Box.

C

Report Series
No 23

Communicating Terminals

May 1981

The Butler Cox Foundation

Abstract

Report Series No 23 — Communicating Terminals

by Tony Gunton

May 1981

For well over a decade now, terminals have been used to communicate with computer systems. During that time, the technology of terminals has improved substantially whilst their cost has fallen steadily. In addition, the number of terminals in use has increased exponentially, and the range of applications for which they are used has widened. Some organisations have already reached the point where they have installed one terminal for every two or every three members of staff.

There is no sign at present that any of these trends are coming to an end. Assuming that they do continue, it is only a matter of time before terminals become as ubiquitous in most organisations as the telephone is today.

In this report, we review the developments in terminal technology and we examine the forces driving the terminal market. We also return to the perennial problem of terminal incompatibility (which we previously discussed in Foundation Report No. 3). In this report, we look at the evolution of communications protocols, and we look also at the progress the standards authorities have made since we published Report No. 3.

Finally, we discuss the role of management services in selecting and providing for communicating terminals. Clearly, the skill with which terminals are introduced into more and more parts of the business will help to determine how much benefit is gained from their use. There can be no doubt that the potential benefits are considerable.

The Butler Cox Foundation is a research group that examines major developments in the fields of computers, telecommunications and office automation on behalf of its subscribing members. The Foundation provides a set of "eyes and ears" on the world for the systems departments of some of Europe's largest organisations.

The Foundation collects its information through its office in London and also through its associated offices in Europe and the US. It transmits its findings to its members in three main ways:

— Through regular *written reports* that give detailed findings and substantiating evidence.

— Through *management conferences* for management services directors and their senior colleagues, where the emphasis is on the policy implications of the subjects studied.

— Through *professional and technical seminars* where the members' own specialist managers and technicians meet with the Foundation research teams to review their findings in depth.

The Foundation is controlled by a Management Board whose members include representatives from the Foundation member organisations. The responsibilities of the Management Board include selecting topics for research and approving the Foundation's annual report and accounts, which show how the subscribed research funds have been employed.

Report Series No 23

COMMUNICATING TERMINALS

by Tony Gunton

May 1981

Butler Cox & Partners Limited Morley House 26 Holborn Viaduct London EC1A 2BP

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

CONTENTS

CHAPTER 1

INTRODUCTION

TERMINAL COMPATIBILITY

Foundation Report No. 3, published in late 1977, discussed the problem of terminal incompatibility. That report showed that the origins of device incompatibility were complex, and it also showed that incompatibility had come about through a combination of historical accident, deliberate marketing policy and technological change. The report was pessimistic, in that it did not foresee an immediate improvement in the position. It did, however, point to several potential sources of relief, of which the two most important were, first, microprocessors, which permitted terminals to be more intelligent, and second, international standardisation efforts, which were showing early signs of success.

This report extends and broadens the analysis contained in Report No. 3, and it does this by looking at certain aspects of communicating terminals other than the problem of incompatibility. As this report shows (and as many managers responsible either for procuring terminals or for providing communications support for terminals will not need reminding), the incompatibility problem persists today in a form, which, regrettably, is hardly less acute than it was three years ago.

The incompatibility problem persists not because no progress has been made. On the contrary, the sources of relief cited in Report No. 3 have delivered some results. Terminals have become more intelligent, and many terminals now available can communicate with more than one computer manufacturer's equipment. Standardisation efforts also have progressed, although perhaps not as far or as fast as might have been hoped for, given the urgency of the terminal compatibility problem.

Unfortunately, at the same time as these advances have been made, the underlying problem has itself been compounded. Terminal intelligence can be used, and indeed has been used, to alleviate the incompatibility problem, but equally it can compound the problem. As Report No. 3 pointed out, an intelligent terminal is difficult to emulate, and it can easily be modified by the manufacturer to inhibit competition. Primarily for this reason, the gulf between the actual and the ideal as far as terminal compatibility is concerned remains as wide as it was three years ago. In addition, terminals are increasingly penetrating into the business world and, as they do so, they are being used for more and more varied purposes.

OBJECTIVES OF THIS REPORT

This report has two main purposes:

— To re-assess the problem of terminal incompatibility.

— To examine the technology of communicating terminals and to assess the future development of that technology.

Also, this report seeks to advise organisations both on the policy they should adopt for procuring communicating terminals and on the influence that this policy might have on communications planning. The report does not contain a catalogue of terminal features because this would, by its nature, contain transient information inappropriate to a Foundation report. In any event, such a catalogue is readily available from several other sources.

SCOPE OF THE REPORT

For the purposes of this report we define a communicating terminal as a device that a human being interacts with (for example, by keying) to gain access to either a remote processing service or a remote information service. Many users will, of course, also use communicating terminal devices to access local services where the communications facilities will be different. Local-access devices present a subset of the problems facing communicating terminals, and for that reason much of the discussion in this report will apply also to local-access devices.

We have excluded remote batch terminals from the scope of the report on the grounds that they are, in effect, used by the remote computer rather than by the human beings who operate them. Remote batch terminals are in any case a highly-specialised class of device.

Strictly speaking, our definition of communicating terminals also includes the ubiquitous telephone, which is sometimes used to perform data entry tasks. However, for data entry tasks the use of the telephone is minimal when compared with the use of conventional keyboards. In this report, therefore, we concentrate on keyboard terminals, whilst acknowledging that the telephone and/or the telephone switching system might have a significant role to play in the future. The role of the telephone and the telephone switching system will be discussed in more detail in Foundation Report No. 26.

INTENDED READERSHIP

The report is intended to be read by those managers who are responsible either for procuring terminals or for providing the communications facilities that support those terminals. It is also intended to be read by those managers (usually management services managers) who establish the policy framework within which such managers operate.

CHAPTER 2

A CLASSIFICATION OF TERMINALS

THE NEED FOR AN ENDURING CLASSIFICATION

To produce conclusions that will have enduring significance, it is necessary first to classify terminals in ways that also can endure. Terminals are normally classified either in terms of their features (keyboard/printer devices, visual display units, etc.) or by the application area in which they are used (for example, banking terminals and point-of-sale terminals). Usually, the former classification is applied to general-purpose terminals and the latter to special-purpose terminals. In the next two sections we discuss the significant categories of terminals within the two headings of general-purpose terminals and special-purpose terminals. Later in the chapter we also discuss the classification of terminals by their mode of use.

GENERAL-PURPOSE TERMINALS

In classifying general-purpose terminals, we have identified five major categories of terminals. These five categories, which we discuss below, are interactive printers (including keyboard/printer devices), visual display terminals, general-purpose printers, word processing terminals and multifunction terminals.

Interactive printers (including keyboard/printer devices)

Interactive printers and keyboard/printers are low-speed or medium-speed devices that are used to interact with a remote computer, such as a timesharing bureau. They incorporate a typewriter-style keyboard (which is usually extended by the addition of several special function keys) and a printer. The printed characters are normally formed by a dot matrix, and the print quality is usually poor. Some of the terminals in this category are portable. Most interactive printers and keyboard/printers do not have any local processing power, but a few have storage facilities (such as cassette tape, cartridge tape or bubble memory) for off-line data entry purposes.

Visual display terminals

Visual display terminals are either medium-speed or high-speed devices, and they are normally used when interaction with a remote computer is necessary, but when there is no requirement for printed output. These devices incorporate a typewriter-style keyboard (which is usually extended with additional numeric keys, editing keys and function keys) and a display screen. The display screen usually consists of a CRT that displays 24 lines, each of which contains 80 characters. This category of terminal can be further divided into the following three sub-categories that classify visual display terminals by the amount of local processing power contained in the terminal:

1. Unintelligent visual display terminals
   Unintelligent visual display terminals have no local processing capability, and they are used mostly as an on-line device to a remote computer. Often, they have only an upper-case character set, a limited set of keys and no editing or other

functional capability. Any processing power available to the user is contained in the computer system to which the terminal is connected.

2. Smart visual display terminals
Smart visual display terminals provide local processing power that permits the user to enter data direct to the display screen from the keyboard without the intervention of the remote computer. The user can edit or modify the data before he transmits it in either a page mode or a block mode. Because smart devices are based on microprocessor electronics, they can also incorporate various other functions and facilities, including:

— Support for peripheral devices such as printers and mass storage devices.

— Communications buffering.

— Limited control of the internal memory.

— Video attributes (such as variable intensity characters, flashing, reverse image and underlining).

— Block graphics.

— A numeric keypad.

— Keyboard control of the cursor position.

— Communication of terminal attributes to a remote computer.

— Tabulating facilities.

— Self-testing facilities.

Smart visual display terminals often permit keyboard control of features that can 'personalise' the terminal. For example, the user may be able to specify the action the terminal will take when the return key is depressed.

3. Intelligent visual display terminals
Unlike smart visual display terminals, intelligent visual display terminals can be programmed by the user, who may, for example, use a high-level language such as a screen formatter. Programs can be loaded either from the terminal's keyboard or from an on-line computer. Intelligent visual display terminals often incorporate more memory than do smart devices, and this additional memory is used to store either pages of data or programs.

General-purpose printers

General-purpose printers are faster than interactive printers, and they use either impact technology or non-impact technology. The quality of their printout is usually relatively poor, but their reliability is good.

Word processing terminals

Word processing terminals are used for entering text, editing text, communicating text and printing text. These terminals can be divided into the five sub-categories of editing typewriters, thin-window typewriters, memory typewriters, stand-alone display word processors and shared-logic display word processors. Each of these five sub-categories is discussed below:

1. Editing typewriters
   Editing typewriters are sold as replacements for IBM typewriters. They consist of a typewriter keyboard and either a golf-ball printer, or a daisy-wheel printer, or an ink-jet printer. These typewriters usually have simple editing facilities such as erase-last-character or erase-last-line. Although few of them at present have communications facilities, there is an increasing trend for editing typewriters to be provided with communications facilities.

2. Thin-window typewriters
   In addition to having a keyboard and a printer, thin-window typewriters also incorporate a one-line (or a few-lines) display. This thin-window display may be constructed either from light-emitting diodes or from liquid crystal displays or plasma panels. Compared with editing typewriters, thin-window typewriters provide more comprehensive editing facilities for the line that is displayed. These facilities permit the user to alter any element of the displayed line before it is printed.

3. Memory typewriters
   Memory typewriters incorporate facilities for storing text either before or after it is printed. The storage facility may be either internal or external to the typewriter, and it may consist of bubble memory, cassettes, mini-diskettes or magnetic cards. The printing of text is usually separated both from the inputting of text and the communicating of text.

4. Stand-alone display word processors
   Stand-alone display word processors consist of a keyboard, a display screen, a high-quality printer and a storage facility. The keyboard of a stand-alone display word processor has, in addition to the standard alphanumeric keys, keys to facilitate the inputting, editing, printing and transmitting of text. The display screen usually displays 24 lines, each consisting of 80 characters. However, the screen may be taller to accommodate an A4-size document, or wider to accommodate displays of 132 characters per line. High-quality CRTs are often used for the display screen.

   The storage facility typically consists of a cassette tape, or a diskette, or a cartridge disc, or a Winchester-type disc. Stand-alone display word processors provide comprehensive editing facilities, such as facilities for moving sections of text from one document to another and for merging two or more documents. Some document filing facilities may be provided, and various output formatting facilities are provided. As with other word processing terminals, most of the installed base of stand-alone display word processors do not communicate remotely. Many stand-alone word processors, however, have the facility, at least as an option, to communicate remotely.

5. Shared-logic display word processors
   From the users' point of view, shared-logic display word processors are functionally the same as stand-alone display word processors. When a department has a need for more than a certain number of word processing devices, it is cheaper to put the processing power in a separate unit and share it between several word-processing terminals. In addition to having this centralised processing power, shared-logic display word processors may also have centralised printing facilities (particularly for high-speed draft-quality printing) and centralised storage facilities in the form of disc packs that hold common files.

Multifunction terminals

Multifunction terminals are ill-defined at present, because there are not yet many products available that do not obviously fall into one of the categories described above. The few multifunction devices that are available integrate word processing with data processing, and hence they combine some of the characteristics of all four categories described above. Some of the multifunction products also include image processing (using raster-scanned displays and printers) and, at least in prototype form, voice input and voice output facilities.

SPECIAL-PURPOSE TERMINALS

Special-purpose terminals employ either special printing mechanisms or special reading mechanisms. Examples of special-purpose terminals are industrial data collection terminals, Post Office counter terminals, ticketing terminals and cash receipting terminals. Three categories of special-purpose terminals deserve particular attention because they may exert a substantial influence on terminals in general. They may do this either because of the particular technology they employ or because of the large numbers of them that are already (or will be) installed. They comprise public-service terminals, point-of-sale terminals and financial terminals.

Public-service terminals

Public-service terminals can be divided into the three sub-categories of telex terminals, teletex terminals and videotex terminals. Each of these three sub-categories is discussed below:

1. Telex terminals
   There are now more than one-million telex subscribers worldwide, and more than half of those subscribers are in Europe. Traditionally, telex terminals have been keyboard/printer devices, although there are now terminals that incorporate visual displays, local storage facilities and editing facilities. Both the mode of communication and the limited character set used by the telex network separate telex terminals from most other terminal devices.

2. Teletex terminals
   Several PTTs are planning to introduce teletex services, and some of their plans include a special-purpose teletex terminal for accessing the new services. Teletex will remove the barriers that separate telex terminals from other communicating terminals. Different PTTs, however, are placing a different emphasis on their teletex services. Some PTTs see teletex as a 'super-telex' service, some see it as a communicating word processor service, and yet others see it as something between the two. Teletex, in brief, is still at a formative stage of development.

3. Videotex terminals
   Although videotex service plans are more advanced than the plans for teletex, videotex is also at a formative stage of development. For domestic use, videotex terminals are based on adapted televisions, and the user interacts with the system either through a numeric keypad (which doubles as a remote controller for the television) or through a touch-tone telephone. For business use, special-purpose videotex terminals are available with a full alphanumeric keyboard. Videotex terminals display 24 (or 25) lines of information of 40 characters, and most terminals conform to CCITT Recommendation V.23. This means they receive information at 1200 bit/s and transmit information at 75 bit/s.

Point-of-sale terminals

Point-of-sale terminals can be sub-categorised into stand-alone devices and large-cluster systems.

1. Stand-alone devices
   Stand-alone point-of-sale devices are used in stores and shops as cash tills, as receipting machines and as store-accounting machines. There may be one or more devices per store, and, if there are several devices in a store, they may be connected together. Stand-alone point-of-sale devices consist of a visual display (either a CRT, or light-emitting diodes, or a thin-window plasma panel), one or more printers, a keyboard and a cash till. The keyboard is usually numeric only, but with a few extra function keys. Point-of-sale devices may also include a variety of additional features such as bar-code readers. Most stand-alone point-of-sale devices can communicate either on-line or off-line with a remote computer.

2. Large-cluster point-of-sale systems
   Large clusters of point-of-sale terminals are most commonly used in supermarkets, where each terminal is connected to a central store controller. The terminals themselves usually have limited processing power and also limited facilities because the central store controller provides, on a shared basis, all but the most rudimentary processing power and facilities. The central store controller usually communicates with a remote computer in an off-line batch mode.

Financial terminals

Financial terminals consist of teller stations and automatic teller machines. Teller stations are used at the front desks in commercial banks and savings banks. They need to be small, and often they incorporate cheque scanners or pass-book readers.

Automatic teller machines provide either limited cash dispensing facilities or more comprehensive services such as the handling of credit transfers, payments and account status enquiries. They incorporate a magnetic stripe reader, a small keyboard, a money dispenser, a printer and a limited visual display. They communicate with a remote computer system in both an off-line mode and an on-line mode.

CLASSIFICATION OF TERMINALS BY MODE OF USE

From the terminal user's viewpoint it may be perfectly adequate to classify terminals as either general purpose or special purpose (as we have done above), especially when that classification is combined with details of the specific features and characteristics the user requires. However, there are two reasons why such a classification is neither adequate nor ideal for the purposes of this report. Firstly, developments both in technology and in applications present an ever-changing environment in which to classify terminals. Secondly, the way the terminal is used, as well as its constitution as a product, will determine both the demands placed on the communications network and some of the criteria for evaluating the device itself. The mode of use of a terminal is therefore likely to be a significant factor, both when considering the suitability of a particular terminal and when considering its impact on other system components (such as the network). Consequently, it is necessary to classify terminals also by their mode of use. Essentially, this classification represents the terminal as it appears to the communications network.

In Foundation Report No. 21, we defined the five types of network traffic as voice communication, bulk data communication, on-line data communication, message communication and video communication. By defining the type of network traffic in that way, we illustrated the different demands placed on a network in terms mainly of bandwidth constraints and timing constraints. Voice communication and video communication are outside the scope of this report, but the remaining three network traffic types can be expanded to form the basis for classifying communicating terminals by their mode of use. This classification is illustrated in figure 1.

Traditionally, different protocols and different network structures have been used for the three types of traffic, as shown in the figure. Each type of traffic in its turn accessed a different set of ports and a different piece of communications software in the computer system that it addressed. To rationalise this situation, computer manufacturers and communications equipment manufacturers have progressively brought the three types of network traffic together. The penetration of terminals into organisations demands such a rationalisation, and it is clearly essential that both the terminal and the network should be able to cope with multiple modes of use if the concept of the multifunction terminal is to succeed.

The complexities of the multifunction approach for both the terminal and the network (and indeed the standardisation efforts) should, however, not be overlooked. Where there is no evident need for a multifunction terminal, separate devices connected to separate networks may remain cost-effective for some years yet.

In this respect, the key distinction is between on-line data traffic and the two other traffic types shown in figure 1. Both bulk data traffic and message traffic operate (or may operate) in an off-line mode where the sender and the receiver are not coupled simultaneously. On-line traffic tends to demand a higher level of intelligence in the terminal and the computer system respectively at each end of the communications link. This higher intelligence is required to control both the flow of information and the logic of the dialogue. Although off-line traffic may be handled in a similar way (except that communication need not be so strictly co-ordinated), it can also be handled successfully by placing the controlling intelligence in the network, which acts as an intermediary. As a result, both the terminal logic and the computer software logic can be made considerably simpler. For the same reason, protocols for off-line communication can be considerably simpler than those required for on-line communication.

The protocols that a terminal uses will largely determine the internal hardware and software structure of the terminal. Communications protocols and standards have been the subject of much international debate in the past few years, and in chapter 5 of this report (commencing on page 35) we provide a detailed review of the current status of the protocols and standards that govern communication.

SUMMARY

Two points emerge clearly from our comprehensive classification of communicating terminals. Firstly, the range of options available is considerable and it is multi-dimensional. Hence, rationalisation, of either the product features, or the network interfaces, or the internal structure of terminals will inevitably be a protracted process.

Secondly, and consequently, terminal design (as far as suppliers are concerned) and terminal selection (as far as users are concerned) must take account of all the dimensions we have identified. Thus, suppliers and users must consider product features, network interfaces and the internal structure of terminals.

Figure 1 Classification of terminals by mode of use

Traffic		Typical usage parameters		
Type and sub-type	Description	Transmission	Network structure	Protocols
Bulk data traffic				
Batch transfer	Transfer of bulk information for processing	Synchronous 2,400 to 48,000 bit/s	Remote job entry procedures over point-to-point lines or sometimes over the PSTN	IBM 2780 IBM 7181
Document transfer	Transfer of bulk information for printing			
On-line data traffic				
Transaction processing	Exchange of messages to complete a transaction	Synchronous 1,200 to 9,600 bit/s	Poll/select over multi-point or point-to-point lines	IBM 3270 ICL 7502
Data entry and data editing	Entry of information for computer processing			Uniscope
Information retrieval	Retrieval of information for examination			
Message traffic				
Information retrieval	Retrieval of information for examination	Asynchronous 300 to 2,400 bit/s	Contention procedures over a switched network (PSTN, packet or private). Sometimes over point-to-point lines	ASCII/Teletype IBM 2741 X.28/X.29/X.3
Interactive processing (timesharing)	Interactive access to computer facilities			
Electronic mail	Sending or receiving a person-to-person message			
Data collection and data distribution	Sending or receiving a message to or from a computer process			

CHAPTER 3

TERMINAL TECHNOLOGY

In this chapter we examine the technology of the constituent parts of communicating terminals. Most terminals are made up of an input device, a display or a printer, and software to control the electronics of the terminal, and we consider each of these in turn. We also discuss the recent trend towards using the terminal's electronics not only for controlling the device, but also for adding more functions to it.

The purpose of this chapter is to review the technologies that are already available, or are becoming available, and to identify those technologies that might be important in the future. The chapter concludes with a review of the trends in the cost of manufacturing terminals.

KEYBOARDS

The keyboard is the most common input device used by terminals, and despite the effort that has been put into developing alternative methods, the keyboard is likely to remain as the most common input device, for the following four reasons:

1. A change to a different device or method for generating text or data would require a substantial training programme spread over several years.

2. Keyboards are used by the data terminals that are now widely installed, and the installed base of data terminals will not be replaced overnight.

3. Keyboard technology is mature, and, consequently, keyboards are low-cost input devices.

4. Evidence from the United States shows that even quite senior executives are prepared to use keyboards.

In this section of the report we review the technology of keyboards under the headings of keyswitch technology, keyboard layout and keyboard interface.

Keyswitch technology

The keyswitch is the fundamental unit of a keyboard, and there are two main styles of keyswitch. The first style of keyswitch uses solid-state non-contact keys that are built into a printed circuit board. Solid-state keys use either capacitive technology or Hall-effect technology (which alters the inductance of the circuit). The main advantage of non-contact keyswitch technologies is that very high reliability can be achieved. A keyswitch life of 200 million depressions is normal for capacitive technology, and 500 million depressions is quoted for Hall-effect technology. However, it is difficult to reconfigure solid-state keyboards because it is necessary either to replace the appropriate memory boards or to re-program the internal memory of the keyboard.

The second style of keyswitch is based on contact technology, where two contacts are joined (or separated) when the switch is depressed. Contact keyboards are very commonly used, particularly in cheaper or low-use terminals. Because the switch itself is inside the key, the switches are relatively simple to repair, and the keyboard can easily be re-configured. Contact keyboards are, at present cheaper than non-contact keyboards, but, with a keyswitch life of between 5 million and 20 million depressions, they are not as reliable as non-contact keyboards. In the future, the cost of contact keyboards will, relative to non-contact capacitive keyboards, become more expensive because of their mechanical content and greater dependence on labour costs. This relative increase in cost, together with the higher reliability of capacitive keyboards, leads us to conclude that, within the next five years, capacitive keyboards will become the norm for computer terminals.

A third style of keyswitch that, potentially, is important is the touch-sensitive switch. This style of switch works by finger contact altering the capacitance of the circuit. The main advantage of the touch-sensitive switch is that it does not suffer from the mechanical unreliabilities both of non-contact switches and contact switches. However, the surface of a touch-sensitive switch needs to be metallised, and if this surface becomes dirty, the switch can become inoperative. Also, a touch-sensitive keyswitch requires complicated electronics in order to overcome the hysteresis effect that makes it difficult to determine exactly when the key is either touched or released. The major disadvantage of touch-sensitive keyswitches, however, is that most typists require a tactile feedback to confirm that the key has been depressed. For this reason, touch-sensitive keyboards are unlikely to be widely adopted for terminals that are used in either a professional or a business environment.

Keyboard layout

Most terminal keyboards use the QWERTY (or AZERTY) layout. This layout was originally invented to avoid clashes between typebars, and to place commonly-used pairs of keys far apart on the keyboard. The problem of clashes between type bars no longer exists, but there are still advantages in separating common pairs of keys, because this separation ensures a balanced division of operation between the two hands.

Some terminals have adopted an alphabetic keyboard layout. Although totally inexperienced typists find it easier to locate individual keys in an alphabetic layout, those with even a little typing experience can find it more difficult to use this layout.

Research has been carried out to design an ergonomic keyboard that makes typing as efficient as possible. For example, the Maltron keyboard is claimed to increase typing speed by up to 50 per cent. An increase of this magnitude makes direct dictation a feasible proposition.

Another important factor that should be considered when designing keyboard layouts is the need to minimise the chance of errors occurring. For example, critical keys such as the break key should not be situated adjacent to the commonly-used keys. At present, different manufacturers position the function keys and numeric keys in different ways. These different layouts lead to confusion when typists change from one layout to another, and there is an urgent need to define a standard keyboard layout for terminals.

Keyboard interface

The keyboard control electronics are usually housed in the display cabinet, although there is a growing trend for the control electronics to be housed in a separate unit. However, the hardware and the software used to code and decode characters is often housed in the keyboard unit itself. Most keyboards are connected to the control unit by a bit-parallel (that is, a character-serial) interface cable which is constructed as either a flat-ribbon cable or a multi-core fat cable.

Several types of terminals, however, now incorporate electronics that assemble the keyboard output into a bit-serial form before it is transferred to the control unit. The method of transferring information from the keyboard to the control unit in a bit-serial form has two advantages. First, the connecting cable can be a thin, flexible cable, which enables the keyboard to be more portable. Second, two-way communication between the keyboard and the control unit can be easily implemented. The return channel to the keyboard can be used to control the illumination of keys or lamps on the keyboard, or it can be used to interrupt the keyboard. Facilities such as these will become increasingly valuable as terminals begin to be used for applications such as electronic mail.

ALTERNATIVES TO KEYBOARDS

Despite the continuing dominance of the keyboard, terminals with other forms of input devices will probably take an increasing share of the terminal market in the not too distant future. The most likely candidates for general-purpose data and text terminals are voice-recognition devices and handwriting terminals. Special-purpose terminals will use an increasing number of other input devices, such as bar-code readers and magnetic-card readers. In this section of the report we examine the technology of several of these alternative input devices.

Voice recognition

Voice recognition devices are already available from several suppliers, as figure 2 shows. Current products, however, are able to recognise only a limited vocabulary of clearly-spoken discrete words, and commercial products that are able to perform continuous voice recognition of a full vocabulary (for example, 20,000 words in English) will not be available for at lest ten years. The products that exist today are expensive, and they are practical only in limited, special-purpose applications, such as baggage handling and postcode sorting.

Because of the problems involved in using keyboards with the extensive Japanese character sets, Japan is putting a significant research effort into voice recognition technology. However, the most advanced Japanese system to date can cope with only a one hundred-word vocabulary in real time. (A comprehensive description of the state-of-the-art of voice recognition technology can be found in the transcript of the eighth Foundation management conference held in Torquay in January, 1981.)

We believe that voice recognition input devices will be important in the future, not least because of Japan's research effort. But there is not sufficient evidence to suggest that voice recognition will ever completely replace keyboards, and in the next ten years voice recognition systems will be used only for specialised applications.

Handwriting input pads

Handwriting is a natural method of creating documents, and handwriting input pad technology is an attractive concept because it can eliminate the transcription carried out when a copy typist uses a keyboard to input a handwritten draft into a computer system. Handwriting input pads operate by using the pressure of a stylus or a pen to reduce the gap between certain sensitive membranes that form part of the electrical

Figure 2 Products available or under development for discrete speech recognition at the end of 1980

Supplier	Status	Unit price ($000)	Number of words in vocabulary	Response time	Percentage accuracy claimed
BTL	D	—	50	Slow	95
Centigram	A	3.5	32	Fast	95
Verbex (Dialog)	A	70	30	Fast	98
Heuristics (1)	A	0.3	30	Slow	Not known
(2)	A	2 to 3	64	Fast	Not known
IBM	D	—	50	Fast	95
Interstate Electronics (1)	A	18	50	Fast	95
(2)	A	2	40 to 100	Fast	98
(3)	A	0.5 to 0.9	16	Fast	98
Threshold	A	10 to 80	50 +	Fast	99
Fujitsu	D	22	Not known	Fast	Not known
Sanyo	D	—	32	Fast	90
Nippon (1)	A	80	120* 1,000	Fast	99
(2)	A	18.2	4	Fast	99
(3)	A	22.7	16	Fast	99
(4)	A	68.2	128	Fast	99

D = Under development
A = Available
* = Connected speech recognition

circuit. The resulting change in capacitance of the circuit elements is measured, and the changes are converted into co-ordinate values.

Handwriting input pad technology can be used relatively easily either to transcribe freehand drawings onto a display screen (as for a graphics application) or to choose between options on a pre-formatted form displayed on a screen. It is, however, much more difficult for a handwriting input pad to recognise individual handwritten characters, because the pad has to recognise the many different ways in which each character can be written. At present, character recognition devices (such as the Quest Micropad) are limited to only an upper-case character set. Obviously, though, to be really useful, handwriting keypads need to be able to recognise connected handwriting. The problems associated with recognising connected handwriting are similar to those of recognising continuous speech, and, like continuous speech recognition devices, continuous handwriting recognition devices will need to store each user's vocabulary and style. Continuous handwriting recognition devices are unlikely to be available as commercial products before the end of the 1980s.

Other input methods

Several other methods of inputting information to terminal devices are either available or under development. Some of these devices are general purpose and others are used in special circumstances. We discuss these devices under the headings of positional devices, optical character recognition devices, and readers and scanners.

1. Positional devices
   Positional devices permit the user to interact direct with the display screen either by touching the screen or by providing a flexible means of positioning the display cursor. These devices include light pens, touch-sensitive pads or screens, joysticks and roller-ball mice.

   Light pens are used to mark a screen location, and they operate by means of a photosensitive cell at the tip of the pen that recognises when the electron beam spot passes the light pen's position. The location software is usually built into the display's electronics, because this makes it easier to synchronise the locating procedure with the raster scan. Touch-sensitive pads are sometimes used instead of light pens, and touch-sensitive screens may be used either to provide for situations where the user is inexperienced or to implement 'soft' keys that the user can program.

   Joysticks and roller-ball mice are both used to move the cursor around the display screen, particularly for graphics applications. They may be used, for example, to position a number in a particular region of the screen, or to point at a screen position. Joysticks operate by the user positioning the stick in one of eight positions around it. A roller-ball mouse rolls across a horizontal surface, and the position of the mouse is determined by using an inductive principle.

   Each of the four types of positional devices mentioned above needs to respond virtually in real time to the user's actions. Consequently, the processing power necessary to control these devices must be located in the terminal, rather than in the host computer system.

2. Optical character recognition devices
   Optical character recognition devices use either arrays of photosensitive cells or scanning technology to recognise printed characters. When the device recognises a character it translates it into a coded form. Optical character recognition technology has the advantage that material that has been typed or printed in a recognisable font can be input. However, the recognisable fonts usually require a change of the printhead on the typewriter that produces the material.

   At one time, optical character recognition technology seemed likely to replace other input methods, particularly for inputting textual material. Its future now seems less certain, because other input technologies, such as voice recognition, handwriting recognition and bar-code readers seem likely to be more effective.

3. Readers and scanners
   Input devices in this category include bar-code readers, laser bar-code scanners,

badge readers and card readers. Bar-code readers consist of a stylus comprising a light source and a photosensitive cell. The stylus is connected by a flexible cable to either a control unit or a terminal. Bar-code readers are used most commonly in a retail environment, but they are used also in an industrial environment (for example, to record the movement of particular items of stock). They are likely to become a universal addition to stand-alone point-of-sale devices, particularly in smaller retail outlets. The major disadvantage of bar-code readers is that every product package has to be marked with a bar code, although 70 per cent of all goods now sold at supermarkets in the United Kingdom are marked with bar codes.

Laser scanners are used in larger retail outlets such as supermarkets, and they read the same bar codes as bar-code readers. The laser scanner is situated in a slot in the check-out desk, and the operator passes the item to be read over the slot with the bar code facing downwards. A low-powered optical laser beam is focused through the slot, and the reflected light from the bar code is measured by a photosensitive device.

Badge readers and card readers read (and write) information recorded in a magnetic stripe incorporated in a plastic card. Plastic cards are used as credit cards and also as keys for automatic teller banking terminals, and they are also beginning to be used both as bus tickets and as keys for vending machines.

As part of their Télématique programme, the French have developed a 'smart' card. (A description of the smart card can be found in the transcript of the eighth Foundation management conference held in Torquay in January, 1981.) Instead of a magnetic stripe, the smart card incorporates a simple four-bit microprocessor and a limited amount of semiconductor memory. Members of the public in France will use these cards to pay for their use of the public terminals connected to the French videotex system.

DISPLAYS

Visual displays were originally developed as an offshoot from CRT technology and television technology. CRT technology has been used almost universally for the production of alphanumeric visual display terminals over a period of nearly twenty years. The reduced prices of the associated logic elements and memory elements, and the constantly increasing manufacturing volumes have enabled the manufacturers of CRT-based devices to reduce their prices steadily at a rate that has often exceeded 10 per cent annually. This trend has been accentuated by the use of standard modules from the consumer television industry, which bring with them the benefits of large production runs and lower costs.

In addition, the quality of CRT displays has improved significantly over the last twenty years and the technology is still improving. As figure 3 shows, the light output obtainable from a CRT has increased by a factor of about 50 during that period. For all the reasons just discussed, CRT technology will remain in common use, and it is likely still to be the dominant display technology in ten years' time.

CRT technology is used in visual displays in three forms — refreshed CRTs, directed-vector CRTs and storage tubes — and in this section of the report we review the technology of each of these forms. However, other display technologies are becoming available, and we also review the technologies of gas plasma panels, and other display technologies currently in use or under development. Finally, we review the technology of colour displays.

Figure 3 Change in the light output of a CRT

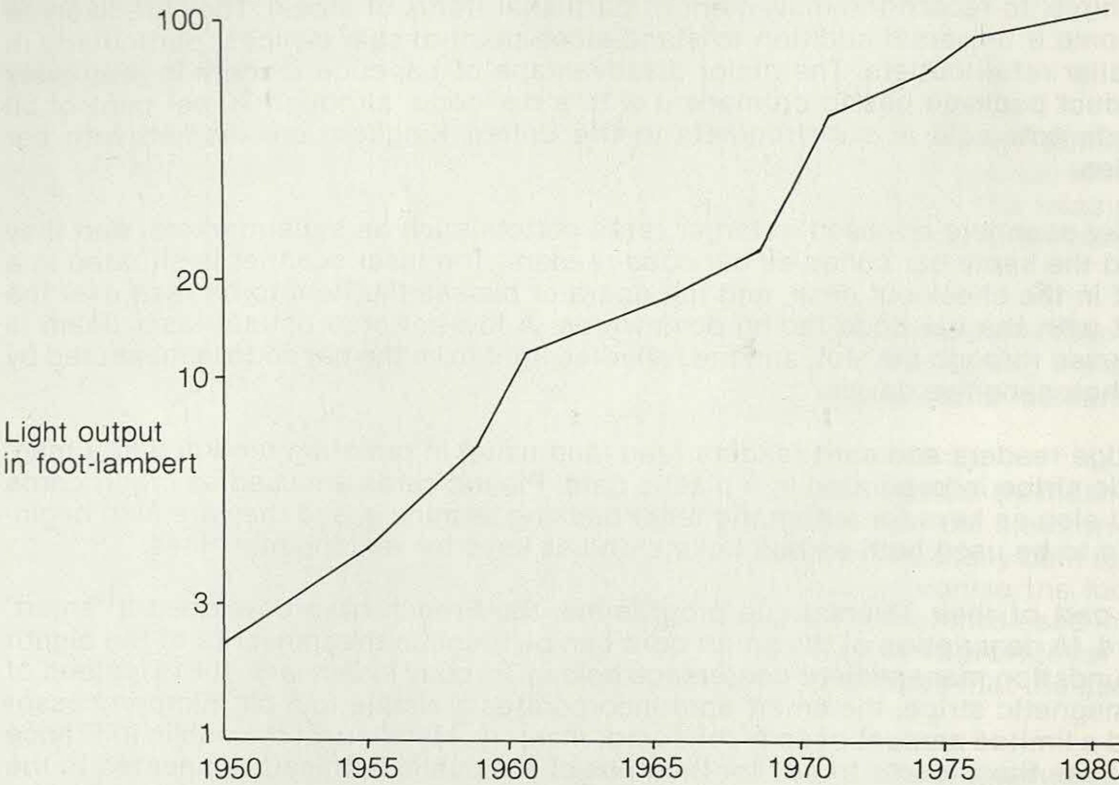

Refreshed CRTs

Television receivers use refreshed-CRT technology, and the image on this type of display is created by a fast-moving beam of electrons striking the phosphor coating that is permanently deposited on the back of a glass screen. The image needs to be refreshed periodically because the glow produced when the beam strikes the phosphor fades after a short time. To achieve the periodic refreshing, the beam is continuously scanned across the screen, and there are four well-established techniques for doing this, as discussed below:

1. Cursive scan
   With the cursive-scan technique, the electron beam is deflected around the outline of each character. This technique is widely used both in large graphical displays and in radar displays. It produces very high-quality images, but it is unlikely to be used in a display that costs less than £20,000.

2. Jiggle scan
   With the jiggle-scan technique, the electron beam is deflected very quickly to scan the character position. The jiggle scan is combined with slower vertical and horizontal scans, and it produces good quality images at a moderate cost. This technique was used with early visual display terminals because it required a slow store access time that could be obtained easily with delay-line techniques. Nowadays, the availability of faster semiconductor stores and the reduced cost of addi-

tional scanning circuitry have largely caused the jiggle-scan technique to be superseded by raster-scan displays.

3. Raster scan
   The raster-scan technique is used both by all television receivers and by most visual display terminals. With this technique, the electron beam is continuously swept across the screen in a succession of scan lines, each of which is displaced a fraction below the one above. Each complete scan from top to bottom of the screen is called a field. With standard television receivers, alternate fields are displaced by half a scan line. Thus, alternate fields are interlaced, and the complete cycle is called a raster. In Europe, the television standard is a 625-line raster, and each field of 312.5 scan lines is repeated every 1/25th of a second.

   The interlacing of alternate fields with the raster-scan technique reduces the flickering effect that periodic refreshing can produce. However, if there is a slight movement between the fields, interlacing may cause some instability of the image. This instability is acceptable in a normal television picture, but it may be more noticeable on a computer terminal, particularly on a terminal that has a high resolution. For this reason, many terminals dispense with interlacing and double the refresh rate.

   The characters displayed on a raster-scan CRT screen are displayed as a dot matrix (typically 5 x 7), and the dot matrix pattern of each character shape that the terminal can display is stored permanently in the terminal in a read-only memory.

   The need to refresh the display every 1/50th (or every 1/25th) of a second is the major disadvantage of the CRTs that are used in terminals. Because the displayed information does not change in successive rasters, the screen contents must be stored. This means that the terminal (or its cluster controller) must incorporate at least one page of storage (called the refresh buffer), in addition to the read-only memory required to store the shape of each possible character.

4. Step scan
   The step-scan technique is a development of the raster-scan technique where the vertical scan is not uniform, but speeds up between rows. The extent to which the vertical scan is speeded up can be adjusted, and so it is possible to adjust the character height independently of the overall height.

Refreshed-CRT display screens typically display 24 lines of 80 characters, but screens with more lines and/or longer line lengths are becoming more common, particularly for use in word processing systems. These devices, however, use logic elements that do not come from the television industry. (Standard television technology imposes a limit of about 25 lines of 80 characters on the display.) Thus, larger screens (in terms of display characteristics) are more expensive than typical 24 x 80 displays.

Directed-vector CRTs

Directed-vector CRTs also use refreshed-CRT technology, but the characters are formed by a series of strokes, or vectors, rather than being formed as a dot matrix. The information that describes the vectors is stored in a read-only memory in the terminal, and image-generation circuitry is needed to construct the characters for display on the screen. Characters formed in this way have a higher resolution than do characters formed as a dot matrix. In addition, simple graphics (such as data-entry forms) can be generated by adding a small number of characters to those stored in the read-only memory. The major disadvantage of directed-vector CRT technology is that it costs more than conventional refreshed-CRT technology.

Storage tubes

A storage tube display, sometimes known as a direct-view CRT, is also a CRT device, but it uses a phosphor coating that retains its luminance for a long time. Storage tube displays therefore do not need to be refreshed to maintain the image. Thus, they do not require a refresh buffer, and, provided that they do not need to display characters, they do not require image-generation circuitry.

Because the image is retained after it is first displayed, it cannot be selectively erased or replaced. Response times are therefore longer than with conventional CRTs, and interaction is not usually possible. Also, large storage tubes are difficult to manufacture, they have inferior contrast characteristics, they are not very good at drawing lines, they cannot display grey scales and they have lower light levels. Despite these drawbacks, however, storage tubes are convenient display devices. They are particularly convenient for graphics applications where the user requires limited interaction facilities (for example, in some computer-aided design applications).

Alternative display technologies

Although CRT technology will remain the dominant display technology throughout the 1980s, it does have some disadvantages. It requires high power consumption and high voltages, and it has undesirable heat dissipation characteristics. In addition, the video display suffers from distortion. Despite these disadvantages, newer display technologies are finding it difficult to displace the CRT. The difficulties stem partly from the limitations of the newer technologies (which stem partly in turn from their early stage of development) and partly from their inability to share the production set-up costs with a mass consumer market. To successfully displace the CRT, a new display technology must have the following characteristics:

— It must be usable in televisions.

— It must operate in colour.

— It must be as cheap as a CRT, whilst also offering significant technical advantages over a CRT.

At present, the most promising alternative display technology is the gas plasma panel. Gas plasma panels operate by passing an electrical discharge through a gas trapped between glass panels, thereby causing the gas to luminesce.

Two types of gas plasma technology are available. The first of these technologies is similar to a conventional CRT in that it needs refreshing regularly from a refresh buffer. The second technology uses a panel that retains its image until it is either modified or deleted. This second technology, however, does not suffer from the disadvantage of storage tubes, because the image can be selectively erased or replaced. This second type of gas plasma display technology therefore has the potential to become a general-purpose display technology.

Plasma display panels have the following advantages compared with conventional CRTs:

- They are small, and a typical 25 x 80 character display of the non-refresh type is about three inches deep. This depth might make plasma display panels more acceptable than conventional terminals in an office environment.
- They require only low voltages.
- They dissipate only small amounts of heat, and therefore they do not need a fan to assist in dissipating the heat.
- They have flat screens, and this feature eliminates the edge distortion that occurs with CRTs.
- They do not require refresh logic and memory.
- They have translucent screens that permit the rear projection of images through the panel.
- They provide an image that is almost entirely free of flicker.

The present plasma display panels, however, have two disadvantages. The first is that they provide only limited colour facilities. Only orange coloured screens are commercially available at present, although several other colours have been developed. Also, Fujitsu claims to have developed a colour plasma panel using red/green/blue technology. The second disadvantage of plasma display panels is their high cost. Large-size plasma displays are at present more expensive than the equivalent CRTs, although as the cost of CRT displays stabilises, plasma panels will be able to compete with CRTs on price.

Other alternative display technologies have been developed, or are under development. Some of these technologies, such as light-emitting diodes and liquid-crystal displays, are unlikely ever to compete with CRTs or plasma panels for use in large-size displays. Other technologies, such as electrochromatic displays, are still at an early stage of development.

There is, however, one application for small displays that may prove to be important. This application is the use of 'soft' keys that can be programmed either to represent users' choices or to allow users to make menu choices. A possible implementation of this concept would be to use the keys also as limited displays. The only display technology that has been used to implement this concept is liquid-crystal display technology. In addition to acting as a limited display, the crystal acts also as a capacitive touch-sensitive keyswitch.

Colour displays

Colour displays are at present based on raster-scan refreshed CRTs. They are similar in principle to standard television tubes, where each point on the screen consists of three dots of phosphor — one each for the primary colours of red, blue and green. These phosphor dots are energised by three individual electron guns. The individual dots can be energised separately, and so a wide variety of colours of varying luminance can be displayed.

The major problem with the use of refreshed-CRT technology for colour displays is that the refresh buffer has to contain additional information to code the colour attributes. This problem is particularly severe for graphics applications, where individual picture elements may need to be separately stored. For example, twelve bits of coded information may be necessary to define the colour and the other attributes of each picture element. If a display is to resolve an array of 512 x 512 picture elements, a memory of about 3 million bits is needed (compared with 2,000 bits for a similar-sized monochrome character-only display). For this reason, colour displays are expensive because the additional memory costs between £5,000 and £8,000.

Simple colour-coded systems, such as videotex systems, avoid the need for large amounts of memory by severely limiting the colour content of their displays. The Prestel system, for example, uses only two values for each electron gun (either on or off), and, therefore, only eight possible colour values are available. Eight colours may well be sufficient for simple alphanumeric applications (the IBM 3790 colour text terminal, for example), but they are not enough for the graphic applications in which most colour displays are used.

However, as memory prices continue to fall, the cost penalty of including large amounts of memory in a terminal will become less severe. In addition, most television industry modules will be oriented to colour devices in the near future. These two factors suggest that, in the near future, colour terminals will be a practical proposition for everyday commercial use.

PRINTERS

Several types of printing technology are employed by the print devices that are a constituent part of most terminals. In this section of the report we classify printers under the headings of line printers, impact printers and non-impact printers.

Line printers

The traditional line-printer technologies of revolving drums, revolving chains and reciprocating trains are still extensively used, but more recently Tally has invented the comb mechanism. This mechanism uses a comb with 132 teeth, each of which has a steel ball on its tip. The characters are formed from a matrix of dots, with seven lines of dots forming a row of characters. The dots are printed by the impact of the balls striking the print ribbon onto the paper. All 132 print positions are struck at the same time, and the appropriate teeth of the comb are moved by an electromagnet. The advantages of the comb mechanism are that it is both simple and reliable. Line printers that use this technology are cheaper and more reliable than fully-formed character printers that operate at similar speeds.

Impact printers

The original type of impact printer was the IBM Selectric golf-ball mechanism, and it has become the standard for conventional typewriter mechanisms. Until 1972, no alternative to the Selectric mechanism was available for printing letter-quality documents. In that year however, Diablo Systems invented the 'daisy wheel' mechanism with their HyType 1 printer. The daisy wheel mechanism provides several advantages compared with the Selectric mechanism, including:

— Better reliability.

— Lower noise level.

— Higher printing speed (55 characters per second compared with the Selectric's 15 characters per second).

— Better control of print and of paper position.

Several mechanisms similar to daisy wheels have been made available recently. For example, Perkin-Elmer's printers use a print cup. The side of the cup is made up of a number of print fingers, and the character shapes are embossed on the top of the fingers. Other manufacturers use daisy wheels with several concentric rings of characters embossed on the wheel.

The most common type of impact printer now available with terminals is the matrix printer. These printers are similar to both the Selectric mechanism and the daisy wheel mechanism in that they print one character at a time. However, they form each character from a matrix of dots rather than by striking an embossed character shape onto the print ribbon. The performance of matrix printers has been improved by the recent introduction of electronic control, and also by the addition of a second print head. As a result, printing speeds of up to 400 or 500 characters per second can now be achieved.

There have been two other recent advances in impact matrix printer technology. The first advance is the use of a single line of needles, rather than a matrix. This single line of needles forms each character by moving along the line repeating a vertical row of dots. The print head on these devices is both cheaper and easier to replace.

The second advance is the use of the so called 'infinite matrix' principle invented by Sanders Technology Incorporated. This principle requires the print head and the paper to be accurately aligned, and it also requires the use of more electronic control. With the infinite matrix principle, characters are formed by repeating rows of print, and on each successive row the print head is slightly displaced from the previous row. The characters can be printed either in draft quality at high speed by using one pass of the head, or in letter quality at a slower speed using multiple passes of the head. Infinite matrix printers can therefore be used where there is a requirement both for printing draft documents at a fast speed and for producing the final documents to the quality of a typed letter.

Non-impact printers

Non-impact printing technologies include ink-jet printers, xerographic printers, thermal printers and electrostatic printers. All non-impact printing technologies share a common disadvantage in that they cannot produce simultaneous multiple copies of the output. Ink-jet printers and xerographic printers use normal paper, but the other technologies require more expensive special papers.

Ink-jet printers operate by ejecting ink jets selectively from a multiple nozzle print head. They are very quiet in operation, and they are able to produce very high quality output. Xerographic printers operate on the same principle as an office copier. They are very fast, and they can achieve speeds of up to 20,000 lines per minute. One advantage of xerographic printers is that special pre-printed stationery is not required, because forms can be overlaid onto the output as it is printed.

Thermal printers use the reaction of special paper to heat, and characters are generated by heating selected probes in the print-head matrix. Thermal printers are small and light, and they are quiet in operation. In addition, they are much more reliable than are most other types of printer (particularly impact printers), and they are inexpensive.

Electrostatic printers operate by applying a voltage to a stationary linear array of conducting nibs. These nibs produce an electrostatic charge on the surface of a dialectically-coated paper. The charge is then developed by the use of a liquid toner to produce a high-contrast image of the required characters. Electrostatic printers are quiet and fast in operation. They also produce a good definition of image, because they use a high density of dots (typically 200 or 300 dots per inch) to form the image. These printers are therefore capable of producing very high-quality text output, which can be printed in a wide range of fonts that often can be selected whilst the printer is operating. Electrostatic printers also have another significant advantage in that, apart from the paper-handling mechanism, they have no moving parts.

THE USE OF SOFTWARE IN TERMINALS

Until quite recently, both the control logic and the functional attributes of terminals were implemented in fixed circuitry. As a result, to install any change involved redesigning the hardware and usually involved replacing either individual components or complete logic boards. However, the advent of the general-purpose microprocessor and the reduced cost of solid-state memories have made it economically feasible to implement terminal control logic and terminal functional attributes in software.

Terminal functions are implemented by programs stored in read-only memory modules, rather than in fixed circuitry. The modules are programmed either by the manufacturer, who produces a standard read-only memory, or by the system builder. In addition, the user may program the terminal controller or the cluster controller, either direct or by down-line loading programs from the host computer. The intelligence that this approach provides to the terminal is made available to the terminal user at three levels (as we discussed in chapter 2 on pages 3 and 4).

The implementation of control logic and terminal functions in software provides several advantages, of which the five most important are:

1. The terminal is provided with a high degree of functional flexibility because new or changed features can be implemented relatively easily in the terminal software. These new or changed features may be required to meet either changing market conditions, or specific customer requirements, or changing communications environments or different interface specifications. One example of the functional flexibility that this approach provides is the speed with which IBM implemented the X.21 interface option on its terminals after that interface had been agreed.

2. The terminal manufacturer can add new capabilities to his product at a low incremental cost.

3. The manufacturing process is simplified because most of the functional variations of a terminal exist only as software.

4. The life of a terminal product can be extended by changing the software several times to stretch the original basic design.

5. The design stage of a terminal product is simplified, and it is faster and cheaper than the design stage of a hard-wired product.

The range of functions that are implemented in software in display terminals is very wide, and it includes:

— Editing features.

— Forms control features.

— Video display features, such as variable intensity, reverse video, and the limited use of colour.

— Video control features, such as the local storage of several pages, scrolling facilities, and the separate control of several independent areas on the screen.

— Communications control features, such as the use of several line protocols.

— Multiple font features.

— The personalised setting of function keys, such as the action the terminal will take when the 'return' key is depressed.

— Simple line-graphics and block-graphics features.

COST TRENDS

The previous sections of this chapter have examined the technology of terminals as it relates to keyboards, alternatives to keyboards, displays, printers and software. In this final section we now examine the underlying cost trends of the various technologies, and the impact those trends are likely to have on terminal devices.

The continuing decrease in the cost of microelectronics has been well documented in Foundation Report No. 15, and, as the electronic content of terminals increases, this decreasing cost of microelectronics will, of course, have an impact on the economics of terminals. Terminal devices, however, consist of more than just electronics, and the costs of some of the non-electronic components of a terminal are increasing rather than decreasing. The net result is that although the average cost to the user of a display terminal has decreased from $2,150 in 1974 to $1,450 in 1980, the rate of decrease is quite modest (often being less than 10 per cent per annum), and it is slowing down. The deceleration in the rate of decrease of the cost of manufacturing a display terminal is even more marked. In fact, since 1978, the manufacturing cost has stabilised at about $680.

In order to measure the trends in the costs of manufacturing terminals we defined a hypothetical terminal device. This hypothetical terminal consists of an Intel 8080 microprocessor, 32k of read-only memory, 32k of programmable read-only memory, circuit board mounts on which the various chips and supporting elements are mounted, a power supply element, a CRT display, a keyboard and a cabinet. With the assistance of several terminal manufacturers we have calculated the cost of manufacturing this hypothetical terminal in 1974 and in 1980. The results of these costings are shown in figures 4 and 5. Figure 4 shows the manufacturing cost in dollars of each of the elements of the hypothetical terminal in 1974 and 1980. It shows that, whilst there have been dramatic reductions in the costs of the elements based on microelectronics (the memory elements and the microprocessor), there have been substantial increases in the costs of the board mounts, the power supply elements, the cabinet and the labour costs.

Figure 5 shows the way in which the percentage of the total cost of the terminal accounted for by each element has changed. For example, in 1974 the power supply

Figure 4 Manufacturing cost (in dollars) of terminal elements in 1974 and 1980

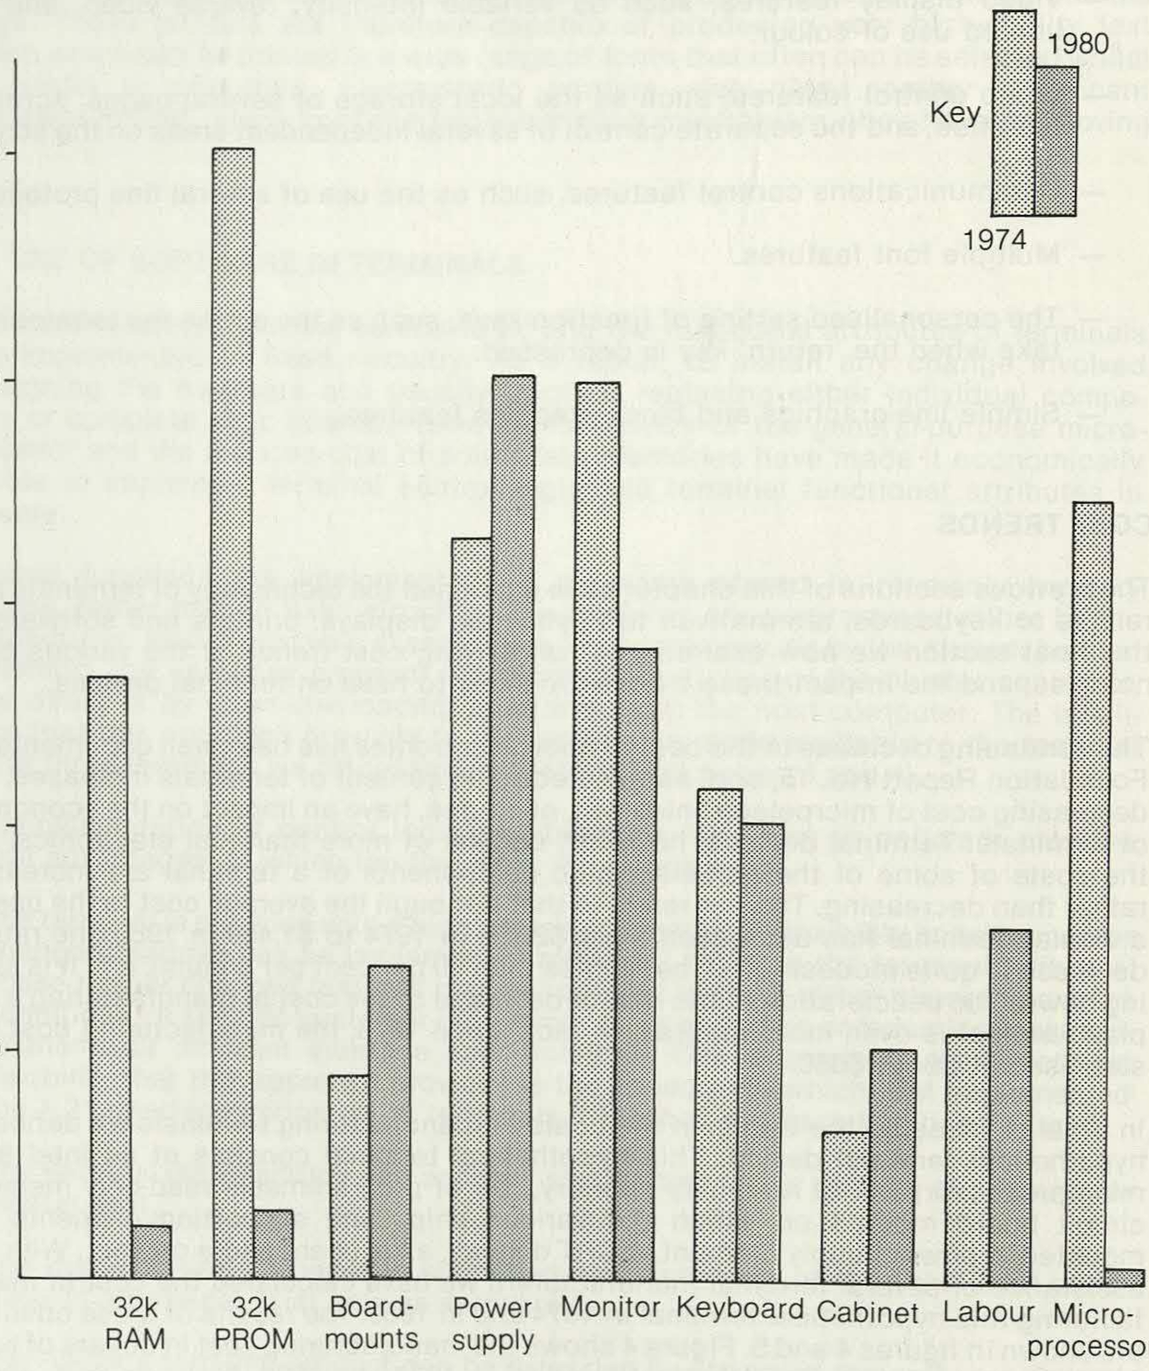

element accounted for 14.2 per cent of the total manufacturing cost, whereas in 1980 it accounted for 29.9 per cent. On the other hand, the microprocessor accounted for 15.1 per cent of the cost in 1974, but it accounted for less than 1 per cent of the total manufacturing cost in 1980.

Figure 5 Percentage of total terminal cost accounted for by each element in 1974 and 1980

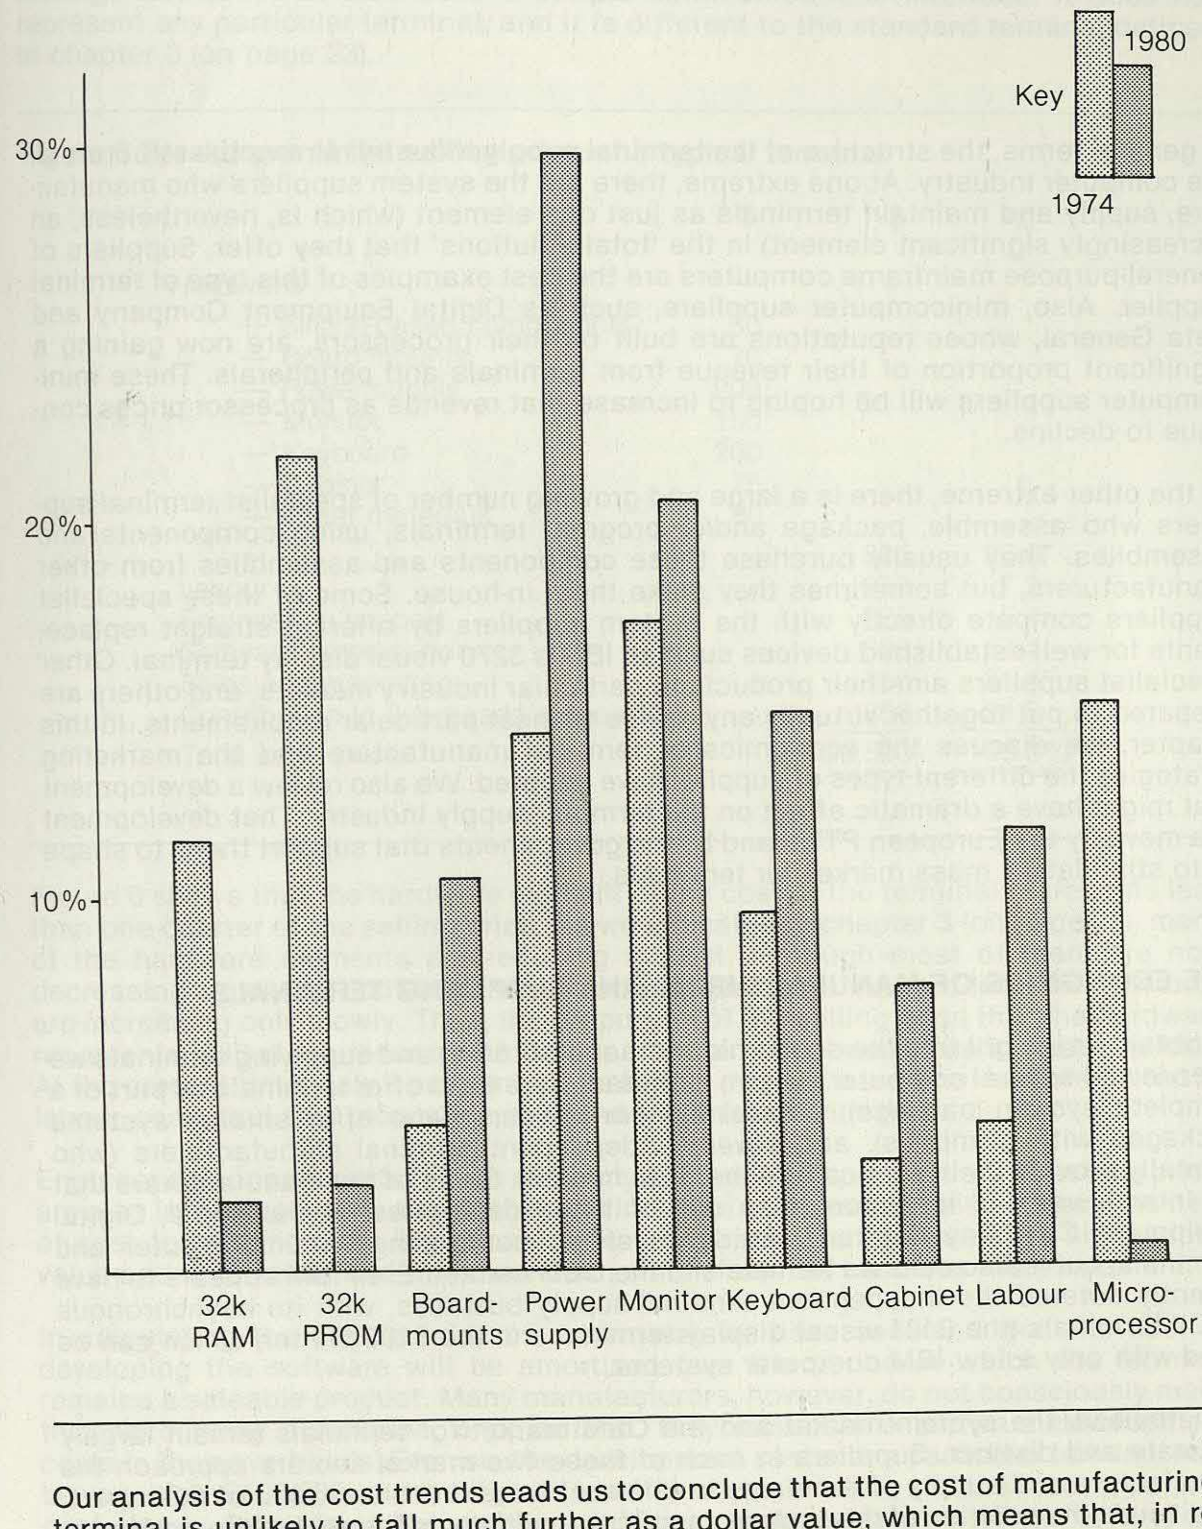

Our analysis of the cost trends leads us to conclude that the cost of manufacturing a terminal is unlikely to fall much further as a dollar value, which means that, in real terms, terminals will become increasingly cheaper. Manufacturers will, however, continue to add additional features and facilities to their terminals.

CHAPTER 4

THE TERMINAL SUPPLY INDUSTRY

In general terms, the structure of the terminal supply industry mirrors the structure of the computer industry. At one extreme, there are the system suppliers who manufacture, supply and maintain terminals as just one element (which is, nevertheless, an increasingly significant element) in the 'total solutions' that they offer. Suppliers of general-purpose mainframe computers are the best examples of this type of terminal supplier. Also, minicomputer suppliers, such as Digital Equipment Company and Data General, whose reputations are built on their processors, are now gaining a significant proportion of their revenue from terminals and peripherals. These minicomputer suppliers will be hoping to increase that revenue as processor prices continue to decline.

At the other extreme, there is a large and growing number of specialist terminal suppliers who assemble, package and/or program terminals, using components and assemblies. They usually purchase these components and assemblies from other manufacturers, but sometimes they make them in-house. Some of these specialist suppliers compete directly with the system suppliers by offering straight replacements for well-established devices such as IBM's 3270 visual display terminal. Other specialist suppliers aim their products at particular industry markets, and others are prepared to put together virtually any device to meet particular requirements. In this chapter, we discuss the economics of terminal manufacture and the marketing strategies the different types of supplier have adopted. We also review a development that might have a dramatic effect on the terminal supply industry. That development is a move by the European PTTs (and by the governments that support them) to shape or to stimulate a mass market for terminals.

THE ECONOMICS OF MANUFACTURING AND SUPPLYING TERMINALS

To obtain an insight into the economics of manufacturing and supplying terminals we interviewed several computer system manufacturers (who offer terminals as part of a complete system package), several system builders (who offer smaller systems packaged with terminals), and several independent terminal manufacturers (who normally provide their terminals to the OEM market). Some of the manufacturers that we interviewed fall into more than one of these categories. For example, Digital Equipment Company Limited provides systems comprising a minicomputer and terminals, but it also sells its terminals to the OEM market. Even IBM appears to have recently entered the independent terminal-supply business, with its asynchronous ASCII terminals (the 3101 visual display terminal and the 3102 printer) which can be used with only a few IBM computer systems.

Nonetheless, the system market and the OEM market for terminals remain largely separate and distinct. Suppliers in each of these two market sectors approach the manufacture and supply of terminals in different ways. In this section of the report, we discuss the position of the system suppliers, and then, in the section beginning on page 29, we describe the differences that a user (or a purchaser) can expect to find in the terminals provided by an OEM supplier.

From our discussions with terminal manufacturers we were able to determine how the sales price of a typical terminal manufactured by a system supplier is arrived at, and this information is shown in figure 6. The figure shows the breakdown of the selling price of a typical terminal that incorporates a screen, a keyboard, some local intelligence, and that also uses a simple communications interface. It does not represent any particular terminal, and it is different to the standard terminal defined in chapter 3 (on page 23).

Figure 6 Breakdown of the selling price of a typical terminal

	$	$	%
Hardware			
— Silicon-based components	50		
— Board mounts	70		
— Power supply	200		
— Monitor	150		
— Keyboard	200		
— Cabinet	75		
— Other materials	80		
		825	23.6
Labour		250	7.2
Engineering support		225	6.4
Software amortisation		600	17.1
Sales and marketing		750	21.4
Contribution to overheads and profit		850	24.3
Total sales price		$3,500	100.0%

Figure 6 shows that the hardware element of the cost of the terminal represents less than one-quarter of the selling price. As we indicated in chapter 3 (on page 23), many of the hardware elements are reducing in cost, although most of them are now decreasing slowly. In addition, those hardware elements that are increasing in cost are increasing only slowly. Thus, the proportion of the selling price that the hardware represents is likely to reduce with time, particularly if labour costs increase quickly. At the present time, only 7 per cent of the selling price of a terminal is represented by labour costs, but this proportion is bound to increase with time.

Engineering support costs, which currently account for 6 per cent of the selling price, are also labour intensive, and so they will increase with time. We also expect this item of cost to become more significant as the differentiation between the higher-cost, value-added terminal and the lower-cost, simpler device becomes more apparent.

In calculating the selling price of a terminal, we have assumed that the cost of developing the software will be amortised over the period in which the terminal remains a saleable product. Many manufacturers, however, do not consciously make this accounting distinction, and, instead, they include the software development costs in their overheads. Even so, the development costs for a particular terminal may be very high (£100,000 was mentioned to us as a typical cost), and the recovery of this cost clearly depends on the number of devices that are sold. At 17 per cent of the selling price (as shown in figure 6), software development is an important element, and it is likely to represent an increasing proportion of the total cost.

A surprisingly large proportion of the selling price of a terminal is attributed to the sales and marketing effort. This cost element will also increase with time, because it is labour intensive. In addition, as the marketplace becomes more competitive, the sales and marketing effort may become even more important and necessary.

POTENTIAL STRATEGIES FOR TERMINAL SUPPLIERS

By examining the breakdown of the selling price of the typical terminal, we have identified the four strategic options that suppliers of terminals can adopt to retain their market positions whilst continuing to make a respectable profit, and we now discuss these four options.

1. Attempting to continue to reduce hardware costs

The supplier can attempt to reduce further the cost to himself of the hardware element of terminals. Although our discussion of the cost trends in chapter 3 (on pages 23 to 25) showed that further cost reductions would have little effect on the overall manufacturing cost of a terminal, there are two ways in which a supplier might achieve some additional cost reductions. The first is to plan for longer production runs, which would enable the production costs, the development costs, and the set-up costs to be spread over a larger number of devices. If this policy proved to be successful for a particular supplier, he would also be able to reduce the selling price of the terminal over a period of time.

The second way of reducing hardware costs is for the supplier to produce some of the components in-house. Some suppliers, in fact, have already followed this policy successfully. One supplier told us that in-house production of selected components could reduce the manufacturing cost of a terminal by as much as 35 per cent. However, for the typical terminal defined above, this reduction would reduce the final selling price by only about 8 per cent. In addition, the supplier would need to exercise great care in selecting the components to be produced in-house. For example, a supplier would gain very little from setting up a microprocessor chip manufacturing plant.

2. Offering standardised software

Figure 6 also shows that one of the largest elements of the selling price of a terminal is the cost of the software. But because the cost of developing software is independent of the number of terminals sold, the software element included in the price of a single terminal can be reduced by selling more terminals. Alternatively, the same software (or a slightly modified version of it) can be incorporated in different hardware packages. Many suppliers of electronic calculators have adopted a similar strategy. They offer several calculators that have slightly different facilities but all use the same chip and the same software.

3. Reducing both the sales costs and the marketing costs

About 20 per cent of the selling price of a typical terminal is accounted for by the sales and marketing costs. There is considerable scope for reducing these costs either by reducing the sales force, or by using a dealer network, or by establishing a retail sales network.

The sales force can be reduced if the supplier concentrates more on the OEM market for terminals. IBM's policy of selling its new asynchronous ASCII terminals through its Direct Marketing Centre is one example of this approach.

If the supplier chooses to use a dealer network, the dealer is responsible for the sales costs, and the supplier has merely to set up a dealer network wherever he wishes to sell the terminal. This approach, has the disadvantage that the supplier loses control of his products, and many suppliers have avoided it for this reason. Nevertheless, several suppliers, (and particularly the independent suppliers) have shown that this can be a profitable policy.

Several suppliers (including IBM, Xerox and Digital) have set up a retail sales network. A large investment is needed to establish such a network in the first place, but experience in the United States shows that the investment may well be justified. For the smaller supplier, an alternative is to take a concession in an existing chain of shops.

4. Adding value to the product

This last strategy option is the one that suppliers most widely adopt. As the expectations of terminal users grow, the supplier can provide increased facilities in his product at very little additional cost to himself. Moreover, he can sell a new model of his product that incorporates these facilities either at the same price as that of the previous model of the product or at a slightly lower price. Almost all of the successful terminal suppliers have implemented this policy, and we expect that the trend of providing improved capability for a small change in the price will continue. However, this policy can increase the software development costs, and it can require additional sales effort.

THE POSITION OF THE OEM SUPPLIERS

The OEM suppliers' policies regarding the manufacturing, selling and maintenance of terminals have several major differences from those of the system suppliers. We summarise those differences in figure 7, and we also discuss them briefly below.

The OEM supplier usually sells his terminals to the market through dealers, whereas the system supplier often sells his products direct to the end user as part of a system. Because the system supplier sells to the end-user market, and because the market he operates in is changing rapidly, he has to offer a rental option for his products. Many of his customers (typically 60 to 70 per cent) will, in fact, rent rather than buy terminals. The disadvantage of this policy is that when the user decides to stop renting the product, the system supplier has on his hands a secondhand product to sell. For this reason, the system supplier is more likely to pay attention to the life-cycle costs of his product than is the OEM supplier.

A similar difference in policy arises when suppliers consider the reliability of terminal devices. The system supplier typically maintains his own terminals along with the rest of the system. By comparison, the OEM supplier maintains his own terminals only if he cannot avoid doing so. Usually, his terminals are maintained either by a system supplier or by an independent maintenance company. When the system supplier calculates his selling price, he must, therefore, take account of the additional cost of maintaining terminals.

However, an OEM supplier's reputation depends much more on the success of an individual terminal product than does a system supplier's reputation. The reason for this is that the sales of terminals form a substantial part of an OEM supplier's business. For this reason, the better-selling OEM terminals are good quality, reliable products.

The major difference that the user sees is in the price of the terminals. If he is able to purchase volume quantities direct from an OEM supplier, rather than through the OEM's dealers, he will achieve a significant cost saving compared with buying terminals from his system supplier. On the basis of the typical terminal costed in figure 6, an OEM supplier could well offer the same terminal for perhaps half a system supplier's price.

Figure 7 The major differences in the policies of OEM suppliers and system suppliers

Policy issue	OEM supplier	System supplier
Supply policy	Usually sells terminals	Very often rents or leases terminals
	Usually sells terminals either through dealers or by heavy advertising	Usually sells terminals through an in-house sales force
	Usually fixes a selling price between 30 and 40 per cent higher than the production cost	Usually fixes a selling price between 3 and 4 times the production cost
Attitude to life-cycle costs	Does not need to consider product life-cycle costs, because, once sold, the terminals are no longer the supplier's responsibility	Needs to consider product life-cycle costs, particularly the cost of maintenance and the likely secondhand re-sale value
Maintenance policy	Does not maintain terminals	Maintains terminals
Installed base	Has installed a large number of general-purpose terminals	Has installed a small number of special-purpose or general-purpose terminals
Significance of terminals to the business	Has a business that depends wholly (or largely) on the sales of terminals	Has a business that depends only to a small extent on terminals, because the supplier sells terminals only as components of systems
Production policy	May produce some components in-house	Rarely produces components in-house
	Utilises some cheaper components	Utilises only high-quality components to keep down the likely support costs

PLUG-COMPATIBLE TERMINALS

The market for plug-compatible terminals shows the same pattern of competitive move and counter-move as do the markets for plug-compatible disc and tape drives and plug-compatible processors. The plug-compatible (PCM) terminal supplier must be able to provide products that:

- Emulate the system supplier's terminal closely enough for the purposes of the particular application.
- Offer features or facilities that give his devices a competitive edge.
- Match counter-moves that the system supplier will make in response to the PCM's product.

To make life difficult for their PCM competitors, the system suppliers attempt to ensure that they offer a product that is a moving target as far as emulation is concerned. As the market leader, IBM is the prime target for PCM competition. (For example, there are more than twenty suppliers of alternative versions of the 3270 visual display terminal in Europe.) IBM, however, is adept at providing products that present difficulties to its PCM competitors. In particular, IBM includes features in its products that it does not announce initially, but then reveals later. For example, IBM announced the enhanced video capabilities of its 3278 terminal when it introduced the 3279 colour terminal and the new version of the 3274 cluster controller to support the 3278. In practice, these capabilities work with existing 3278 terminals when the new controller is used. PCM terminals might not have this facility, although most do.

Many of the PCM's problems are becoming more difficult, rather than easier, to come to terms with. With advanced line protocols and the newer network architectures, it is more difficult to emulate the appropriate devices correctly. As a measure of the complexity, we summarise in figure 8 the approximate size of the software needed to emulate various terminals.

Some of the 30k bytes of software required to emulate the SNA 3276 terminal is needed to handle the complex microcode necessary for the SDLC line protocol. However, much of the software is also needed to handle other aspects of SNA (such as error recovery, encryption of data and network management features). Because it is so difficult to handle all of these features (particularly in a stand-alone terminal device), many PCMs limit themselves to emulating just the terminal functions.

It is clear, then, that when a user evaluates a terminal system that attempts to emulate another manufacturer's equipment, he should be careful to ensure that it incorporates all the features he needs.

In some circumstances (such as when terminals need to communicate with more than one type of host computer), there will be no alternative to emulation. This situation arises in particular industries (for example, in airlines). In these industries, the problems of emulation are amplified, because the amount of software needed to handle concurrently several different communications protocols is more than the sum of the individual amounts of software needed to handle each protocol individually.

PCMs also face a dilemma when they design the user interface of a terminal that will emulate several different terminals. Different terminals may, for example, use function keys or other terminal features in different ways. The designer has to choose between carrying these differences through to the user (at the risk of confusing the user), or designing a terminal that is a hybrid of those emulated, and mapping all the variants onto this hybrid. The latter approach is likely to produce the best result, but it is more challenging technically.

Figure 8 Size of software required to emulate various terminals

Terminal type	Size of emulation software (in bytes)
Teletype 33	4k
Dec VT52	4k
IBM 2741	5k
IBM 2780/3780	9k
Burroughs TC/TD	12k
IBM 3270 BSC	13k
Univac V100/200	16k
ICL 7502	16k
IBM 3275 BSC	16k
IBM 3276 SNA/SDLC	30k

(Source: *Data Communications,* November, 1980)

THE MASS-MARKET APPROACH

Either directly or indirectly, most PTTs have been in the terminal business for some time as suppliers of telex terminals and, in some cases, as manufacturers of them. AT&T, through its subsidiary, the Teletype Corporation, also supplies general-purpose keyboard terminals worldwide, and more recently it has launched a range of business terminals.

Of far greater potential significance than this direct participation by the PTTs in the market for business terminals are their recent attempts to create a mass market for terminals. These attempts are centred around the PTTs' videotex developments. A key element of the concept of viewdata as originated by the British Post Office (now British Telecom) was the opportunity to appeal to the mass domestic market. Since then, the videotex concept has been developed further by the French PTT with its plans to distribute to telephone subscribers a directory enquiry terminal in place of printed directories. Pre-trial batches of terminals have already been produced, and orders for 250,000 terminals have now been placed with two French manufacturing companies. Originally, a target ex-factory price of $100 was set for the directory enquiry terminal. We estimate that manufacturers will now incur an ex-factory cost of $105 (excluding any contribution to overheads and profit) for the planned large-volume production runs.

In terms of features and quality, the directory enquiry terminal will probably fall short of what is now expected of a business terminal. However, the dramatic cost reductions achieved as a result of the mass-market volumes could be passed on into terminals designed for business use, in the same way that television CRT technology has been used in visual display terminals. Also, lower levels of quality may well prove quite acceptable in terminals that are intended for infrequent, casual use. Low-priced devices will, of course, make it easier to justify large numbers of terminals for casual use, and thus will reinforce the existing trend towards using terminals less intensively.

Apart from France, where this bold experiment is being conducted, the only country in which videotex terminals are available in significant numbers is the United Kingdom. At 1980 prices, the additional price for videotex in a full-size colour television was $1,000. In other words, videotex doubled the normal price. We estimate that volume production (hundreds of thousands of terminals) will bring down the additional cost of building-in the videotex adaptor and the integral modem to $115, including overheads.

Several purpose-designed Prestel business terminals with a full alphanumeric keyboard are also available. These terminals, typified by STC's Novatel and Pye TMC's Visa, sell at between $1,000 and $1,500, a price range similar to the price range for teletype-compatible visual display terminals ('glass' teletypes). Again, volume production of these business Prestel terminals would have a significant impact on their price. Figure 9 shows a breakdown of the ex-factory costs for both small-volume and medium-volume production runs.

The cost of 'downmarket' visual display terminals is now approaching the small-volume production run cost for a Prestel terminal shown in figure 9. However, such visual display terminals do not, at this stage, compare with the medium-volume cost shown in the figure, which, assuming that the sales and support costs are minimal, could result in a sale price as low as $550. The sales and support costs could be minimised if Prestel business terminals were distributed and maintained by established outlets, such as television rental companies and electronic equipment retailers.

In addition to the potential difference in sales volume, the difference in ex-factory cost between the Prestel business terminal and our typical terminal also reflects the deliberate adoption of minimum-cost technology for Prestel terminals. The use of minimum-cost technology may well be an essential element in any supplier's strategy that aims to achieve a mass-market sales volume for terminals. Once the mass market has been established, extra features can be added later when high-production volumes permit this addition at minimal additional cost. The way in which the market for pocket calculators was first created, and is now being developed, provides an excellent example of this kind of strategy.

Thus, although many imponderables remain, videotex and similar developments could quickly revolutionise the market for general-purpose terminals. The French PTT has no doubts about this, at least none that it has expressed publicly. The French PTT expects to have four million directory enquiry terminals installed by 1985 and 30 million by 1990. The directory enquiry terminal or derivations of it, may not compare in terms either of engineering, or features, or reliability with its established competitor — the general-purpose visual display terminal. Like the microprocessor, however, its low cost seems certain to make the videotex terminal an inevitable part of the business scene.

The videotex terminal may also make a dramatic impact on the pace of development of office systems. For example, an organisation wishing to equip one in four of its 20,000 employees with a visual display terminal would, at the medium-volume production cost, be faced with a bill of $2.75M for Prestel-based terminals compared with $4.5M for glass teletypes. In our view, therefore, it is potentially as dangerous to ignore the videotex terminal as it is to ignore the microprocessor.

Figure 9 Ex-factory costs (in dollars) for Prestel business terminals

Cost item	Small-volume production	Medium-volume production
Modem/interface	80	40
Logic	95	30
Power supply	20	10
Monitor	65	50
Keyboard	40	30
Cabinet	20	10
Labour	80	40
Direct costs	400	210
Contribution to overheads	400	210
Ex-factory price	800	420

(Source: Butler Cox & Partners Limited)

CHAPTER 5

COMMUNICATIONS PROTOCOLS AND STANDARDS

We now turn from the terminal as it appears to the user and the buyer, to the terminal as it appears to the network and the staff who design and operate the network. We look, in other words, at the protocols and standards that govern communication.

We begin with a general discussion of the structure of protocols, and we then examine the way in which specific communications protocols have evolved. This evolution has been dictated largely by the computer suppliers and the terminal suppliers, who have dominated data communication from the outset. The PTTs, of course, supply the transmission facilities used for data communication, but only recently have they begun to influence the methods used for data communication. We continue, therefore, with an assessment of the impact of new PTT services, and we then review the impact that several new communications products will have on communications standards. Finally, we review the activity of the various standards authorities.

PROTOCOL STRUCTURES

Increasingly, communications protocols are defined in terms of levels, with each level containing a specific set of functions. Each level of protocol requires a corresponding layer either of hardware or software in the terminal. Consequently, it is necessary to understand protocol structures in order to be able to understand the workings of a terminal, and, in particular, to be able to appreciate the dimensions of the compatibility problem.

To illustrate protocol structures, we have adopted the International Standards Organisation's (ISO) Open-Systems Interconnection (OSI) reference model. We have done so not because the OSI model is universally appropriate (and indeed it is not yet finalised), but because it does present a comprehensive picture, and so it can demonstrate the full dimensions of the terminal compatibility problem. In reality, many present and future applications involving communicating terminals will not implement all of the levels of protocol in the OSI model. For some applications, not all of the levels are needed, because two or more separate levels can be combined into one level for simplicity and efficiency. For our purposes, we have combined level 3 of the OSI model (the network level) and level 4 (the transport level) into one level, which we call the 'transport' level.

Members of the relevant ISO committee have proposed that these two levels should be combined. It can be argued that a separate transport level is needed because existing network levels (such as in CCITT's X.25 packet interface) only partially serve the required purpose. It follows that it would be better to define a transport level that displaces the network level completely, rather than to patch a new transport level over the top of the existing network level.

Figure 10 shows our transport level (3/4) and the remaining five levels in the OSI model, in which level 1 is the lowest level. Against each level in the figure we give a summary of the function of the level and also an indication of the most likely effect of limited compatibility at that level. These effects will result if a particular level of protocol in a terminal is not fully compatible with its counterpart layer in the terminal or computer system with which it is communicating. For the lower protocol levels, the same effects will result from incompatibility with the communications network to which the terminal is attached. In figure 10, we have shown the effects of limited compatibility. Full incompatibility, clearly, will mean that nothing useful will be achieved at all.

Figure 10 Protocol structure and the effect of incompatibility

Protocol level		Function	Effect of limited compatibility
Number	Name		
1	Physical	Transmits a bit stream between devices	Poor diagnostics
2	Data link	Transfers records across a communications link	Poor performance
3/4	Transport	Provides an end-to-end transport service for records	Poor performance and ineffective error handling
5	Session	Co-ordinates the dialogue between the communicators	Difficulty in accessing or changing between services
6	Presentation	Converts information into the appropriate format for output	Poor presentation of information
7	Application	Selects or provides the required service	Clumsy or confusing operating procedures

It is worth noting that limited incompatibility at the lowest three levels (physical, data link and transport) primarily affects performance, which can easily be quantified (for example, in terms of line costs). Limited compatibility at the three highest levels (session, presentation and application), and possibly at level 3/4 (transport), may result in a poor user interface. Here, the cost of incompatibility is not so easily measured (except in extreme cases where an application fails totally), but the cost could easily prove to be far more damaging than poor performance would be. For that reason, and because it is easy to overlook the deficiencies in the higher levels of protocol until the terminal is in use, higher level protocols deserve special attention when terminals are evaluated for purchase.

So far, standardisation has had most impact at the lower levels (X.25, for example, encompasses the lowest three levels in the OSI model). Standardisation of the higher protocol levels will be more difficult to achieve, both because requirements are more diverse and because many software elements may be involved. For example, programming languages and their compilers may make assumptions about the way terminals work.

These assumptions might be about the way a screen is formatted, or about the use of a function key, etc. Applications programs also may well make similar assumptions, and often these programs will make explicit references to features of the particular terminal for which they were designed. Effective standardisation of the presentation level may therefore depend on programming language standards, and indeed on the programming standards adopted in a particular installation.

THE EVOLUTION OF PROTOCOLS

Three main stages are apparent in the evolution of protocols for communicating terminals — asynchronous protocols, basic-mode synchronous protocols, and bit-synchronous protocols. In this section of the report we first describe these three types of protocols, and then provide advice on how an organisation should select a protocol for a particular application.

Asynchronous protocols

The term 'asynchronous', when applied to data communication, refers to the method used to frame data characters when they are output onto a transmission line. This framing is effected by sending a start bit and a stop bit that respectively precede and follow each character, and for this reason asynchronous transmission is sometimes referred to as 'start-stop' transmission. However, the term 'asynchronous' is commonly used to describe the teletype protocol (used by unbuffered keyboard terminals that use the ASCII character code), and we use it in this sense in this report.

As we have implied in the previous paragraph, asynchronous protocols were designed for the unintelligent terminals that are typically used in a conversational mode with timesharing systems. Asynchronous protocols incorporate limited error-checking facilities at both the physical level and the link level. At the physical level, error checking takes the form of a character parity bit and a longitudinal checksum. At the link level, messages are retransmitted when a negative acknowledgement is received. This retransmission is often combined with the technique known as 'echoplex', where characters are echoed back by the receiving device to the sending device, so that they can be displayed for a visual check by the terminal user. Asynchronous protocols are now widely used also by more intelligent buffered devices, which are often screen-based devices.

Asynchronous protocols use contention procedures. In other words, to initiate a dialogue a person enters a message (after dialling to establish a physical connection, if necessary) then waits for a reply. The dialogue then continues on a turn-and-turn-about basis. Asynchronous protocols can therefore be used only on point-to-point lines (typically with terminals connected locally to a computer system), or alternatively, with remote terminals that dial-up over the public telephone network.

Asynchronous terminals can also communicate via packet-switching networks, by connecting to a packet assembly/disassembly (PAD) facility in the node of a packet-switching network. The PAD effectively handles the conversion between asynchronous protocols and packet-switching (X.25) protocols. Using a packet-switching network, communication can take place only between an asynchronous terminal and a packet device (normally a computer system). Communication cannot occur between two asynchronous terminals. CCITT Recommendations X.3, X.28 and X.29 specify the rules for the terminals, for the network and for the host computer systems that implement this arrangement (as illustrated in figure 11).

Figure 11 CCITT's X-series Recommendations

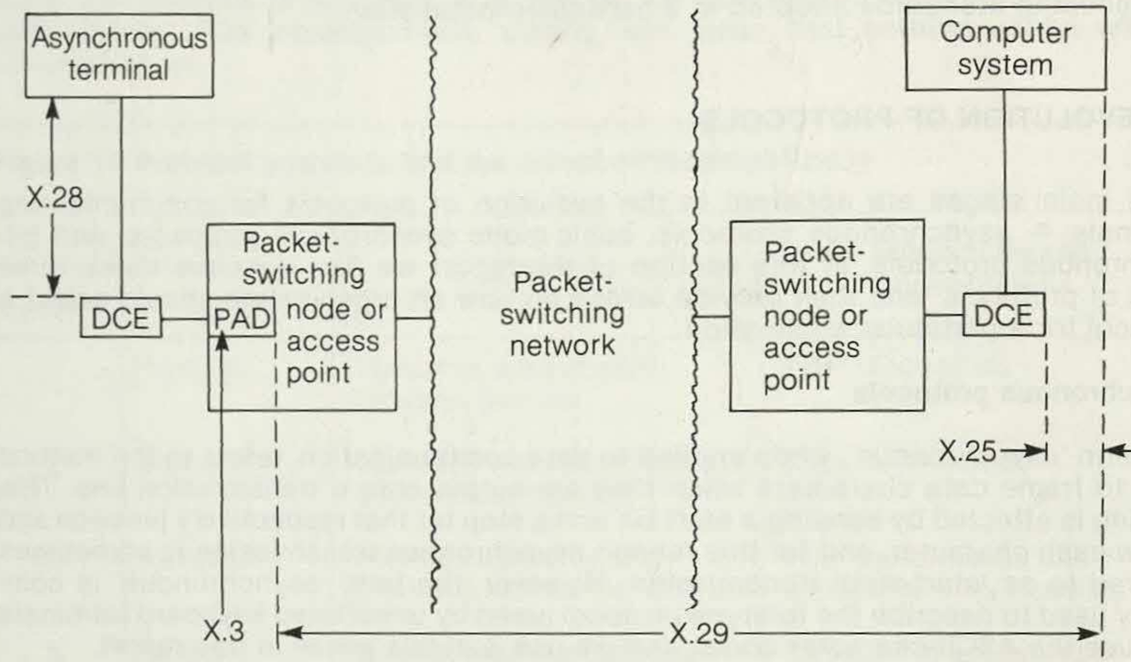

DCE = Data circuit-terminating equipment
PAD = Packet assembly/disassembly
X.3 = PAD facilities in a public data network
X.28 = Terminal/DCE interface for asynchronous terminals accessing a PAD
X.29 = Procedures for exchanging control information and user data between a packet-mode device and a PAD

Basic-mode synchronous protocols

'Basic mode' is the term the International Standards Organisation uses in its standard that describes the procedures for this type of protocol. In fact, that standard postdated IBM's introduction of Binary Synchronous Communications (BSC), sometimes referred to as Bi-synch, which has similar characteristics to the basic-mode standard defined by ISO. IBM's BSC comprises a family of protocols, the best known of which are 3270 BSC for visual display terminals and 2780 BSC for remote batch terminals. Because of IBM's dominant market position, these standards are closer to being de-facto standards than ISO's basic-mode is to being a de-jure standard. ISO's basic-mode standard has not been effective, because all the computer manufacturers have developed their own, mutually incompatible variants of the ISO standard.

The common features of basic-mode synchronous protocols are as follows:

— They are intended to be used for synchronous transmission between buffered devices, where each message is sent as a continuous stream of 8-bit characters, framed by synchronisation characters.

— The method of transmission permits higher line speeds (up to 9600 bit/s for terminals and occasionally higher).

— Error control is both more sophisticated and more effective than with asynchronous working.

— Poll/select techniques, which tightly control the receipt and the transmission of messages, permit multipoint lines to be used to reduce line costs.

Like asynchronous protocols, basic-mode protocols have code-dependent features. This means that certain characters are reserved for transmission control and for message framing functions. Although an 8-bit character code is assumed, binary data can be accommodated by a technique known as character-stuffing. This technique is used when an 8-bit field that occupies a character position in the binary data might be confused with a control character. In this situation, a special additional character is inserted before the field to give warning that the situation exists. This additional character is then stripped out when the message is received.

Bit-synchronous protocols

Bit-synchronous protocols, typified by ISO's High-Level Data Link Control (HDLC) standard and IBM's Synchronous Data Link Control (SDLC) protocol, differ from basic-mode protocols in the way in which they provide data transparency. Each message is enclosed in a standard envelope, and information within the envelope may assume any form, except that of the framing sequence. Where such a sequence occurs, a single bit is added to the message, and this bit is then stripped out when the message is received. The handling of this bit stuffing and bit stripping and also the handling of the envelope procedures and the transmission control procedures requires a certain amount of processing power in the communicating devices, but it is insignificant in terms of today's processing costs.

One further advance that is reflected in bit-synchronous protocols is a move from half-duplex working to full-duplex working. With half-duplex working, the communicating devices transmit alternately, whilst with full duplex working, both parties may transmit at the same time. These two modes of working should not be confused with the use of half-duplex or full-duplex lines. Full-duplex lines can be used to advantage with a half-duplex protocol like 3270 BSC, to avoid the need to turn the modem round between messages and also to eliminate some waiting time.

In practice, most, if not all, current implementations of bit-synchronous protocols for terminal control use half-duplex logic, both because full-duplex logic is difficult and expensive to program and because most terminal applications do not require full-duplex working. The logic of most terminal dialogues is, naturally enough, half-duplex. On the other hand, high-performance packet-switching networks (both public and private) use full-duplex protocols, and computer manufacturers also will use them for inter-computer or inter-node traffic.

The HDLC standard specifies several basic control functions and various classes of procedures within which certain of these functions can be used by one or both of the communicating devices. This enables the sophistication and the power of the devices to be matched with the procedures used. 'Unbalanced' communication procedures approximate to the basic-mode poll/select procedures used to control terminals on a multipoint line, while 'balanced' communication procedures are used for either high-throughput processor-to-processor communication or node-to-node communication.

Choosing a protocol

The evolution through the three types of protocol reflects a number of changes both in the application of terminals and in technology. In particular, it reflects:

— The declining cost of terminal logic and memory.

— The availability of, and the need for higher line speeds.

— The development of on-line data processing and the emergence of new patterns of traffic.

Figure 8 (on page 32), shows that the amount of memory needed to emulate different protocols is not small. In turn, this means that the additional cost of implementing the more advanced protocols is not yet insignificant — and it may never become so, since it involves more than just the cost of logic and memory in the terminal. Where applications genuinely benefit from the additional capability provided by advanced protocols, there clearly will be advantages (or else reductions in cost elsewhere, such as in communications lines) to set against the additional cost of the terminal equipment required to implement the protocols. Where that is not the case, the only motivation for an organisation to move to the more advanced protocols may be either pressure from the organisation's computer supplier or an awareness of the long-term benefits that the organisation might gain. This raises the question, of course, of whether there are indeed any such benefits for an organisation to pursue.

To answer that question it is necessary to consider again the classification of terminals by mode of use, which we gave in Chapter 2 on pages 7 and 8. Although we have presented the three types of protocol in terms of an evolution (which, chronologically speaking, they did undergo), the merits and demerits of each protocol type depend to a large extent on the type of traffic to be carried. A summary of the merits and demerits of the three protocols are set out below:

1. Asynchronous protocols
   Asynchronous protocols are best suited to those applications in which their limitations regarding error control and speed are not significant. In other words, asynchronous protocols are best suited to the class of applications we have termed 'conversational'. Asynchronous communication is a de-facto standard that is widely supported both by terminal suppliers and computer suppliers. In addition, it is supported by packet-switching networks. It can also cope (although with some difficulty) with on-line transaction processing, but it is not suitable for batch traffic.

2. Basic-mode protocols
   Different versions of basic-mode protocols have been developed for on-line traffic and for bulk-data traffic, as exemplified by IBM's 3270 BSC, and 2780 BSC protocols. These protocol structures were not designed to integrate those two forms of traffic, although some suppliers have made efforts to do so (for example, ICL's extended Basic Mode, XBM). The use of a basic-mode protocol is probably not necessary for most conversational applications, except where the conversational traffic is interleaved with on-line data traffic.

3. Bit-synchronous protocols
   Bit-synchronous protocols have been designed quite specifically to handle all forms of data traffic concurrently in an effective way. Packet switching, which is a specific implementation of these protocols, combines both the fast transit times

characteristic of on-line data and the throughput capacity demanded by bulk data with the switching capability that is often needed for conversational working.

Two barriers prevent these advanced protocols being widely adopted. The first, as already mentioned, is their adverse impact on terminal costs, which is most significant for the cheaper, asynchronous terminals used for conversational working. (Hence the provision of PAD facilities in packet-switching networks.) The second barrier is the need to convert existing applications that may have only just settled down to basic-mode protocols, where the result of the conversion may be marginal short-term gains and uncertain long-term gains. The size of the problem that these barriers represent can be gauged by the difficulty IBM has found in persuading its customer base to abandon the BSC protocols and take up the SDLC protocol.

THE IMPACT OF NEW PTT SERVICES

Until the late 1970s, the pace and the direction of developments in communications protocols was largely dictated by companies in the computer industry, and they continue to exert a major influence. In the last few years though, influence has been exerted from outside the computer industry, and particularly by the PTTs, who are in a position to exert considerable influence on the market through their service and their tariff policies. Most PTTs have introduced (or plan to introduce) public data networks and these networks will influence protocols up to the transport level (level 3/4). They also plan to introduce services that will influence the protocol levels above the transport level.

We have already mentioned packet switching, including PAD facilities for the connection of asynchronous terminals. It is probable that by 1983 or 1984 the majority of business premises in Europe will be within local-call distance of a packet-switching node or access point that will permit communication with any device that is attached to a packet-switching network throughout Western Europe and North America.

In addition to developing packet-switching networks, some PTTs are introducing circuit-switching networks (using the X.21 interface), teletex services and videotex services. We now review these three significant developments.

X.21 circuit-switched networks

Circuit-switched data networks using CCITT's X.21 standard are being introduced in several European countries, with the Nordic countries and West Germany leading the way. X.21 standards have the following advantages over the prevailing interface standards (V.24 and V.28) used for connection to the telephone network:

- They provide a common interface for different transmission speeds.
- They provide faster connect times than dial-up networks (0.2 to 0.5 of a second compared with 3 to 15 seconds for the public telephone network).
- They provide better error-handling facilities.
- They provide new features such as closed user groups and automatic call redirection.

The X.21 standard can be used both on switched public networks and on private leased lines. A bridging standard (X.21 bis) has been defined to permit V.24 terminals to access X.21 networks, but X.21 bis excludes auto-dialling. At present, X.21 defines a single-circuit interface, but a multiplexed interface is also being defined for high-traffic devices such as computer systems.

X.21 and its variants will probably be revised as public networks are developed. In the next few years, therefore, there will be only a limited impact outside those countries where circuit-switched data networks are being installed.

Teletex

Teletex is both a logical extension and a modernisation of the concept of the telex service. The European PTTs, who are its main instigators, have been considering it for six years, but there is also considerable interest in it in North America. A demonstration of international use took place between West Germany, Canada and Sweden in mid-1980, and in November 1980 CCITT ratified the international standards for teletex. The three countries mentioned above expect to offer public services in 1982, and most other European countries seem likely to follow suit between 1982 and 1985.

Teletex, as defined by CCITT, differs from telex in that it uses memory-to-memory communication, rather than terminal-to-terminal communication (as illustrated in figure 12). This approach has several significant implications, including the following:

— Beyond the teletex interface the terminal can take any form, and, if desirable, it can therefore incorporate multiple functions and a variety of input and output technologies.

Figure 12 Schematic of a teletex service

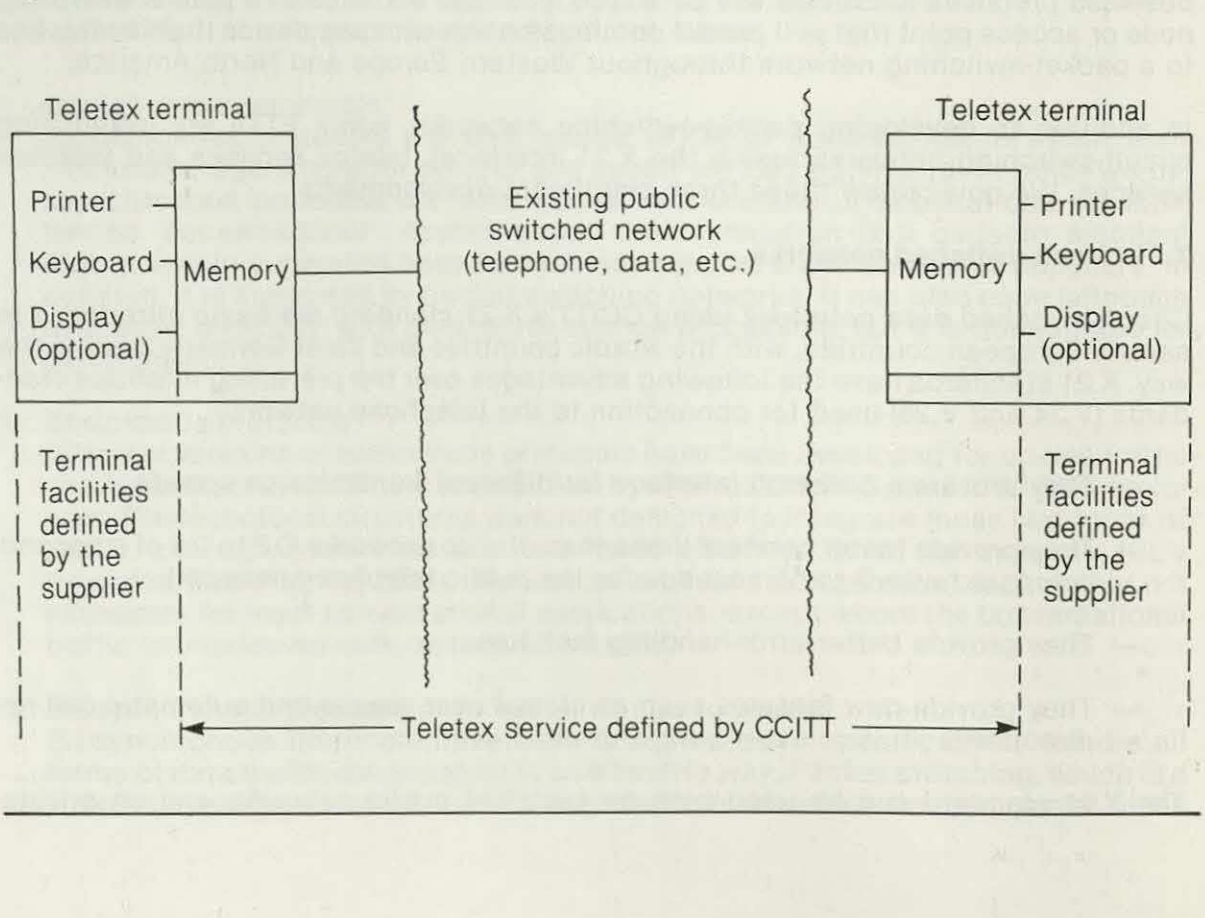

— Local use of the teletex terminal is independent of any communication that is taking place.

— There is no need for identical terminals at the sending and the receiving points.

The CCITT teletex standard is intended to be network-independent, and it envisages the use of either the public telephone network, or a packet-switching network or a circuit-switching network. Interworking with telex will be provided from the start of service, but, for this purpose, teletex terminals will need to be able to restrict the character set and the format to suit the telex network.

A transmission speed of 2400 bit/s is recommended for international teletex use—although national implementations may use different speeds. At a speed of 2400 bit/s an A4 sheet can be transmitted in less than five seconds.

One of the major objectives of the teletex standard is to define a character set that can be used with communicating terminals. The need for such a standard character set has been demonstrated by the difficulties different word processing terminals encounter when trying to communicate with one another. The teletex character set provides for a full Latin character set, and it also includes special national characters. (The teletex character repertoire is in fact the same as the text repertoire defined for videotex.) Subsets of the character repertoire may be selected for particular purposes. Terminals will transmit keyboard charcters only, plus an indication of the character subset in use, to enable the receiving device and/or its operator to check that the subset matches the printer. Teletex terminals are therefore expected to be capable of receiving all the character codes in use.

The teletex standard also requires that terminals should be capable of displaying a full A4 format (that is, up to 56 lines of 77 characters or 40 lines of 92 characters).

A terminal that met all the requirements that are explicitly or implicitly implied by the teletex standard would be an expensive device, and it would probably compare in price with a screen-based stand-alone word processor. But simpler devices that handle just the basic teletex functions will probably be available in the near future. In particular, the PTTs' approach of defining a memory-to-memory standard will undoubtedly encourage suppliers both of word processors and intelligent terminals in general to add a teletex interface to their products at a low marginal cost. The ability of teletex to interwork with telex terminals, and the close similarity there is between the videotex and the teletex character sets are also factors that will encourage the use of teletex. Consequently, we believe that teletex will have a significant impact on the market for communicating terminals. Some practical problems remain to be resolved, however, both for the PTTs and the suppliers. For the PTTs there are problems of administration and policy, and for the suppliers there are problems of design. The skill with which the PTTs and the suppliers address those problems will determine how quickly the teletex market develops. Our estimates of the development of the teletex market are shown in figure 13.

In summary, we believe that teletex has much to recommend it. We believe therefore that, as soon as the market permits, organisations should include compatibility with teletex in their specification of requirements for communicating terminals (or store-and-forward exchanges) that they are likely to install for communicating text either inter-company or internationally, or both. The use of teletex for intra-company communication is less attractive, because it would disenfranchise most of the terminals that are already installed. They would not be able to cope with the character set and other features implied by the teletex standard.

Figure 13 Development of the teletex market

Year	United Kingdom	West Germany	France	Western Europe
1982	0.5	2	1	4.5
1984	2	5	4	14
1986	10	22	15	56

The table shows the cumulative shipments (in thousands) both of teletex terminals and teletex-compatible word processors by the end of 1982, 1984 and 1986. Data terminals and microprocessors with a teletex capability will also be installed in the timeframe shown.

Videotex

Activity in the videotex area has so far been centred on the new capabilities that videotex systems offer, rather than on their inter-connection with other services. The videotex standards recommended by CCITT have not defined the methods to be used for interworking between videotex services and other services such as teletex or facsimile. As far as terminals are concerned, the overlap between videotex and other services occurs both at the format level and at the character set level. This assumes, of course, that the PTTs provide 'gateways' between the different networks used by the different services.

Interworking between videotex and telex is to be limited to the videotex format and the alphanumeric characters from the telex alphabet. As we have already indicated, the character repertoires for teletex and videotex are almost identical, and figure 14 shows the fallback representation for the small number of characters that do not match. Interworking between videotex and teletex will also be limited to the videotex frame format.

THE IMPACT OF NEW COMMUNICATION PRODUCTS

The design of communications interfaces for terminals is influenced not only by the PTTs (who provide the facilities for transmission between sites) but also by the suppliers of the private communications equipment installed on those sites. We summarise below the impact of three key product developments — proprietary network architectures, local-area networks and electronic private automatic branch exchanges.

Proprietary network architectures

The proprietary network architectures marketed by the computer suppliers, of which IBM's Systems Network Architecture (SNA) is the prime example, may affect not only the transport protocol adopted by terminals, but also the higher levels of protocol. For example, two features of SNA — the multi-system networking facility (for accessing more than one host system) and the character string facility (for compressing data for printing) — both depend on supporting logic in the terminal. In the terms of the ISO reference model described on pages 35 and 36, these two features are respectively

Figure 14 Recommended representation of those videotex characters that are not part of the teletex repertoire

Videotex character	Fallback representation on teletex
←	<
→	>
↑	¦
↓	¦
‘	, (apostrophe)
’	,
“	,,
”	″
── (horizontal bar)	– (hyphen)
Any block graphic	/ (solidus)

(Source: CCITT)

session level and presentation level protocols. Some degree of compatibility at these protocol levels is therefore necessary if SNA terminals are to access non-SNA hosts. Compatibility is also required if non-IBM terminals are to exploit features such as these when they access SNA host computers.

These examples show clearly how proprietary network architectures have increased the demands placed on those competitive suppliers who wish to sell terminals into a computer supplier's installations. It also demonstrates the need for standards for open-systems interconnection, as envisaged by ISO.

Local-area networks

The term 'local-area network' is used to describe a variety of products with differing aims and differing performance characteristics. However, all of those products have the following three features in common:

1. They use distributed switching techniques, rather than the centralised matrix switching used both by telephone exchanges and most current data switches.
2. They are intended principally to support local communication on one site, although some products may also be connected inter-site.
3. They use high internal operating speeds, usually in excess of 1M bit/s.

Two categories of local-area network are of potential significance for communicating terminals. The first includes those local-area network products that 'explode' systems that hitherto have been integrated (such as shared-logic word processing systems). These products therefore permit greater flexibility both in configuring the system and in locating the elements that make up the system. Logica intends to use the Cambridge Ring technology in this way.

The second category of local-area network product that is of potential significance for communicating terminals includes those products that are intended to be used as general-purpose networks to inter-connect a variety of devices of different types. Ethernet, which was originally developed by Xerox alone, but is now being developed jointly with Digital and Intel, falls into this category of product. (A detailed discussion of local-area networks can be found on pages 26 to 32 of Foundation Report No. 21.)

The suppliers of products in the first category may, for practical and marketing reasons, adopt a similar policy to that adopted by the computer suppliers for their proprietary network architectures. Such a policy would encourage communication outside the network itself only with other similar networks or with products from the same supplier. On the other hand, the market success of products in the second category will depend on their ability to service a wide variety of devices that are already in place. These products will therefore be compelled to support existing protocols such as those used by BSC terminals and by asynchronous terminals. Similarly, they will need to provide gateways for inter-site communication over either leased lines or packet-switching networks. Their competitive edge will derive from their performance and their functionality.

The position relating to standards for local-area networks is at present confused. There are two basic variants of the logic for controlling access to a local-area network. The first is a token-passing scheme (as exemplified by the Cambridge Ring), and the second is a broadcast-contention scheme (as exemplified by Ethernet). The IEEE in the United States has endorsed both schemes.

Cost is the main advantage in standardising the interface between a terminal and a local-area network. A standard interface will make it possible to mass-produce the logic circuitry required, so that it can be built into the terminals and the other devices that use the local-area networks. Also, standards both for the methods of operating a local-area network and for the diagnostics produced by a local-area network would simplify the network management task.

Electronic private automatic branch exchanges

Electronic private automatic branch exchanges, particularly those that use digital switching, are better adapted than are their analogue, electro-mechanical counterparts to switching non-voice traffic as well as voice traffic. Electronic private automatic branch exchanges support communicating terminals in two ways:

1. By providing tranparent through-circuits, which are set up by means of the telephone but are used by conventional terminals.

2. By acting as a store-and-forward device for special-purpose terminals. For example, IBM offers keyboard terminals and visual display terminals for use with its 3750 exchange. These terminals use multi-frequency (analogue) signalling to communicate with the exchange. Data stored by the exchange may then be retrieved by a computer system over either a BSC or a SDLC connection.

Electronic exchanges will probably not exert a significant influence on the market for communicating terminals. The suppliers of private automatic branch exchanges, however, are likely to offer a range of devices (such as enhanced telephones, keyboards or display terminals) that they design for use with their exchanges.

THE STANDARDS AUTHORITIES

The internationally agreed standards for public data networks promise to exert a major influence on communication standards up to the transport level. However, the fact that protocols such as the teletype (asynchronous) protocol, and IBM's 3270 BSC and 2780 BSC protocols have established themselves as de-facto standards is a measure of the failure of the standards authorities to get to grips with the problem of terminal incompatibility. This failure is not necessarily a reflection of the competence of the standards committees. Rather it reflects the conflicting interests and viewpoints both of the organisations that implement the agreed standards and of their representatives on the standards committees. Those organisations include the PTTs, the equipment suppliers, the universities, and government agencies such as the Institut National de Recherche en Informatique et en Automatique in France, the Gesellschaft für Mathematische Datenverarbeitung in West Germany, the National Physical Laboratory in the United Kingdom and the National Bureau of Standards in the United States. We now review the positions of the main interest groups.

The PTTs

The PTTs are predominantly in the business of communications, and so they can be expected to favour standardisation of the communications protocols. Generally they do so, but the extent to which they are effective in making standards is limited by two factors. The first is that their monopoly position, combined with their limited experience, may inhibit them from tackling protocol problems at the levels above the transport level. The second is that, as has happened with videotex, they may be competing among themselves for worldwide markets.

The suppliers of computer equipment

The computer suppliers operate in an increasingly competitive market, and one of their major pre-occupations is to retain their existing customers. Incompatibility in communications protocols is one weapon that they can use to do this. Essentially, the computer suppliers are in the business of selling equipment, and so they are bound to evaluate communications standards in terms of the effect those standards will have on equipment sales. From the computer suppliers' point of view, some degree of compatibility may well be desirable because it can both expand the market and reduce diversity. Thus, it will keep manufacturing costs down to a reasonable level. On the other hand, too much compatibility may represent a severe threat to a supplier's hard-won competitive position.

The suppliers of communications equipment

Excluding the major computer suppliers, for whom communications equipment is part of a total system, the suppliers of communicating equipment and services depend on communications for their business, just as the PTTs do. Unlike the PTTs, however, they can be pragmatic and selective in their approach to the market. They, if anyone, are likely to provide the bridges between today's widely-used communications protocols and the international standards that are being assembled so painstakingly and also so slowly.

The North American packet-switching carriers provide an illustration of the way in which the communications suppliers can provide these bridges. The operators of the Datapac network in Canada and the Tymnet and GTE-Telenet networks in the United States are now testing software to support 3270 BSC terminals across all three packet-switching networks. Polling of the terminals is intercepted at the network node nearest to the host computer, and the polling message is propagated at the node nearest to the terminal controller. (IBM offers a 'black box' that users can install privately to perform a similar function, but it ought to be far more cost-effective to implement this function in a public network.)

Sectional interests

The three groups we have just discussed all have a direct commercial interest in communications standards. In addition, there are several influential groups that represent either particular users' interests or national interests. Probably the most significant of these groups is the National Bureau of Standards (NBS), which has the task of defining standards for all United States government computer work. The government represents such an important segment of the market for United States computer companies that NBS standards can very quickly become de-facto standards. Fears have been expressed that the NBS, which is working to a shorter timescale than ISO's, will be the first body to produce comprehensive communications protocol standards, and so will pre-empt ISO's work on open systems interconnection.

National standards bodies

In the belief that ISO's open systems interconnection reference model will take too long to formulate, both the United States and West Germany have decided to develop what they term National Interim Standards, which are based on the OSI reference model. However, these national standards do not guarantee compatibility in the future, even though they are based on the OSI reference model. In fact, the reference model is so general at this stage that many interested parties can and do interpret it in different ways to suit themselves. IBM, for example, is applying pressure to have the reference model aligned with its own SNA. Needless to say, such an alignment will not be achieved with the willing consent of IBM's main competitors.

CCITT, meanwhile, is pressing ahead at a faster pace than ISO, in an attempt to consolidate its success in establishing standards for public data networks. These standards will affect the lower levels of protocol in the ISO model — the levels up to the transport level. Here again, there is no guarantee that CCITT Recommendations will match the standards that ISO subsequently specifies.

SUMMARY

The two main conclusions that can be drawn from this analysis of the standards world are as follows:

1. OSI standards are unlikely to be of practical value before the mid-1980s, and even then, there may be a transitional period of several years whilst the conflicting standards that have been adopted in the meantime are brought into line.

2. The PTTs, through CCITT, are likely to be effective in establishing communications standards, but only up to the transport level. There are likely to be an increasing number of suppliers, including the PTTs themselves as well as the specialist communications companies, that provide conversion services or conversion products. Those services or products will enable devices that work to existing standards to be used on public data networks.

CHAPTER 6

SUMMARY OF FINDINGS

In this report we have described communicating terminals from three different points of view:

1. From the viewpoint of the user, who sees a terminal with a range of features and facilities, and equally from the point of view of the buyer, who acquires such terminals for others to use.

2. From the viewpoint of the communications network, which sees a device demanding a certain type and level of communications service.

3. In terms of the logical structure of the hardware and the software elements involved in the task of communication.

In this chapter, we summarise the findings from our research in terms of these three different views of communicating terminals. When an organisation acquires terminals it must reconcile all these viewpoints, at least to some degree. We concentrate in particular on those aspects of our findings that affect corporate policies for terminal selection and terminal support, rather than on those aspects that affect individual purchasing decisions.

THE USER/BUYER VIEWPOINT

The use of terminals

Several general changes are discernible both in the ways that communicating terminals are being used and in the purposes for which they are being acquired. Most importantly, as the penetration of terminals into the business environment increases, so terminals are being increasingly used on a casual and infrequent basis by inexpert, rather than expert, users. This change in use implies a need for a different type of terminal that provides:

— A simpler keyboard layout, with fewer function keys, etc.

— Different forms of feedback (for example, audible or visible feedback as well as tactile feedback).

The change in the use of terminals may also imply that a different means of introducing terminals into an organisation is required, because it may be neither possible nor helpful to train users intensively. More probably, careful attention will need to be given to operating procedures, and particularly to 'help' facilities, which will need to be supportive without being either intrusive or onerous.

As a result of this trend to the less intensive use of terminals, there is both an opportunity and a latent demand for terminals to be used for more than one purpose. The shorthand term 'multifunction terminal' has been coined to describe this multipurpose use of a terminal, but the different purposes for which a terminal can be used can, in fact, take several different forms. At the simplest level, the terminal may be used to access two similar services (for example, a timesharing bureau as well as an information database). This form of use could be termed 'multi-access'. Provided that the methods of access are compatible, multi-access demands little more than a line-switching capability in the networks that are used to access the services.

At the next level of complexity, terminals may be used both for data processing and text processing. One example of this type of use is a terminal that is used both to retrieve information from the corporate database and to compose messages to be sent electronically. Whether or not this form of use affects the facilities and the features of the terminal depends on the type of processing undertaken. For example, a casual user's needs may be satisfied with the limited editing facilities available on a basic data terminal, whereas a secretary will require many of the features found on word processors. In both cases, the demands placed on the network may extend beyond a line-switching capability to include also an involvement in network addressing-schemes and in establishing sessions (that is, in sign-on procedures, etc.). We suggest that the term 'multifunction' best suits this extended mode of use.

The next level of multi-purpose complexity is a 'multi-media' capability, but we could find no clear evidence of user demand for this. A multi-media capability implies a need to combine, at the terminal, facsimile information or voice information with data or text. Facsimile information and voice information are essentially analogue in content, whereas data and text are coded in a digital form. This difference makes it difficult to combine the different forms of information. The first place where these different forms of information will merge is probably at the output station. For example, non-impact printers are capable of printing facsimile information as well as character output streams, and raster-scan displays can be used for displaying both analogue and digital information. Multi-media terminals that combine voice information with non-voice information at the terminal are, we believe, many years from realisation.

Looking further ahead to a time when communicating terminals will be as normal a part of the office environment as telephones are now, it will be possible to distinguish four categories of terminal device, rather than the existing simple divisions into either expert and inexpert devices or regular-use and occasional-use devices. Figure 15 describes these four categories, and in each of these we have, as appropriate, included the telephone and other devices.

Terminal technology

As far as terminal technology is concerned, the four key points that we identified in our research are:

1. The QWERTY/AZERTY keyboard will continue probably throughout the 1980s as the dominant method of input both for data and text.

2. The CRT, deriving from the television industry, will continue to develop, and it is unlikely to become a constraint at the time when higher-capacity displays (for example, A4 page-size formats) or higher-resolution displays (for example, graphics) are needed. Consequently, CRT technology will continue as the dominant display technology, and it will continue to use raster-scan techniques, rather than vector techniques. Raster-scan techniques are cheaper than vector techniques, and they are also capable of handling graphics, for which there are already some signs of a growth in demand.

3. Business terminals are being provided with increasing amounts of processing

Figure 15 Future categories of terminals

Terminal category	Standard features	Optional features	Used by
Basic terminal	Telephone plus a few special-feature keys, or a videotex business terminal	Extra keys, videotex screen, keyboard, loudspeaker, tape recorder, ansafone and interfaces to: — local-area networks — digital public networks — V.24 devices — IEEE devices — Ethernet	Managers, staff in small offices and professionals who use oral communication extensively
Clerical terminal	Visual display terminal with a telephone and sufficient local processing power for diary, electronic mail and word processing functions	As for the basic terminal, plus extra processing power and storage, and facsimile (group 3) capability	Clerks
Techno-crat terminal	Programmable computer with a telephone interface, a colour display, many function keys, facsimile receipt capability, a multiplexed communication interface, and several linked filing systems with additional links to departmental and corporate files as appropriate	Extra processing power and storage, facsimile scanner, and a microfilm link	Managers in high technology companies, scientists and engineers, and secretaries
Special-ist terminal	Designed for each set of special requirements. Current examples include newspaper page-make-up terminals and terminals used by telephone switchboard operators		Designers, artists, draughtsmen, process control operators, telephonists, etc.

power and memory. This processing power and memory is being used both to provide additional functions (such as a graphics capability or additional buffering) and to process information locally.

4. From an ergonomics point of view, terminals are improving, as designers turn their attention from the fundamentals of performance to the capability of the terminal as a complete system.

The terminal market

The terminal market, like terminal technology, is evolving steadily, although there are intimations of revolutionary changes to come. The market position can be summarised as follows:

1. Terminals will steadily become cheaper, although suppliers will generally aim to supply terminals with more functions, rather than supplying cheaper devices.

2. There will be little scope for individual organisations to save money by buying terminals in bulk, although an organisation that purchases terminals from an OEM supplier rather than from a system supplier can achieve significant savings.

3. Terminal sales will continue to grow quickly, with suppliers making more sales direct to users rather than to buyers such as management services departments.

4. Minimum-cost-technology devices, such as videotex terminals, may be the forerunners of more sophisticated terminals that will use the same minimum-cost technology, but will be sold at prices well below those current today for similarly equipped devices.

THE NETWORK VIEWPOINT

One key decision that an organisation has to make when it constructs the infrastructure to support communicating terminals is what protocols that infrastructure should support. Clearly, some flexibility can be built into any network structure, and this means that the choice of terminals is not entirely constrained by the network. That flexibility will not be limitless, however, and it will cost money. Therefore, an organisation must make some decisions (even though they are only provisional) that attempt to forecast the demand that terminals will place on the network. In the absence both now and for the foreseeable future of any widely applicable and internationally recognised standards, such decisions must be based partly on expediency (by taking into account the protocols currently in use) and partly on judgements of the direction that terminal communication within the organisation will probably take in the future. Such a judgement could be based, for example, on the relative importance of different modes of use, which in turn would depend partly on the nature of the business and partly on the policy for information systems. At the simplest level, the choice of network structure, the choice of terminals, and the choice of protocols all depend on the type of communications traffic. For example, heavy conversational traffic favours asynchronous terminals working either over the public telephone network or over a packet-switched network, whereas heavy on-line data traffic favours basic-mode synchronous protocols.

In the course of time, developments both in corporate information systems and in the economics of communications, will make it necessary for organisations to change their protocols and network structures. In figure 16, we summarise the migration paths that organisations might follow as different types of communications traffic come together.

THE STRUCTURAL VIEWPOINT

The full seven-level ISO reference model for open-systems interconnection is not a prerequisite for terminals to communicate. As an ideal concept, though, it is undoubtedly worth pursuing, particularly when it is coupled with the cheap processing power that can now be incorporated into terminals. However, while the standards authorities are pursuing the concept of open-systems interconnection, users

Figure 16 Possible migration paths

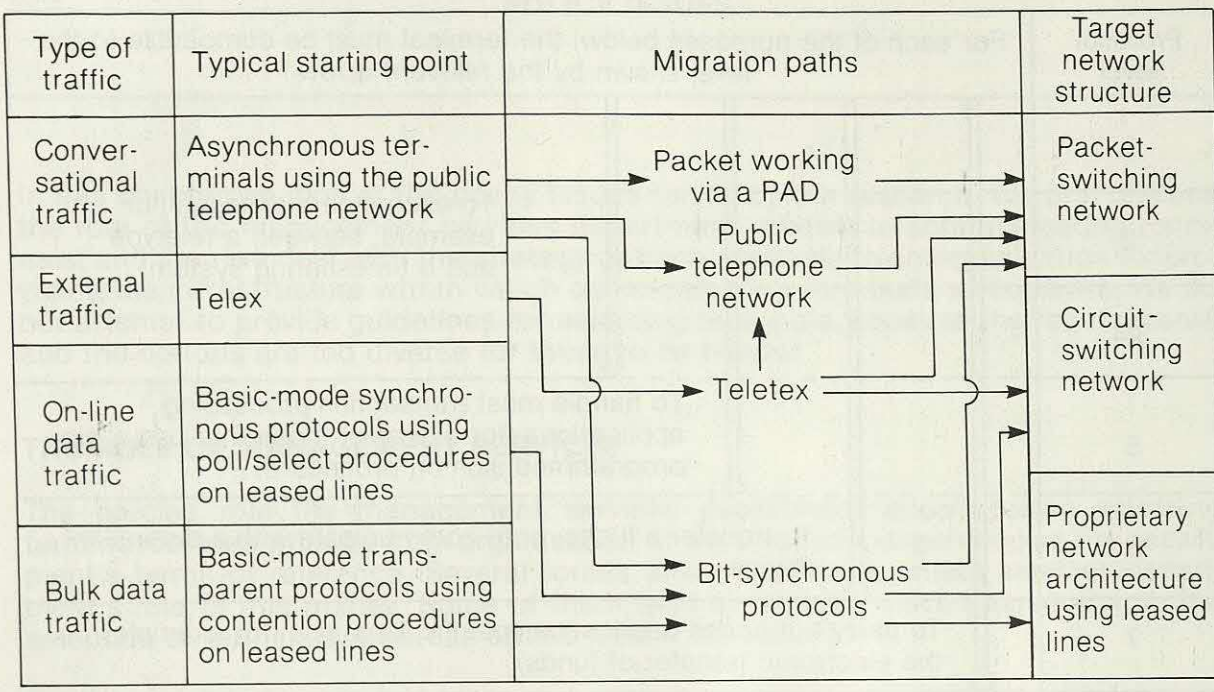

will need to mix and match the various partial (and sometimes conflicting) communications protocols that are available now.

The level of compatibility that is required by both a terminal and the device (or service) to which it is connected varies according to the terminal's purpose. As a guide, figure 17 shows, in terms of the ISO model, the different levels of compatibility that are required for these various purposes. For example, a simple terminal, such as a teletype, is aware only of the procedures up to level 3/4 that are used to establish a physical connection with either a timesharing system or a PAD. The other levels of protocol are, in effect, implemented by the user himself (for example, by signing on, by entering keyword commands, etc.). This example is a slight over-simplification because the terminal and the timesharing computer must also have code compatibility, but it serves to illustrate the principle. Clearly, it is wasteful to have compatibility beyond the point that the particular purpose demands. Conversely, it is important to ensure that compatibility goes as far as is needed (as, for example, when an organisation acquires plug-compatible devices).

In essence, figure 17 illustrates that the user has a choice as to where he executes some of the levels of protocol. When accessing a timesharing system, for example, the terminal user himself is responsible both for the session protocol (the sign-on procedures, etc.), and also, to a certain extent, for the presentation protocol (the lay-out of commands, etc.), and the applications protocol. Equally, software-executed protocols may, if it is convenient, reside elsewhere than in the terminal. As an example, when the terminal user depresses a function key, the terminal logic may implement an element of the presentation protocol by causing a command (such as 'GO TO') to be sent over the network.

Figure 17 Levels of compatibility required for various communications purposes

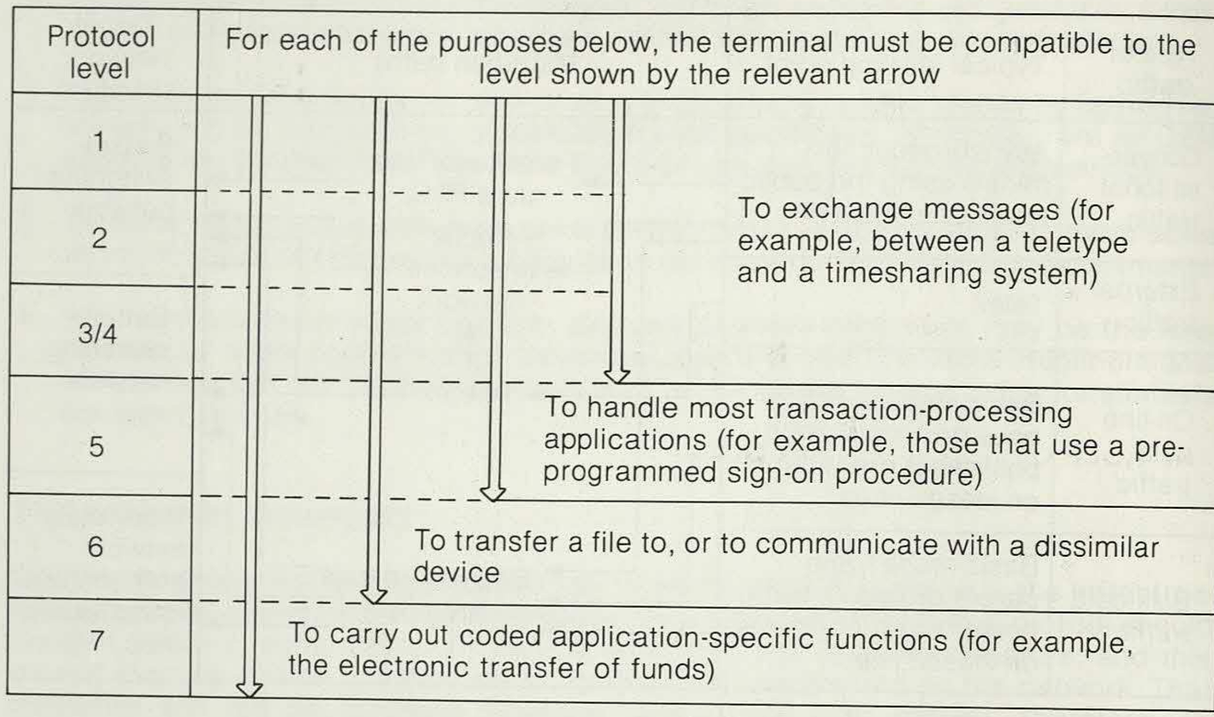

Alternatively, the terminal might send a code that simply indicated that the user had pressed a particular function key. The mapping routines in the network or in the interface software in the host computer would then interpret the code as appropriate. The place where this mapping is carried out will determine the ease with which a terminal can be used for different purposes. It will also determine how easy it is for different terminals to access the same service in broadly the same manner.

The ideal arrangement, in functional terms, is probably to use a virtual terminal protocol that is implemented either within the communications network or at the access points to the communications network (such as a PAD). All application programs would address the virtual terminal. With terminals that are not designed for the virtual terminal protocol, mapping from the virtual terminal to the real terminal is carried out by protocol conversion routines. With terminals that are designed to support the virtual terminal protocol direct, the mapping is carried out within the terminal itself.

In the following chapter (on pages 58 and 59) we identify the virtual terminal approach as one of the strategy options that user organisations might pursue with respect to communicating terminals, and we describe the difficulties of implementing such a strategy at present.

CHAPTER 7

POLICY ISSUES

In this chapter we look at the policy issues raised by our research. We first discuss the role of the management services department relative to communicating terminals, and then we deal with the strategy options available to an organisation for providing the infrastructure within which communicating terminals will operate. We do not attempt to provide guidelines for selecting terminals, because the requirements and the options are too diverse for these to be helpful.

THE ROLE OF MANAGEMENT SERVICES

The precise role the management services department adopts when selecting terminal devices will vary from organisation to organisation, depending on the department's terms of reference. Several forces affect the management services department's role in this matter. Some of them tend to demand more intervention in the selection of terminals, whereas others lessen the need for it.

On the one hand, more and more end users are coming into contact with terminal devices, and those end users are learning for themselves (or through their unions) what features and facilities they should and should not expect in a terminal. Experienced users such as these clearly can and should exert a strong influence on the selection of terminals.

On the other hand, the increased penetration of terminals into the business environment means that terminals are being used by inexpert users and, perhaps more significantly, by occasional users. Any lack of user expertise can be remedied by training and by experience, but occasional use brings with it problems that may be more difficult to solve. In particular, occasional users may become sufficiently familiar with the operating procedures to find conventional 'help' facilities tedious. However, they may not be sufficiently familiar with the operating procedures to dispense completely with the help facilities. Sometimes, as we mentioned in chapter 6 on page 49, it may be appropriate for inexpert users to use a different type of terminal to that used by expert users. Very simple operating procedures, such as those employed by videotex terminals, will also be attractive both to inexpert and occasional users. Also, inexpert users will probably not be well informed about either the options they have available or the features that they need. Consequently, they will tend to depend heavily on the management services department for expert advice.

In addition to advising on terminal selection, the management services department will often also provide or manage the communications infrastructure, and it will therefore set the network interface standards to which terminals must conform. It is clearly desirable that such interface standards should not impose a major constraint on the choice of terminals, and, for that reason, many organisations will seek to maintain the maximum flexibility in the network. Even so, we believe that some constraints on the choice of terminals are unavoidable, and we discuss the main options and decision factors later in this chapter.

There are two further roles that the management services department might play, and they are both likely to gain in importance over the next decade. The first concerns the skills an organisation requires when it deals with OEM terminal suppliers. We have already discussed the different market approaches and pricing policies that the OEM suppliers and the system suppliers have adopted. An organisation can achieve cost savings of 50 per cent or more by purchasing terminals from an OEM supplier. But, in order to assess the risks and the trade-offs involved in dealing with an OEM supplier an organisation requires greater skills (both commercial and technical) than it requires in dealing with a system supplier.

The second role that the management services department can play is as a broker of secondhand terminals. This role would enable organisations to re-distribute terminals from those end users whose requirements have outgrown the terminals, to other end users who can make good use of those terminals.

THE ROLE OF THE COMMUNICATIONS INFRASTRUCTURE

The role of the communications infrastructure is to provide a service both to the terminals and the computers that need to communicate with one another. However, decisions about the shape of the infrastructure cannot be entirely subordinate to decisions about the terminals and the computers. In practice, the decision about the infrastructure must be a compromise between, on the one hand, the need to keep future options open whilst continuing to support existing equipment, and, on the other hand, the cost of providing the infrastructure. One key decision an organisation must take concerns the distribution of intelligence between the terminal, the network and the host computer systems with which terminals communicate across the networks. By 'terminal' we mean here both the terminal devices themselves and the other equipment associated with the terminal at the same location. Figure 18 summarises the arguments for and against placing intelligence in one or other of these three elements of information systems.

By developing the analysis shown in figure 18, we have identified four distinct strategy options that user organisations might pursue, although, in practice, hybrid solutions may be appropriate. Each option implies a certain approach to the communications network design as well as to terminal procurement. These four strategy options are described below:

1. The highest common factor approach

This strategy recognises that there is an absence of manufacturer-independent standards of broad scope, and so it adopts those low-level standards that can be relied on for terminals. The communications network and other facilities (such as service computers attached to the network) can be used to provide the flexibility and the wider functionality that the terminals may lack. A typical example of this approach is the use of asynchronous terminals that are supported by a packet-switching network. These terminals can cope both with conversational work and some on-line work, whereas the packet-switching network both provides the switching capability and compensates (at least to some degree) for the error-handling limitations and the speed limitations inherent in asynchronous protocols. (Bulk traffic could also be carried over the packet-switching network, but local distribution would be separate from that of the asynchronous terminals.)

The main advantage of this approach is that it exploits proven technology, and this means that it is straightforward to implement. In addition, it takes advantage of the wide choice of inexpensive asynchronous terminals that are now available. On the other hand, this approach tends to exclude the more sophisticated terminal devices, particularly those that are commonly used for on-line applications.

In the longer term, a strategy based on packet switching will benefit from the developments set in motion by the widespread implementation of public packet-switching networks. More and more terminals and computer systems with full X.25 capability (that is, with a bit-synchronous protocol) will come onto the market over the next few years. In addition, many new public services, such as teletex and videotex, are also likely to use packet-switching networks to distribute information.

Figure 18 Distributing communications intelligence within a communications system

Element in which intelligence may be placed	Arguments for placing intelligence in the respective elements	Arguments against placing intelligence in the respective elements
Terminal	Local autonomy and flexibility Reduced communications costs Increased ease of integrating different forms of traffic (e.g. batch and on-line)	Lack of manufacturer-independent standards
Network	Consistent session-level interface for multi-access terminals Reduced terminal costs Increased ease of handling corporate requirements (e.g. encryption and addressing schemes)	Increased difficulty in using public networks like the PSTN for fall-back
Host computer	Consistent applications-level interface	Increased difficulty for multiple access Limitation of terminal choice

2. The single supplier approach

The decision to standardise on a single system supplier for terminals goes hand in hand with the decision to adopt that supplier's proprietary network architecture. The main disadvantage of this strategy is that it limits the choice of terminals, although the existence of a thriving PCM terminal industry, particularly for IBM devices, provides some flexibility. At present, the cost penalties for an organisation of a limited choice may not be great. However, they could quickly become substantial if an organisation cannot exploit cheap mass-market terminals when they become available.

Despite its disadvantages, this strategy has several points in its favour:

— It is relatively safe from a technical point of view.

— It ought to be easy to manage.

— It promises to provide a consistent and a controlled environment for terminal users.

3. The intelligent network approach

By 'intelligent network' we mean a network that is capable both of supporting a variety of communications protocols and of applying any required protocol conversions. This strategy, therefore, goes one step further than the highest common factor approach because it attempts to accommodate a range of existing protocols, rather than selecting the common elements that all the protocols either share or can be made to share.

The intelligent network approach has three disadvantages. Firstly, unless it is severely limited in scope, it can be an expensive option to pursue. AT & T reputedly discovered the difficulties of this approach when it attempted to implement a general-purpose solution in the form of its planned Advanced Communications Service. Secondly, the strategy offers no guarantee of long-term stability. The experience of the PCM terminal industry demonstrates that a continuous programme of enhancement to an intelligent network may be necessary to keep the protocol conversion routines up to date. Thirdly, protocol conversion is essentially a compromise. It may therefore prove unsatisfactory where the differences that have to be resolved are fundamental rather than superficial. Where conversion is carried out solely for the purpose of transport (for example, by a PAD in a packet-switching network), the problems are likely to be solvable. On the other hand, where the purpose of conversion is to make one terminal look in every respect like another terminal, more significant design problems may emerge.

4. The virtual terminal approach

Except where it is used as a bridge between existing standards and an enduring standard, protocol conversion, as used by the intelligent network approach, must inevitably remain a short-term expedient. A virtual terminal protocol could be such an enduring standard. The purpose of a virtual terminal protocol is to specify a general representation of a particular class of terminal device (for example, visual display terminals).

The difficulty, of course, is that a generally applicable virtual terminal protocol does not yet exist and does not yet appear to be in prospect. There is a dilemma that needs to be resolved before a viable virtual terminal protocol can be designed. On the one hand, a virtual terminal protocol must be wide enough both to cover the needs of a reasonable proportion of users and to allow an organisation to justify the management effort and the systems effort necessary to introduce the protocol. It must also be sufficiently wide to enable terminal suppliers to justify the development effort necessary to support the protocol.

On the other hand, it is because a virtual terminal protocol needs to be narrow enough to be defined precisely without too many compromises and inefficiencies that there is scepticism about ISO's chances of success with their open-systems interconnection standards, which must incorporate a virtual terminal protocol in some form.

Equivalent standard protocols will also be needed for file transfer procedures and, possibly, for job transfer procedures.

Even with a viable virtual terminal protocol, there are still some difficulties that remain to be solved. For example, systems software may remain coupled with the peculiarities of the supplier's own terminals. It will then be necessary either to amend the systems software or to introduce conversion routines at the communications interface. These conversion routines would bring the supplier's protocol into line with the virtual terminal protocol. As another example, the virtual terminal protocol may become obsolete because it fails to anticipate future technological developments.

Despite these formidable difficulties, the virtual terminal protocol approach, given the right circumstances, is clearly feasible. The multi-access reservation systems being developed in several European countries have to provide communication between standard terminals and several different host computer systems. Implicitly or explicitly, the protocol used by these multi-access systems is a virtual terminal protocol, even though it is designed for a limited range of applications. One of the keys to the definition of a general-purpose virtual terminal protocol may well be to define (and limit) the range of applications that it must accommodate.

The biggest difficulty, undoubtedly, will be to find a satisfactory vehicle to specify and promulgate the virtual terminal protocol. The ideal vehicle would be an international standards body, such as ISO, but we believe that a more narrowly constituted body (such as a user group from a specific industry) has a greater chance of achieving worthwhile results within a reasonable time. Such a body must, of course, also carry sufficient commercial influence for its voice to be listened to by the terminal suppliers.

CHAPTER 8

CONCLUSION

Management services departments will find it increasingly difficult to exercise detailed supervision over the acquisition of terminals, just as they find it difficult to control the purchase of microcomputers. (In any case, the microcomputer itself is, at least potentially, a communicating terminal.) It is, therefore, important for an organisation to establish a framework in which decisions on terminal acquisition can be made in a rational way, based on established criteria and guidelines.

For large organisations, the corporate network may be the key element in such a framework because it imposes some limitations of choice on those terminal users who wish to gain access to corporate services available via the network. For other organisations, a master strategy may not be possible. Instead, those organisations require several tactical expedients that take account of changing technology, of changing marketing strategies by the suppliers, and of changing public and private communications services. The principal reason for an organisation to adopt a pragmatic strategy is that there are no comprehensive and dependable manufacturer-independent standards, around which an organisation can build a long-term strategy.

A pragmatic strategy may take advantage either of particular suppliers' protocols or of public network standards, depending on an organisation's current investments and its perceived needs. However, because of the standards vacuum, few organisations will be able to avoid either taking risks or making compromises. In this report we can, of course, only draw attention to the problems and point to those partial solutions that are now available.

Communicating terminals will be so important to the development of information systems that, in our view, few organisations of any size can afford to be without a strategy of some kind, no matter how imperfect that strategy may be.

The Butler Cox Foundation

Report Series No 23 Communicating Terminals

Please record your interest in follow-up work on the subject of this report, highlighting those aspects of the topic of most interest to your organisation.

Completed by...

Job title...

Organisation ....

Address...

....

Please return to: The Butler Cox Foundation
Butler Cox & Partners Limited
Morley House
26 Holborn Viaduct
London EC1A 2BP

The Butler Cox Foundation

Report Series No 23 Communicating Terminals

Please note below any points on which you would like the author to expand at the professional and technical seminar associated with this report.

In order to take account of your comments we need to receive them from you by 12 June 1981.

Completed by..

Job title..

Organisation..

Address...

....

Please return to: The Butler Cox Foundation
Butler Cox & Partners Limited
Morley House
26 Holborn Viaduct
London EC1A 2BP

The Butler Cox Foundation

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

Butler Cox & Partners Limited
216 Cooper Center, Pennsauken, New Jersey 08109, USA
Tel (609) 665 3210

Italy
Sisdoconsult
20123 Milano – Via Caradosso 7 – Italy
Tel 86.53.55/87.62.27

Benelux, Germany, Switzerland
Akzo Systems bv
Postbus 350, 6880 AJ Velp (Gld), The Netherlands
Tel 85-649474

France
Akzo Systems France
Tour Akzo, 164 Rue Ambroise Croizat,
93204 St Denis-Cedex 1, France
Tel (1) 820.61.64

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