Report Series
No. 46

Network Architectures for Interconnecting Systems

June 1985

The Butler Cox Foundation

THE BUTLER COX FOUNDATION

REPORT SERIES NO. 46

NETWORK ARCHITECTURES FOR INTERCONNECTING SYSTEMS

ISSUED JUNE 1985

Abstract

There is a growing need to link different computer systems in order to transfer data, and to link processors in order to share resources. The computer suppliers, however, have developed communications protocols that are incompatible. This Report provides a management perspective of the issues involved in computer system interconnection; it is aimed at information systems managers and their telecommunications and systems planning departments.

The report examines the relative importance of proprietary network architectures and open system standards in interconnecting hitherto incompatible computer systems. We expect that open systems standards will have an important role to play as the bases of some proprietary network architectures, and of gateways between different proprietary network architectures. However, proprietary network architectures, and in particular SNA, will always provide more functionality than is specified in the current version of open system standards. We conclude that the simplest solution to the problem is to standardise, wherever possible, on the network architecture of a single supplier.

Research method

The research for this report was carried out in late 1984 and early 1985, and was led by *Philip Aspden,* a senior consultant with Butler Cox specialising in telecommunications.

In preparing the report, Philip Aspden was assisted by three other members of the Butler Cox consultancy staff — *Roger Camrass,* the Director of Telecommunications Studies, *Doug Taylor,* a senior consultant specialising in data communications, and *Edward Vulliamy,* a senior consultant also specialising in telecommunications.

The research comprised over 40 in-depth interviews with network users, computer suppliers, PTTs and government officials. The communications managers of 14 large organisations (eight in the United Kingdom, three in France, one in Belgium, one in Italy and one in Sweden) were interviewed. Case histories based on a selection of these face-to-face interviews illustrate the diversity of current approaches to computer networking. We should like to thank all the organisations that participated in the research and, in particular, those that have allowed their experiences to be published.

We also reviewed many papers on proprietary network architectures and open system interconnection (OSI) standards, and drew on previous research carried out by Butler Cox into the corporate network market as part of the Strategic Studies Programme. Finally, we also made use of Butler Cox's accumulated expertise gained from carrying out a range of communications strategy consultancy assignments.

The main findings of the research are highlighted in the report synopsis.

Additional report copies

Normally members receive three copies of each report as it is published. Additional copies of this or any previous report (except those that have been superseded) may be purchased from Butler Cox.

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

NETWORK ARCHITECTURES FOR INTERCONNECTING SYSTEMS

CONTENTS

NETWORK ARCHITECTURES FOR INTERCONNECTING SYSTEMS

REPORT SYNOPSIS

As computer systems become more important to the running of an organisation's business, so the need to interconnect these systems becomes greater. More than ever before, there are now compelling organisational, commercial and functional reasons for interconnection and for comprehensive networks. In practice, interconnection is difficult, mainly because of incompatible communications standards. Today's problem for organisations is to decide on the best strategy for tomorrow's networks, amid the conflicting claims and pressures of proprietary architectures and emerging international standards. Organisations now need to define a communications infrastructure that will accommodate all the requirements of their future network services and applications.

Selecting the best route forward is the subject of this report, which looks in turn at the basic problem, at the solutions available today, and at the options in planning for tomorrow.

Approaches to the interconnection problem

The basic problem of incompatible communications standards certainly applies in the many organisations that use computer systems from several suppliers. But even when an organisation standardises on products from a single supplier, connectability is not guaranteed.

In broad terms, the problem is to choose between two approaches, one leading to integrated communications using either a proprietary network architecture or a single communications standard, and the other based on the use of several different standards, with interconnection via commercial 'off-the-shelf' products or bespoke solutions as required.

To a large extent the choice will reflect the business style of the organisation: those organisations whose business units need to work closely together will favour the integrated approach. The experiences of specific organisations are outlined as case histories in the Appendix to the report. They show that for established proprietary network architectures such as SNA and DECnet, implementation can be fairly straightforward. On the other hand, some organisations have developed their own communications software, although this is no longer an attractive option because few organisations now have the necessary skills and resources.

The need to use several standards, rather than an integrated approach, can arise for a variety of reasons. Organisations may use different suppliers' systems for different purposes, may wish to avoid dependence on a single supplier, and/or may have inherited different systems through company mergers. In addition, the growing demand for electronic communication between organisations and their customers, suppliers and business partners implies a multi-standard approach.

Both approaches, integrated and multi-standard, have their benefits and disadvantages. To adopt multiple standards can limit functionality, reduce the scope for sharing resources, and lead to increased costs. For those reasons, many organisations are moving to integrated networks and, in particular, to proprietary network architectures. On the other hand, the multi-standard approach enables the best supplier for a particular job to be selected; is consistent with decentralised management; and avoids the danger that an organisation will be locked into a single supplier's products and prices.

Proprietary network architectures

Proprietary network architectures are seen by many user organisations as their best interconnection solutions for the foreseeable future. Such architectures have been introduced by most major computer suppliers in recent years, following the lead of IBM with SNA in 1974. In essence, compared with earlier communication products, current versions of proprietary network architectures provide better utilisation of resources, improved network management and greater resilience.

The two most developed architectures are IBM's SNA and DEC's DECnet — the former intended primarily to give remote access to centralised resources, the latter promoting a distributed computing environment. As the most widely used network architecture, SNA is a de facto standard, a proven product, and can be fairly straightforward to implement. Significant further development is being pursued. DECnet also has been, and will be, extensively developed. The report considers the current status, and the likely development, of these and other leading proprietary network architectures.

OSI standards

Proprietary network architectures are here today; OSI standards are the hope for tomorrow. The aim of the OSI initiative is to permit interconnection between computer systems, regardless of supplier, through the use of internationally accepted standards that would be implemented by all computer manufacturers. A seven-layer reference model defines a structure for communication and identifies services and protocols that support the structure. This structure means that standards for each layer can be developed independently.

The impetus behind OSI standards is considerable, with support coming from suppliers, PTTs, users and governments. Nevertheless, several problems still face the OSI initiative. The process of obtaining international agreement is difficult and lengthy. Different suppliers will produce different versions of the same standard, requiring the establishment of internationally compatible protocol testing procedures. And, such is the scale and complexity of OSI development work, it is very difficult for users to contribute.

We expect that electronic-mail products based on the full OSI model will be available in 1985-86, with products for mainstream data processing following two years later. These products will be OSI-based proprietary network architectures, and OSI-based gateways between existing proprietary network architectures. However, the initial products will provide greater functionality than is specified in the early OSI standards. OSI-based gateways will be of crucial importance, but they will have teething problems, and it would be risky for organisations to adopt early OSI gateway products for critical applications.

Future trends

Very few user organisations possess the skills and resources to develop their own bespoke OSI-based network infrastructures cost-effectively. Against this background we expect, with the exception of SNA, an increasing proportion of proprietary network architectures to be based on OSI. Nevertheless, we expect SNA to remain the dominant proprietary network architecture. In the short term, the choice of SNA will offer the greatest possibilities for interconnection, together with the side-benefit of access to the IBM-compatible supply industry.

Organisations with requirements for external links should consider using the services of value added network operators.

Choosing to standardise on a proprietary network architecture is a strategic issue, having a fundamental impact on future data processing and office automation directions. The overriding consideration should be to aim at simplicity rather than technical elegance in the network. The trend is firmly towards a single workstation per desk, able to access a variety of in-house and external systems and services. This demands an appropriate 'open' network infrastructure. Foundation Members should continue to support the OSI initiative, while appreciating that it will be five years or more before OSI as such can provide a real alternative to SNA.

CHAPTER 1

THE REQUIREMENT FOR COMPUTER SYSTEM INTERCONNECTION

Most large organisations now have several types of data network installed throughout the business. The earliest networks were installed to provide remote access to expensive mainframe computers. They were based on asynchronous data communications standards and were used predominantly for timesharing applications. Later networks used synchronous data communications standards (typified by IBM's binary synchronous — BSC — standard) to provide faster data flows and better error control. Often, asynchronous terminals were attached to these networks via terminal controllers, using synchronous communications for the links between the controllers and mainframe computers.

The next stage in the evolution of corporate data networks was marked by distributed computing based on minicomputers where local processors were linked together and/or to a central mainframe. One advantage of this type of network was that much of the transaction processing load could be shifted to a local minicomputer, so reducing communication volumes and costs.

Today, data networks increasingly are based on network architectures, either proprietary to a particular supplier or purpose-built by pioneering user organisations. Most computer suppliers have now introduced their own layered network architectures, often as a means of rationalising their existing incompatible data communications standards.

Nevertheless, many organisations still use a mixture of network types, mainly because the early networks have proved to be extremely durable. (Many organisations still use star-shaped networks of asynchronous terminals, for example.)

As computer systems become more important to the running of the business, so the need to interconnect computer systems becomes greater. Even if all the equipment is from a single supplier, interlinking terminals and computer systems attached to one network type to those attached to another network type can be difficult. It is even more difficult to interlink systems if the equipment is from several suppliers.

THE INCREASING NEED FOR COMPUTER SYSTEM INTERCONNECTION

The business needs for computer system interconnection arise in three ways:

—To meet organisational requirements.

—To meet commercial requirements.

—To meet functional requirements.

Organisational requirements

The extent of organisational decentralisation will affect the need for computer system interconnection. An oil company might wish to give greater autonomy to its business units, for example. At the same time, the corporate headquarters needs to be able to respond quickly to changes in the marketplace and so requires up-to-date information from individual business units. One way of achieving this is to link the headquarters and business-unit computer systems.

The need for organisational flexibility also has an impact on the need for computer system interconnection. For example, the organisation might wish to:

—Sell off either a loss-making concern or a unit that is not central to the business.

—Merge business units for synergistic reasons.

—Set up a joint venture with another organisation.

—Move into overseas markets.

—Move into new product areas.

Where computer systems are important to the activities of the business units, any of the above changes could lead to new requirements for computer system interconnection or integration. An organisation therefore needs to build its computer systems so that they can easily accommodate new requirements brought about by business changes.

Commercial requirements

Increasingly, many organisations see their computer systems as competitive weapons that can help them either to reduce the costs of existing goods and services or to provide new services. One of Butler Cox's consultancy clients is a manufacturing company that owns its retail outlets. This manufacturer/retailer is seeking to reduce costs by linking more closely the retail demand for its products with the manufacturing function. This in turn requires that the retail computer system be connected with the production computer systems. The aim is to reduce stock levels and work-in-progress. The company can also react faster to changes in the market.

Another example of commercial requirements leading to a need for computer system interconnection is provided by a major supermarket chain. Each week this organisation receives tens of thousands of invoices from its suppliers, most of which have been produced by the suppliers' computers. On receiving the invoices the supermarket chain has to enter them into its own computer system for processing. The obvious inefficiency in this procedure has stimulated interest in transmitting invoices electronically between suppliers and retailers.

Functional requirements

The need for computer system interconnection may also arise because users of terminals or computer systems originally installed for one purpose wish to use them for additional purposes. A user may wish to use his terminal to access more than one mainframe computer, for example, or a processor originally installed as a standalone device may also need to access a mainframe computer. A rapidly growing requirement of this type is for business microcomputers to be interconnected with corporate mainframe computers, both for information retrieval and for downloading information for local processing.

Although many of the large numbers of business microcomputers currently being installed in large organisations will be used initially as standalone devices, information systems managers anticipate an increasing demand for microcomputer-mainframe links. In Foundation Report No. 43 (Managing the Microcomputer in Business) we said that microcomputer users have to go through a learning process before they perceive the need for such links. We predicted in that report that user demand for microcomputer-mainframe links would, in most organisations, become a significant requirement from mid-1985 onwards.

Another example of the need for additional links from existing equipment is provided by the growing demand to link word processors to each other and also to mainframe computers (to extract financial data for inclusion in a report, for example). Some senior executives cannot understand why it is difficult for their secretaries' word processing systems to do this.

DIFFICULTIES OF ACHIEVING SYSTEM INTERCONNECTION

Despite the growing requirements for interconnecting computer systems, several factors make this interconnection difficult to achieve in practice. Even if an organisation chooses to standardise on computer products from a single supplier, connectability between the different products is not guaranteed. In practice, many organisations use computer systems from several suppliers and, in general, these suppliers use incompatible communications standards. This problem has been exacerbated because many organisations have adopted the IBM microcomputer although they do not use IBM for mainstream data processing. In addition, many of the other popular business microcomputers (Apple, Commodore, Sirius, etc.) are not provided by established data processing companies, and it will not be easy to link these machines to corporate mainframe computers.

Another reason for the difficulties in interconnecting computer systems is the lack of internationally agreed communications standards for interconnection, regardless of supplier. There are de facto standards, however (IBM's 2780 standard for batch file transfer, for example) and transport-level international data communications standards such as X.25. But the use of these standards is not without problems. One surprise to emerge from the setting up of a value-added network service in the United Kingdom was that there are at least 28 versions of IBM's 2780 standard. And, with international standards, different suppliers' versions of X.25 may not be compatible because X.25 offers the implementer options at certain points.

Interconnecting computer systems from different suppliers is also inhibited in some organisations by the lack of suitably qualified staff. Establishing and maintaining such links can be a very time-consuming business, requiring skilled (and scarce) data communications staff. A major problem with multi-supplier networks is that, if the network fails, the user organisation is responsible for establishing which supplier's equipment is at fault. The fault could be in the computer (or terminal) equipment at either end of the link or in third-party equipment (a modem or the PTT circuit, for example). Moreover, the demands on these scarce networking staff are increasing as data networks become larger and more complex, and as a wider choice of publicly available transmission options becomes available.

THE NEED FOR A CORPORATE COMMUNICATIONS INFRASTRUCTURE

The diversity of networks installed, the increasing need for computer system interconnection, and the difficulties of achieving this interconnection are all forcing organisations to consider their total data communications requirements, and to define the communications infrastructure that is needed to meet those requirements. The aim is to provide a communications framework within which future network services and applications can interconnect within the organisation, and in some cases with external services and applications as well. The requirements for such a corporate communications infrastructure might include any or all of the following:

- Providing user access from any terminal to any other terminal, printer or data storage device linked to the network.
- Providing user access from any terminal to any application on any mainframe linked to the network.
- Providing the ability to interconnect an application on one computer with an application on another computer.
- Checking for transmission errors and, where possible, correcting them.
- Removing incompatibilities in character coding and screen formats.
- Providing protocol and speed conversion as necessary.
- Restricting access to a user's data according to the user's wishes.
- Monitoring continuously all parts of the communications system (mainframes, lines, modems, terminals, etc.) and providing warnings of any malfunction in any part of the system.
- Providing flexibility and resilience so that, when parts of the network are inoperative, alternatives can be used.

Faced with the above range of ideal requirements, organisations have either:

- Accepted the fact that they will use several incompatible communications standards, with interconnection being provided where necessary via publicly available products or purpose-built solutions; or
- Adopted an integrated set of communications standards by standardising either on a single communications standard or on a proprietary network architecture.

The major computer suppliers now encourage their customers to adopt the proprietary network architecture approach because, for the suppliers, these are the key networking products that provide a framework within which future interconnecting products can be designed. An additional pressure to adopt a proprietary network architecture comes from the fact that the latest release of an operating system often supports only the latest network architecture software.

Nevertheless, a third option is emerging, and that is to base the organisation's communications infrastructure on open system interconnection (OSI) standards. During the 1970s it was recognised that the large number and diversity of communications protocols provided by the different suppliers was not in the best interests of user organisations. Moreover, it was believed that the confusion in this area was limiting the growth potential of the data communications industry. As a result, the OSI initiative was launched by the International Standards Organisation. The aim is to allow interconnection between computer systems through the use of a single internationally accepted communications standard that would be implemented by all computer manufacturers.

Whichever approach is adopted, a corporate network infrastructure is expensive to build and takes a long time to implement. It involves a long-term commitment to a particular networking approach, and the effects of decisions taken today could still be felt in ten or fifteen years' time.

The purpose of this report is to clarify the four main issues that need to be considered to ensure the long-term viability of a corporate networking policy:

- What approaches to computer system interconnection are organisations adopting today, and what lessons can be learnt from their experiences?
- What is the practical relevance of OSI standards?
- How will suppliers' proprietary network architectures develop in the future?
- What are the implications of the above for a corporate network strategy?

Each of these issues forms the subject of a chapter of the report. The scope of the report includes both intracompany and intercompany communication, but excludes links to computer-aided design or manufacturing applications. The issues are addressed within the context of interconnecting autonomous computers used for business, commercial and administrative purposes.

CHAPTER 2

APPROACHES TO COMPUTER SYSTEM INTERCONNECTION

Previous research by Butler Cox found that differing business and commercial pressures had led different types of organisations to adopt different approaches to computer system interconnection. Two main approaches can be identified, and we have called these the integrated communications approach and the multi-standard communications approach.

Using an integrated approach, an organisation bases its communications infrastructure on an integrated set of standards by choosing either a single standard or a proprietary network architecture. Nevertheless, organisations adopting this approach recognise that a limited amount of non-standard communications is inevitable.

Organisations that adopt a multi-standard approach accept that they will use several incompatible communications standards, including in some cases proprietary network architectures. Equipment using different standards is interconnected where necessary by using either commercially available products or purpose-built solutions.

The integrated approach is more likely to be adopted by organisations where there are strong commercial pressures for the business units to work together. Typical examples would be organisations in the banking, finance, travel, distribution and large-scale production and manufacturing sectors, such as food, pharmaceuticals and engineering. The multi-standard approach is more likely to be adopted by conglomerates where the individual units operate in different types of business; and by large, diverse companies whose activities are nevertheless linked by a common thread. This latter category would, for example, include petrochemical companies that have separate businesses for extraction, refining, chemical production and distribution, but whose common thread is oil.

The research for this report confirmed our earlier work. In the organisations we studied, 70 per cent of the group having strong pressures for business units to work together had adopted an integrated approach, whilst only 15 per cent of the remainder had done so.

We now discuss the main findings of our research under the headings of integrated and multi-standard communications, where we also comment briefly on the growing trend for computer systems in different organisations to be interconnected. Finally, we draw out the lessons that can be learnt from the experiences of the organisations we studied. (The experiences of some of the organisations are presented in detail as case histories in the Appendix to this report.)

INTEGRATED COMMUNICATIONS

Overall, 60 per cent of the organisations we studied had adopted an integrated communications approach. We identified several reasons for this, but the main one was to satisfy a business requirement to transfer data between computers quickly and easily. For a large United Kingdom supermarket chain such as Sainsbury, there was a need to transfer data quickly and easily between distribution depots, supermarkets and headquarters, but at the same time to have systems that could adjust quickly to changes in the market. These requirements led Sainsbury to implement a well-established communications standard (SNA). This company believed that a multi-standard approach would not have provided the required functionality. More importantly, it would have taken scarce development resources away from the more important task of application development.

Another example of the need for close interworking between separate sites, this time in different countries, is provided by SKF. This Swedish-based company has several factories, each of which specialises in manufacturing a limited range of roller bearings. SKF's markets are spread across the countries in which it manufactures, and a considerable amount of marketing and technical data is required to support the product shipments. In order to support these data flows, SKF has adopted an integrated communications approach.

Some organisations have adopted an integrated approach to ensure the availability of critical applications. As computers become more central to the activities of organisations, computer systems have to be more reliable. These systems can be disrupted by events such as machine failure, fire in the computer room or strikes by key staff. In such an emergency, critical applications must be switched from one computer centre to another. Achieving this is much easier with integrated communications.

Finally, an integrated communications approach can arise when the information systems department has adopted a policy of dealing with a single supplier. Frequently, this means adopting a single communications standard.

Of the organisations we studied that had adopted an integrated approach, 60 per cent had implemented a proprietary network architecture, 25 per cent had developed their own high-level communications software, and 15 per cent used earlier (pre-proprietary network architecture) communications standards. The trend is for large organisations to move away from earlier communications standards.

Implementing proprietary network architectures

For established products such as SNA and DECnet, implementing a proprietary network architecture is now regarded as a fairly straightforward process, as confirmed by the SKF and Sainsbury case histories. However, our researches indicated that the characteristics of existing applications, hardware and software may present some difficulties in migrating to a network architecture. Complex or unusual communications requirements also may cause difficulties or delays in implementing proprietary network architectures, however well-established they are.

Implementing a less-well-established product is not so straightforward, however, as the West Midlands County Council case history illustrates. This organisation's experience of implementing ICL's IPA shows that it is important to move forward cautiously. It also shows that the implementers must be chosen carefully, and that management ideally should allow them the freedom to experiment (though in many organisations commercial pressures will prohibit this).

A common finding with proprietary network architectures is the need to train staff adequately if a proprietary network architecture is to be used to its full potential. For one multi-host SNA network it was reported that only about five per cent of the 1,000 terminals on the network accessed more than one application. In this case, it was believed that more users would take advantage of the facility to access any application on the network if they were made more aware of the capabilities of SNA.

Our research also highlighted other important aspects of using well-established communications standards. For an international company like SKF with a need to integrate its computer systems across national boundaries, one of the benefits of SNA is that the same products are available in different countries.

Nevertheless, in opting for SNA, Sainsbury chose not to procure equipment exclusively from IBM. Most organisations who have implemented SNA, in fact, purchase equipment both from IBM and IBM plug-compatible manufacturers. Policies for purchasing SNA products have varied enormously, from one organisation that bought most of its equipment from IBM to others that procured competitively as many products as they could.

Implementing purpose-built high-level communications software

Two of the case histories (SNCF and Société Générale de Banque) are good examples of the situation where equipment from several computer suppliers is used, and an 'open' communication system is required.

SNCF (the state-owned French railway company) is encouraged not to use computers from only one supplier. Users need to access, from the same terminals, mainframes from Bull, IBM and Sperry. In the absence of internationally agreed communications standards, SNCF decided to write its own high-level communications software.

In the case of Société Générale de Banque (SGB) the proprietary network architectures available at the time the communications infrastructure was installed could not provide the required functions. For example, SGB wanted to set up a private X.25 network, and at the time neither SNA nor DECnet supported X.25. This bank's data processing requirements were too diverse to rely on one computer supplier. SGB's primary requirement was to transfer rapidly and easily all banking transactions validated each day in each of about 1,000 branches to one of the regional processing centres for consolidation of both regional and local databases. Because of this, SGB decided to develop its own 'open' communications software based on OSI principles.

In all the cases where commercial organisations developed their own high-level communications software, they had available at the time a sufficient number of experienced networking staff. These organisations admit that because of the present general shortage of such staff, they would be unlikely now to set out to write their own high-level communications software.

MULTI-STANDARD COMMUNICATIONS

Many of today's computer systems were developed at a time when the need to communicate between systems was not paramount. The ability to meet local computing and cost requirements was more important. As a consequence, many organisations now use computers from several suppliers and, if they wish to interlink different computers, they have no choice other than to use the communications standards of several suppliers. Moreover, the cost of modifying existing applications often prohibits migration to equipment from a single supplier.

Multi-supplier environments, and often as a result multi-standard communications, have arisen for a variety of reasons, including:

- —No supplier can satisfy all the organisation's needs. For example, several organisations we spoke with used IBM computers for administrative applications and DEC computers for technical applications.
- —Some organisations wish to avoid becoming dependent on a single supplier in case that supplier goes out of business, or because such a policy might result in the organisation having little or no influence over its supplier. Some public bodies are encouraged not to adopt a single-supplier policy.
- —Some organisations wish to take advantage of competitive procurement.
- —Some organisations wish to take advantage of leading-edge products from innovative suppliers.
- —Multinational organisations want to use the products of the leading local supplier for their local operations.
- —An organisation may have been formed from the merger of two or more companies.
- —Some organisations need to connect to external companies for information exchange.

Of the organisations we studied that had a multi-standard communications environment, 70 per cent had interconnected equipment with different communications standards. In the early days of computer systems interconnection, they wrote their own purpose-built software to interconnect such equipment. This approach was found to be expensive and now has been superseded by the use of commercially available products. Our research showed that there are five main ways of interconnecting computer systems with incompatible communications standards:

- —Transferring files via magnetic tape.
- —Emulating interactive terminals.
- —Emulating remote batch terminals.
- —Emulating terminal cluster controllers.
- —Installing gateways between proprietary network architectures.

Magnetic tape transfer is still commonly used and is a perfectly acceptable way of transferring non-time-critical files.

Terminal emulation

An intelligent workstation (or a business microcomputer) can be used to emulate several different types of terminal. Foundation Report No. 43 described ICl's Conductor software which runs on an IBM PC, permitting the PC to communicate as an IBM 3270 terminal, as a videotex terminal or as a DEC VT100.

Terminal emulation can be an expensive form of interconnection, however, if it requires a separate physical connection to each computer system. Circuit costs can be reduced considerably if the workstation is linked by a single connection to a switched network. One organisation we investigated had set up a private packet-switched network allowing terminals to access ICL and IBM mainframes. Some terminals could emulate both ICL and IBM interactive terminals.

Microcomputer links to mainframes are often achieved by the microcomputer emulating a standard terminal type such as an asynchronous terminal or an IBM 3270. Our research showed that, in late 1984, about 50 per cent of the organisations we studied had successfully connected a microcomputer to a minicomputer or mainframe. This percentage is increasing rapidly, and it will not be long before most organisations will have achieved such links. Some of the organisations said that, at present, their microcomputer-mainframe links were not being used for real applications, but the organisations were anticipating the need for such links from microcomputer users. Most of the links involved the IBM PC.

Remote batch terminal emulation

The most frequently used remote batch communications standard is emulation of IBM's 2780 terminal. This is an old (pre-SNA) standard, and it is often used to transfer files between IBM and non-IBM computers (Hewlett-Packard, DEC, Wang, ICL, etc.). Occasionally it is also used to connect two non-IBM computers.

The 2780 standard is rather limited because it permits file transfer only between directly connected processors. One organisation, we were told, discourages the use of this standard because users demand additional software to enhance its functionality.

Several organisations reported using the 2780 standard to transfer data files between Wang and IBM computers, although the transfer of revisable word-processing files between equipment from these two suppliers is still some way off.

Terminal cluster controller emulation

In order to reduce communication costs, some organisations have installed a device that emulates a supplier's terminal cluster controller and have attached to it inexpensive (frequently asynchronous) terminals that emulate the supplier's terminals. An important extension of this approach is where the device also acts as a terminal controller either for another network or for a standalone minicomputer.

Some organisations with several geographically separated sites have adopted a distributed processing approach with minicomputers linked to a central mainframe for batch processing and a limited amount of interactive work. In such an environment, each minicomputer emulates a remote batch terminal and a terminal controller. The advantage of this approach is that it improves reliability (because the organisation is not dependent on one central computing facility) and reduces the communications requirements, thereby reducing communications costs. On the other hand, it is likely that the staff at remote sites will have limited computing expertise and there is therefore a requirement, as the Sainsbury case history demonstrates, for remote and reliable management of local software changes and equipment monitoring.

Gateways between proprietary network architectures

There are two main approaches to interlinking proprietary network architectures — a dedicated gateway processor or an integrated gateway. The former provides a single point of entry from one proprietary network architecture to another, and so any terminal or application on the first network will be able to access the second network via the single gateway processor. An integrated gateway processes applications in parallel with handling communications between processors and terminals on its own network and the second network.

Gateways reduce communications costs whilst providing enhanced facilities for users because any terminal may be used to access processors in either network. This type of product is relatively new and so there is little practical experience of their use, although two of the organisations we investigated reported satisfactory trials of DECnet/SNA gateways. Because of the growing numbers of implementations of proprietary network architectures, we expect gateways to become important elements of a multi-standard communication infrastructure.

Intercompany communications

Intercompany communications inevitably demand a multi-standard communications approach. Our research showed that an increasing number of organisations are establishing electronic links with customers, suppliers and business partners. Among the industries most advanced in this respect are the banks, with SWIFT; the airlines, with SITA; and the travel industry, which has installed electronic links between tour operators and travel agents.

The pressure to reduce product cycles (a pressure that originally encouraged the development of integrated internal computer systems) is now encouraging organisations to speed up and computerise links with external suppliers and distributors. This is particularly true for the food industry, where pilot electronic links are being established between producers and distributors. These links are being provided by value added networks and, because of the more liberal telecommunications environment in the United Kingdom, they are likely to develop more quickly in that country than elsewhere in Europe.

Early experience is showing that electronic intercompany communications can reduce costs. More importantly, this type of communications promises to provide some businesses with a competitive advantage.

Because of the increasing importance of intercompany communications, one of the case histories describes the Clearing House Automated Payment System (CHAPS). This example illustrates two of the key elements of intercompany communications:

- —The importance of providing a reliable service. Information must be transferred accurately, and the fallback position when errors are found, or in emergency situations, must be automatic.
- —Competition between the users of the service must not be precluded.

THE LESSONS LEARNT

We have found that there are advantages and disadvantages to both of the main types of approach to computer system interconnection. The main disadvantages of adopting a multi-standard approach are:

- —The communications functionality may be limited because there may be restrictions on the applications and databases that can be accessed by a particular terminal.
- —The scope for sharing resources is reduced, which may in turn affect the ability of the network to function when some critical computer facilities are not operational.
- —Management costs are more expensive, because staff need to be familiar with more than one communications standard, and the different networks may need to be managed from different centres.
- —Duplicate channels may be needed between the same two places.

—Costs are increased because the user organisation has to take more responsibility for interconnection issues — deciding which of several parties involved in a link is responsible for a fault, for example. This means that a higher calibre of networking staff is required than with the integrated communications approach, where the equipment is designed to interconnect.

These disadvantages have encouraged many organisations to migrate from a multi-standard approach to an integrated communications approach which, in the majority of cases, means adopting a proprietary network architecture. Our research indicated that this trend will continue; several of the organisations we studied suggested that they would soon be opting for a proprietary network architecture. We did not hear of any organisation that was planning to move in the reverse direction.

However, there are two main benefits to a multi-standard approach. First, the organisation is able to select the best supplier for a particular job. Second, a multi-standard approach is consistent with a decentralised management environment. In such environments, individual business units should be autonomous and have the freedom to install the computer systems they want. For these organisations, the need for computer systems to interconnect is of lesser importance.

There are also drawbacks to adopting an integrated communications approach. The organisation may be locked into one particular supplier's products (and prices), for example, while more appropriate products may be available from other suppliers. In addition, most organisations are unlikely to be able to persuade a supplier to change its mind if the supplier decides to pursue a product strategy that is not aligned with the organisation's requirements. These drawbacks are not as severe for organisations standardising on the major suppliers' products, because competitive equipment is available from plug-compatible and look-alike suppliers.

CHAPTER 3

THE OPEN SYSTEMS INTERCONNECTION INITIATIVE

Although very few organisations today explicitly base their communications strategy on OSI standards, there is widespread support for the OSI concept. At a conference for United Kingdom members of the Butler Cox Foundation held in April 1985, about 80 senior information systems and telecommunications managers debated the motion: "This house believes that OSI is a distraction". The motion was defeated by a substantial majority.

In this chapter we review the progress to date with the development of OSI standards and assess their future prospects. First, we describe the aims of the reference model and the layers defined by it.

AIMS OF THE OSI INITIATIVE

The aim of the Open Systems Interconnection (OSI) initiative is to allow interconnection between computer systems, regardless of supplier, through the use of internationally accepted communications standards that would be implemented by all computer manufacturers. The term 'open systems interconnection' originated within the International Standards Organisation (ISO) and the formal responsibility for the development of OSI standards rests with the ISO, although other international standards organisations are also involved. ISO's membership comprises the standards bodies from more than 80 countries (ANSI — American National Standards Institute, BSI — British Standards Institution, AFNOR — Association Française de Normalisation, DIN — Deutsche Industrie-Norm, etc.).

Because of the increasing involvement of the PTTs in the provision of data communications networks and services, the Comité Consultatif International Télégraphique et Téléphonique (CCITT) has also taken a great interest in the development of OSI standards. The CCITT is responsible for developing recommendations for the design and operation of telecommunications equipment and services. Its members are nominated by national governments, and comprise the national agencies responsible for offering public telecommunications services as well as representatives of equipment manufacturers.

The third main body concerned in Europe with the development of OSI standards is the European Computer Manufacturers Association (ECMA). This association has about 30 members, including the European divisions of the major United States computer manufacturers.

ISO, CCITT and ECMA recognise that they must move in step on the development of OSI standards, and great efforts are made to achieve this. Documents originating from one body are distributed to the other two for comment. In addition, some standards experts contribute to the work of more than one organisation, and this encourages the development of consistent standards.

Activity on OSI standards began in 1977, and significant progress has been made in the past two years. There is a great deal of emotional and political support for the OSI initiative from the major suppliers, PTTs, many governments and some major user organisations. Although the OSI initiative has not yet resulted in specific products, the momentum that has built up will ensure that OSI has a long-term future.

THE REFERENCE MODEL FOR OSI

The complexity associated with exchanging data between computer systems meant that a methodology was required for organising the work into manageable parts. The first task was therefore to create a framework for coordinating the development of data communications standards and for placing existing standards in perspective. The overall framework became the international standard (ISO 7498) entitled "Information processing systems — open systems interconnection — basic reference model". As the title indicates, this standard is only a model, and provides insufficient detail for suppliers to build OSI products.

The OSI reference model defines a structure for communication between computer systems and identifies a set of communications services and protocols that support the structure. A system in this context is defined as a single information processing unit, comprising one or more computers together with the software, peripherals, terminals, human operators, physical processes and means of data transfer that go with them.

Work on absorbing existing standards into the reference model and filling the gaps has now been progressing for several years. Although significant progress has been made over the past few years towards the original aim of open systems interconnection, there is still some way to go.

The concept of layering

An important concept of the reference model is that of layering, which can be understood by considering the way in which two people communicate. Three communications levels (or layers) can be distinguished:

— The cognitive level.

— The language level.

— The transmission level.

The cognitive level encompasses concepts such as understanding and knowledge. For the communication to be meaningful, both people must have a shared understanding and some common knowledge about the topic to be communicated. The language level, however, is not concerned with the content of the communication. The problem at this level is how to put the concepts that need to be communicated into words that will be understood by the recipient. If someone speaking only English makes a telephone call to a hotel in Japan to book a room, and if the receptionist answering the call speaks only Japanese, then it is unlikely the reservation will be made, even though the receptionist will understand perfectly well the concept 'reservation'.

At the third level, the transmission level, neither the concept nor the language has any relevance. The only issue is what medium should be used to communicate the information (telephone call, telex, letter, etc.). The choice will depend, among other things, on the urgency of the communication.

Two important principles should be noted:

— The layers are independent in the sense that the language used is not concerned with the content of the communication, and the transmission medium is independent of the content and language of communication.

— The layers form a hierarchy, with the cognitive layer at the top. Each layer uses the functions of the layers below it.

These two principles are important elements of the structure of the OSI reference model, which we now examine in more detail.

The layered structure of the model

The OSI reference model comprises a seven-layer hierarchic structure. Splitting the model into seven layers simplifies the task of developing standards because it allows standards for each layer to be developed independently. The ISO working groups used formal techniques to determine the number of layers and the functions that should be allocated to each layer. Nevertheless, the number of levels and the allocation of functions to the layers is, to some extent, arbitrary.

Two types of standard are produced for each of the seven layers — services and protocols. The first type relates to the functionality contained within and beneath a particular layer, and represents in abstract terms the services provided to the layer above. The other type of standard specifies the protocol(s) to be used within this layer when two open systems intercommunicate. The differences between the two types of standard can be illustrated by considering the services and protocols of one particular layer of the seven-layer model — the transport layer. The transport protocol provides a mechanism for the reliable exchange of data between processes in different computers. Thus, the transport protocol ensures that data is delivered in the correct sequence, with no loss or duplication. The transport services relieve the higher levels of the need to manage the physical communication facilities.

The seven layers are known as the physical, data link, network, transport, session, presentation and application layers. Before describing each layer, it is worth considering briefly what lies outside the seven layers, because this is also important to the concepts of OSI. Above the top layer (the application layer) there is an application process (an operator at a terminal, for example, or a program that is using the OSI protocols to communicate with an application in another open system). Below the bottom layer (the physical layer) there is the physical medium used to connect the computer systems. This physical medium could be a conventional cable, an optical fibre or the ether for radio communications.

DESCRIPTION OF THE LAYERS OF THE REFERENCE MODEL

The OSI reference model is represented schematically in Figure 3.1. The figure illustrates the hourglass effect of intercommunication, where the scope of the upper and lower layers are broad to accommodate, respectively, future requirements and new transmission options, and the scope of the middle layers is narrow to achieve universal communicability. We now describe the functions assigned to each of the seven layers.

Figure 3.1 Hourglass effect of intercommunication

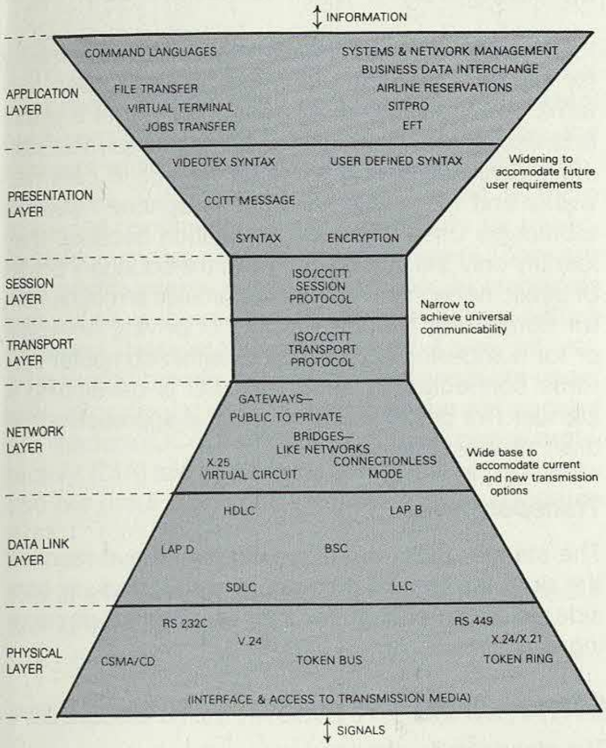

Physical layer

This layer provides the means to transmit raw data (in the form of a stream of binary digits) across a continuous communications path. The standards in this layer can be divided into four areas:

—Mechanical – the dimension of plugs and the allocation of pins, etc.

—Electrical – voltage levels on wires.

—Functional – the meaning of defined voltage levels on specific wires.

—Procedural – the rules that apply to the various sequences in which events may occur.

Data link layer

This layer shields the higher layers from the characteristics of the physical transmission medium, and provides reliable transmission that is free from basic errors. Because errors may occur during transmission, the data link layer provides error detection and correction facilities.

Network layer

The network layer is different from the two lower layers in two very significant respects. It has to be independent of the communications medium employed and it has to provide end-to-end connectivity between OSI systems across interconnected networks of various kinds (circuit-switched analogue, circuit-switched digital, packet-switched and so forth). The network layer provides the necessary routeing functions across the network(s) and provides the higher layers with transmission independence (apart from quality of service).

Transport layer

The task of the transport layer is to provide the quality of service that will satisfy the users' needs. The three main elements of quality of service are:

—To determine the level of error detection required. The requirements for error control may be less stringent for text transmission than for data transmission.

—To reduce costs, by multiplexing multiple transport connections onto a single connection, for example.

—To increase the speed of delivery. The transport layer can use multiple connections in parallel to increase the throughput.

Session layer

The session layer provides interaction management services for the presentation layer by selecting the type of interaction to be used (two-way simultaneous, two-way alternate, or one-way).

Presentation layer

The presentation layer is concerned with managing the problems of encoding and representing data that needs to be communicated between dissimilar systems.

Application layer

The application layer provides the interface between the communications environment and the application process using it. Sometimes the interface is an indivisible part of the application process itself, in which case the process is perceived as residing partly within the application layer. Alternatively, the application process can use an interface that is constructed by selecting a subset of the application service elements made available by the application layer. In this case the application process is wholly outside the reference model.

So, what is the procedure for developing ISO standards, how long does it take, and what stage have the OSI standards reached?

PROGRESS WITH THE DEVELOPMENT OF OSI STANDARDS

There are four stages in the development of all ISO standards: working documents, draft proposals (DP), draft international standards (DIS) and international standards (IS). Progress from one stage to another takes place following approval by ISO members, as registered by ballots of members. If a stage is rejected, further work is carried out on the standard to overcome members' objections. The DIS stage is usually a document that is unlikely to be changed significantly before it reaches the IS stage. It takes at least a year to progress from the DIS stage to the IS stage, but most manufacturers regard the DIS stage as sufficiently stable to enable them to start developing products based on the draft standard. Occasionally the DP stage is sufficiently stable to allow product development to begin.

Standards for the OSI reference model itself and for the individual layers have reached different stages, and the status of each as of early 1985 is reviewed below. The standards relate specifically to directly connected computer systems, although some reference is also made to indirectly connected communications, as in the case of store-and-forward messaging.

The reference model

Although the reference model is now an international standard, it is recognised that there will be a continuing need to extend it, or to provide further descriptions of some of the concepts defined by it. Currently the reference model applies only to computer systems that are directly connected. An extension to the model is being prepared to cover indirectly connected communications. Further addenda are being considered to include naming and addressing, and security.

Physical layer

The main physical level standard is the well-known V.24, which has been implemented in products for many years. V.24 specifies local connection of terminals to computers, and connections with modems for operations over switched or leased lines. X.21 and X.21 (bis) are extensions of the basic V.24 standard.

Data link layer

The ISO data link standard is HDLC (high-level data link control). This standard has also been stable for many years and is implemented in many products.

Network layer

The network service standard is well advanced, and in late 1984 it was out for ballot to become a draft international standard.

Many different protocols are required in the network layer because of the variety of real networks over which network services can be provided. The development of network protocol standards is not as advanced as the network service standard.

ISO has placed on the network layer the responsibility for global intercommunication between open systems. Among other things this means that a global addressing (or numbering) system is required. Existing global numbering systems for public data networks and the public switched telephone networks cannot be used without modification because they identify only the access points at the boundary points of public networks. They do not provide a mechanism for numbering boundary points of private networks or for numbering access points within computer systems connected to networks. ISO is developing a standard for global addressing that is approaching the draft proposal stage.

Transport layer

The standards for the transport layer have reached the draft international standard stage, and are considered stable enough for suppliers to start developing products.

Session layer

The standards for the session layer have also reached the draft international standard stage.

Presentation layer

The presentation layer standards are well advanced, and the draft international standard stage was expected to be reached in mid-1985.

Application layer

Standards for the application layer are, by definition, application-specific, and potentially there are a very large number of them. At present, the standards bodies are concentrating on four main types — file transfer access and management (FTAM), job transfer and management (JTM), terminals, and electronic mail.

File transfer access and management is a basic standard for providing limited file transfer features. It has reached the draft proposal stage and was on course to reach the draft international standard stage in early 1985. Work will start on extensions to the basic standard as soon as it has stabilised.

Full service standards and basic protocol standards for job transfer and management are at the draft proposal stage. These standards should reach the draft international stage in mid-1985. Work on the full JTM protocol standard has started and this could reach the draft proposal stage in 1985.

Five terminal classes are under consideration — basic, forms, graphics, image, and mixed-mode. The most advanced are the standards for the basic class and these were scheduled to reach the draft proposal stage in early 1985.

The CCITT has agreed electronic mail standards (the X.400 series) for messaging systems and internal document structure. The document-structure standard permits a limited amount of editing by the recipient without destroying the format of the document. ISO, CCITT and ECMA are working on more advanced document-structure standards for mixed-mode documents (text, voice, and image). There are also developments under way to specify standards that will allow the recipient to have improved editing features. These standards are some way off yet, even though broadly parallel developments are under way, in particular at IBM. IBM's document content architecture (DCA) and document interchange architecture (DIA) are addressing the same general area and are more advanced than the formal standards work.

This is an example of one of the major problems still facing the OSI initiative.

PROBLEMS STILL FACING THE OSI INITIATIVE

An innovative product or type of service can be incorporated into a supplier's range of products more quickly than a standard can be developed and incorporated into products. Moreover, the development of a standard requires international agreement, which may not be easily reached where there are conflicting national interests, and where suppliers may be trying to protect their own interests.

OSI standards inevitably will always lag behind the best current practice. Nevertheless, 'de jure' standards can be created by committees when sufficient suppliers and users accept that it is mutually beneficial for them (as in the case of X.25 standards).

But problems will persist even when several suppliers provide products based on the same standard because there will inevitably be difficulties in interlinking the different suppliers' versions of the standard. In this context, experience with teletex and X.25 does not set a very good precedent. Protocol testing will go some way towards alleviating these problems, but it will be expensive to set up and will probably require financial support from governments in the early stages. It will also be difficult to ensure the international consistency of protocol testing. Here again, government initiatives will be required to set up the appropriate administrative machinery.

Another problem concerns the limited and perhaps unrepresentative experience on which the standards makers can draw. Some of the OSI standards are 'future' standards in that they are being developed in advance of any widespread experience of the modes of computer use with which they are concerned. The aim of the standards is to permit modes of interworking between computer systems that, except in the case of computer systems from the same supplier, are currently impossible to achieve.

The final problem arises from the enormous scale and complexity of the OSI standards development work. People with the appropriate skills are in short supply. The standards committees are frequently criticised for not including enough user representatives. In practice, it is very difficult for users to make an impact on standards committees. Standards documents are very detailed and are often written in a specialised style and, as a consequence, the time needed to come to grips with the subject matter generally exceeds the time that user representatives can give. Hence, standards committees tend to be dominated by representatives of suppliers who work full-time on standards matters.

SUPPORT FOR THE OSI INITIATIVE

The impetus behind the OSI initiative is considerable — it would not have reached its present position if that were not the case. This support has come from suppliers, PTTs, users and governments. We now review in more detail the commitment of each of these groups to the OSI initiative.

Suppliers

During 1984 most major suppliers of computer equipment made public statements supporting OSI standards. Early in 1984 twelve European computer companies (AEG, Bull, CGE, GEC, ICL, Nixdorf, Olivetti, Philips, Plessey, Siemens, STET and Thomson) announced their support for the OSI initiative.

In June 1984 Fujitsu, ICL and AT&T issued a statement of intent to support OSI standards. In September 1984, IBM announced that it was developing software that would provide mainframe support for selected functions in the OSI transport and session layers. IBM also has announced that it will provide an OSI/SNA gateway when there is a market requirement. Other major American suppliers (including DEC, Sperry and Honeywell) also told us that they will be supporting the OSI initiative.

We believe there are two main reasons for the suppliers' support for OSI standards.

First, the corporate network market cannot support a dozen or so different proprietary network architectures. Some of the smaller suppliers will not have the resources to continue with the development of their network architectures, and they may have to withdraw from the proprietary network architecture market, and possibly from the large-scale data processing market as well. However, if enough suppliers adopt the OSI standards then the size of the market for OSI-based network architectures might become significant. As a consequence, a small supplier would enhance its prospects of long-term survival by adopting OSI protocols.

The second reason arises from the incompatibility of proprietary network architectures. It would be less expensive for suppliers to use OSI protocols for interlinking proprietary network architectures than to develop a large number of bespoke gateways.

IBM, on the other hand, is positioning itself to take advantage of a market for OSI products, if one develops.

PTTs

During the past ten years or so the PTTs have become increasingly involved in providing data communication services, and they expect the revenues from these services to grow substantially. They have already demonstrated their willingness and ability to implement international data communication standards with their support for X.25, and the successful implementation of X.25 standards is an important factor in the credibility of the OSI initiative. The PTTs are continuing their support for international standards by their involvement through the CCITT in the development of OSI standards.

Users

During the research for this report we met with a few large organisations that had developed their own high-level communications software but are now planning to migrate to the higher levels of the OSI model. Many more organisations have made public statements supporting the OSI initiative, but, in the absence of higher-level products from suppliers, have not so far been able to convert this support into action.

Just over half the user organisations we interviewed supported the OSI initiative. These organisations looked forward to the day when their choice of computer could be based on performance rather than, as now, on its ability to interconnect with their existing systems. Others looked forward to OSI standards being used as bridges between proprietary network architectures. Several remarked that the success of X.25 was due to the strong support provided by the European PTTs. They believed that the success of the OSI initiative would depend equally on strong support being provided by the PTTs.

Despite this user support for the OSI initiative, only a few organisations with large corporate networks were planning to migrate to products using the higher levels of the OSI model. Those who, at best, were lukewarm to the OSI initiative supported their views with some or all of the following arguments:

- Most of those who had chosen a single proprietary network architecture thought that the majority of their interconnection worries had been taken care of. They believed that the OSI initiative had little relevance to them.
- Many believed that standards are determined by market forces, not by formal standards committees.
- Innovation will always mean that standards are out of date. Expressed another way, standards will provide only a minimum set of requirements.
- The network management features of early implementations of the full OSI model will be inferior to those of the leading proprietary network architectures.
- Suppliers will seek to protect their market base, and this will inhibit them from implementing the OSI standards.
- Many organisations expressed concerns about the incompatibility both of different suppliers' versions of the same standard and of possible variations between countries. They cited the cases of the X.25 and teletex standards as examples of these difficulties.

Governments

The European Community has consistently supported the OSI initiative. It is concerned about the survival of the European IT supply industry in the face of competition from the United States and Japan. By encouraging European suppliers to adopt OSI standards, it hopes that a sufficiently large market for OSI products will be created, thereby securing the long-term viability of the European IT industry. In addition to encouraging IT suppliers to co-operate in implementing OSI standards, the European Community seeks to:

- Use its own and its national governments' procurements to support OSI products (through the information exchange systems for the Esprit programme, for example).
- Develop administrative procedures that ensure that each European country adopts OSI standards in a consistent fashion.
- Harmonise protocol testing and certification procedures across the community.

In the longer term, the European Community is likely to provide support for demonstrator projects that show how different products from different suppliers can interconnect.

Towards the end of 1984, the OSI initiative also received support from the Japanese government. The Japanese Industrial Standards Committee announced that it would formally adopt OSI as a Japanese industrial standard. This might lead to the Ministry of International Trade and Industry (MITI) adopting OSI standards.

FUTURE PROSPECTS FOR OSI

The OSI initiative has already stimulated a great deal of interest within the information technology community, largely because a wide range of users and suppliers have accepted that the initiative's aim is valid. In addition, the seven-layer model and the standards for the individual layers have provided an intellectual framework for the discussion of network architectures. The success of the X.25 transmission standard also has indicated the potential for the OSI initiative.

In the short to medium term, we expect that electronic-mail products will be the first to be based on the full seven-layer OSI model. These will be available in 1985 or 1986. OSI-based products covering all seven layers of the model for mainstream data processing will be available two years later. These products will be of two types:

— Proprietary network architectures based on OSI standards.

— OSI-based gateways between proprietary network architectures.

CHAPTER 4

PROPRIETARY NETWORK ARCHITECTURES

Many user organisations believe that, for the foreseeable future, systems interconnection issues are best addressed by basing the organisation's network strategy on a supplier's proprietary network architecture. During the past few years, most major computer suppliers have been busy defining their own proprietary network architectures (see Figure 4.1), following the lead of IBM who introduced SNA in 1974. (They are also all involved to a greater or lesser extend with the OSI initiative.) From a supplier's point of view, a proprietary network architecture provides a framework within which future products can be designed so they can interconnect with each other.

Figure 4.1 Computer suppliers' network architectures

Supplier	*Network architecture*
Burroughs	Burroughs Network Architecture (BNA)
Digital Equipment Corporation	Digital Network Architecture (DNA) (implemented as DECnet)
Hewlett-Packard	Distributed Systems Network (implemented as AdvanceNet)
Honeywell/Bull	Distributed Systems Architecture (DSA)
IBM	System Network Architecture (SNA)
ICL	Information Processing Architecture (IPA)
Sperry	Distributed Communication Architecture (DCA)

A proprietary network architecture provides greater functionality than earlier communications products, including:

- —Better utilisation of resources. A network architecture allows the user to access a wider range of resources (mainframes, discs, printers, etc.) than those connected directly to a mainframe. Implementing a proprietary network architecture provides the opportunity of using those resources more effectively.
- —Improved network management. In large networks, the condition and utilisation of the links are important measures for day-to-day management purposes. Proprietary network architectures provide much improved management information compared with earlier communications software.
- —Improved resilience. A network architecture can provide automatic re-routeing if particular links are inoperative. In addition, if a mainframe resource is not available, it may be possible to divert priority work to another mainframe.

This chapter reviews the historical development and future directions of proprietary network architectures, focusing particularly on IBM's SNA and DEC's DNA/DECnet. These two are chosen for detailed discussion because they are the two most developed architectures. They also represent the two main streams of development in this field. The current status of several other major suppliers' proprietary network architectures is also reviewed.

IBM and DEC were the two computer manufacturers that pioneered the concept of network architectures, and their different approaches reflected their positions in the marketplace. IBM, with its large mainframe market, was seeking to rationalise its remote-access hardware and software. The original emphasis of SNA was therefore on remote access to centralised resources. The network was controlled by the mainframe, to which all paths led. DEC, on the other hand, was in a different market. It was predominantly a minicomputer manufacturer and was seeking to develop ways to interconnect its computers to create an environment of enhanced resources. DECnet therefore promoted a distributed computing environment.

THE DEVELOPMENT OF SNA

SNA is recognised as a de facto standard and a yardstick by which to judge other proprietary network architectures. Moreover, it is a proven product, and relatively straightforward to implement, as some of the case histories in the Appendix illustrate. SNA was made public in September 1974, but the first releases, known as SNA-O and 1, provided only limited functionality. The main purpose, however, was to establish the concept of an all-encompassing, wellstructured architecture. In the early releases a network consisted typically of a host computer and a channel-attached communications controller that was linked by leased lines to several remote cluster controllers. Terminals attached to the cluster controllers could access multiple applications on the host computer but, in the early releases of SNA, the only terminals available were for banking and retail applications.

The next major release (SNA-2) in 1976 allowed for remote communications controllers (but never more than one local and one remote controller in tandem), the direct attachment to the host of local cluster controllers and the support of public switched communication lines. By this time general-purpose terminals were also available.

SNA-2 did not greatly extend the scope of the architecture, but SNA-3 (for which products were first delivered in 1977) did. With SNA-3 it became possible to interconnect host computers via their local communications controllers, and for terminals attached to one host to access applications on other network hosts. Each of the hosts was responsible for controlling a set of resources (discs, terminals, etc.), known as a 'domain'. The domains co-operated on a non-hierarchical basis, but within a domain centralised management was adopted, as in earlier versions of SNA.

Another important feature of SNA-3 was the fact that, once a terminal was in 'session' with an application in a domain other than its own, the data traffic to the application bypassed the terminal's host domain.

SNA 4.1 introduced the important feature of parallel sessions between application programs on different hosts. This enabled as many sessions to be established as were necessary to acquire the required bandwidth between application programs. SNA 4.1 also provided better network management tools, including the network communication control facility (NCCF) and the network problem determination application (NPDA) packages.

SNA 4.2 introduced the possibility of using more than one physical path between two network nodes and of several end-to-end network connections between two network nodes, independent of the network in between. These facilities provided better load balancing and network availability. This release also enabled an unlimited chain of communications controllers to be connected to the host. Communications controllers could also be interconnected to form a mesh network.

In 1983 IBM announced software products that made it possible to interconnect individual SNA networks without affecting the characteristics of each network. This is a useful feature both for those organisations that might wish to incorporate the SNA network of a newly acquired subsidiary and for intercompany communications. The software products handle directory services and access security so that only designated resources within a network can be accessed from another network.

The evolution of SNA is summarised in Figure 4.2. For almost a decade, this architecture has been the cornerstone of IBM's teleprocessing product strategy.

Figure 4.2 The evolution of SNA

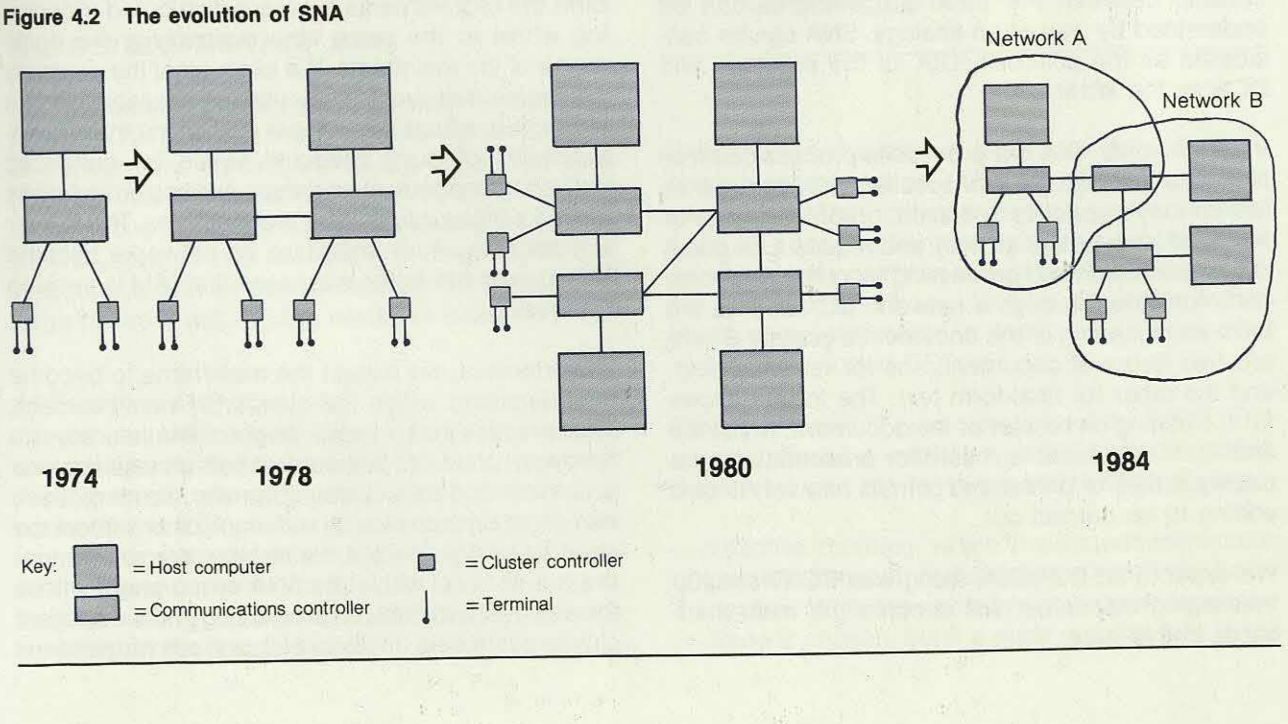

It is now the most widely used network architecture, with more than 10,000 SNA hosts. The evolution of SNA since 1974 has made it a powerful and flexible product. Its networking capability has progressed from a limited tree structure to a fully meshed network of host computers, with alternative routeing and the ability of each terminal to access any host on the network. Also, when SNA was first released, it supported only analogue leased lines. Today it also supports circuit-switched and packet-switched services. Finally, a widely praised SNA network management feature has been developed.

FUTURE DIRECTIONS FOR SNA

Despite its success to date, the pace of development of SNA is not likely to slacken. We believe that the most significant future developments will be in the areas of:

— Integration with the DCA/DIA architectures.

— Local area networks.

— Distributed intelligence.

— The connection of non-SNA devices.

— Support for OSI standards.

SNA, DCA, DIA integration

With organisations installing ever-expanding networks, IBM will want to exploit the SNA infrastructure for office automation applications. For office applications involving document distribution, IBM will be seeking to integrate SNA with two other architectures — Document Interchange Architecture (DIA) and Document Content Architecture (DCA). The relationship between the three architectures can be understood by way of an analogy. SNA can be considered as the postman, DIA as the envelope and DCA as the letter.

In other words, DIA is a process-to-process communications architecture that specifies how documents (which may consist of text, data, graphics, image or voice, or a mixture of all five), and requests for document distribution and processing functions, are to be communicated through a network. DCA defines the form and meaning of the document's content. There are two forms of document: one for revisable text, and the other for final-form text. The former allows for full editing on receipt of the document, whilst the final-form document is meant for presentation on a display screen or printer and permits only very limited editing to be carried out.

We expect that DCA/DIA, along with CCITT's X.400 messaging standards, will become the main standards in this area.

As yet, IBM has only partially implemented DCA/DIA in mainframe software products and intelligent workstations and minicomputers. We expect that IBM will be integrating these office information architectures with the SNA environment as part of its overall strategy for the office system marketplace. Moreover, many of the products will be mainframe-based in order to sustain the sales of IBM mainframes.

Local area networks

IBM has also started to address the problem of linking together the various office (and factory) systems it has developed. To do this it has announced three different local area networks. These are:

— A general-purpose token-passing ring network for linking computers and workstations. This local area network is used to exchange information and share resources. It is the most strategic of the company's three local area networks. It is also meant to provide the backbone connection for the three networks. So far only the cabling specification for this network has been announced.

— LocalNet PC, a low-cost network that allows PC users to share files and printers and send messages from one PC to another. Of the three networks, this is the only one completely specified.

— An industrial local network. This will integrate factory floor data collection.

We expect that, over the next few years, IBM will be seeking to develop these local networks and to integrate them with the SNA environment.

Distributed intelligence

IBM has the problem of aligning SNA to accommodate the requirements for more distributed processing whilst at the same time maintaining the dominance of the mainframe. We expect that the company will tackle this problem by making available certain strategic products that will run only on the mainframe. Examples of such products would be enhanced network management packages and implementations of IBM's office information architectures. The former will become more important as networks become bigger, and the latter is an area that IBM is seeking to develop.

Nevertheless, we expect the mainframe to become less dominant within the overall SNA environment. This process has already begun. IBM has several hardware products that can act both as a standalone processor and as a cluster controller. Some of these can communicate directly with each other without the need to be linked via a mainframe. We expect that the future trend within the SNA environment will be towards more distributed processing (more intelligent cluster controllers, for example), and communications between distributed processors and between terminals to bypass the mainframe.

The connection of non-SNA devices

Originally SNA was designed to accommodate synchronous terminals, but many organisations still have large numbers of asynchronous terminals that need to be connected to the corporate network. IBM has now recognised this fact and has announced a powerful cluster controller that, among other things, allows asynchronous terminals and binary synchronous 3270 terminals (as well as standard SNA terminals) to be connected to an SNA network. This development is seen by IBM as a key step in its move towards integrating office systems within the SNA environment. Makers of protocol converters, however, will no doubt see it as a move by IBM to take over their markets.

Support for OSI standards

IBM supports the use of OSI standards for communicating between an SNA network and other network architectures. Moreover, it is developing software that will provide mainframe support for OSI Levels 4 (transport layer) and 5 (session layer). Those wishing to use these products when they become available will have to write their own higher-level software. In our view, these moves by IBM represent an attempt to position the company to take advantage of an OSI market, if the opportunity arises. We do not believe that SNA will evolve towards the OSI standards because, for the foreseeable future, the OSI standards provide considerably less functionality than SNA. Moreover, OSI standards are some way from being implemented for mainstream data processing applications.

Figure 4.3 The evolution of DECnet

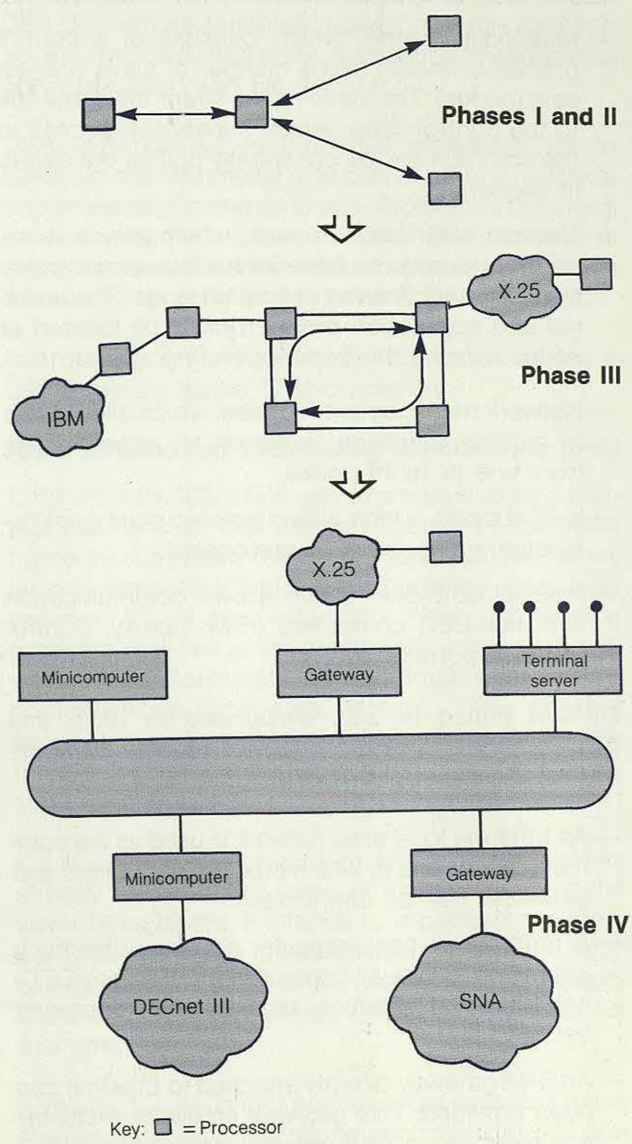

THE DEVELOPMENT OF DNA

DEC announced its Distributed Network Architecture (DNA) in 1975. It was first implemented as a series of software and hardware products under the name DECnet. The historical development of DECnet is summarised in Figure 4.3, which shows that it has evolved from a small network providing only limited functionality between adjacent nodes to a potentially large network with a wide range of features.

DECnet Phase I implemented only a small subset of the global architecture by allowing similar computers using the same operating system to work together. A program on one computer could make a logical connection to a program on another computer so that interactive communications could take place between them. These connections operated only on a point-to-point basis — for the computers to communicate, they had to be directly linked.

In 1978 a new version of DNA was defined, and a new implementation, DECnet Phase II, was announced. This still operated on a point-to-point basis but added file transfer, remote resource access, remote batch command and submission, and down-line loading to the Phase I facilities. By the end of the Phase II developments, a total of eight operating systems supported DECnet.

DECnet Phase III was announced in 1980. This version offered a maximum of 255 network nodes, and introduced the following additional features:

—Adaptive routeing, which makes communication possible between nodes that are not directly connected. Wherever possible the least expensive route is chosen, but if a node or a link fails in the least-cost path, an automatic change is made to the next best least-cost path.

—Multi-point linking, which consists of a control processor (node) and a number of slave processors (nodes). The slave nodes share the same link to the control node, which supervises access to the common link by continually polling the slaves in sequence.

—Network-command terminals, which allow a terminal at one node to have interactive access to a system or application at another node. The terminal and application program must be located at nodes running the same operating system.

—Network management facilities, which allow users to monitor message traffic on all network lines from one or more nodes.

—X.25 support, which allows point-to-point communication using the X.25 protocols.

—Internet software, which allows communication with non-DEC computers (IBM, Sperry, Control Data Corporation etc.).

DECnet Phase IV was announced in 1983, and products were released in 1984. The main elements of this, the latest phase of DECnet, are:

—An Ethernet local area network is used as the communications bus to which processors, routers and gateways can be connected.

—A terminal on one computer can now communicate directly with all computers running any DECnet Phase IV product, regardless of operating system.

—An SNA gateway directly attached to Ethernet has been provided. This gateway emulates a cluster controller on an SNA network, and allows a DEC VT100 terminal to emulate the basic IBM 3270 terminal. It also allows a 3270 terminal to appear to the DEC environment as a VT100. In addition, the gateway provides remote job entry into the SNA environment, and interlinking between DEC's office systems and IBM's office information architectures (DCA/DIA).

—An X.25 gateway directly attached to Ethernet.

—Terminal servers attached to Ethernet.

—Compatibility with DECnet Phase III. Phase IV nodes can be part of a Phase III network.

—The maximum number of nodes has been increased to 1,023.

A fundamental change in direction took place in the transition from DECnet III to DECnet IV in that a local mesh network has been replaced by a communications bus to which processors and gateways can be attached. The major benefit of this change is that most of the communications functions were removed from processors and placed in specialised routers and gateways. This will help counteract the criticism that the communications functions in DECnet, particularly for interactive work for remote users, used too much central processor power. Other benefits of using a communications bus are that fewer network links are required (but with some loss in network resilience), and flexibility is increased because the terminals accessing a particular node need not be close to that node.

FUTURE DEVELOPMENT OF DECNET

An important theme of the development of DECnet has been the ability to link it to non-DEC environments in order to satisfy user requirements for systems interconnection, and this trend will continue in the future. We also believe that the protocols in future versions of DECnet (perhaps DECnet V) will be based on OSI standards.

The most important non-DEC environment is IBM's SNA, because of the large IBM customer base. DEC perceives the IBM environment as providing the large number-crunching computers as well as the large databases, with DEC providing the intermediate and small systems. DEC already has available a gateway into IBM's SNA. We believe that the facilities of this gateway will be enhanced to provide access to any new significant IBM features as they become available (links with IBM's token passing ring, for example). In addition, we believe that DECnet's connectability will be extended by providing links to other types of local area networks, to private branch exchanges and to ISDN facilities.

DEC, along with most other computer suppliers, has declared its active support for the OSI initiative. Nevertheless, some of those we spoke with during our research questioned the extent of the company's commitment to OSI. They thought that DEC might be reluctant to open up its market by adopting OSI standards, because this might result in a net loss of market share. Despite this, we expect that DECnet will, in the future, be based on OSI standards for all seven layers, but will be enhanced to provide extra proprietary features. The communications bus will also provide both a gateway based solely on OSI standards for communication with other proprietary network architectures, and a gateway to DECnet IV. The latter, together with continued support for DECnet IV, will protect DEC's customer base.

OTHER PROPRIETARY NETWORK ARCHITECTURES

We now note the main features of Sperry's DCA (Distributed Communications Architecture), Honeywell/Bull's DSA (Distributed Systems Architecture) ICL's IPA (Information Processing Architecture), Burroughs' BNA (Burroughs Network Architecture) and Hewlett-Packard's AdvanceNet. We describe the services provided by these proprietary network architectures, the types of computer on which they are implemented, and the supplier's attitude to the OSI initiative.

Sperry's Distributed Communication Architecture (DCA)

Sperry's DCA was introduced in 1976 following the recognition of the need for a company-wide unifying set of communications standards interconnecting the company's different products. The architecture has been implemented in a set of products that allows terminals attached to terminal controllers to access any mainframe via linked communications controllers (known as DCPs — distributed communications processors).

The advent of the OSI model in 1979 caused Sperry to re-evaluate its position, and it decided to adopt the terminology of the model, and to incorporate in DCA the standards for the individual OSI layers as they become available. The lowest three layers of DCA are now very similiar to the equivalent OSI standards for these layers, and in late 1984 a pilot DCA implementation of the OSI transport layer (DIS 8073) was being tested in Norway. DCA products for this layer are expected to be available sometime in 1985. Sperry anticipates that DCA session-layer products will be available in 1986.

Sperry recognised from the outset that its proprietary network architecture would need to allow the connection of non-DCA products, both for incorporating older equipment and for co-existing with other proprietary network architectures.

For example, Sperry has developed products that will allow IBM cluster controllers, IBM remote job entry terminals and Sperry terminals emulating IBM terminals to access SNA and binary synchronous networks via a network of DCPs. The company is also developing products that will allow IBM terminals to emulate Sperry terminals on a DCA network.

Honeywell/Bull's Distributed Systems Architecture (DSA)

Honeywell and Bull are both strong supporters of the OSI initiative. DSA has the same seven-layer structure as the OSI model and the intention is to evolve DSA towards OSI standards. DSA provides peer-to-peer communication for the whole range of Honeywell/Bull processors. For the large (mainframe) processors, the lower four layers of DSA are implemented on a front-end processor and the upper three layers on the mainframe. For the smaller minicomputers, all seven layers are implemented on the processor itself. The main services provided by DSA are remote terminal access, remote database access, remote batch entry, application-to-application communication and file transfer.

DSA incorporates international standards up to Level 3 of the OSI model, and part of Level 4 has been implemented in some products. Prototypes of Level 5 are being tested. Later in 1985, Bull expects to demonstrate an X.400 (electronic mail) application incorporating parts of the Level 5 OSI standard.

Interconnection with IBM and ICL networks via a DPS 6 minicomputer is also possible.

ICL's Information Processing Architecture (IPA)

Until recently, ICL's IPA was implemented as a hierarchical network that allowed access to the mainframe by dumb terminals (via a terminal controller), by minicomputers and by a proprietary local area network. ICL is now extending IPA by introducing products that will allow computers, local area networks and terminals to communicate via a contention-bus local area network. IPA's network services now include remote session access, file transfer, remote batch-file entry and interprocessor communication.

ICL is a strong supporter of the OSI initiative, and has already implemented some of the standards of the lower three layers. It intends to implement the standards for the higher layers as soon as they are sufficiently stable — some of ICL's latest networking software already incorporates Level 4 standards, for example.

Although ICL no longer provides IBM-lookalike mainframes and communications controllers, it still provides software that allows ICL minicomputers and intelligent terminal controllers to emulate IBM cluster controllers. In addition, ICL intends to provide an SNA/IPA gateway for its contention-bus local area network.

Burroughs Network Architecture (BNA)

BNA provides networking capabilities on a peer-to-peer basis between Burroughs' range of mainframes and minicomputers. Terminals link directly into these computers, there being no communications controller. Network services provided by BNA include remote terminal access, file transfer, job transfer, remote file access and interprogram communications.

BNA is a layered architecture, although the functionality of each BNA layer does not correspond exactly with the OSI model, especially for the upper four layers. Burroughs is committed to implementing OSI standards, however, and the company has given priority to implementing the OSI standards for remote terminal access and file transfer.

Gateways to other suppliers' computer systems (IBM and ICL, for example) can be provided by a communications processor attached to a BNA network. For access to an SNA network, the gateway can emulate a terminal cluster controller and a remote job entry terminal.

Hewlett-Packard's AdvanceNet

Hewlett-Packard's AdvanceNet products provide OSI-based networking capabilities between the company's different computer systems: Hewlett-Packard 3000 (general commercial applications), Hewlett-Packard 1000 (technical and automation applications) and Hewlett-Packard 9000 (CAD/CAE applications). Interprocessor links can be either by point-to-point circuits, X.25 networks or Ethernet-like (IEEE 802.3) local area networks. AdvanceNet features include network file transfer and virtual terminal remote database access. In addition, Hewlett-Packard plans to provide multi-vendor local area network facilities.

The company's minicomputers can interconnect with IBM equipment in a variety of ways. For example, Hewlett-Packard 3000s can emulate either an SNA cluster controller for interactive terminal access to an SNA network, or an IBM 8100 for remote job entry, with remote access being provided by an AdvanceNet network. In addition, any Hewlett-Packard minicomputer can be accessed by an IBM PC emulating a Hewlett-Packard terminal.

TRENDS IN THE PROPRIETARY NETWORK ARCHITECTURE MARKETPLACE

We have identified three main trends in the proprietary network architecture marketplace:

- —There will be a greater emphasis on distributed intelligence.
- —There will be an increasing need for interconnectivity between different proprietary network architectures.
- —An increasing number of suppliers will adopt the OSI standards.

Distributed intelligence

In an IBM environment the mainframe is dominant, and SNA is only slowly coming to terms with distributed processing power. DECnet, on the other hand, has always catered for a distributed intelligence approach. The overall trend will be to allow for more distributed processing within a supplier's proprietary network architecture. ICL, for example, is moving from a hierarchic networking approach to an approach similar to that of DEC, where the main local communications are via a communications bus, to which processors and gateways are attached.

Connectivity between network architectures

Products for interlinking the computer systems of different suppliers have been available for some time. We expect the market for these products to increase because of the growing requirements for systems interconnectivity. The most important part of this market is for products that provide links into the SNA environment. In general, IBM does not actively develop links for other suppliers' products, although it publishes sufficient material for them to be developed by other computer suppliers or by third parties.

Most, if not all, other computer suppliers see links between their proprietary network architectures and SNA as being important, and are developing products to provide such links. Often, these links provide more than simple emulation of an IBM terminal. Increasingly, they allow IBM terminals to emulate a terminal in the foreign proprietary network architecture as well.

The main methods of linking network architectures are:

- —Through terminals that can access more than one network architecture.
- —By emulating terminal controllers, so that asynchronous terminals, or microprocessors or terminals from supplier A, can emulate supplier B's terminal on supplier B's network.
- —By emulating communications controllers. Sperry, for example, has emulated the teleprocessing system of an IBM mainframe and communications controller so that a network of Sperry's communications controllers appears to the IBM mainframe as a communications controller from another part of the SNA network.
- —Through value added networks that provide protocol conversion facilities.

In the longer term (five years or more), we expect that OSI-based gateways will provide the links between proprietary network architectures. Until then, the above methods will continue to be important because of the increasing requirement for interconnectivity. The method chosen will depend on three main factors — the functionality required, the costs, and the types of computer system to be linked.

Increasing use of OSI standards

The OSI initiative already is supported by all computer suppliers. As the OSI standards are agreed, suppliers will adopt them for one or both of the following purposes:

—They will base their proprietary network architectures on OSI standards, as they become available, but will provide proprietary enhancements over and above the basic set of standards.

—They will develop gateways, using the basic OSI standards, for linking with other suppliers' network architectures.

Nevertheless, we expect that SNA will continue to be developed largely independently of OSI standards. This is partly because the cost of bringing it into line with OSI would be enormous, but mainly because SNA has a decade's head start.

CHAPTER 5

IMPLICATIONS FOR CORPORATE NETWORK POLICY

Our research confirmed that corporate policies for data networks will increasingly be based on computer suppliers' proprietary network architectures — and these will be essentially of two types: SNA and OSI-based.

At present, the OSI standards provide less functionality than existing proprietary network architectures. Even so, many suppliers (other than IBM) will base their proprietary network architectures on the OSI reference model, but from the outset will provide more functionality than is specified in the OSI standards. In this respect, the situation will be the same as has existed for many years with ANSI and Codasyl standards for Cobol and database management systems.

The implication is that there will be no such thing as a universal standard OSI product. Suppliers basing their network architectures on OSI principles will always provide more (or different) functionality than is specified in the OSI standards, both in response to perceived market demands and because new technology can be incorporated into their products faster than agreement can be reached on the relevant standard.

Thus, there may be significant differences in individual OSI-based network architectures, which means that the interconnection of two such networks will not be a trivial task. Furthermore, cost-effective interconnection may mean reducing the level of inter-network functionality, even when OSI-based gateways are used.

We expect OSI-based gateways to be of crucial importance for interlinking different proprietary network architectures, but we anticipate considerable teething problems in making gateways from different suppliers interwork reliably. Because of this, we advise Foundation members to avoid early OSI gateway products for critical data processing applications. By the early 1990s, however, we expect that OSI-based gateways will have matured sufficiently for them to be widely used.

Nevertheless, we expect SNA to remain the dominant proprietary network architecture, even though its development will continue to be largely independent of the OSI initiative. Furthermore, we expect that an increasing proportion of organisations will base their network policy on SNA, for the following reasons:

— For the foreseeable future, SNA will continue to be the most highly developed proprietary network architecture. In particular, it has good network management features that have been widely praised by users and that will become more important as networks become larger and more complex.

— It is a de facto standard. In the absence of a fully fledged ISO standard, SNA is the next best thing.

— Pressures to interconnect systems are increasing. For users with both IBM and non-IBM equipment, the advantages of adopting SNA are significant because of the range of worldwide services and support IBM can provide. Moreover, once an organisation chooses IBM as its company standard, there is a low probability that it will change to standardising on products from another (non-IBM-compatible) supplier.

— IBM has important developments in hand that will enhance the functionality of SNA. Examples include the integration of DCA/DIA with SNA, greater use of distributed processing, and the acceptance of asynchronous terminals into the SNA world.

The OSI initiative will have a significant impact on other suppliers' proprietary network architectures, however, and SNA increasingly will have to co-exist in an environment in which most (and probably all) other architectures are based largely on OSI standards.

Most Foundation members now have substantial and well-established data networks, more often than not based on one or more proprietary network architectures. And many of those not using such an architecture at present are actively considering doing so. The alternative to choosing one or more proprietary network architectures is for an organisation to develop its own OSI-based network infrastructure. Very few organisations will have the skills and resources available to make this a viable and cost-effective approach, and it is not an approach we would recommend.

If your organisation is strongly committed to a particular supplier's networking approach, we see no good reason for changing to another approach, unless there are overwhelming pressures to do so. If, however, you have adopted a multi-supplier (which often means a multi-standard) networking approach, we believe that you should seriously consider reducing the number of suppliers to one or two.

In deciding which network architecture to adopt, you need first to consider the requirements for interconnection — both in terms of workstations accessing host applications and for computer system interconnection. If most of your in-house interconnection requirements will be to interlink workstations and computer systems from the same supplier, that supplier's proprietary network architecture will be the best choice. If the requirement is to interlink workstations and computer systems from different suppliers, a network architecture that maximises the interconnection possibilities needs to be chosen.

At first sight, it may appear that choosing an OSI-based architecture will maximise the potential for computer system interconnection. However, each supplier will be providing its own proprietary features in addition to the OSI standards, and, as a result, there will always be greater functionality within a network architecture than between them. But, because SNA is, and will remain, the dominant network architecture, most other suppliers will want to ensure that their own network architectures can interconnect with SNA. There will therefore be as much, if not more, interconnectivity between SNA and other (OSI-based) architectures as there will be between different OSI-based architectures. The implication is that, for the foreseeable future, the greatest interconnection possibilities will be provided by choosing the dominant proprietary network architecture — SNA.

Choosing to standardise on SNA has a further advantage, in that there is a thriving plug-compatible supply industry for IBM products. This means that an organisation is not constrained to purchase equipment only from IBM. The case histories in the Appendix show that some organisations have chosen to standardise on SNA, but also to purchase equipment from plug-compatible suppliers. Choosing to standardise on another supplier's network architecture and products may well restrict your choice of computer equipment to that supplier alone.

If some of your computer system interconnection requirements are for links outside the organisation, you should consider using the services provided by value added network operators. The range of industry-specific VAN services on offer throughout Europe is set to grow rapidly, although outside the United Kingdom they will more usually be provided by the PTT rather than by licensed network operators.

The overriding consideration, however, should be to adopt the simplest technical solution that matches your particular interconnection requirements. Data networks are complex, and the temptation to adopt the most elegant solution for its own sake should be vigorously avoided.

Choosing to standardise on one or more proprietary network architectures is a strategic issue because it will have a fundamental impact on future directions both for data processing and office automation strategies. The trend is firmly towards a single workstation per desk, able to access a variety of in-house and external systems and services. This objective cannot be achieved without an appropriate network infrastructure. Furthermore, standardising on a particular network architecture implies a long-term commitment to the supplier and its networking approach. Replacing one network architecture with another will involve a major upheaval, even if both architectures are based on OSI principles.

Our research showed that there is considerable support from user organisations for the OSI initiative. At present, this support is at the conceptual level; very few organisations have yet taken practical steps as a result of the OSI initiative. The underlying belief is that OSI's greatest potential lies in forcing computer suppliers (even IBM) to adopt a more 'open' approach in their networking products. For this reason alone, Foundation members should continue to support the OSI initiative, even though it is likely to be at least five years before OSI as such can provide a real alternative to SNA.

APPENDIX

CASE HISTORIES

The case histories have been selected to illustrate the different approaches to computer systems interconnection that are being used by different types of organisation. All of them indicate that the organisations concerned had either completed, or were about to embark on, a significant change in the way they approach computer systems interconnection. We believe this is generally true. Most large organisations are actively developing their corporate investment in computer systems and computer networks in order, for example, to improve their responsiveness to market forces, or to improve the internal management of networks.

Two of the case histories (SKF and West Midlands County Council) describe the implementation of a proprietary network architecture. For established products such as SNA and DECnet this is now regarded as a fairly straightforward process. However, the constraints of existing applications and systems may cause some difficulties in migrating to a network architecture. A common element of these experiences was the need to train staff adequately if a proprietary network architecture is to be used to its full potential. The SKF case history also describes a successful attempt to link SNA and DECnet.

The case histories also highlight several important points concerning the use of one supplier's protocols. For an international company like SKF that needs to integrate its computer systems across national boundaries, one of the benefits of SNA is that the same products are available in different countries. In opting to standardise on SNA protocols, J Sainsbury nevertheless chose to procure equipment competitively from IBM and plug-compatible manufacturers. The anonymous case history illustrates the difficulty of interconnecting different suppliers' versions of the same communications standard, in this case X.25.

Two of the case histories (SGB and SNCF) are good examples of the situation where, although more than one computer supplier is used, an 'open' system is required. In the case of SGB the proprietary network architectures available at the time could not provide the required functionality. In the absence of OSI products, both organisations had to develop their own communications software, a task that can be undertaken only by large organisations.

The case histories also provide insights into various aspects of network management. SKF pointed to the complexity of some of IBM's network management software. For SGB, the realisation that telecommunications costs were 10 per cent of all non-salary IT costs had encouraged the bank to look more closely at network management issues. For Air France, the expense of managing several distinct networks had encouraged it to migrate to a more integrated (and hence easier to manage) network.

The case histories also illustrate that user attitudes to the OSI initiative can range from utterly sceptical to highly supportive. Some organisations see no value at all in the OSI initiative, whilst others are drawing up plans that eventually will lead to them implementing products covering the higher levels of the OSI model.

AIR FRANCE

Air France is the French flag carrier providing international passenger and freight airline services. We met with M Gruson, telecommunications manager, to hear how Air France's telecommunications requirements were being met by using the protocols of several suppliers.

Development of networking

Air France has a long history of computer networking, beginning in 1963 with eight remote terminals linked to an IBM 7000 Series computer. Since then many networks have been installed, the principal ones being:

- —A Sperry network used for processing reservations, fare calculations and flight check-in. The network comprises controllers in Nice and about 12,000 terminals worldwide linked to the host computers by telecommunications loops using a specially designed protocol.
- —An IBM SNA network used for payroll, accounting and freight applications. This network comprises

three mainframes and local communications controllers in central Paris (Vilgenis), and remote communication controllers at Orly Airport and Charles de Gaulle Airport. There are 1,700 terminals in France and a few overseas as well.

—A network of some 15 Bull Level 6 minicomputers linked to about 100 terminals situated at Orly and Charles de Gaulle Airports. These minicomputers and their terminals are used for applications such as aircraft maintenance and staff records.

—A DECnet network linking three VAXs used for office automation applications.

Links between networks

Data files were first transferred between the IBM and Sperry networks in the early 1970s, using a specially defined airline protocol called P1024. This protocol is used also to link the IBM and Sperry systems with SITA, the international airline messaging system.

In 1978 a special terminal was developed for the Sperry network. This terminal could also be used to access IBM applications. By 1982 some terminals attached to the Level 6 minicomputers could also access applications on the Sperry systems. The IBM system is linked to the DEC and Bull systems using IBM's 2780 protocol. The main links between the various networks are shown in Figure A.1.

Current concerns and future developments

The growth in the number of applications has given rise to several concerns:

—Users want to access more and more applications. Terminals designed for one application cannot always access another application, with the result that the number of terminals has grown significantly.

—The growing number of networks has led to duplicated parallel circuits. This approach is particularly expensive for international circuits.

—Each network has its own management centre. These centres are in different places, and this increases the labour costs of managing the networks.

—The amount and diversity of equipment at each site make system maintenance expensive.

These concerns led Air France to reconsider its overall network strategy. As a result the company decided to reduce the diversity of its systems and migrate towards a more integrated network based on OSI concepts. The stages in this migration are as follows:

—Set up a primary X.25 network based on six main nodes, as shown in Figure A.2, and connect all existing terminals and computers to these nodes through satellite network processors.

—Extend the primary network by installing small satellites on the telecommunications loops and adopting the X.25 protocol instead of the existing specially designed protocol.

Figure A.1 Main Air France network links

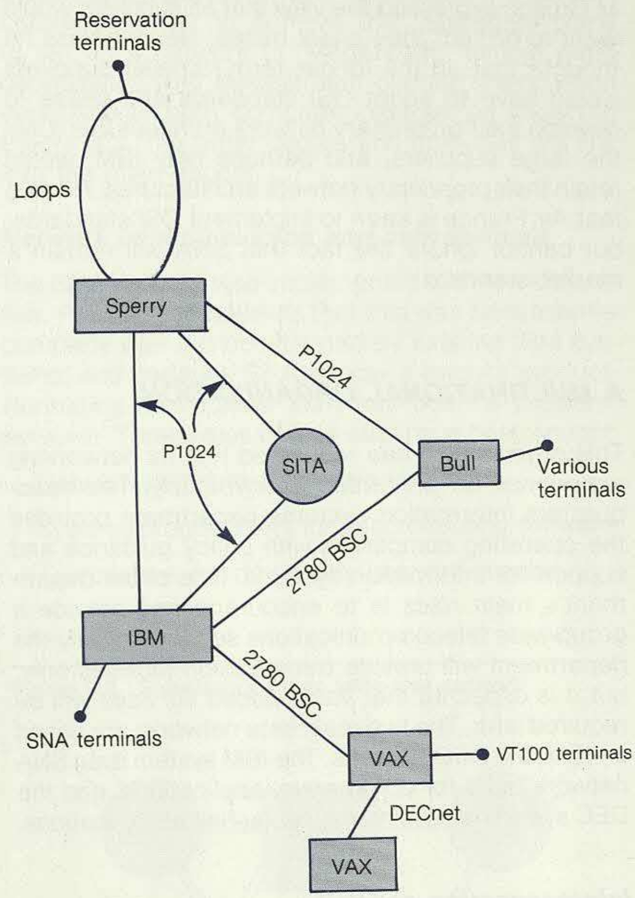

Figure A.2 Air France's proposed primary X.25 network

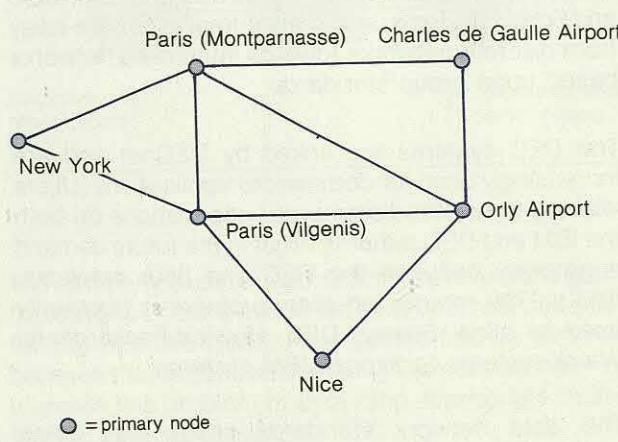

Attitude to the OSI initiative

M Gruson expressed the view that all suppliers would want to protect their client bases. Nevertheless he thought that, in the longer term, smaller suppliers would have to adopt OSI standards and cease to develop their proprietary network architectures. Only the large suppliers, and perhaps only IBM, would retain their proprietary network architectures. He said that Air France is keen to implement OSI standards, but cannot ignore the fact that SNA will remain a market standard.

A MULTINATIONAL ORGANISATION

This organisation has requested that its networking experience be presented anonymously. The headquarters information systems department provides the operating companies with policy guidance and support for information systems. One of the department's main roles is to encourage and provide a group-wide telecommunications service. Initially the department will provide transmission facilities only, but it is expected that value added services will be required later. The two main data networks are based on IBM and DEC systems. The IBM system is an SNA network used for commercial applications, and the DEC system is used mainly for technical applications.

Interconnection policy

Group headquarters is seeking more management information in standardised form from the operating companies, but at the same time is giving them more autonomy. At the computer network level this is leading to a requirement to interconnect IBM, DEC and Sperry systems, and is encouraging a move away from discrete networks towards interlinked networks based upon group standards.

The DEC systems are linked by DECnet and are increasingly used for commercial applications. Users wishing to access commercial applications on both the IBM and DEC systems might in the future demand a gateway between the DEC and SNA networks. IBM's 2780 remote job entry protocol is frequently used to allow Sperry, DEC, Hewlett-Packard and Wang systems to access IBM systems.

The data network standards policy was being reviewed at the time we met with this organisation. A single-supplier policy is not a possibility because no one supplier could provide the range of products required. The review is considering which network architectures should be used and the extent to which the networks should be based on international standards. In the long term the company wants to use OSI standards.

The SNA network

The organisation's SNA network was constructed over two to three years without any serious problems, and those concerned ascribed this to the fact that SNA is now a mature product.

The system allows any terminal to access any application on the network, although only about five per cent of the 1,000 terminals on the system make use of this facility at present. The belief is that more people would take advantage of this feature if users were made more aware of the capabilities of SNA.

IBM's 2780 communications standard is used to link Sperry, DEC, Hewlett-Packard and Wang systems into the SNA network. Users are requesting specially developed software to enhance the functionality of this communications standard and, as a result, they are now discouraged from using it.

The DEC network

In late 1984 the organisation had a 57-node worldwide network of DEC machines, only a minority of which are used by headquarters staff. Most of the rest belong to the various autonomous business streams, and a few belong to third parties. Because third parties are linked into the network, great care had to be taken with network numbering. Each of the autonomous companies manages its own part of the network. Central network management was not considered because it would not be consistent with the organisation's management style.

The experience with interactive terminal work across the DEC network was disappointing. Response times were poor and this type of working is now prohibited. The main application across the network is now file transfer.

Microcomputer-mainframe links

Several types of microcomputers (IBM PCs, Decmate word processors, DEC PCs, etc.) have been linked into VAX minicomputers using VT100 emulation packages. Some of these microcomputer types have also been linked into IBM systems using IBM 3270 emulation.

Attitude to OSI standards

The headquarters information systems department is following closely the development of OSI standards. It is a member of the national Standards Users Association. OSI communications are seen as a competitor to SNA in the long term. Concern was expressed about the weakness of OSI standards in the network management area. This area is seen as being increasingly important as networks become larger and more complex.

J SAINSBURY PLC

J Sainsbury plc operates a chain of about 250 retail supermarkets in the United Kingdom selling food and drink, and household goods. Sales for the financial year 1984/85 were £3 billion, and the company had an eight per cent share of the United Kingdom trade in food and drink. During the research for this report we interviewed Alan Jacobs, Sainsbury's director of data processing; Paul Davies, communication design manager; and Barry Mason, office automation manager.

Network policy

In 1982 the company was using a mixture of ICL and IBM data processing equipment. At that time, Alan Jacobs, who had just joined the company, recognised that the difficulties of interconnecting computer systems from two suppliers would consume scarce development resources that could better be used to extend the company's computer systems. The main priority was to provide the flexible systems that senior management required in order to react quickly to changes in the marketplace. The company decided to standardise on IBM-compatible equipment and adopted SNA as its communication standard.

IBM-compatible equipment was chosen because of IBM's large client base and the resulting large pool of available software. It was also thought that IBM-compatible equipment would provide the greatest scope for the future. The company chose SNA protocols for the network because this was the primary communications standard that IBM would be developing.

The network and network applications

Sainsbury has adopted a mixed supplier policy for its IBM-compatible hardware. Of the four current mainframes, two are secondhand IBM 3033s and two are Atlas 10s from ICL. For security purposes these mainframes were located at two sites, London and Stevenage, and were subsequently linked by a British Telecom Megastream circuit.

In early 1985 the main computer centre was in London. It had two mainframes and three communications controllers supporting about 1,000 terminals in 250 supermarkets, 25 depots and at head office. The two main types of work processed at the London computer centre are ordering, distribution, stock control and buying applications (Priority 1) and development work (Priority 2). The London centre has been configured so that either mainframe or any communications controller can handle all Priority 1 work.

The Stevenage centre currently processes accounting applications but, in due course, it will operate fully in parallel with the London centre. Its communications facilities are being extended so that, in an emergency, it can handle all the Priority work from the London centre. In order to achieve this, Sainsbury plans to set up a network of 'Meganodes' at four main depots/sites with Megastream circuits linking each depot/site to both computer centres, as shown in Figure A.3. The other depots and branches will link into the main depots.

Network commissioning and management

The SNA network was implemented with few difficulties. Paul Davies believes that this was because the company was not constrained by existing IBM systems, and because SNA is now a mature product. Recruiting appropriate staff has been a problem, however. Three types of SNA staff have been sought:

—Network planners and implementers.

—Systems programmers for teleprocessing software.

—Technicians for day-to-day network support.

Figure A.3 Sainsbury's future wide-area network

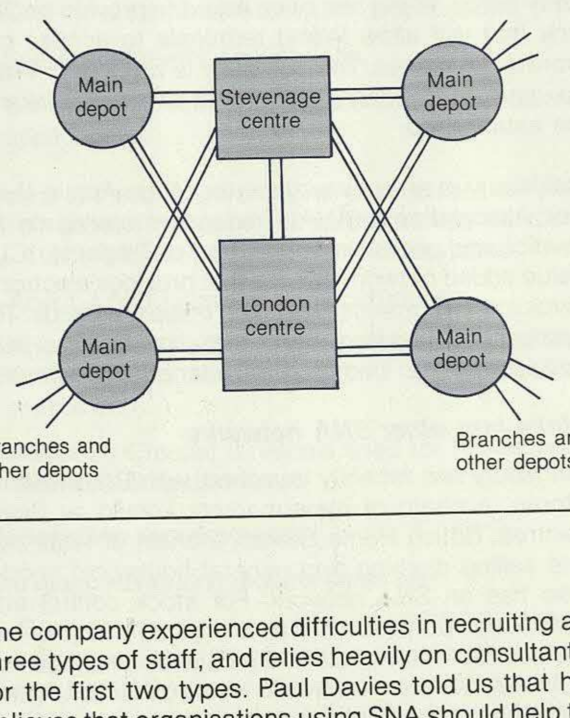

The company experienced difficulties in recruiting all three types of staff, and relies heavily on consultants for the first two types. Paul Davies told us that he believes that organisations using SNA should help to alleviate this problem by providing appropriate training for existing personnel.

Mr Davies admits that Sainsbury still has some way to go in learning how to use SNA to its full potential, particularly with regard to network management. Some of the network management packages took a long time to implement, and their full value cannot be realised because the company does not use IBM modems.

Links between the data network and non-SNA systems

The company has stressed to computer suppliers that SNA compatibility is an essential requirement. The most significant example of this requirement is provided by the ICL System 25 minicomputer that will control the laser scanning equipment in each of the 250 supermarkets. The minicomputers must emulate terminal controllers in an SNA network but, more importantly, they must be supported remotely because supermarket personnel are unlikely to be familiar with computer systems.

The need for remote support stems from the relative complexity of the System 25 data processing operation and the company's need to respond quickly to changes in market conditions. This latter need may require new information to be loaded very rapidly into each supermarket system and, in the event of problems, the centre must be able to support and reconfigure the branch system. ICL has been developing the software to meet these requirements; by June 1985, 70 systems were operational.

As a further example of the company's SNA compatibility policy, Wang has been asked to provide an SNA link that will allow Wang terminals to access corporate databases. The company is waiting for Wang to upgrade its existing equipment so that the link can be established.

Sainsbury is an active supporter of the Article Number Association's (ANA) Tradacoms standards for invoice and order exchange, and of Tradanet, ICL's value added network service that provides electronic invoicing and ordering based on these standards. The company is making sure that its systems take account of the Tradacoms/Tradanet requirements.

Links into other SNA networks

Sainsbury has recently launched, with British Home Stores, a chain of hypermarkets known as Savacentres. British Home Stores, a chain of retail outlets selling clothing and general household goods, also has an SNA network. For stock control and distribution purposes, there is a requirement to link the two SNA networks. There have been difficulties with IBM's SNA network interconnection (SNI) package concerning the addressing of nodes in the two networks, and different versions of the teleprocessing software operating in the two networks. The link has now been established through the introduction of a 'black box' protocol convertor between the two networks.

Attitude toward OSI standards

With regard to communications standards, Alan Jacobs said that although SNA might not be ideal, it is available in product form and it works. Before Sainsbury could consider adopting OSI standards, they too would have to be available in well-established products.

A/B SKF

The SKF group is the world's leading manufacturer of roller bearings. Its headquarters are in Gothenburg, Sweden; it employs 45,000 people worldwide; and it has 75 factories (30 manufacturing roller bearings) in 17 countries. Less than 20 per cent of total group turnover is derived from production in Sweden. We spoke to Bengt-Åke Svensson, the group information systems manager responsible for worldwide telecommunications, about the group's IBM SNA approach to networking (although SKF's experiences illustrate the way in which the company has adopted the protocols of more than one supplier).

The data network and its application

Each of the SKF factories specialises in manufacturing a limited range of bearings in large quantities for shipment to all SKF markets. With 25,000 product variants, a considerable amount of marketing and technical data is required to support these shipments, and an international data network has been set up to handle the data flows. This network uses IBM's SNA protocols, and at the end of 1984 it consisted of eight nodes in five European countries, supporting about 3,000 terminals (see Figure A.4). A typical data centre has an IBM 3083 mainframe and an IBM 3725 communications controller. In addition, some of the data centres have DEC computers.

Development and management of the data network

In 1977, SKF decided to review its network requirements. The alternatives considered were an SNA network, a network based on X.25 protocols, and a network based on Comten processors. The company had a history of using IBM computers for administrative work (and DEC machines for technical work), and so it was no surprise that IBM was chosen.

Bengt-Åke Svensson told us that, from a technical point of view, it was not particularly difficult to install and commission the SNA network. The network nodes and the telecommunications links were already in place. The SNA software worked and IBM's support, including training, was good. However, Mr Svensson stressed the importance of determining the network management procedures before the network is installed, especially for an international network. He emphasised the need to:

— Decide on approval procedures (both at the national and international levels) for changes in network software, links, etc.

Figure A.4 **SKF's international data network**

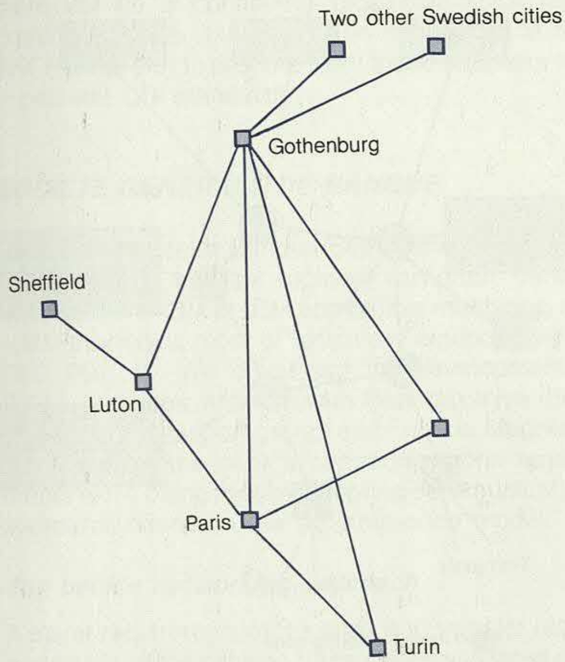

—Agree the charging structure.

—Determine how transnational support is to be provided in an environment in which several languages are spoken.

Management of the network is carried out both locally and centrally. Each data centre has a teleprocessing section that reports to the local company. The group information systems department monitors the whole network, and reports to a committee of executives that decides what functions the network should provide.

Mr Svensson reported that SKF had generally found the IBM network management packages very useful. The network problem determination application (NPDA) software had caused difficulties, however, because it was complicated to use. Mr Svensson thought more training was required before users could gain the full benefits of this software.

Early in 1984, the company successfully installed a DECnet/SNA gateway. This is a standard DEC product that allows a DEC computer to emulate an IBM cluster controller and remote job entry terminal, and allows a DEC VT100 terminal to emulate an IBM 3270 terminal.

Mr Svensson said that his two most immediate tasks were to extend the data network to Philadelphia, and to use IBM terminals to access applications on the DEC computers. On the latter task, he is waiting for appropriate DEC products to be announced. Ultimately, he would like to link DEC terminals across the SNA network.

Attitudes towards OSI standards

Mr Svensson had no immediate concern about OSI standards, although he thought that national differences in the implementation of OSI standards (if any) would be important to anyone running an international OSI network. His view was that the support of the European PTTs would be critical in getting OSI standards implemented. In the longer term he thought that OSI communications standards had the potential for linking proprietary networks.

SNCF

SNCF is the state-owned French railway company. Its networking experience illustrates how communications requirements are being met by adopting 'open' standards. We met with SNCF's communications manager, M Desconclois.

SNCF's mainframe systems are of three types:

—IBM for ticket and reservation applications.

—Sperry for freight applications.

—Bull for administration, payroll and maintenance applications.

Access to these mainframes, all of which are in Paris, is from about 4,700 terminals spread across the country. Three different access devices are used:

—About 4,000 Logabax microcomputers.

—About 700 Olivetti teletype terminals. These are being phased out and replaced by the Logabax machines.

—About 50 Crouzet terminals used for high-speed-train (TGV) reservations.

Networking requirements

The main networking requirements are:

—Open communications. Because SNCF is a state-owned company it is encouraged not to use computers from only one supplier. Users want to access all three types of mainframes from the same terminal.

—High reliability and high capacity for the network.

—Promotion of distributed processing to support and encourage decentralised management.

Because of the lack of high-level communications standards, SNCF decided to write its own high-level communications software. The company also decided to build a private packet-switched network, for reasons of cost and reliability.

Packet-switched network architecture

SNCF's private packet-switched network is depicted in Figure A.5. The network consists of:

—Twenty-two packet-switching nodes. These are monitored by a Bull DPS 6 minicomputer that checks the nodes are working and collects operational information, but exerts little control.

—Sixty terminal concentrators that are monitored by a DEC PDP 11 machine. The terminal concentrators decide on priorities in case of simultaneous demands on terminals.

—A PDP 11/70 manages the terminals and provides a store-and-forward and store-and-retrieve messaging system.

The network is very much terminal-orientated. If mainframe-to-mainframe communication is required, then the mainframe emulates a terminal on the network. Because of the absence of high-level communication standards, SCNF developed its own software for these levels. This consists of a transport layer called station transport (ST) and a virtual terminal protocol (PTV) for the session and presentation layer. The PTV allows each mainframe to communicate with a single type of 'virtual' terminal. At the application layer there are three separate applications — batch file transfer to the mainframes, distribution of output from the mainframes to terminals, and database query. Figure A.6 shows how these layers are implemented on the various mainframes.

Future developments

SNCF has a considerable number of new projects under development, the major ones being:

—A freight train assembly management application, which will be implemented on DPS 6 minicomputers at each of the 40 major railway stations. The network will be used to pass on information as freight trains move across the country. A Bull DPS 8 mainframe will coordinate the information centrally.

—Personnel and administrative applications, which will be based on a Bull DPS 7 in each region. These applications will use the network to provide headquarters with data.

—Systems to allow the public to access train timetable information via the free videotex terminals being provided by the PTT for telephone directory services. It is expected that three million of these terminals will be installed by the end of 1986.

—At present, SNCF has separate IBM, Bull and Sperry networks, each running under a proprietary network architecture. Although these networks will remain, the transmission facilities wil be moved progressively onto the packet-switched network for transmission purposes.

—Teletype and dial-up access to the network will be provided for all terminal types.

Figure A.5 SNCF's private packet-switched network

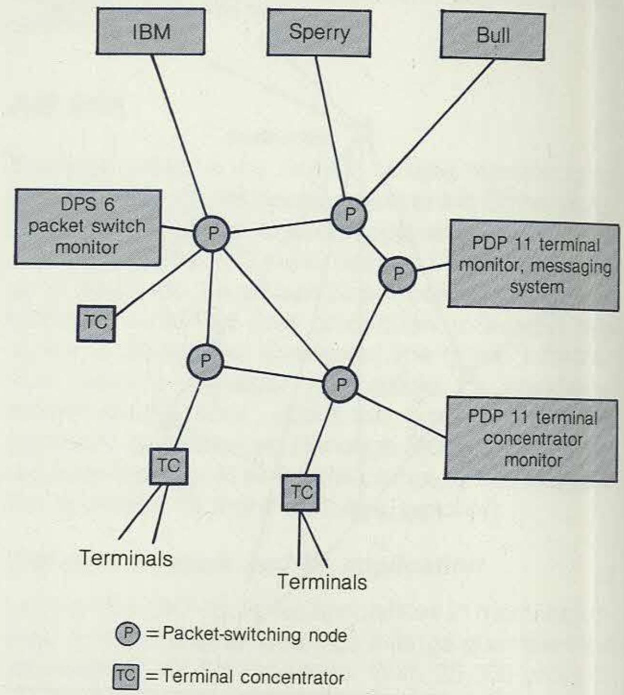

Figure A.6 Implementation of software layers at SNCF

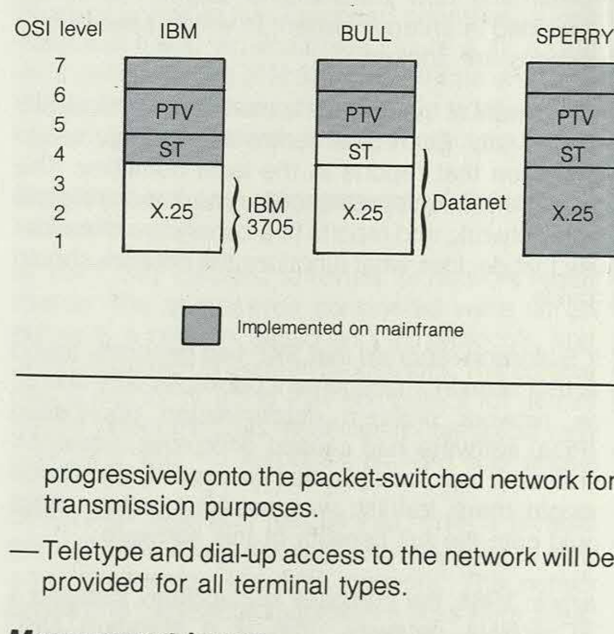

Management issues

In the absence of international standards, SNCF chose to develop its own high-level communications software for a multi-vendor network. At the time of the decision, SNCF had the specialist staff required to do this work. Now those staff are not so readily available and so SNCF has decided to shift the emphasis towards the greater use of proprietary network architectures.

The company is encouraged to procure computer equipment from several suppliers, and this means there will be a continuing requirement for open communications standards. The implication is that SNCF will prefer to procure from those suppliers that implement OSI standards.

SOCIETE GENERALE DE BANQUE

Société Générale de Banque (SGB) is the largest bank in Belgium. It has six regional computer centres equipped with IBM or IBM-compatible machines, and 1,200 branches, most of which are equipped with a DEC PDP 11. We discussed the development of SGB's computer network with Paul van Vyve (head of the system support group) and Francis Magnée to find out how the bank's communications requirements were being met by adopting communications standards based on the OSI reference model.

The bank's networking approach

A major requirement of the bank is to transfer rapidly and easily all the banking transactions validated each day in each branch to one of the regional processing centres for consolidation of the customer database. In addition, the local databases at the individual branches must be brought up to date before the opening of business the next day.

In order to meet this requirement, the bank decided to install a private packet-switched network and to develop its own communications software for the higher levels of the OSI model. The main reasons for this choice were that:

—Such a telecommunications network would be able to support all present and future communication needs.

—The approach would provide an 'open' network with access protocols based on internationally agreed standards.

—The bank's data processing requirements were too diverse to rely on one supplier. The type of communications architecture must be independent of any manufacturer to allow multi-vendor systems to be interconnected. Choosing the IBM or DEC communications architecture would have led to a system that would make it difficult to interconnect with other suppliers' equipment.

—At the time of the decision, the Belgium PTT did not provide a public packet-switched service.

—The communications architectures offered by IBM and DEC at the time of the choice did not meet all of SGB's requirements.

—Choosing an open network would provide the possibility of combining some of the individual networks already installed by SGB into a single data network, and would therefore reduce line costs.

The packet-switched network

SGB installed a private packet-switched network for linking the PDP 11s in the branches with the data processing centres. The network is based on Northern Telecom's SL-10 hardware, and is represented schematically in Figure A.7. The network has nine main SL-10 nodes and seven peripheral SL-10 nodes. One of the nodes is designated as the network control centre (NCC). The NCC gathers network alarm data and permits remote trouble-shooting throughout the network. It also collects performance statistics, traffic information and accounting data.

SGB's communications software

The X.25 network provides the lower three layers of the OSI model. SGB developed its own version of Layer 4 (the transport layer), which it has called the network communications manager (NCM). The functions of the NCM are as follows:

—Opening and closing links between end-systems.

—Multiplexing and de-multiplexing between links.

—Transferring data.

Figure A.7 SGB's packet-switched network

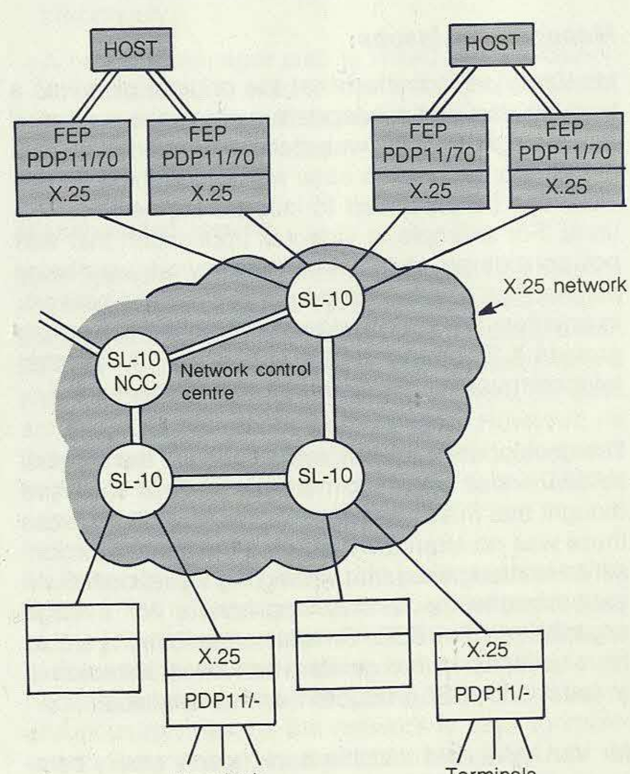

—Controlling the flow of data and data errors on an end-to-end basis.

SGB has also developed software for several functions corresponding with OSI Levels 5, 6 and 7. The main ones are:

—Batch file transfer between an application running on a DEC machine and an application running on an IBM machine. This facility is used to update the regional and local customer databases.

—A file transfer protocol to broadcast new versions of programs to the branch computers.

—An interactive application that allows a terminal connected to a PDP 11 to access a database controlled by an IBM mainframe.

Gateways

Each IBM host is connected to the X.25 network through batch and real-time front-ends which implement the gateway functions needed for the communication between the branch processors and the mainframes in the regional computer centres. Both front-ends have a direct channel interface to the host and connect through high-speed links (64k bit/s) to the X.25 network. The batch front-end that supports the data collection and database update application emulates a 3803 magnetic tape controller, whereas the real-time front-end implements an SNA gateway emulating a 3705 communications controller.

Management issues

Mr Van Vyve admitted that the original plan was a bold initiative, but subsequent events have confirmed that the right decision was made. An important benefit of the approach is the ease with which the X.25 network can be expanded to incorporate new applications. For example, a videotex application that was not considered at the time the network was being planned has now been incorporated into the network. In addition, SGB is considering using the network to provide X.25 bearer facilities for some of the other telecommunications networks run by the bank.

The development and implementation of the network software took about 50 man-years. Paul Van Vyve thought that this investment was worthwhile because there was no alternative offering the same functionality. He recognised that writing high-level communications software is only appropriate for a large organisation like SGB. Smaller organisations would have to implement proprietary network architectures, at least until OSI products become available.

Mr Van Vyve said that the bank is only slowly coming to terms with the need to manage and maintain the network. A recent study has indicated that 10 per cent of all non-salary IT costs are devoted to telecommunications. This study has encouraged the bank to take the management of the network more seriously. With such a large network, maintenance is also a problem. The bank has recognised that it needs to employ a new type of person — local telecommunications maintenance engineers. The plan is to employ four maintenance technicians centrally and about 20 technicians locally.

Currently, network users are charged on the basis of the number of terminals they operate. A new billing procedure is to be introduced that will more accurately take into account the network resources used. Paul Van Vyve expects that this new procedure will encourage the users to demand a more bespoke service.

Future developments

SGB is considering ways to develop its communications software. The three main options are to:

—Continue to write and maintain the bank's own software.

—Employ a software house to write and maintain the software.

—Migrate to OSI products, as they become available.

SGB expects ultimately to adopt the third option.

Besides the main network described here, SGB has four other networks. In addition, an office automation network will be created by linking 40 PDP 11/44s and 11/24s via DECnet. All of these networks are candidates for using the X.25 bearer facilities of the main network. Some of the links in one of the (SNA) networks are already provided by the main SGB network.

SGB also plans to integrate voice and data at the transmission level. Paul van Vyve said that this might be implemented after 1988.

WEST MIDLANDS COUNTY COUNCIL

The West Midlands County Council (WMCC) provides local government services for about 2.6 million people living in an urban area centred on Birmingham, England. The main services provided are police, fire, strategic planning, transportation and engineering, waste disposal, consumer services and passenger transport. We talked with Mike Spencer (Assistant County Treasurer, Computers) and one of his group leaders, Martin Evans, about the Council's experience of an early implementation of ICL's open system local area network (OSLAN). This case history illustrates how the protocols of one supplier are being used to satisfy communications requirements.

The network strategy

In 1981 the WMCC produced a computer network plan. Immediately prior to the formulation of the plan, the following systems were planned:

—A network of two ICL 2966s and one ICL 2988 mainframes with 150 terminals linked directly or remotely. There were plans to increase the number of terminals to 450 by the end of 1984. The systems included payroll, passenger transport administration, traffic accident and police incident analysis.

—A centralised typing pool (based on Wang equipment) with links to remote sites.

—A network of Tandy microcomputers linked by Clearway, a proprietary asynchronous local area network.

—A proposed Sigma network for computer-aided design applications in architectural planning and transport engineering. Subsequently it was decided to opt for a network of Apollo microcomputers linked by a token-passing local area network.

The original aim of the computer network plan was to allow a standard terminal to access several mainframe applications. The intention was to identify an 'ideal terminal' that was inexpensive and would also offer a range of functions, including database access, diary management, videotex and messaging services, plus specialist features such as a high-resolution screen for graphics work. It was soon realised that such an ideal terminal was not available, and the plan was modified to cater for four main types of terminal:

—For senior managers, a videotex terminal.

—For middle managers, a combined terminal and microprocessor.

—For specialist users, a terminal with high-resolution graphics and word processing.

—For general use, a conventional terminal.

Apart from the videotex terminals, any terminal might require access to several mainframe applications. The most important requirement was for access to local databases (payroll, personnel, trading standards, etc.) or to remote databases such as legal databases at the House of Commons in London or at the European Commission. The specialist terminals were included in this requirement because, for example, a standard letter might be prepared on a Wang terminal that would need to access a mainframe file for the mailing list. In another instance, a report prepared on a Wang terminal might need to include financial information extracted from a mainframe file.

In order to achieve its strategic aim, WMCC decided to implement a local area nework to interconnect its systems. The Council worked with ICL to implement one of the first installations of ICL's open system local area network (OSLAN), which is based on Ethernet.

The OSLAN configuration

Processors are attached to the OSLAN by open system link units (OSLUs), each of which has eight ports. The configuration of the OSLAN is illustrated in Figure A.8 (overleaf), and is described below:

—The ICL mainframes are linked to a communications controller, which in turn has separate CO3 (a proprietary ICL communication standard) and asynchronous links to the OSLAN. CO3 communications are more efficient and so, unless otherwise required, all asynchronous communications pass automatically through an asynchronous/CO3 protocol converter attached to an OSLAN node.

—Each OSLU can in theory support synchronous communications with eight DRS 20/50 microcomputers, each of which can control several of the smaller DRS20/10 microcomputers. In practice, each OSLU was able to support three DRS 20/50s and five asynchronous devices.

—The Wangnet is linked to an OSLU via protocol converters.

—Other non-ICL equipment (Apollo terminals, Tandy microcomputers) are linked into the network asynchronously.

—A network manager also is linked to the OSLAN. At present this is used to load the OSLU software, but future versions will collect a limited amount of network usage statistics.

Management issues

Mike Spencer admitted that the Council had taken a risk in deciding to opt for an Ethernet-based system before standards were available. He is satisfied with the functionality of the network, but has reservations about its costs. At £1,000 per OSLU port, he does not think OSLAN can be cost-justified. However, he thought that the price of an OSLU is likely to decrease to a level at which further interfaces to the OSLAN can be cost-justified.

Mr Spencer was very happy with other aspects of the OSLAN. The network cable is inexpensive and is easy to lay. The network can be extended easily, because it does not have to be taken out of service when nodes are added. The interconnection of two asynchronous devices via the network is easy to implement. The network is resilient: if an OSLU or an attached device goes down, the OSLAN continues to operate.

Figure A.8 West Midlands County Council's OSLAN

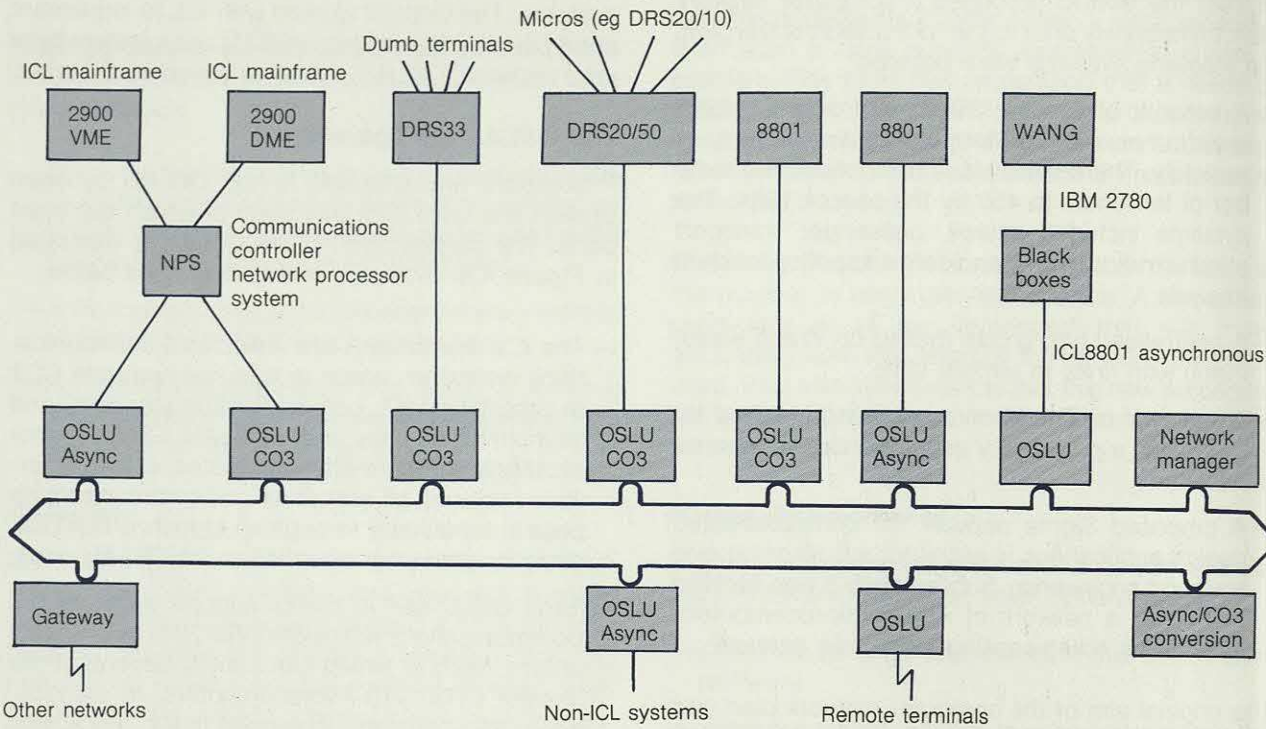

Looking back on the installation phase of the project, Mike Spencer said that two key lessons had been learnt. First, it is important to move forward cautiously. Second, the implementers must be chosen carefully, and management must let them experiment.

Impact of the OSLAN on IT procurement policy

Prior to installing the OSLAN, the Council tried to implement a rigid IT procurement policy. Only the ICL CO3 communications standard could be used, for example. This policy did not always find favour with end users, who wanted the equipment that best suited their requirements. The installation of the OSLAN permitted the WMCC to relax its procurement policy because a wide range of equipment could be attached to the OSLAN.

Attitudes towards OSI standards

Both Mike Spencer and Martin Evans were sceptical about the success of the OSI standards initiative. Specifically, they were concerned with the slow speed of the ratification process.

They also believed that innovation could easily make 'standards' out of date. In their view standards would only cater for a minimum requirement. They cited the example of the 'Fortran standard' and how it had changed over the years. They also suggested that suppliers would want to protect their existing client base, and this might inhibit the implementation of standards.

Future developments

The WMCC is close to implementing its aim of allowing any terminal to access several mainframe applications. However, the graphics terminals cannot access the videotex system, and database access from the Wang terminals is not particularly user-friendly. Apart from these exceptions, the functionality provided by the OSLAN is almost up to expectations. The Council's main concern is that the high cost of extra OSLUs is preventing it from adding further devices.

Although there are no definite plans to do so, it is possible that other mainframes might be linked to the OSLAN. If ICL produces an IBM/CO3 conversion package, then some IBM devices might be linked to the OSLAN as well.

CLEARING HOUSE AUTOMATED PAYMENT SYSTEM (CHAPS)

This case history illustrates the way in which a value added network service is providing protocol conversion so that different computer systems in different companies can be interlinked. The service operates in the City of London banking environment and provides an improved financial system for the same-day clearing of large sterling payments. The case history also highlights the need for a highly resilient service as a network becomes a central element of the business activity. The network also has to be flexible, because the participants operate in a competitive environment.

The London Town Clearing System

The Town Clearing System is a manual operation in the City of London used by banks to make same-day large sterling payments, reliably and with confidence. It operates only within the boundaries of the City of London, and there is currently a lower limit of £10,000 per transaction. Clearing systems in the United Kingdom are operated by a group of large banks collectively known as 'clearing banks'. Each clearing bank maintains a special clearing account at the Bank of England, and this account is used at the end of each day to make or receive net payments to or from other clearing banks. Non-clearing banks can offer payment services but these must be via a clearing bank.

In the early 1970s the clearing banks began to consider whether the time had come to create an up-to-date automated payment system. A totally manual system had obvious capacity limitations, and limiting the service to the City of London also caused irritation. The success of the Clearing House Interbank Payment System (CHIPS), an automated system in New York, and of SWIFT (Society for Worldwide Interbank Financial Telecommunications) eventually persuaded the major United Kingdom banks to create an automated same-day payment service for the City, and ultimately for the whole country.

Design of the CHAPS system

The first design for CHAPS was based on a large central computer with the banks linked into the system by standard connections. As the development work went ahead, it became clear that security issues and the lack of flexibility in a centralised system would cause problems in a competitive environment. Eventually CHAPS I was abandoned, and CHAPS II was instigated.

The second design for CHAPS consisted of gateways linked by British Telecom's packet-switching service. Each clearing bank links into a CHAPS gateway. Some gateways are owned by a single bank and others are shared by more than one bank. The non-clearing banks link direct into a gateway owned by a clearing bank or via a clearing bank's computer system. Each gateway has customised software that converts the clearing bank's electronic payment into the CHAPS standard, and standard gateway software (GWS) that controls the movement of the electronic payment (see Figure A.9).

Figure A.9 The CHAPS network

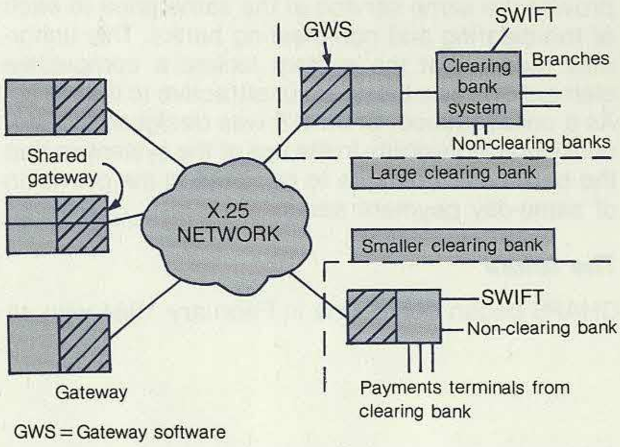

A CHAPS payment

A CHAPS payment is initiated when a bank's customer asks for a same-day payment to be made. If the payment is to a customer at a non-clearing bank, then the settlement bank is chosen from a routeing table stored within the system.

The gateway of the originating bank acknowledges the CHAPS payment, stores audit trail information, encrypts the message and then sends it. The receiving gateway authenticates the message, stores its own audit trail information and acknowledges receipt of the payment to the originating gateway. The gateway is then responsible for passing on the payment to the recipient.

Security and reliability

CHAPS handles substantial sums of money, and it was essential that very high levels of security and reliability be built into the system. Standard built-in features include:

—The packet-switching links are duplicated.

—Every payment message is authenticated.

—Every payment message is encrypted.

—Audit trails are kept at several stages of each message's journey.

The gateways are based on Tandem NonStop computers which are designed to continue operation after multiple failures. In addition, provision is made for a settlement bank to switch to an alternative gateway if its gateway fails. In the event of a complete failure of a gateway, the recovery procedures enable a rapid switch to be made to the fallback gateway with a minimum of retransmission whilst still maintaining the integrity of the total nework.

Competitive aspects of CHAPS

One of the drawbacks of CHAPS I was that it would provide the same service at the same price to each of the clearing and non-clearing banks. This uniformity meant that the system lacked a competitive element and was therefore unattractive to the banks. As a consequence, CHAPS II was designed to allow considerable flexibility in the use of the system so that the banks can continue to compete in the provision of same-day payment services.

The future

CHAPS began operations in February 1984 with all the clearing banks participating. Most clearing banks now have their branch networks connected to CHAPS so that automated same-day payments can be sent and received in all parts of the United Kingdom. During 1984 some 50 non-clearing banks arranged connections with their clearers so that they could also take advantage of the facilities provided by the CHAPS system. CHAPS is now playing an important role in helping to maintain London's position as a major financial centre, and has laid the foundation for further automation in the City — improvements in automated corporate cash management services, for example.

THE BUTLER COX FOUNDATION

Butler Cox

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

Objectives of The Foundation

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

The Foundation not only provides access to an extensive and coherent programme of continuous research, it also provides an opportunity for widespread exchange of experience and views between its members.

Membership of The Foundation

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

The Foundation research programme

The research programme is planned jointly by Butler Cox and by the member organisations. Half of the research topics are selected by Butler Cox and half by preferences expressed by the membership. Each year a short list of topics is circulated for consideration by the members. Member organisations rank the topics according to their own requirements and as a result of this process, members' preferences are determined.

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

The report series

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

Previous reports

No. 1 Developments in Data Networks*
No. 2 Display Word Processors*
No. 3 Terminal Compatibility*
No. 4 Trends in Office Automation Technologies*
No. 5 The Convergence of Technologies
No. 6 Viewdata*
No. 7 Public Data Services*
No. 8 Project Management
No. 9 The Selection of a Computerised PABX*
No. 10 Public On-line Information Retrieval Services*
No. 11 Improving Systems' Productivity
No. 12 Trends in Database Management Systems*
No. 13 The Trends in Data Processing Costs
No. 14 The Changing Equipment Market*
No. 15 Management Services and the Microprocessor
No. 16 The Role of the Mainframe Computer in the 1980s*
No. 17 Electronic Mail
No. 18 Distributed Processing: Management Issues
No. 19 Office Systems Strategy
No. 20 The Interface Between People and Equipment
No. 21 Corporate Communications Networks
No. 22 Applications Packages
No. 23 Communicating Terminals
No. 24 Investment in Systems
No. 25 System Development Methods
No. 26 Trends in Voice Communication Systems
No. 27 Developments in Videotex
No. 28 User Experience with Data Networks
No. 29 Implementing Office Systems
No. 30 End-User Computing
No. 31 A Director's Guide to Information Technology
No. 32 Data Management
No. 33 Managing Operational Computer Services
No. 34 Strategic Systems Planning
No. 35 Multifunction Equipment
No. 36 Cost-effective Systems Development and Maintenance
No. 37 Expert Systems
No. 38 Selecting Local Network Facilities
No. 39 Trends in Information Technology
No. 40 Presenting Information to Managers
No. 41 Managing the Human Aspects of Change
No. 42 Value Added Network Services
No. 43 Managing the Microcomputer in Business
No. 44 Office Systems: Applications and Organisational Impact
No. 45 Building Quality Systems
No. 47 The Effective Use of System Building Tools

**These reports have been superseded.*

Future reports

No. 48 Measuring the Performance of the Information Systems Function
No. 49 A Strategy for Systems
No. 50 Making Data Available to Users
No. 51 Secure Access to Computer Systems
No. 52 Organising the Information Services Function

Butler Cox & Partners Limited
Butler Cox House, 12 Bloomsbury Square,
London WC1A 2LL, England
☎ + 44 1 831 0101, Telex 8813717 BUTCOX G

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

France
Butler Cox SARL
Tour Akzo, 164 Rue Ambroise Croizat,
93204 St Denis-Cedex 1, France
☎ (1)820.61.64, Telex 630336 AKZOPLA

United States of America
Omni Group Limited
115 East 57th Street, NY 10022, New York, USA
☎ (212) 486 1760

Australia
Mr John Cooper
Business House Systems Australia
Level 28, 20 Bond Street, Sydney, NWS 2000
☎ (02) 237 3232, Telex 22246

Italy
SISDO BDA
20123 Milano – Via Caradosso 7 – Italy
☎ 498 4651, Telex SISBDA 350309

The Nordic Region
Statskonsult AB
Stortarget 9, 5-21122 Malmo, Sweden
☎ 46-401 03 040, Telex 127 54 SINTAB