area networks, which are considered in the next chapter, may also be linked together across several adjacent sites to form a wide-area network. The distinction is therefore somewhat arbitrary.

COMPUTER SUPPLIERS' NETWORK ARCHITECTURES

ISO's open systems interconnection model is intended to provide a standards framework for all data communication systems. The data network architecture defined by a particular computer supplier provides a similar framework for that supplier's products. A proprietary network architecture establishes a framework within which the supplier's future products can be designed so that they can interconnect in a distributed processing environment. Each architecture has been based on a hierarchy of layers (in the same way as the ISO model) and the architecture defines the protocols for interfacing one layer to the next. The software for each layer can therefore be developed independently of the software for other layers. Most of the major computer suppliers have introduced (or at least have announced) their particular network architectures during the past five years. IBM, however, introduced SNA in 1974.

Each architecture is implemented as a family of hardware and software products that perform the functions of the architecture's layers. Figure 7 overleaf lists some of the proprietary network architectures that have been defined.

Benefits from a network architecture

In addition to the benefits to the computer supplier of having a common and consistent set of communication protocols for his products, a network architecture also provides direct benefits both to the system designers who use the products, and to the end users of the systems that are based on them. The

© Reproduction by any method is strictly prohibited

Figure 7 Computer suppliers' network architectures

Supplier	Network architecture	Protocol layers	Key features
Burroughs	Burroughs Network Architecture (BNA)	2	PC, X.25
Data General	Xodiac	4	PC, PS, X.25
Digital (DEC)	Digital Network Architecture (DNA/Decnet)	6[a]	PC, FS, X.25, SNA
Hewlett-Packard	Distributed Systems Network (DS-1000, DS-3000)	4[b]	PC, PS, FS
Honeywell	Distributed Systems Architecture (DSA)	7	PC, FS, X.25
IBM	Systems Network Architecture (SNA)	4	PC, DL, X.25
ICL	Information Processing Architecture (IPA)	4	X.25
NCR	Communications Network Architecture (CNA). NCR/TAM and NCR/DTN	3	PC, X.25, SNA
Prime	Primenet	4	PC, PS, FS, LN, X.25
Univac	Distributed Communications Architecture (DCA)	2	Not known

Key:
a = For phase III of Decnet
b = Based on X.25
PC = Supports inter-program communication
PS = Supports peripheral sharing
FS = Supports file sharing
LN = Supports high-speed local networks
X.25 = Supports X.25 interface to packet-switching networks
DL = Supports downline loading of programs
SNA = Supports links to SNA networks

two most important benefits are a reduction in the programming effort that is required to implement systems, and a greater flexibility in the use of terminals and processors.

Reduction of programming effort

A typical data communication application can involve database software, teleprocessing software, input/output software, communications-controller software and terminal communication protocols, and all of these may need to be written specifically for that application. The trend towards distributed processing will increase the communication element of the total software. A comprehensive network communication system should provide much of this software, so avoiding the need for bespoke programming for each application. In the past, each type of application tended to use a particular teleprocessing access method, so that programmers had to be familiar with a range of teleprocessing software products. For example, IBM provided 35 different access methods and 15 different communication protocols. A comprehensive network architecture should reduce the variety of communication methods.

Flexibility in terminal and processor usage

The network system should allow any terminal to be used for more than one type of application, and to access programs on any number of host processors. This multiple access avoids the need for several terminals at one location. The network system should also permit a wider variety of processors and terminals to communicate with each other.

A proprietary network architecture usually is designed to accommodate only the particular supplier's processors and terminals. But suppliers are beginning to extend the scope of their architectures and products to accommodate other suppliers' equipment. Non-IBM terminals can now connect to SNA networks via an X.25 network, for example, and Digital's Decnet can interwork with SNA networks.

Network architecture features

A network architecture is defined mainly by the specification of the functions performed by each layer and by the protocols used to communicate between layers. As a message passes through the network from application to application (or from user

The Butler Cox Foundation
© Reproduction by any method is strictly prohibited

terminal to host processor), it passes through a succession of software modules each of which performs the functions of the respective layers. Each module adds (or removes) the header data that is required to communicate across the network at that layer. As well as providing products that implement the basic data communication function of the architecture, the computer suppliers also provide supporting products that assist in the operation and management of the network as a whole. These products also make use of the layered structure of the architecture.

Foundation Report No. 21 contained a description of the layers used in the two most well developed architectures — IBM's SNA and Digital's Decnet. (Strictly speaking, Decnet is the collective name for the products that implement the Digital Network Architecture.) That report also contrasted the different natures of the two architectures. Although both architectures use data packet techniques, SNA has been developed mainly to connect remote terminals with host processors whereas Decnet caters mainly for inter-processor communication. The reader is referred to the earlier report for a more detailed discussion of the differences between the two architectures. Report No. 21 did not describe the way in which the architectures had been implemented in particular products. To appreciate fully the case histories in this chapter of users' experiences with SNA and Decnet it is necessary to have a basic understanding of such products. For those readers not familiar with SNA or Decnet we now provide a brief description of the respective products.

IBM's Systems Network Architecture

Figure 8 illustrates the way in which SNA products can be used to interconnect IBM processors and terminals. The three main protocol layers of SNA are the application layer, the function management layer and the transmission sub-system layer.

At the application layer, the software is processed in the host computers and cluster controllers. The formatting of data for a particular type of terminal (character, page or multi-page device) is handled by the application software in either the host processor, or in the controller (provided that it has sufficient intelligence).

The function management layer corresponds approximately to the session layer of the ISO reference model, and it provides facilities to connect terminals to the appropriate host computers and programs. It initiates distributed functions such as file transfers, and maintains the independence of separate applications even though they share the same resources. In addition, it provides facilities to support distributed functions from the host machines. Function

Figure 8 IBM network using SNA products

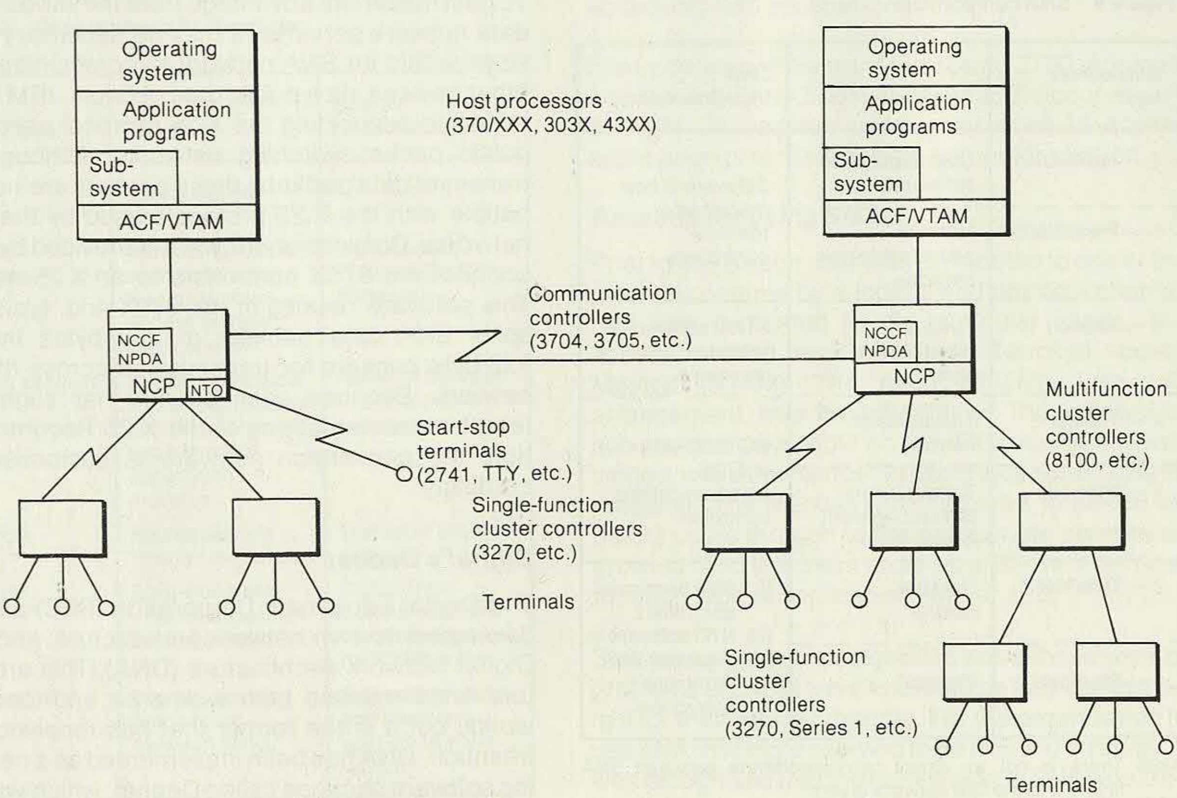

© Reproduction by any method is strictly prohibited

management layer software is implemented in the host processors and forms the major part of the operating system for the network. The product used for a single host network is the Virtual Telecommunication Access Method (VTAM), and the Advanced Communication Function version of VTAM (ACF/VTAM) is used for a multi-host network.

The transmission sub-system layer is responsible for the routeing and movement of data between processors and terminals. It contains three sub-layers — transmission control, path control and data link control. The first two of these are implemented as the Network Control Program (NCP) held in the 370X communication controllers. NCP provides for the physical management of the network by controlling the lines and terminals attached to the 370X. It also includes error recovery procedures. The data link sub-layer is implemented by using devices designed to operate with IBM's Synchronous Data Link Control (SDLC) protocol. Terminal devices that do not support SDLC may be connected to an SNA network by using an additional software product — the Network Terminal Option (NTO) — which resides in the 370X controllers. NTO makes non-SNA terminals, such as teletypes or IBM 2741 terminals, appear to the host program as SNA terminals. (Users may also add their own conversion software for non-SNA terminals not catered for by NTO.)

Figure 9 shows the various components of SNA and how they relate to the different layers in the ISO reference model. SNA products offer users more than a data network. They comprise a complete distributed system. Within each protocol layer of SNA there are several products that support a range of network facilities.

Figure 9 SNA components

OSI network layer	*SNA layer*	*SNA implementation*
7 — Application	User application software	Software in host machines or terminal controllers
6 — Presentation	Network addressable unit services	
5 — Session	Functional management layer	VTAM software in host machines
4 — Transport	Transmission control	NCP software in 3705 communications controller
3 — Network	Transmission path sub-system control layer	
2 — Data link	Data link control	(i) SNA terminals/controllers (ii) NTO software in 3705 non-SNA terminals/controllers
1 — Physical	Physical	

Note: There is not an exact correspondence between SNA layers and the OSI network layers.

1. Monitoring and diagnostic aids

The Network Communication Control Facility (NCCF) supplements ACF/VTAM by automating various operational procedures that are used to physically control the network. A further aid, Network Problem Determination Application (NPDA) can be used in conjunction with NCCF to collect (and then display) statistics (including data about errors) of the performance of network components such as controllers, modems and lines.

2. Co-existence with non-SNA applications

A large existing network that uses the Binary Synchronous Communication (BSC) protocol and asynchronous devices cannot be converted immediately to an SNA network. A migration path has to be provided to allow the conversion to be spread over a period of time during which only part of the network will be using SNA protocols. IBM provides the Partitioned Emulator Program (PEP) to allow non-SNA applications (using say, 270X terminal protocols) to share the 370 X controllers and host machines forming the SNA network.

3. Interfaces to public data networks

To gain maximum advantage from the various public data network services, it may be necessary for devices within an SNA network to communicate with other devices via a public data network. IBM is committed to supporting the X.25 protocol used in the public packet-switching networks. Although SNA transmits data packets, these packets are not compatible with the X.25 protocols used by the public networks. Conversion software is provided by IBM to connect the 370X controllers to an X.25 network. This software resides in the 370X and, typically, it splits SNA data packets of 256 bytes into two 128-byte packets for transmission across the X.25 network. Because each country has slightly different implementations of the X.25 Recommendation, the conversion software is customised accordingly.

Digital's Decnet

The Digital Equipment Corporation (DEC) also has developed its own network architecture, known as Digital Network Architecture (DNA). This architecture encompasses both wide-area and local networks, but it is the former that has received most attention. DNA has been implemented as a networking software package called Decnet, which was first introduced in 1975. Decnet allows any DEC com-

© Reproduction by any method is strictly prohibited

puter to interwork with any other DEC device with files, programs, and other resources being shared as required. The initial phase I version of Decnet allowed only similar machines that used the same operating systems to work together. A phase II version, released in 1978, catered for computers using different operating systems. It also included extra facilities such as improved security and the sharing of resources between machines. A further advance was the release of phase III in 1980. This most recent version of Decnet includes additional facilities to assist in the management and control of the network. Examples include adaptive routeing to allow for the failure of any link, and various error reporting and monitoring routines that produce network operation statistics.

Figure 10 lists the various components of the phase III version of Decnet, and shows how they relate to the different layers in the ISO reference model. Users have a choice of implementing either the X.25 protocol for the lowest three layers in the network, or DEC's own datalink protocols used in the Digital Data Communications Message Protocol (DDCMP) software.

Decnet users can communicate across public X.25 packet-switching networks as well as across their own private networks. There are also facilities for interworking with other networks such as SNA, although Decnet and SNA are not directly compatible.

Figure 10 Decnet phase III structure

Network layer	*Function*	*Decnet implementation*
User	User supplied functions	—
Network application	Network functions, such as remote file transfer and access for user applications	Data Access Protocol (DAP) software
Network services	Establishes logical link for communication between any two network application modules	Network Services Protocol (NSP) software
Transport	Routes packets through the network	Transport protocol software
Data link	Error-free transmission of packets between nodes	Digital Data Communications Message Protocol (DDCMP) software
Physical link	Communications mechanism between nodes	Software modules and hardware specific to each device

EXPERIENCE IN INSTALLING AN SNA-BASED SYSTEM

In this section of the report we relate as a case history the experience of one of the United Kingdom's nationalised industries in installing an SNA-based system. This industry has six computer sites, the two principal sites being about 100 kilometres apart. At each of those principal sites there is an IBM 370/158 (or similar) mainframe computer, providing both interactive and batch processing services to a large number of users throughout the country.

Data communication application

Users can access the principal computer centres in three ways:

- —Local users are connected to their nearest centre via the public telephone network or leased lines.
- —Remote users are connected to one of the centres via multiplexed leased-line circuits which terminate at their local concentrator or computer.
- —Any user may, if necessary, be connected to the other principal centre via 48k bit/s leased-line circuits that link the two centres. There are duplicate circuits, along different routes, to provide a standby service if one circuit should fail.

During normal office hours the data traffic is mainly interactive, generated by 3270-type terminals accessing IMS applications.

File updates and file transfers (using 2780 terminals and the Job Entry Subsystem — JES) occur mainly at night. In addition users may need to access application programs at both computer centres.

Reasons for using SNA

Prior to 1978 each user was connected to one of the principal centres by a logical and physical channel that was dedicated to his particular needs. The access method used was Basic Terminal Access Method (BTAM). The main difficulty with this arrangement was its inflexibility. The number of sub-channels available on the 370 mainframes was limited, while the number of users was continually increasing. The manual work required to switch (or patch) users through to the appropriate centre was growing. And the users required a different terminal for each type of application (IMS or JES).

SNA was a product designed to solve these sorts of problems and IBM were known to be fully committed to it for their future products. For the organisation to use SNA therefore seemed to be a sound strategy. It was recognised that the conversion to SNA would be a major one, particularly as it involved replacing BSC

© Reproduction by any method is strictly prohibited

protocols with SDLC. The change was justified early in 1978 in terms of cost and also in terms of other benefits expected for the operation of the data network.

Implementation of the system

The move from BSC required extensive changes to the terminal access methods, to the communications controllers and to the user terminals. Because of these changes, it would have been impossible to implement SNA for all applications at the same time. The implementation strategy adopted was to use the communications controller software that can emulate a 270X transmission control unit, thereby enabling the communications controller to handle BSC batch and interactive traffic. This strategy made possible a gradual transition to SNA, based on two major stages:

—Establishing a single-host SNA network at one of the principal computer centres.

—Extending the SNA systems to embrace the other principal computer centre to provide a multi-host network.

The main tasks in the first stage were:

—To replace BTAM by VTAM on the host IBM 370/158.

—To introduce NCP and Emulation Program (EP) on the two 3705 communication controllers and to enlarge their memory size.

—To convert all the user terminals from BSC to SDLC operation, concentrating initially on the IMS users.

The timescales for the two stages of the implementation are shown in figure 11. In the second stage, the main tasks were:

—To enlarge the memories of the 3705 controllers at the first centre, and also to provide them with faster CPUs and special line sets to connect them to the 48k bit/s links.

—To install two new 3705 controllers at the second computer centre and a third temporary 3705 to act as a concentrator.

—To replace the VTAM software at the first centre by ACF/VTAM, and to replace the BTAM software at the second centre by ACF/VTAM.

—To replace the NCP software at the first centre by ACF/NCP and to introduce ACF/NCP on the new 3705s at the second centre.

—To introduce a new release of IMS (release 1.1.5).

—To convert all the second centre IMS user terminals from BSC to SDLC operation.

To make the development and testing of the network easier, the 3705s for the second centre were temporarily installed at the first centre. In addition, the circuits on the 48k bit/s inter-site links were reorganised to provide wideband circuits for the conversion task.

Problems encountered in implementing and using SNA

The first stage of the implementation proceeded according to plan, but the second stage took longer than expected. The delays were caused by five main factors.

1. The complexity of the task

The second stage consisted of several tasks that apparently could proceed in parallel, but in reality were interdependent. Delays in any one task affected the other ones.

2. Use of the 48k bit/s circuits

Hardware and software difficulties were encountered in interfacing the 3705s to 48k bit/s lines. At the time (late 1979) there was no prior experience to draw on to solve the problems.

3. The need to maintain a fully operational service

The new network architecture was installed in an environment that attempted to maintain a fully operational service for existing users, and even attempted to increase the number of users. This meant that each step in the implementation had to be abandoned if it were not completed on time, to allow normal service to be resumed at the scheduled times. For example, all major changes had to be completed overnight and the network had to be operating satisfactorily before the prime day shift commenced the following morning. Hence, the evening's work had to be aborted if it was falling behind schedule, and this led to a duplication of effort compared with that which would have been needed if the work had been carried out continuously. Also, the implementation had to be planned to proceed in rather smaller steps than would otherwise have been necessary.

4. The lack of enthusiasm from the users

The main beneficiaries of the SNA conversion were the data processing operations staff. The changes brought about by the introduction of SNA did not directly benefit the end users, and so they were reluctant to abandon their existing service which, from their point of view, was perfectly adequate. The users' terminal had to be upgraded and the users had to be trained to access the required application. But initially the terminals were being used only for a single application as before, and so the users perceived no benefit. Only towards the end of the project did the users start to benefit from the increased facilities. Perhaps the users' enthusiasm for the

The Butler Cox Foundation
© Reproduction by any method is strictly prohibited

Figure 11 Implementation schedule for a large SNA network

Stage 1 *(one centre using SNA)*	Date	*Stage 2* *(both centres using SNA)*
Decision to install SNA at first site	1978 April	
Install and test VTAM and NCP. Upgrade 3705 controllers. Convert and test two 3276 VDU controllers to BSC/SDLC mode	June July August	
First live SDLC user	September	
	December	Decision to link second site and introduce remote 3705 concentrator at second site
Convert all existing IMS users at first site to SDLC. All new users connected using SDLC	1979 March April May	Remote 3705 and ACF/NCP installed and tested at first site. Reorganise multiplex link users to second site
	May June July August	Upgrade first site 3705 controllers
	September October	Remote 3705 moved to second site
	October November December	Test remote 3705 on 48k bit/s link and ACF/NCP at first site
	1980 January	ACF/VTAM implemented at first site
	February	Remote 3705 implemented
	March	ACF/VTAM and NCP implemented at second site
	March	First live user of SDLC at second site
	April May June	Remote users at first site transferred to second site 3705
	May to September	All second site IMS users converted to SDLC
	September	Two sites linked direct via 3705 controllers rather than 3705 remote concentrator
	November	Remote 3705 returned to IBM

change to SNA would have been greater if they had received a less satisfactory service from the operations staff under the old system.

5. IBM software support

This industry was a relatively early European user of SNA and it had an ambitious application. For these reasons, IBM was called upon to provide quite substantial amounts of software support. On one occasion it was necessary for an SNA expert from the United States to be on site for two weeks in order to clear a particular problem. Today the SNA products are more mature, and it is unlikely that such extensive software support would now be necessary.

The costs of installing SNA were clearly considerable. The actual amounts are confidential, but further references to the costs of using SNA are made later in this chapter on page 25.

Benefits derived from SNA

In addition to the expected benefit of flexibility of the network to meet the increasing user needs, this user has found the network monitoring facilities provided by SNA to be extremely useful. The Network Communications Control Facility (NCCF) and, to a lesser extent, the Network Problem Determination Aid (NPDA) have been valuable in three areas in particular:

— Allowing the entire network to be monitored from a single point.

— Isolating communication faults.

— Providing network operation statistics.

This SNA user believes that it would not be possible to operate its large network without a facility like

© Reproduction by any method is strictly prohibited

NCCF. Also, the user 'help-desk' facility provided by IBM has been particularly valuable as a trouble-shooting aid.

EXPERIENCE IN INSTALLING A DECNET-BASED SYSTEM

Our second case history of a proprietary network architecture user is that of a large automobile manufacturer. Although this organisation has extensive IBM mainframe computer installations, it has standardised on the use of DEC minicomputers for controlling manufacturing processes. It was therefore natural to select Decnet as the network architecture for linking those minicomputers together.

Data communication application

The minicomputer network installed in one of the manufacturing plants is shown in figure 12. The computers and terminals are connected to each other via channels on the Videodata cable system installed throughout the building. This system is manufactured by Interactive Systems Inc., which is a subsidiary of 3M. Each device is connected to the Videodata cable via a radio-frequency modem. (Further details of this cable system are contained in the case history in chapter 5, beginning on page 32.)

Reason for using Decnet

As explained above, the choice of Decnet was largely determined by the fact that the data communications need was to link together DEC minicomputers. Decnet offered the network facilities required by the organisation, and its use avoided the need to develop a networking system in-house.

Implementation of Decnet

Decnet was introduced on the site early in 1980, and no significant problem was encountered with its implementation. The system was installed very quickly. For example, a program to link two applications together took only three weeks to write and implement, and further application linkage programs have been based on this initial program.

A small number of software problems arose from the interaction between application programs and the Decnet software, but were all solved fairly quickly. This organisation has considerable experience of using DEC equipment and software and, for this reason, was probably better able to solve the problems than the average Decnet user might be. Nevertheless, DEC's support has been very good when it has been required.

Figure 12 Decnet network based on a Videodata cable system

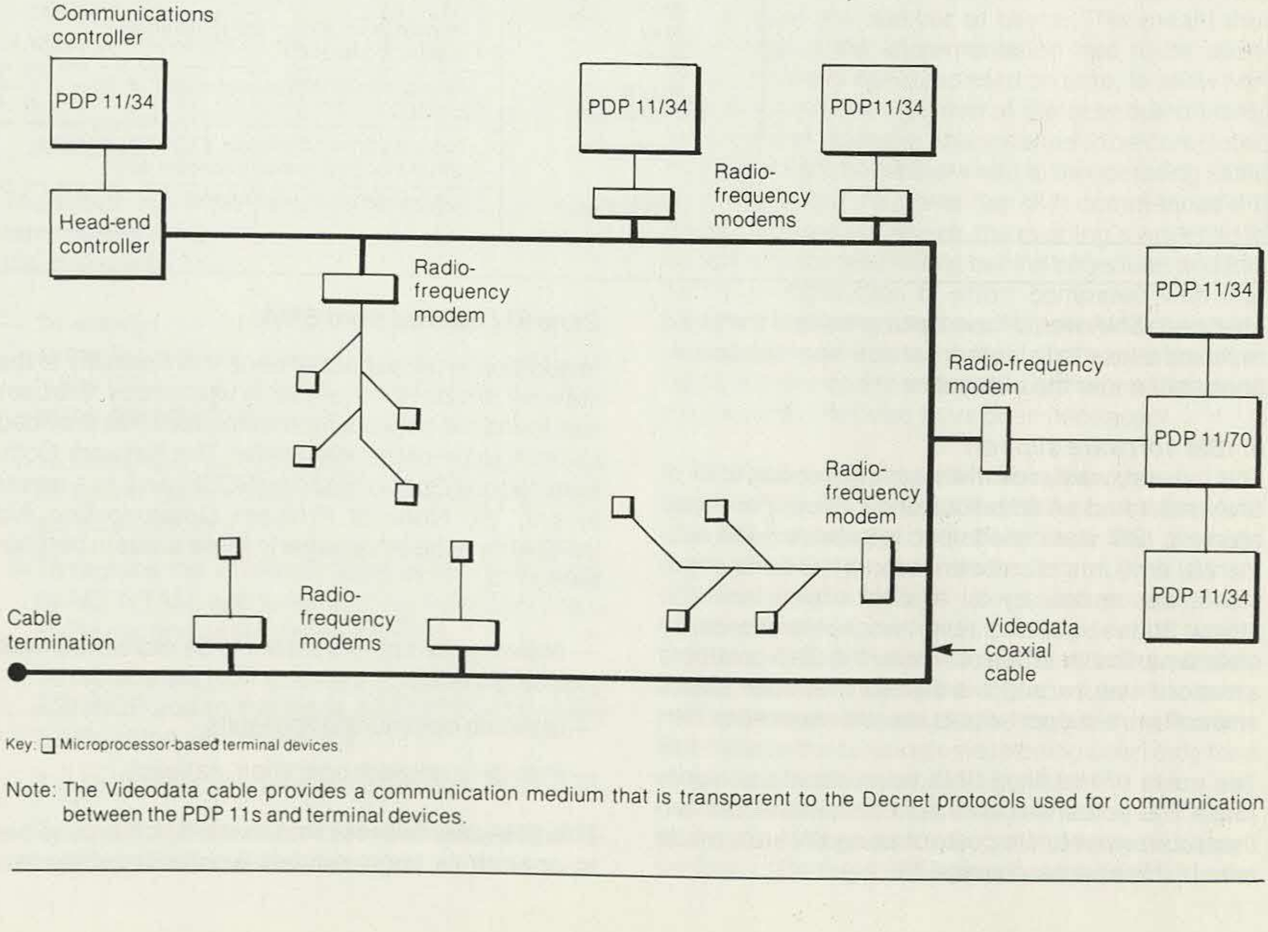

Note: The Videodata cable provides a communication medium that is transparent to the Decnet protocols used for communication between the PDP 11s and terminal devices.

The Butler Cox Foundation
© Reproduction by any method is strictly prohibited

Benefits of using Decnet

This user of Decnet has found the diagnostic aids to be very useful, but says that it is necessary to have a comprehensive understanding of Decnet to be able to locate faults. Messages from the Decnet software tend to be too cryptic. For example, a message that says "LINE DOWN" does not give any indication of why the line is down.

The network management statistics provided by Decnet were found to be good at the byte and packet levels. However, this user has had to write his own software to gather statistics about the higher transmission layers. Statistics relating to the number of times a particular application had accessed another within a given period, for example, were not provided by the Decnet software.

Future plans for using Decnet

The current systems have been implemented using phase II of Decnet, and this organisation is now committed to implementing phase III. The long-term strategy is to move towards a totally open network by successively adopting the higher layers of the ISO open systems interconnection reference model as they become standardised.

Costs of using Decnet

For this particular user, the costs of installing Decnet were less than they would be for other users because the Videodata cable was already installed. The cost of the initial software package for use on one minicomputer was £3,000 ($6,000). This cost included the licence for using Decnet, so that the cost of using the software on additional minicomputers was only £1,200 ($2,400) per machine. The initial package also included all documentation and up to 90 days' support from DEC.

The cost of interfacing each machine to the cable depended on the machine. The average cost was of the order of £600 ($1,200) per machine, including the modem.

OTHER EXPERIENCE OF USING SNA AND DECNET

The other large users of SNA and Decnet that we interviewed also provided details of their experience with these network architectures. They all confirmed the points made in the two case histories related in the previous sections of this chapter.

We now summarise the respective strengths and weaknesses of SNA and Decnet as perceived by all the users we interviewed.

Strengths of SNA

The strengths of SNA as perceived by its users were as follows:

- —SNA is now a proven product that provides powerful facilities for distributed processing.
- —SNA makes the network very flexible by avoiding the need to dedicate terminals or channels to a single application, and by making it easy to add extra terminals.
- —The SNA network management and control facilities are very good. Diagnostic aids at the modem interface level are available when IBM modems that support the NPDA software are used.
- —Direct cost benefits can be achieved by using multi-purpose terminal devices for several applications instead of using separate terminals dedicated to single applications.

Weaknesses of SNA

The weaknesses of SNA as perceived by its users were as follows:

- —SNA depends on the VTAM software in the host computer to route all calls. For example, one user with two offices only a few kilometres apart, each with its own 3705 controller, still has to route all calls between the two offices via the host mainframe situated some 300 kilometres away. In addition to the extra communication traffic involved, this arrangement also requires the host processor to have rather more processing power than it would need if terminals were able to communicate directly. SNA is therefore most suited to terminal-to-central-processor communication rather than terminal-to-terminal communication.
- —The 3705 controllers are expensive and their technology is now rather dated. Plug-compatible manufacturers have found it difficult to develop a 3705-compatible device, and one of the users we spoke to said he would be wary of using a new plug-compatible product, because of experience with earlier products. The best-known plug-compatible products are the 3805 and 3809 controllers marketed by ITT, both of which use their own version of SNA software. IBM is expected to announce a replacement for the 3705 later in 1982.
- —SNA is expensive to use. A 3705 costs between $4,000 and $7,000 per month (depending on its size, the number of lines and the number of channel adaptors). The ACF/VTAM software with the Multi-system Network Facility (MSNF) feature costs about $1,350 per month per machine. The costs of NCP, NCCF and NPDA are respectively $190, $110 and $55 per month per machine.

© Reproduction by any method is strictly prohibited

—SNA is made to appear more complex than it really is. Both the documentation and training courses need to be improved. Nevertheless, the time taken for staff to become familiar with SNA should not be underestimated. It can take a year for operational staff to become confident in its use.

—SNA can connect with an X.25 network only from a host machine.

Strengths of Decnet

The main strengths of Decnet as perceived by its users were as follows:

—Decnet was easy to install, and good installation support was available from DEC.

—Decnet has the flexibility of SNA, but is more suited to inter-processor communication on a peer-to-peer basis than on a master-to-slave basis.

—Decnet's own low-level protocols contained in DDCMP are very efficient for communication over long transmission paths. Up to 256 data packets may be transmitted before an acknowledgement is required, compared with only eight packets for SNA.

—Decnet caters for a wide range of terminal speeds from 110 bit/s up to 1M bit/s (using parallel transmission techniques).

Weaknesses of Decnet

The weaknesses of Decnet as perceived by its users were as follows:

—The statistics from the network management system are not as comprehensive as those from SNA. Also, the messages and prompts are not as informative as they might be.

—Although Decnet offers dynamic route allocation for traffic, it is not able to share traffic between equally attractive routes.

DEC insists that prospective users of the Decnet software complete a 'network profile' document that defines all the parameters of the network to be used. The intention is to assist the user in obtaining the most cost-effective network, and the document forms the basis of the support that DEC will provide. It is, however, difficult for a user to forecast his requirements, even with DEC's help.

EXPERIENCE WITH OTHER WIDE-AREA NETWORK PRODUCTS

Several of the organisations we interviewed as part of the research for this report had installed proprietary network control systems, and one organisation had installed a private packet-switching network. A summary of their experiences with these products is set out below.

Network control systems

The private data network users that we interviewed could be divided into two broad categories:

—Those whose modems and multiplexors were supplied by a single manufacturer.

—Those whose data communication equipment was supplied by more than one manufacturer.

Organisations in the former category were more likely to have installed a proprietary network control system because such a system is, in general, designed to work only with one supplier's modems and multiplexors. (Other suppliers' modems may be connected to the network control system, but the full diagnostic facilities of the system are then not available.)

Those organisations that had installed a network control system claimed that it would be difficult to maintain and operate their networks without it. Such a system enables them to extend a network to other locations without having to employ extra skilled staff at each location to solve the network operating problems that inevitably arise. They can also reconfigure the network more readily should the need arise. Users of network control systems believe that the cost of the systems is more than justified by the facilities provided.

Nevertheless, those organisations that had not installed a network control system were still managing to cope. Some organisations indicated that they would, however, use network control systems if they rationalised their networks in the future.

The network control systems that were being used had some disadvantages. They are more appropriate for star-shaped networks than for the mesh networks for which IBM users might adopt SNA to obtain more comprehensive monitoring and control facilities. Criticisms of network control systems included their poor interface with operations staff. It can take three months for the data network operators to become familiar and confident with the facilities offered by such a system. Also, the information provided about traffic and fault statistics is too voluminous to be used as it is presented. Further analysis of the data is necessary before any real benefit can be gained from it.

Surprisingly, network control systems may still be needed with private packet-switching networks, because existing packet network products are deficient in network control facilities. For example,

The Butler Cox Foundation
© Reproduction by any method is strictly prohibited

the monitoring facilities of a packet network apply only at the data link level rather than at the physical transmission level. Individual modems and multiplexors cannot be controlled by the packet switch itself.

The decision to use a network control system depends also on the organisation's policy for centralising or decentralising the data communications function. We return to that topic in chapter 6.

Private packet-switching networks

Packet-switching networks for use over a wide area are expensive to develop. It is difficult to justify their cost unless there is a very high volume of traffic that can take advantage of the economies offered by the technology. Such networks have been developed principally as public networks for PTTs (or government bodies). But, once the initial development costs have been paid for, packet-switching products can be offered for sale to private users at more reasonable prices. A few large organisations have installed (or plan to install) their own packet-switching networks, and we interviewed one such user.

This user decided to acquire a packet-switching network to replace a very extensive star network of leased lines that connected terminals to a central computer facility. The existing network was inflexible in that it was difficult for terminals to access alternative applications, and the large number of modems created a substantial network management problem. Packet-switching technology offered much more flexibility, and enabled a network to be established that could be used to distribute various kinds of information (including electronic mail). The estimated volume of traffic justified the purchase of a private network rather than the use of the public packet-switching network. The estimated cost saving was approximately 30 per cent compared with the public network tariffs.

A further reason for choosing a private network was that it could handle existing protocols (such as IBM 3270 and 2780), whereas protocol converters would be required to connect to the public service. At the time this user was interviewed, the supplier of the packet-switching network had been selected but the choice had not been made public. The chosen product is based on a very modular structure, with separate microprocessors being used for each protocol conversion.

Our interviewee had visited another user who already had installed a packet-switching network to satisfy himself that the change to packet switching would not introduce additional transmission delays. Also he believes that the packet network will effectively decouple the end-user applications from the communications facility. This organisation recognises, however, that to replace a simple star network by a packet network will require a substantial change in philosophy and procedures within its central communications department. The new network will provide a complete communications service rather than just physical linkages.

© Reproduction by any method is strictly prohibited

CHAPTER 5

LOCAL AREA DATA NETWORKS

The three previous chapters have been concerned with data networks used primarily for inter-site communications. In this chapter, we now consider networks intended primarily for local use within a site.

As with wide-area networks, the potential user of a local network has several options open to him. In addition to a network formed from several point-to-point links where the switching function is provided by software in the processors, the local network can also be based on:

—A data circuit switch.

—The PABX, using either modems on the normal telephone extensions, or the special data facilities that are available on some of the more modern computer-controlled PABXs.

—A dedicated data network that uses a common communication channel to link together a variety of terminal devices and processing units.

Data circuit switches tend to be used in isolation at the hub of star networks, and the use of PABXs for data traffic was considered in Foundation Report No. 26 — Trends in Voice Communication Systems. Hence, in this report, we concentrate on the last of these options, the so-called local area networks.

During the past two years or so, interest in local area networks has grown rapidly. Before we review the available types of local area networks we first examine the reasons why potential users are so interested in them. We then provide three case histories of experience with local area networks. These particular networks use three very different transmission techniques. Finally, we review the experiences of some other local network users.

THE NEED FOR LOCAL AREA NETWORKS

The burgeoning interest in local area networks stems from the proliferation within organisations of intelligent terminals, microprocessors and minicomputers, all needing access to the same data files, and needing to share common resources such as printers.

The demands that are placed on networks to interconnect these devices are becoming more stringent. Such a network may be required to transfer files between processors, for example, or to enable an intelligent terminal to access concurrently several different processors or files. Requirements such as these call for error-free transmission at higher speeds than that required for communication between a dumb terminal and a central mainframe. The network also needs to be more flexible, either to allow additional devices to be added easily or to enable existing devices to be moved quickly from one office to another.

In addition, terminals acquired originally for different kinds of applications (such as information retrieval, word processing and process control) need to be interlinked for some purposes. Such devices are not likely to have been acquired from the same supplier and they will use different communication protocols. Hence, the local network has to provide some form of protocol conversion.

It needs also to be user friendly so that unskilled operators can easily set up the calls. And, if any faults develop, the network should provide facilities to assist the users both to locate the source of the fault and to distinguish between faults in the network itself and those within the attached devices.

The emerging local area networks are seen as a promising method of meeting these demands. In addition, by providing better communication facilities, they provide a big stimulus to the trend towards distributed computing.

LOCAL AREA NETWORK PRODUCTS

Local area networks take many forms, and new variants are introduced to the market at frequent intervals. It is therefore difficult to define a particular network in a way which is comprehensive and which delineates it clearly from other products. But local area networks tend to have the following technical characteristics in common:

—They use packet-switching techniques as opposed to circuit-switching techniques.

—Their communication links are shared by a multiplicity of terminals, rather than each link being dedicated to an individual terminal. (Ring, bus and mesh topologies are used rather than a star formation.)

The Butler Cox Foundation
© Reproduction by any method is strictly prohibited

—Their switching and control functions are distributed throughout the network rather than being concentrated at one point.

—They transmit data at rates in excess of 100k bit/s, and as high as 20M bit/s.

Within these basic characteristics there are many options for building local area networks. Figure 13 shows some of the ways that the different characteristics may be implemented in relation to the different layers of the ISO reference model.

Figure 13 Local area network options

OSI reference model		*Options for local area networks*
Level	*Description*	
—	Topology	Star, daisy chain, ring, bus or mesh
0	Physical	Coaxial cable, twisted pair, multi-core cable or optical fibre
1	Signalling	Baseband or carrier (broadband)
1	Physical link	Dedicated, frequency-division multiplexing or time-division multiplexing
2	Transmission interface	Polling or contention
3	Transmission interface	Packets
4	Transmission interface	Logical records
5/6/7	Services offered	File transfer, electronic mail, terminal support, resource sharing, etc.

These options may in principle be combined in many different ways but, in practice, certain combinations are favoured for technical and economic reasons. There is no completely satisfactory method of classifying local area networks, but some of the more useful criteria are:

—The physical nature of the links between devices (coaxial cable, twisted pairs, multi-core cable or optical fibres).

—The topology of the network (mesh, bus or ring).

—The data transmission method used — baseband or modulated carrier (commonly referred to as broadband).

—The method adopted for sharing the common communication channel — polling, contention resolution techniques such as carrier-sense multiple-access with collision detection (CSMA/CD), token passing or frequency-division multiplexing techniques.

—The nature of the devices to be connected (microcomputers, minicomputers or mainframes, or terminals linked to a host computer).

A further possible classification is whether the network is open or closed. A closed network is one that interconnects a single supplier's products; an open network is one that caters for those of several suppliers. The computer suppliers' network architectures described in the previous chapter are a particular variety of closed networks that can be used within a site as well as over a wide area.

We have chosen to classify local area networks as baseband bus networks, broadband bus networks and ring networks. The three case histories related later in this chapter represent these three main categories. Before we describe the characteristics of each type, we first describe the general features of local area networks.

General features of local area networks

The local area networks used by the organisations that we interviewed all made use of a shared wideband channel (known as the backbone) which was installed throughout the site, rather than a star or mesh network of dedicated lines.

Computers are usually connected to the backbone through a controller and a transceiver, which matches the electrical signal to the backbone. The controllers are currently too expensive to be installed in commercially available terminals and, for this reason, terminals are connected to the network through a multiplexor which is itself connected to the backbone in the same manner as a computer. The multiplexors are conventional devices that support data transmission rates up to 19.2k bit/s. They also allow the terminal user to establish a network call, and sometimes they provide protocol conversion.

The key elements in any local area network are the network controllers. They provide the network intelligence, and can be used to provide interfaces with various processes and programs and also extra communication services such as data encryption. In a microcomputer network (known as a micronet), the network controllers usually are integrated with the microcomputers to minimise costs.

The functions that are necessary to operate a local area network are shared between the network controllers. The four main functions are:

—Network management.

—Call establishment.

—Addressing.

—Transmission.

© Reproduction by any method is strictly prohibited

Network management functions include network configuration, monitoring and control. In small networks these functions are often omitted, but in large networks a great deal of effort may be devoted to them.

Most local area networks do not establish calls in the conventional sense but, in some networks, calls are established and managed by a central switching unit (as in a PABX).

The internal addressing scheme used by a local area network needs to be converted to a user-oriented naming system. This conversion is usually performed by a name server that maps names to network addresses. The mapping scheme can be changed, so that the user does not need to concern himself with changes in network addresses.

The transmission functions ensure that adequate signal levels are maintained on the network. Large networks require amplifiers (for analogue networks) or repeaters (for digital networks), and broadband cables require equipment at the head end to provide frequency conversion. Also, polled networks require a central polling unit. In addition, a gateway function is required when traffic is transmitted from one network to another — even if both networks are of the same kind.

Baseband bus networks

As the name implies, baseband bus networks consist of a common bus (normally a coaxial cable). Baseband signalling (which provides a single communication channel) is used, and devices are connected to the cable by means of suitable interface units.

The devices share the communication channel by using a contention-based access protocol such as CSMA/CD. With such a protocol, each interface unit checks whether the communication channel is free before it transmits data to another unit. If two units try to transmit simultaneously they both recognise the conflict, and they both wait for a short (but random) period of time before attempting to retransmit the data. In this way, it is unlikely that they will then 'collide' again. A further refinement is achieved by allocating different priorities to different devices. The priorities are set by including a fixed delay in the unit in addition to the random time element.

Once a unit has seized the communication channel the full bandwidth is available to it, and data may be transferred at the bus transmission speed. Baseband bus networks are therefore particularly suited for applications that have a high peak-to-average data transfer rate. One drawback of baseband bus networks is that their performance does not degrade gracefully as the sum of the demands from the devices approaches the maximum capacity of the cable.

Baseband bus networks are the cheapest local area network option for any required data transmission rate. They are used therefore both for the cheapest local area networks — the micronets — and for the fastest networks such as Hyperchannel (which is used to interconnect mainframe computers).

Micronets allow several microcomputers to communicate with one another and with shared peripherals such as discs or printers attached to the microcomputers. The cost of a micronet interface unit varies from about $200 to more than $2,000 (about $600 is typical). Most micronets can be used only with one type of microprocessor and usually they do not support terminals that are connected directly to the network.

Ethernet, developed by Xerox, is a well-known example of a baseband bus network which has been configured to interconnect minicomputers. Figure 14 illustrates a simple Ethernet configuration. Ethernet uses coaxial cable and operates at 10M bit/s. Each cable may be up to 550 metres in length and several cables may be linked into a single 'ether'.

The Ethernet standard is being promoted by Intel, DEC and Xerox, and has been adopted by other suppliers including Hewlett-Packard, Nixdorf, ICL and Case. Ungermann-Bass has developed a terminal support network, Net/One, which is compatible with Ethernet, and Net/One has been adopted by Case as the basis for its Casenet product line.

The most widely installed type of baseband bus local area network, however, is the Datapoint ARC Interprocessor Bus (IB). This product was designed specifically to interconnect Datapoint computing and office automation products. The ARC IB is not a true bus because its topology consists of a series of interconnected star networks. Nevertheless, its logical behaviour is that of a bus. ARC IBs can be interconnected by low-speed lines, which may be controlled by Datapoint's Infoswitch PABX.

Broadband bus networks

In broadband bus networks, the coaxial cable bus carries several simultaneous communication channels. The data is transmitted as one or more modulated carriers, and radio-frequency modems are used to interface devices to the bus. Within any one of the communication channels, contention-based protocols may be used in the same way as with baseband bus networks. Such protocols are used to connect terminals to the Sytek Local Net System 20 (at 128k bit/s) and to interconnect computers via the Local Net System 40 (at 2M bit/s). A con-

The Butler Cox Foundation
© Reproduction by any method is strictly prohibited

Figure 14 A typical Ethernet local area network

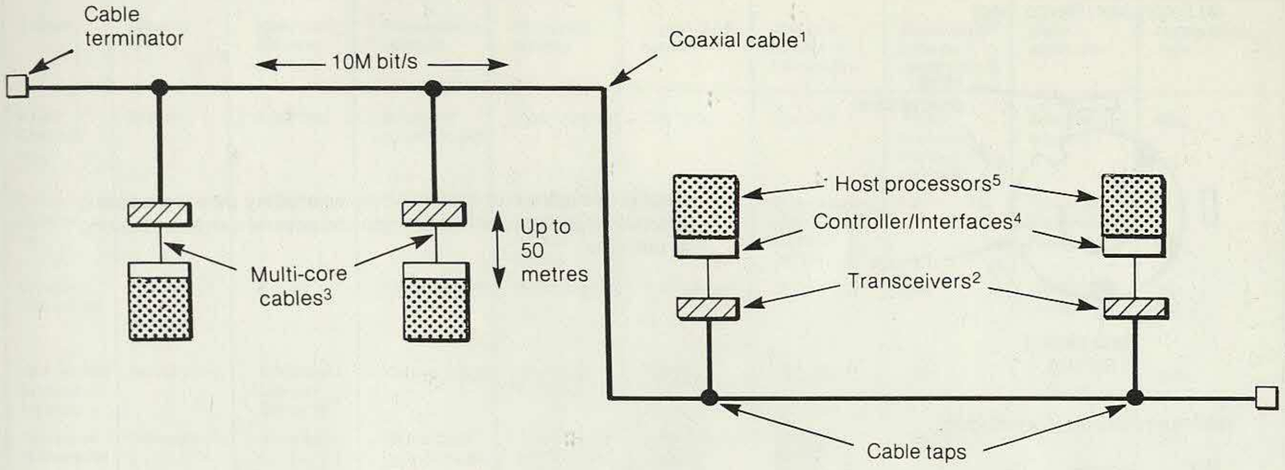

Notes:

1. Each coaxial cable segment can be up to 500 metres long. Segments may be connected by repeaters and gateways up to a total length of 2,500 metres.
2. Transceivers isolate the controllers and host processors from the cable. They comprise line receivers and drivers and packet collision detectors.
3. The screened multi-core cable contains power supply for the transceivers and circuits for transmitted and received signals and collision detection signals.
4. The controllers implement the link level of the ISO reference model.
5. Software that implements the higher levels (3 to 7) of the ISO reference model is required in the host processors. Up to 100 (approximately) host processors may be connected to one Ethernet.

tention protocol is also used in the Wang local area network (Wang Band, which operates at 12M bit/s).

The flexibility of the broadband bus approach allows a variety of dissimilar services to co-exist on a single physical cable. Wang Net is a broadband cable system that uses circuit-switching techniques to allocate one frequency to a pair of attached devices for the duration of a connection. The cable used for Wang Net can be shared by a Wang Band network, video channels, speech channels and so forth.

Broadband local area networks tend to be more expensive than baseband networks for a given data rate. This type of local area network has been used more widely for communication between terminals and host computers than for communication between computers.

Ring networks

Local area networks based on the ring principle transmit data from one station to the next around a continuous bus. Communication is therefore always in one direction (that is, it is simplex).

In most ring networks data is transmitted in 'minipackets' whose size is restricted to about 20 bytes. The small packet size leads to high communication overheads but also to fast transmission times. A wide variety of methods are used to control access to ring networks. The most common methods include empty slot (used with the Cambridge Ring), register insertion (used with the Hasler SILK) and token passing (used with the Apollo Domain). Figure 15 overleaf shows how these three methods differ. In some ring networks, notably in the Cambridge Ring, the minipackets return to the transmitter with a positive or negative acknowledgement that indicates whether or not the packet has been read by the receiving station.

Baseband signalling is generally used with ring networks, but a wider variety of physical cabling is used than for bus networks. (Twisted pairs, coaxial cable and optical fibre cable are all used.)

A ring network can continue to operate only if each station in the ring remains operational. It is never possible to guarantee the reliability of electronic devices and, for this reason, some manufacturers (for example, Hasler and Racal) have provided their rings with alternative routes. In this way, a faulty station or cable segment can be by-passed.

A ring using minipackets is able to offer two types of service to attached devices:

—A byte stream, which is an unstructured stream of minipackets.

—A packet mechanism that imitates those used in buses.

© Reproduction by any method is strictly prohibited

Figure 15 Access control methods for ring local area networks

(a) *Empty slot (Pierce loop)*

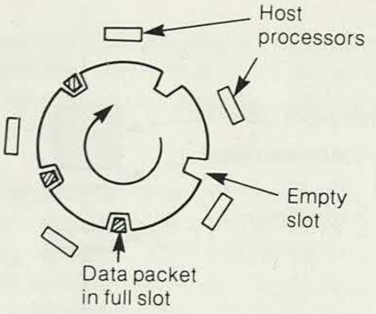

Host processors can transmit only when they detect a passing empty slot. Slots are of fixed length and several can be on the ring at one time.

(b) *Token passing (Newell loop)*

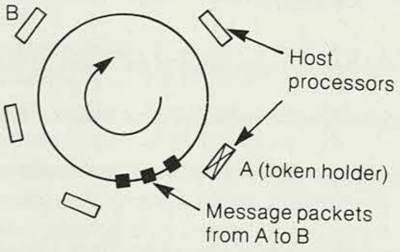

Host processors can transmit a message to (or remove a message from) the ring when they hold the 'token'. (The token is a special bit pattern transmitted round the ring.) The host holds the token until it receives back the packets it transmitted. The receiving host reads the packets intended for it as they circulate past. There is only one message on the ring at any one time, but it can be of variable length.

(c) *Register insertion (Lui loop)*

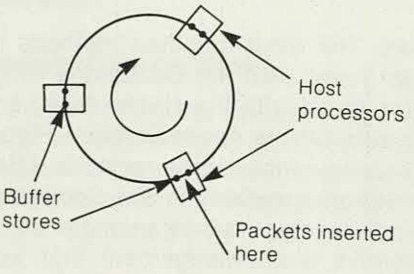

When a host processor wants to transmit a message, it delays incoming messages in the input buffer while it inserts its own message after the end of the previous message received.

The Cambridge Ring is the best known of this type of local area network. This product is now being manufactured by several suppliers and in a variety of forms to suit different kinds of application. The Cambridge Ring operates at 10M bit/s (though higher rates would be easy to achieve) and can be used to interconnect minicomputers and microcomputers.

Other ring networks include Xionics' XiNet, the CM/1 ring for IBM Series 1 computers, and the recently announced Hasler SILK, which offers voice switching as well as data transmission facilities.

Summary of products

We conclude this section on local area network products by presenting in figure 16 a representative list of the products currently on offer. (The figure does not include products based on data circuit switches or PABXs.) For each product the figure shows key characteristics such as the maximum distance spanned, the transmission medium, the accessing scheme and the type of configuration (ring or bus).

EXPERIENCE WITH A BROADBAND NETWORK

The broadband cable system in this case history has been installed in two manufacturing locations by the automobile manufacturer that was the subject of the Decnet case history on pages 24 and 25. The initial cable was installed before the Decnet system. This broadband system is perhaps not a completely typical local area network, because the system provides point-to-point links instead of a flexible network of interconnections. The experience related is therefore more relevant to the use of broadband transmission cable rather than to the full facilities that are offered by true local area networks. Nevertheless, we believe that the experience of this manufacturer is relevant because of the longer

The Butler Cox Foundation
© Reproduction by any method is strictly prohibited

Figure 16 Some local area network products

Supplier	*Network name*	*Maximum[1] distance spanned*	*Transmission medium*	*Accessing scheme*	*Transmission speed*	*Maximum number of connections*	*Approximate[2] price per connection ($)*	*Major application*	*Configuration type*
Apollo Computer Inc.	Domain	3,250 feet	Baseband coaxial cable	Token passing	10M bit/s	Over 100	39,000 (includes Apollo computer)	Scientific and education	Ring
Corvus Systems Inc.	Omninet	4,000 feet	Twisted-pair wire	Carrier-sense multiple access (CSMA)	1M bit/s	64 microcomputers and vendor's peripherals	500 to 1,000	Office automation and education	Bus
Datapoint Corporation	Attached resource computer (ARC)	4 miles	Coaxial cable	Token passing	2.5M bit/s	255	NA	General purpose	Bus
IBM General Systems Division	Series/1 ring	5,000 feet between Series/1s	Coaxial cable	CSMA/CD variation	2M bit/s	16 series/1s	3,825	Data processing	Ring
Interactive Systems/3M	Videodata	40 miles	Broadband coaxial cable	Time-division multiplexing	100k bit/s per channel	248 per channel	600 to 900	General purpose	Bus
Logica Ltd.	Polynet	NA	Baseband coaxial cable	Empty slot	10M bit/s	NA	1,600	Data processing	Ring
Nestar Systems Inc.	Cluster/One	1,000 feet	16-wire cable (flat ribbon, twisted-pair, or round-shielded bundle)	CSMA/CD	240k bit/s	65 Apple micro-computers	395	Small business	Several types
Sytek Inc.	Localnet (System 20 and System 40)	5 to 20 miles	Broadband coaxial cable	CSMA/CD	128k bit/s to 2M bit/s	16,000 to 64,000	500	General purpose	Bus
Ungermann-Bass Inc.	Net/One	4,000 feet	Baseband coaxial cable	CSMA/CD	4 or 10M bit/s	200	500 to 1,000	Office and data processing	Bus
Wang Laboratories Inc.	Wangnet	2 miles	Broadband coaxial cable	CSMA/CD and FDM	9.6k bit/s to 12M bit/s	512 to several thousand, depending on band and equipment	400 to 1,500	General purpose	Bus
Xerox Corporation	Ethernet	1.5 miles	Baseband coaxial cable	CSMA/CD	10M bit/s	100 to 1,000, depending on traffic	1,200	Office	Bus

Notes: 1. Repeaters will be required for the long distances.
2. The price per connection is not directly comparable between the networks because of the differences in the devices connected to the networks.
NA = not available.

(Source: *Data Communications*, December 1981)

period of time over which its broadband cable system has been in use.

Data communication application

Initially, the broadband system was used to link some 200 terminals to a central computer at one of the largest manufacturing plants in the United Kingdom. There was no requirement to switch the terminals to any other computer or to each other. There was a need, however, to be able to re-locate terminals easily and to add further terminals, or possibly other communication traffic (including video and speech) to the system.

Reasons for using broadband cable

The main reasons for using a broadband cable were as follows:

—A conventional star network would have required considerably more cabling in the very large manufacturing plant.

—A star network of dedicated cables would also have been much less flexible when terminals needed to be added or moved, because of the cabling changes required.

—The coaxial cable used was simple to install. It could be attached to the roof supports without the need for ducting or conduits.

—The bandwidth of the cable could cater for additional communications traffic (such as video or speech channels) if the need should arise. It seemed sensible to include spare capacity to avoid the need for further large expenditure on another network for the plant in the future.

The system selected was the Interactive Systems Inc. (ISI) Videodata system. At the time the decision

© Reproduction by any method is strictly prohibited

was made (1977), Videodata was the only product that met the requirements of this company. The company had seen a Videodata system installed in one of the Ford Motor Company's plants in the United States and so they were not the first pioneers in its use. The commercial arrangement, however, was that the company should purchase its system from Ferranti in the United Kingdom. Ferranti were to be the United Kingdom distributor of the Videodata system for industrial applications.

Implementation of the Videodata system

The initial system, installed by the end of 1979, required about six kilometres of cable. The installation was straightforward, as have been extensions to the system and re-routeing of the cable. A second smaller system (four kilometres of cable) was installed at a second factory at the beginning of 1980, where it provides the physical link for the Decnet system described in the previous case history.

Problems encountered in using the Videodata cable

Several technical problems have been encountered with the use of the Videodata cable, and the main ones are listed below:

—When a branching unit is introduced into the cable it needs to be terminated correctly, in order to avoid interference which would affect the total performance of the cable. Similar problems can occur if a modem is not terminated correctly.

—Frequency drift in the radio-frequency modems causes interference on adjacent channels.

—The system was supplied without any diagnostic aid, making it very difficult to locate faults. One problem concerned a faulty modem that was transmitting permanently. It proved particularly difficult to locate the modem that was the source of the problem. The company has therefore developed its own equipment to locate faults and to test new equipment. (Ferranti have subsequently developed similar equipment.) Nonetheless, it would still be desirable to have diagnostic aids built into the cable modems. For example, the staff within the manufacturing division of the company attempt to service the modems themselves, but they have to call on the central computer service department to locate faults.

—There is no easy way to switch to a standby system if a fault develops in the cable head-end. As a result, all processing associated with the head-end must be interrupted until the change-over is completed.

—In practice, the modems must not be more than 30 to 60 metres away from the cable.

—Because of the layout of the cable network, all terminals connected to a branch of the cable beyond an amplifier are put out of action if the amplifier develops a fault.

Experience has shown that faults are more likely to occur after the system has been disturbed in any way. For this reason, the system is left switched on for as long as possible. The company has also experienced faults just after the cable has been extended, although this problem could be due to the supplier sub-contracting the work to people who are less familiar with the system.

Benefits from using Videodata

Despite the technical problems, the Videodata system has proved to be a very flexible network. Re-routeing and extending the cable have proved to be easy tasks, provided that the modified installation is properly tested before the system is put back into use. Such modifications usually have been completed within a weekend — which would not have been possible, for example, with a network based on twisted-pair cables. Also, and despite the lack of diagnostic aids within the system, the company believes that faults have been easier to locate and repair than with a conventional twisted-pair network.

Twisted-pair cabling would have cost approximately twice as much to install as the Videodata cable, and it would have cost more to re-locate such cables.

Costs of using Videodata

When the first system was installed in 1979 the cabling, including amplifiers but excluding modems, cost about $100,000 to install (that is, about $17,000 per kilometre of cable). The modems cost about $1,000 each, giving a total of about $300,000 for the complete system.

The second system, installed in 1980, cost about $45,000 for the cabling and amplifiers (that is, about $12,000 per kilometre of cable). The total cost including the modems was about $85,000.

Ferranti charges a fixed fee of £3,500 ($7,000) per annum for maintaining the system plus £200 ($400) per repair and £250 ($500) per man-day of support.

EXPERIENCE WITH A BASEBAND BUS LOCAL AREA NETWORK

The local area network that is the subject of this case history has been installed within the central Computing and Technical Services Department of a large public utility in the United Kingdom. This department acts as a consultant to various regional computing centres, and is involved in formulating

The Butler Cox Foundation
© Reproduction by any method is strictly prohibited

the plans for the computing facilities within the utility. It is concerned also with all the various kinds of computer application, from customer billing to process control and specialised scientific computing.

Data communication application

Users both at the main computer centre (where the department is situated) and at the regional computer centres of the utility need to use their terminals to access several different computers. Also, the utility's laboratories need to retrieve data files for local access by a variety of terminals and microprocessors. The regional centres are considering the potential of local area networks to act as the communications medium for electronic mail between communicating word processors. The local area network installed at the central department is being used to gain experience of such a network before any decision is made about which networks would be most suitable for the regional offices and laboratories.

Reasons for choosing a baseband bus network

The local network had to provide a switching function, because any user might want to connect his terminal to any other device on the network. The central department was aware of the present uncertainty about standards for local area networks. It knew that this uncertainty could mean that the chosen network would have to be modified when standards were agreed. Thus the cost of the chosen network had to be low enough to enable it to be written off if the network could not be modified to meet the agreed standards. The requirement for a low-cost network that could provide a switching function led the utility to choose a baseband network rather than a broadband network.

A bus topology was chosen because the central department was concerned about the reliability of a ring-based network. The product selected was the Ungermann-Bass Net/One, which was available from a local distributor, Thame Systems. Net/One is an Ethernet-type network operating at 4M bit/s. The central department was assured by Thame Systems that Net/One could be upgraded to operate at 10M bit/s (thereby making it compatible with Ethernet) for a cost of £250 per attached device.

Installation of the Net/One network

Most of the terminals were installed in offices that previously had not contained any type of computer equipment. Many of the offices either did not have a suitable power supply outlet or the outlet was not located in a convenient place for the terminal. The central department therefore decided to install a new power circuit to service all the offices. Also for each office, a cable spur had to be installed between the Net/One network interface unit (NIU) on the bus and the terminal. The main power cable and the network bus cable were both installed in a cable tray attached to the ceiling of the corridor along which the offices were situated. Holes had to be drilled through the wall of each office for the cable spurs from the main cable tray.

Problems encountered in using Net/One

Apart from the physical problems of installing the cable and power supplies referred to above, four types of problems were encountered with the Net/One hardware and software.

1. Net/One hardware problems
The hardware delivered by Thame Systems was at first very unreliable. The equipment failed at frequent intervals and, apparently, there was only one engineer who could diagnose the faults quickly. Eventually, Thame was persuaded to conduct an exhaustive investigation of the faults. As a result, all of the problems were solved within about a week and the equipment has since been reliable.

2. System problems
The utility's staff had no previous experience of Net/One and, because of this lack of experience, they often were unsure whether system faults were due to Net/One or to their own systems. The staff were therefore reluctant to call for support from Thame Systems if they believed the problems might be caused by their own equipment. Also they found it difficult to configure the network NIUs to match each type of terminal, because a complete understanding of the terminal-to-computer interface was required before this could be done. One particular type of terminal could not be attached to the network because the staff did not have a terminal specification.

3. Operational problems
An operational problem can occur because of the Net/One log-on procedure. A terminal user first has to log-on to the network and then log-on to the computer he wishes to access. It is then very easy for the user to log-off from the network without logging-off from the computer. This is not considered to be a serious problem in the development environment of the central department currently using the network, but it would be undesirable in normal commercial or administrative user-application areas. The network software is to be modified to remove this problem.

Each time the system is switched on the NIU programs have to be downloaded from the network operation station via the bus. This operation takes a relatively long time, and so the system is not switched off unless it cannot be avoided.

© Reproduction by any method is strictly prohibited

4. Lost data

A design fault was found in the NIUs whereby data could be lost when a virtual circuit was being used.

In addition to the above hardware and software problems, two serious problems have arisen concerning the original expectations of the product:

—The cost of upgrading the system is considerably more than the utility had been led to expect that it would be. For example, within a period of three months the cost of increasing the transmission speed to 10M bit/s had escalated from the price mentioned originally (£250 per station) to about £1,750 per station.

—The utility has not been provided with the software development tools for the network that it had been led to believe would be available. This lack of development tools means that it is more difficult for the central department to effect some desirable software modifications. For example, Net/One is provided with software that makes it suitable only for asynchronous terminals, but the central department would like to incorporate previously developed software for computer-to-computer communication. Also, the department foresaw that some of the regional centres might ask it to develop network software for the different kinds of communications traffic that they would have.

Benefits from the use of Net/One

The problems listed above might give an unduly pessimistic view of the utility's experience with Net/One. Nonetheless, once the initial installation problems had been solved, the network has been easy to use and it has been reliable. The utility is aware that it suffered from the inevitable problems of being a pioneer user of the product. Thame Systems, as a distributor of a new product, also suffered from an initial lack of experience of the product, but overall Thame has provided good service and training facilities.

Future plans

The utility foresees that it may well use different types of local area network for different types of traffic, possibly with gateways linking the different networks. In the future, it may use Net/One only for computer-to-terminal data traffic.

The utility hopes that its Net/One installation will have a life of about four or five years, but it is concerned because Thame Systems has upgraded the product it now offers to operate at 10M bit/s (and therefore to be compatible with Ethernet). The utility either has to pay for its own system to be upgraded at considerable expense (unless it can negotiate a better price than Thame is currently quoting) or it has to rely on Thame Systems supporting a product that has now been superseded.

EXPERIENCE WITH A BASEBAND RING LOCAL AREA NETWORK

A major chemical company that uses many DEC PDP 11 computers has been experimenting during 1981 with a local area network but, unlike the public utility, has chosen a ring system in preference to a bus-based network. This case history recounts the reasons for this choice and the experience with the ring.

Data communications application

Like the public utility, this company wanted to gain experience of the capabilities and problems of using a local area network to interconnect several processors and terminals. The aim was to share resources between the terminals and computers and allow complete freedom of access between them. The majority of the computers to be linked to the network were DEC PDP 11s. In addition, several Computer Automation computers had to be linked to the network, together with a variety of terminal types.

Reason for choosing a ring network

The company chose to experiment with a ring network of the type developed at Cambridge University. There were several reasons for choosing a ring network in preference to a bus:

—The simplicity of a ring seemed more attractive than the relative complexity of an Ethernet-type of bus.

—The ring network uses twisted-pair cables, which for this particular company were easier to install than coaxial cable.

—Suitable hardware was available at the time the decision was made. Also, the company expected the price of the ring interface units to fall because a supplier (Ferranti) was planning to produce LSI chips for the units.

—A few similar networks had already been installed and were in operation. In addition, if any problems arose, the designers of the ring network were readily accessible at Cambridge University. Furthermore, the company would have access to any further facilities for the network as the University researchers developed them.

Implementation of the network

At the time the decision to experiment with a ring network was made in 1980, commercial products were not available. The company therefore commissioned a small software company (Top Express) in

The Butler Cox Foundation
© Reproduction by any method is strictly prohibited

Cambridge to supply the software for the network. (Top Express employed some of the University researchers who had helped to design the ring network.) The hardware was specified by the company in conjunction with Top Express and was supplied by another small Cambridge-based company, Orbis. The equipment was installed in two phases early in 1981, with the installation (including cabling) being completed in a total of only nine hours. The software also was installed in two phases, with the low-level protocols being installed first. The installation of the software also went very smoothly, and was completed in a total of eight hours.

The initial network included:

—A monitor station and error logger.

—A PDP 11/34 interface unit.

—A Computer Automation interface unit.

Problems in using the ring network

There were only two hardware faults in the first six months of operation. One of these faults did not interrupt the operation of the ring, but the other one was a failure of a ring transformer. (This fault was detected and reported by the next station on the ring.) There were also several software errors, and two of those errors were still outstanding at the time we conducted the interview. The software problems that had been solved had not proved difficult to diagnose, and the support provided by Top Express was very good (apart from a period when some of the staff were on leave during the University's summer vacation).

Benefits from using the ring network

The company has found the network very easy to use and simple to extend. Extensions are made by breaking the ring, installing a new station and updating the name server. Nevertheless, the company has found that extensions require rather more work than it had anticipated.

The additional facilities that the company anticipated are being made available. There is now a gateway facility that allows the ring to be connected to an X.25 network, for example, together with gateways to link two or more rings together.

The network is providing the experimental experience that was the main objective of the project.

Future plans

The experimental use of the network had not been completed at the time we carried out the interview. If the experiment continues to be successful, the company will consider purchasing a ready-made proprietary system rather than involving itself in further development work.

Costs of using the network

The total cost of the network was £17,000 ($34,000), of which £13,000 ($26,000) was for the hardware and £4,000 ($8,000) was for the software. The hardware would cost about half of the original price if it were to be purchased now. The software was comparatively inexpensive because of the experimental nature of the product.

OTHER EXPERIENCE WITH LOCAL AREA NETWORKS

Two other users of local area networks also provided us with information about their experience with them. One user had installed a Net/One network and the other a Cambridge Ring.

The Net/One user had, in co-operation with a software house, extended the facilities of the network by adding the HDLC and X.25 protocols to the two lowest levels of protocol. This network provides access to a Comet electronic mail system and to a shared printer. Once the cable was in place, the network was easy and quick to install, and modifications to the network also have been easy to carry out. The hardware has proved to be very reliable, but the lack of network management facilities is seen as a definite disadvantage. The only major problem to arise had been caused by the wrong type of cabling being installed, but this was an administrative mistake rather than a technical issue.

The Net/One cannot cater for the forecast volume of traffic because the style of communication (byte echoplex) used by the particular application does not make efficient use of the network capacity. A change to dual networks using a single broadband cable is a possible solution to this problem, but the user has not yet decided on an organisation-wide strategy for local area networks. Within the same organisation, for example, there is also a Nestar micronet that serves 80 Apple computers.

The Cambridge Ring is used by an educational establishment which chose the ring in preference to an Ethernet-type of network. The network was commissioned in mid-1978, and now has 96 terminals connected to it. It allows the end users to run batch and interactive programs either on an ICL 2900 mainframe or on DEC Vax and PDP 11 computers, and it provides gateways to other remote computers. Microcomputers connected to the ring can retrieve programs and files from the main processors for local processing.

Both of the above users have had little difficulty in installing their local networks. However, both networks were experimental in nature and the users were willing to contribute their own technical resources to the task of setting them up.

© Reproduction by any method is strictly prohibited

CHAPTER 6

DATA NETWORKS — THE KEY ISSUES

In the previous four chapters, we reported the results of our research into the experience of users with different types of data network. In this chapter we identify the most important data network issues from the user's viewpoint. These key issues are discussed under three headings:

—Planning and choosing data networks.

—Operating and controlling data networks.

—Managing the data communication services used by an organisation.

PLANNING AND CHOOSING DATA NETWORKS

Private data networks have tended to grow in a haphazard way. As new computer applications have been implemented, additional links or services have been added to the networks. The cost of these additions has been justified by the benefits from the enhanced system. Such an ad hoc response to user demands clearly has disadvantages, although it does not require any overhead expenditure. The primary issue now facing those responsible for data communications within an organisation is whether the organisation should have a long-term plan or strategy for the evolution of its data networks. We believe that a long-term plan is necessary but, before that plan can be devised, certain other issues first need to be resolved:

—The organisation must decide how far ahead it can forecast the demand for data communications, and the plan must then allow for the inevitable uncertainty in such forecasts.

—The organisation must decide if it requires a completely open network that enables any type of equipment from any supplier to be connected to it. Alternatively, it might decide to restrict its choice of equipment to that of one supplier, or to that which conforms to certain standards.

—The organisation must decide if it will construct a single corporate data network or separate networks for different applications. Additionally, it needs to decide if the data network(s) should integrate with other corporate networks (such as those for voice and text).

—The type (or types) of data network to be installed must be determined, and the use that will be made of public data services must also be taken into account.

—The network supplier policy must be determined. In particular, the organisation must decide whether the network(s) will be constructed from equipment provided by a single supplier.

We now consider in turn the primary issue and the other issues listed above.

Long-term plans for data communications

The lack of a co-ordinated plan for data communications within an organisation can lead to several problems, including:

—Inefficient use of transmission facilities, because the possible advantages of sharing them between applications have not been obtained.

—Incompatibility both of network and terminal facilities. Compatibility can be provided only by replacing existing equipment.

—Inability to respond quickly to requests either for communication services for new applications, or for changes to existing services. This is particularly true when new PTT-provided facilities are needed. The installation lead times for such facilities often are several months (and sometimes longer than a year).

All of the organisations we interviewed agreed that there should be a strategy, or at least a long-term plan, for data networks. Furthermore, the data network plan should not be subservient to the other business plans and strategies, but should contribute to them. This is particularly true for those businesses where the mode of communication with customers can be changed by the use of on-line computer systems made feasible by new public data services or by cheaper data networks and terminal equipment.

Some businesses (such as airlines) have an obvious need for on-line systems. But, even in those organisations where the need is not so obvious, there is now an increasing awareness that data networks are a vital part of their operating infrastructure.

© Reproduction by any method is strictly prohibited

Forecasting data communications requirements

The need for a long-term plan implies also the need to forecast the likely data communication requirements. This is no easy task, because data networks create what is sometimes called the 'motorway' effect. A new data network is likely to unleash a latent demand from users that is not apparent until the network is in use. Moreover, on-line computer applications often are not a simple replacement for an existing manual or batch system, and they may condition the user's behaviour to such an extent that a totally unexpected pattern of usage emerges.

Because of these uncertainties, the networks should be provided on the assumption that the forecasts may be significantly lower than the actual usage. It is still worthwhile attempting to produce forecasts that are as accurate as possible, but the further into the future the plan extends the less accurate the traffic forecasts will be.

The main requirement when producing a long-term forecast is to identify the nature of the communications traffic and its geographical distribution, rather than to predict exact traffic volumes in terms of bits per day. Thus, the organisation needs to determine, for example, when and where office automation projects are likely to generate electronic mail traffic, and when professional workstations are likely to require access to corporate data files.

This inherent difficulty in forecasting the demand for data communication facilities emphasises the need to monitor the actual use that is made of the network. The sooner the actual traffic can be compared with the predictions, the sooner will any potential inadequacy in the network be revealed.

The user organisations we interviewed were also in agreement on the need for detailed short-term implementation plans. Networks are now more sophisticated than they used to be, and their implementation warrants the same attention to detail as that given to the implementation of the computer applications themselves. These plans are particularly important when an existing network is being enhanced. In such a situation, the plans need to consider the impact of the changes on existing network users so that continuity of service can be maintained. They need also to make allowance for the human factors associated with the change. Existing end users may need to be trained and persuaded of the advantages of the new network, and operations staff may need time to become familiar with the new monitoring and control facilities.

Open or closed networks

The advantages of having a completely open network are self-evident but, in the absence of well-defined and agreed communication interface standards, the network itself must provide the means to interconnect the various communication protocols. Protocol conversion will become an increasingly expensive component of the network as the variety of terminal devices increases.

Current communication standards define only the lower layers of the ISO reference model. The CCITT's X-series of recommendations (the best known of which is X.25) has been a most useful and powerful influence on network layers up to the data transport layer. But the absence of universally agreed standards above that layer has forced each supplier to adopt its own specifications. (Suppliers also have commercial motives for defining their own protocols.)

Although several of the users we interviewed intend to pursue the goal of a completely open network, they recognise that it is likely to take ten (or more) years before agreement is reached on the higher-level protocols of the ISO reference model. Their strategy is to base their current networks on the X.25 Recommendation, and then to adopt the standard for each new layer as it is defined. Provided that these organisations monitor continuously the standard-making activities of the CCITT and ISO, they should not be faced with the need to make sudden and possibly catastrophic changes to their networks. Also, they must consider carefully (again in the light of the current standards debate) whether to acquire any new equipment that does not conform to their existing protocols. The tendency inevitably will be to limit the range of devices that may be connected to the network so as to minimise both the problems of incompatibility and the costs of protocol conversion.

The lack of an agreed standard for local area networks is another problem for potential users. Regrettably, the latest reports from the IEEE sub-committee that is attempting to define a standard suggest that agreement on a unique standard is unlikely, and that there will be several options within the eventual standard.

Integrated or separate data networks

The ultimate goal of an organisation's data communications strategy may be a single all-purpose network. But in the short term the diversity of data traffic is likely to be handled better by several different networks. For example, inter-site mainframe computer applications may be served best by a private leased-line network based on the computer supplier's proprietary network architecture, whereas the intra-site traffic generated by various office systems may have its own local area network. But, to operate efficiently, each type of network requires

© Reproduction by any method is strictly prohibited

the users, and particularly the operators, to have a thorough understanding of it. For this reason, the number of different networks should be kept as small as possible. Full integration of the diverse types of data traffic will be achieved over a period of time by providing gateways to interface the separate networks with each other.

At the physical transmission level the different networks will, however, share the same physical trunk connections so as to minimise line rental costs. Indeed, at the lowest layers of the ISO reference model data traffic may share transmission lines with voice traffic. In the future, there will be scope also for integrating data and text (and possibly facsimile as well) at the higher network levels.

The potential for integrating different modes of communications clearly has implications for the way in which data communications should be managed, and also for its relationship with other communications functions within the organisation.

Types of data network

The user experience that we have related in this report does not suggest that there is any one type of network (either for inter-site or intra-site communication) that is preferable to all others. Although initial problems had been encountered with many of the networks, they had all eventually proved to have no serious long-term disadvantage in terms of performance. From the user experience we conclude that the criteria that should be used when an organisation chooses a network type are the nature of the applications it is to serve, and the availability of suitable networks for those applications. Other factors to consider are the cost of the network relative to other offerings and the constraints imposed by the existing computers that are to be connected to the network.

For some particularly vital communication applications it is advisable to base the network plans on two distinct types of network, so that one type can provide a fall-back should the other one fail. Some organisations may, for example, choose to use a public data network as the fall-back for a private dedicated leased-line network.

The user experience that we researched showed that for public data networks there was no reason for preferring packet switching to circuit switching. The main reasons for choosing one or the other were availability and geographic coverage (in particular, the ability to connect to foreign networks).

An organisation that chooses to use a public network should be aware that it will have less control than it has over its own private network. Some public network users have experienced problems, particularly when the initial demand for a new service has exceeded the PTT's expectations. The PTTs are discovering that data services require a higher level of support than traditional telephone services, and also a higher level of supplier involvement. Before they can satisfy completely the demands of their public data network customers, the PTTs have first to acquire knowledge of and expertise in running those networks.

Network supplier policy

No standards apply to the specification of network control systems and, to gain the most benefit from the network monitoring and control aids that are available, organisations have to purchase all of their data transmission equipment from a single supplier. Many organisations are uneasy about dependence on a single supplier and, in any event, there may be no one supplier that can meet all the organisation's needs. One organisation we interviewed restricted the range of modems and multiplexors used to a single supplier, but the data switch it needed had to be purchased from another source because the modem supplier did not have such a device in its product line. This organisation was keen to restrict the range of equipment to a single supplier for another reason. If equipment from different suppliers is used then additional time may be lost in tracing faults.

Data transmission equipment is now very reliable, and some users are questioning the need for maintenance contracts. Other users, however, had found it necessary to apply considerable pressure to some of the newer suppliers in the market in order to obtain an adequate maintenance service.

OPERATING AND CONTROLLING DATA NETWORKS

We now turn from the planning of data networks to their day-to-day operation. The primary issue for operating and controlling a data network is whether to have a central control group for the whole network, or separate network operations groups for each site. Other key issues that users have found to be important are:

— Traffic management and access control.

— Change control.

— The distinction between network development and routine network operation.

— Education and training of network operation staff.

— Security and privacy.

As in the previous section, we now discuss each of these issues in turn.

The Butler Cox Foundation
© Reproduction by any method is strictly prohibited

Centralisation of network operations

Most organisations would not question the need to centralise and co-ordinate the planning of their data networks. There is, however, not such a clear-cut argument for centralising the day-to-day operation and control of the network. Nevertheless, many organisations have adopted a centralised approach, because of the relative shortage of people possessing the high levels of technical skill associated with the more sophisticated network products. In a centralised approach, however, the necessary technical aids must be available to the central staff to enable them to monitor and control the network from one location. The specialised network control systems described earlier in chapter 4 and some of the computer suppliers' network architectures, such as SNA, do offer this type of facility. SNA is particularly flexible in that the network control and diagnostic software can be accessed by a terminal connected to any node of the network.

Another advantage of centralised network operations is that all contact with suppliers (whether of private equipment or public services) can be channelled through a single department.

But centralised control is against the general trend towards decentralisation, both of computer processing and of company organisation as a whole. An inherent danger with centralised control is that the network will place restrictions on the users' activities. To be truly useful a network must meet the needs of all of its users.

Traffic management and access control

We referred earlier to the difficulties of forecasting data traffic. This difficulty is often due to the motorway effect as unexpected users (and uses) emerge once the network is operational. To prevent the network becoming overloaded, the data traffic should be monitored on a routine basis, and the actual traffic should be compared with that forecast. Also, access to the network should be controlled from the start of operation. (The needs to monitor traffic and to control access are, of course, additional arguments for centralising the operation of the network.)

If the network operations department does not attempt to measure and control the usage of the network, it is likely to be faced with apparently ad hoc requests for extra capacity at short notice. It will only be able to react to demands for its service, rather than anticipate them.

Change control

For reasons similar to those advocated above for measuring and controlling the usage of the network, we believe that organisations also should adopt a standard procedure for changing the network or the way in which devices are connected to the network. One large user we interviewed is adopting a procedure whereby all new requests for data communication services are submitted as logical requirements. The central network service function then determines the most economic and practical way of meeting the logical needs.

A formal procedure for requesting network changes should avoid the problems caused by several separate (and substantial) system changes being planned to occur simultaneously.

Development and routine operation of networks

In the past, many organisations found that their data processing systems appeared to be in a continual state of development. This situation was caused by a failure to delineate clearly between the development of a new system and the routine operation that should follow its implementation. A similar situation now exists in some organisations regarding the development of new networks. If there is no clear boundary between the development (and implementation) of a new network and its routine operation, then the causes of operational faults (and the responsibilities for curing them) will not be clear.

This does not mean that network development and operations activities should be divorced completely from each other, any more than computer system development should be isolated from computer operations. There should still be a dialogue between the two. If a network is to be run efficiently, it should be designed with the needs of the operators in mind. One organisation we interviewed gives its network operations staff the power to veto any proposed network changes which they believe will result in their not being able to operate the network efficiently.

Education and training

There are an increasing number of elements in any data communication path between a terminal and a processor. This implies that those involved in the design and debugging of networks (network engineers and computer systems staff) require a comprehensive knowledge of the complex communication software used to control the path. One user we interviewed arranges for network operations staff to be involved in the implementation of all new computer applications that require communication facilities. The aim is to provide the network staff with a better understanding of the applications so they will be in a stronger position to solve any communication problems that may occur with them.

Security and privacy

Most of the users we interviewed had considered ways in which they could prevent their data net-

© Reproduction by any method is strictly prohibited

works being used to gain unauthorised access to their computer systems. None of the experience quoted suggested that any particular methods were better or more necessary than others. Most organisations did not use data encryption techniques; instead, they were content to rely on user-identity numbers and passwords. The decision whether to use data encryption seemed to be a function of the perceived degree of security associated with the application rather than of any evidence that identity numbers and passwords were inadequate. For example, applications involving financial transactions are more likely to use encrypted data transmission than are manufacturing applications.

MANAGING DATA COMMUNICATION SERVICES

Probably the most important issue associated with managing the data communications facility within an organisation is that of how to justify and pay for it. In this respect, the data communications facility has the same problems as any other company-wide service, such as the internal telephone network or the data processing department. The facility is there to provide user departments with whatever they want, but it must operate at a realistic cost. We believe that, at the least, a notional charge must be made for any service provided, even if no funds actually change hands.

Billing user departments for data communication services is an attractive method of controlling the services. It encourages user departments to introduce only those data communication applications that are justified, while at the same time encouraging the data communications department to provide the network as economically as possible. But a long-term view needs to be taken by the data communications department when it sets the tariffs for internal services. To establish a general facility will almost certainly involve greater initial expenditure than that required to provide only the particular services that are needed. If the user departments are asked to bear the full cost of a general service, then they may choose to use the public networks, or they may install their own data networks because they believe that the centrally provided service is too expensive. Moves such as these will exacerbate the costing problem of the central service because they will reduce the potential for economies of scale.

The data communications department therefore needs to market its services (except perhaps in the most autocratic of establishments). It should convince the potential users that it can provide a better service than any alternative network, and it should respond to their needs rather than try to control them. The data communications department needs also to educate senior management and users about the long-term benefits of a professionally run internal central service.

The department does, however, need to work in co-operation both with the computer services and end-user departments to agree an internal set of standards for the equipment that is permitted to connect to the network. In addition, procedures need to be agreed for:

- —Extending the network.
- —Adding new users to the network.
- —Operating the equipment connected to the network.
- —Reporting faults, and bringing more serious and persistent faults to the attention of management and suppliers.
- —Adopting new software in devices connected to the network.

A broader issue is whether it is sensible to divorce entirely the management of a data communications network from the management of other telecommunication services. Several user organisations mentioned that the biggest problem facing an organisation wishing to integrate speech, data, message and other communication networks may not be a technical problem at all. The problem is more likely to be concerned with the traditional organisation of the company.

Historically, telecommunications facilities often have been the responsibility of an administrative, secretarial or site services department, whereas data communication facilities have tended to be the responsibility of the data processing function. The existing telecommunications staff have perhaps provided only an interface between the data communications specialist and the PTT for ordering the PTT-supplied lines and equipment.

The users we interviewed did not express strong views about the urgency of resolving the organisational conflicts. Nevertheless, we believe that organisations should begin as soon as possible the process of bringing the separate telecommunication responsibilities together. Only in this way can organisations prepare themselves to take advantage of the opportunities that will be provided by the wider availability of products integrating voice and data, and by the provision of integrated services digital networks by the PTTs.

Several of the organisations that we interviewed had recognised the inadequacy of their data communications control, and have instituted studies to redefine the role of the data communications function within the organisation.

The Butler Cox Foundation
© Reproduction by any method is strictly prohibited

CHAPTER 7

GUIDELINES FOR DATA COMMUNICATION MANAGERS

In chapter 6 we identified a wide range of issues that users of data networks need to resolve. In this final chapter of the report we summarise the main lessons that emerge from our research by presenting a series of guidelines that data communication managers can use when they set about resolving those issues. The guidelines are presented under three headings:

—Long-term plans and strategy.

—Medium-term plans.

—Operation of the network.

The guidelines are not intended to be complete prescriptions. We have examined the issues from only one perspective — that of the experience of a relatively small sample of users. Hence, these guidelines need to be used in conjunction with an assessment of how an organisation expects its own data communication needs to grow, and also with information about how the technology, products and services are expected to develop in the next five to ten years.

In many ways the guidelines are similar to those that apply to any data processing applications. This similarity reflects the increasing sophistication of data networks and the trend for data networks to be perceived as a company-wide service.

LONG-TERM PLANS AND STRATEGY

1. Organisations need to recognise that the data communications function will change from one that provides ad hoc services for particular applications, to one that provides a utility service to the company as a whole. At present many organisations do not plan ahead for their data communication facilities other than to provide short-term palliatives for day-to-day problems and needs. The impending change in the basic nature of the corporate data communications function will force organisations to place a greater emphasis on long-term plans and strategies for providing the service. In the future, data communications will require at least as much management attention as other company-wide services, such as the internal telephone network.

2. Responsibility for the provision and efficient use of data communication services (via private or public networks) should be given to one clearly defined part of the organisation. Ideally this should be the part that is responsible also for the provision and use of other telecommunication facilities. Integrating responsibilities in this way will anticipate the eventual integration into a single physical network of speech, data, text, facsimile, and other forms of communication.

3. The long-term data communications strategy should attempt to set out a blueprint of the ultimate network facilities required by the organisation. In particular, it should identify the paths along which to migrate towards an open network environment. Like all long-term plans, the data communications strategy will need to be reviewed at intervals to take account of changes in the business. The strategy will also need to be reviewed in the light of the new products and services that are expected to be available, and in the light of the progress made towards agreeing standards for the higher levels of the ISO reference model.

4. The data communication strategy should not be formulated simply on the basis of the forecast needs of the data processing department and the rest of the organisation. Instead, the strategy should be developed in conjunction with the overall business plans, which in turn should take account of the business opportunities made available by developments in data communications.

5. The long-term plan should consider the way in which user departments will be charged for their use of the data communications service. A long-term view should be taken in determining the internal charges. For example, the charging structure should not be designed to recover completely the initial investment costs during the period in which they are incurred.

6. The long-term data communication requirements should not be specified as detailed traffic forecasts. Instead, they should be expressed in terms of the different types of data traffic (speed, type of data call and geography). Defining the requirements in this way will allow networks to be specified in terms of wide-area and local area components, and in terms of the most appropriate points for gateways between components.

© Reproduction by any method is strictly prohibited

7. The long-term strategy should set out the way in which additions and amendments to existing networks should be authorised and controlled. A centralised control mechanism is desirable (though not essential) but, whatever the mechanism, organisations should institute a set of standard procedures for:

   —Authorising changes to the network.

   —Authorising the attachment of any new devices to the network.

   —Authorising changes in the software of attached devices, if it could affect the communication functions.

   —Authorising new users to access the network.

   —Dealing with faults, including the steps to be followed in tracking and clearing persistent faults.

   The operation of such procedures may, of course, be delegated to the separate operating units as appropriate. Nevertheless there should be a central co-ordinating mechanism that ensures that several major changes to the network do not occur simultaneously. (Foundation Report No. 21 contained a blueprint for a corporate communications strategy that suggested a possible sub-division of the responsibility for data and other networks.)

8. Although the ultimate goal may be a completely open network, it will be necessary to control the range of equipment (in terms of suppliers and protocols used) that may be connected to the existing networks. This control is particularly important if proprietary network control systems are to be installed. Once again, this does not imply that every purchase of terminals or data transmission equipment should be referred to a central body for approval. But the central body should at least list the approved equipment, protocols and interfaces that can be supported by the networks. Such a list must be defined in sufficient detail to avoid incompatibility. For example, to specify that "all SNA products" can be attached is too general.

9. The organisation's data network standards should be reviewed regularly in the light of new international standards as these are agreed by CCITT, ISO and other bodies.

MEDIUM-TERM PLANS

1. When specific additions or changes to existing networks are being considered, detailed traffic forecasts are essential. However, those responsible for implementing the additions or changes must treat the predictions with caution. These forecasts should be used as a standard against which the actual traffic can be compared once the change has been made, so that gross mistakes in the planned capacity of the network can be rectified as soon as possible.

2. Detailed plans are required for the successful and prompt installation of a data network. Due allowance should be made in the plans for:

   —Late delivery of equipment and lines.

   —Problems caused by inexperience (both by the supplier and by the organisation itself) in using any new product or service.

   —The need to provide an uninterrupted service to existing users of the network.

   —The need to educate the end users and convince them of the advantages of the amended or new network facility.

3. The user experience related in this report shows that there are no obvious reasons for rejecting any particular type of wide-area or local area network. In all the case histories in this report, the users reported eventual satisfaction with the networks they had chosen. The user experience shows that the new local area networks have been installed with fewer difficulties than might have been expected. One general lesson to emerge from the case histories is that organisations should be aware of the dangers of being a pioneer user — particularly if the chosen network is a public data network. The most important criteria for selecting a data network facility are to choose a network approach that best suits the application and to choose a stable supplier.

4. Some organisations will, nevertheless, choose to be pioneer users of a new product produced by a relatively new supplier. In this situation, the organisation should join forces with the other pioneer users of that product so that, if necessary, the supplier can be pressured into providing a better service or eliminating any persistent faults in the equipment.

OPERATION OF THE NETWORK

1. Operational staff should be involved, if at all possible, in planning and implementing all new network applications. At the planning stage, the network operations staff can contribute their views on any problems that the proposed change or additions might cause. During the implementation stage, operational staff can become familiar with the new applications, so that they are better prepared to deal with any network operation faults that may occur.

The Butler Cox Foundation
© Reproduction by any method is strictly prohibited

2. Because all network products have some drawbacks, a thorough analysis of all the features of each potential product or service should be made before a particular facility is selected. Specific features that should be examined are:

   —Network monitoring features.

   —Fault location aids.

   —Network management statistics.

   —Documentation and training.

   —The degree of user friendliness of the facility.

3. Procedures should be well defined for handing over the network from the development and implementation stages to the stage of routine operation. These procedures should be similar to those used for data processing applications.

4. Organisations should seriously consider using a centralised network control system for wide-area networks.

© Reproduction by any method is strictly prohibited

CONCLUSION

The purpose of this report was to assist data communication users to select the most appropriate form of network products and services, in the light of user experience with the different types of data networks and products that are available today.

The experience documented in this report has shown that users of all types of data communication products and services have reported success in using them (although there may have been initial teething problems). It has shown also that no one type of product or service is preferable. Rather, it has shown that organisations need to evaluate the range of products and services that are now available against their own particular applications requirements. To assist data communication managers to perform this evaluation, the report concluded with a set of guidelines that can be used when organisations set about resolving the key issues identified in the report.

A common thread that emerged from the user experience was that data communication costs are accounting for an ever-increasing proportion of the total data processing budget. Organisations are now more conscious of the benefits to be gained from searching both for cheaper networks and for more efficient ways of using their existing networks. The user experiences documented in this report will provide valuable input to that search.

The Butler Cox Foundation
© Reproduction by any method is strictly prohibited

APPENDIX

The main CCITT technical recommendations relevant to data network standards are as follows:

- X.3 The packet assembly/dissembly (PAD) facility in a public data network. This defines how a non-packet mode terminal communicates with a packet switch.
- X.20 The general-purpose interface between start-stop terminals and a data network.
- X.20 bis The V.21-compatible interface between start-stop terminals and a data network. (V.21 is the CCITT recommendation for 200-baud modems used on the PSTN.)
- X.21 The general-purpose interface between synchronous terminals and a data network.
- X.21 bis The interface between synchronous terminals designed to operate with V-series synchronous modems and a data network. CCITT recommendations for synchronous modems include V.27 (4,800 bit/s), V.29 (9,600 bit/s), V.35 (48k bit/s).
- X.25 The interface between a packet-mode terminal and a packet-switched data network. X.25 incorporates the lowest three layers of the OSI reference model — physical, data link and network control.
- X.28 The interface between a start-stop terminal and the PAD facility in a data network.
- X.29 The procedures for communication between a packet-mode terminal and a PAD across a packet-switched network.

© Reproduction by any method is strictly prohibited

Butler Cox & Partners Limited
Morley House, 26-30 Holborn Viaduct, London EC1A 2BP
☎ 01-583 9381, Telex 8813717 LNCO

Belgium & The Netherlands
SA Butler Cox NV
Avenue Louise – 479 – Louizalaan,
Bte – 47 – Bus,
Bruxelles 1050 Brussel
☎ (02) 647 15 53, Telex 61963 BUTCOX

France
La Fondation Butler Cox
Tour Akzo, 164 Rue Ambroise Croizat,
93204 St Denis-Cedex 1, France
☎ (1) 820.61.64, Telex 610789 ASFRA

United States of America
Butler Cox & Partners Limited
216 Cooper Center, Pennsauken, New Jersey 08109, USA
☎ (609) 665 3210

Switzerland and Germany
Butler Cox & Partners Limited
Morley House, 26-30 Holborn Viaduct, London EC1A 2BP
☎ (London) 583 9381, (Zurich) 302 0848

Italy
Sisdoconsult
20123 Milano – Via Caradosso 7 – Italy
☎ 86.53.55 / 87.62.27, Telex 311250 PPF MI

The Nordic Region
Statskonsult
PO Box 4040, S-171 04 Solna, Sweden,
☎ 08-730 03 00, Telex 127 54 SINTAB