Report Series
No 16

The Role of the Mainframe Computer in the 1980s

December 1979

The Butler Cox Foundation

Abstract

Report Series
No 16

The Role of the Mainframe Computer in the 1980s

by Tony Brewer and
Tony Gunton

December 1979

While mainframe computers themselves get cheaper, the cost of developing and maintaining the applications which run on them seems to be climbing irresistibly. For many organisations, the backlog of applications waiting to be implemented is a cause for serious concern.

This is reason enough on its own to question whether "the mainframe approach" to computing is the right approach for the 1980s.

Another reason is the move to distributed computing which, after being in the air for some time, now shows signs of making a significant impact on the computing way-of-life of many organisations.

This report evaluates the role and potential progress of the mainframe from the point of view both of users and suppliers, and reviews the technological developments that are taking place. The report's purpose is to forecast the role of the mainframe in the 1980s and to suggest how this might affect management services' plans.

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

The Foundation collects its information in Europe and the US, where it has offices through its associated company. It transmits its findings to members in three main ways:

- As regular *written reports,* giving detailed findings and substantiating evidence.

- Through *management conferences,* stressing the policy implications of the subjects studied for management services directors and their senior colleagues.

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

The Foundation is controlled by a Management Board upon which the members are represented. Its responsibilities include the selection of topics for research, and approval of the Foundation's annual report and accounts, showing how the subscribed research funds have been employed.

Report Series No. 16

THE ROLE OF THE MAINFRAME COMPUTER IN THE 1980s

by Tony Brewer and Tony Gunton

December 1979

Research team: Rob Moreton
Bob Shepherd

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

CONTENTS

CHAPTER 1

INTRODUCTION

The rather curiously named "mainframe" might be more accurately termed the mainstream of data processing, because most sizable organisations rely predominantly on the mainframe for their computer services. The mainframe has been threatened by the minicomputer and more recently by the microcomputer, which apparently offers computing that is more cost-effective and, perhaps as importantly, promises to deliver it right to the end-user's doorstep. This has led some commentators to conclude, overhastily, that the mainframe is a dinosaur, and that its demise cannot come too soon for the good of all concerned.

The purpose of this report is to take a more rational look at the mainframe and to try to forecast its role in the 1980s. We do this by discussing the requirements of those three groups who are most closely involved with the mainframe today — the data processing managers who buy and run mainframe computers, the end-users who rely on them for service, and the suppliers who sell and support them. We concentrate on these three groups because data processing is by its very nature a conservative industry, so that change is more likely to result from pressures from within the industry than from new opportunities outside it.

We begin in chapter 2 by providing a working definition of the mainframe and by discussing its current status and the pressures acting on it. Then, in chapters 3, 4 and 5, we look at the mainframe from the respective points of view of the data processing manager, the end-user and the supplier. In chapter 6, we discuss the developments that are taking place in hardware architecture and some of the most significant findings that are emerging from research work.

Finally, in chapter 7, we attempt to draw all the threads together and to provide some guidance both for management services managers and data processing managers on how the mainframe might figure in their plans for the 1980s.

CHAPTER 2

THE MAINFRAME – DEFINITIONS AND CHARACTERISTICS

DEFINITION OF A MAINFRAME COMPUTER

The term "mainframe" had its origins in the early days of computing, when it was used to describe the racking that held the processor and the core store of a computer. The word was then generalised to include all of the directly-connected components of a second- or third-generation computer. It took on a further lease of life in the early 1970s, when it was used to discriminate between the larger business computers and the smaller special-purpose minicomputers. The term is still in very general use but it is applied in many different contexts and its meaning has now become rather blurred. We have met the following different ways of attempting to define the term mainframe:

1. Definition based on architecture
   A mainframe computer, defined by its architecture, would have an internal word size of 32 bits or more. It would have peripheral units connected with both the main memory and the processor via standard channels, it would have a main memory size several times greater than the average applications program size, it would have a multi-function operating system, and it would be constructed using either transistor technology or integrated circuit technology.

   However, the latest products available from the so-called minicomputer suppliers would also meet the requirements of this definition, and so would certain very powerful, special-purpose machines which not everyone would describe as mainframe computers.

2. Definition based on size and power
   A mainframe computer, defined by its size and power, might now have a main memory storage capacity of at least 256K characters, a back-up storage capacity of at least 100M characters and an instruction time of less than 10ms. However, as figure 1 shows, time very quickly renders this type of definition obsolete.

3. Definition based on organisation
   A mainframe computer, defined by its place in the user's organisation, would be a central computing facility. This definition was suggested to us by an eminent academic, who pointed out that the traditional definitions, which are based on hardware characteristics, no longer have any meaning.

   This definition does have a precise meaning, but it conflicts with the accepted meaning of the word, and so it may not be a very useful definition.

4. Definition based on use
   This can hardly be regarded as a viable definition, when viewed against the rapid growth of application of minicomputers and, to a lesser extent, microcomputers to business problems. In addition, the uses to which the computer is put often conflict with the purposes for which it was designed. Thus, the IBM 360 and 370 ranges were not originally designed to handle transaction processing, but they are now widely used for this type of work. Similarly, minicomputers designed for on-line work tend to get involved in batch processing also.

Figure 1 The changing capacity of computers

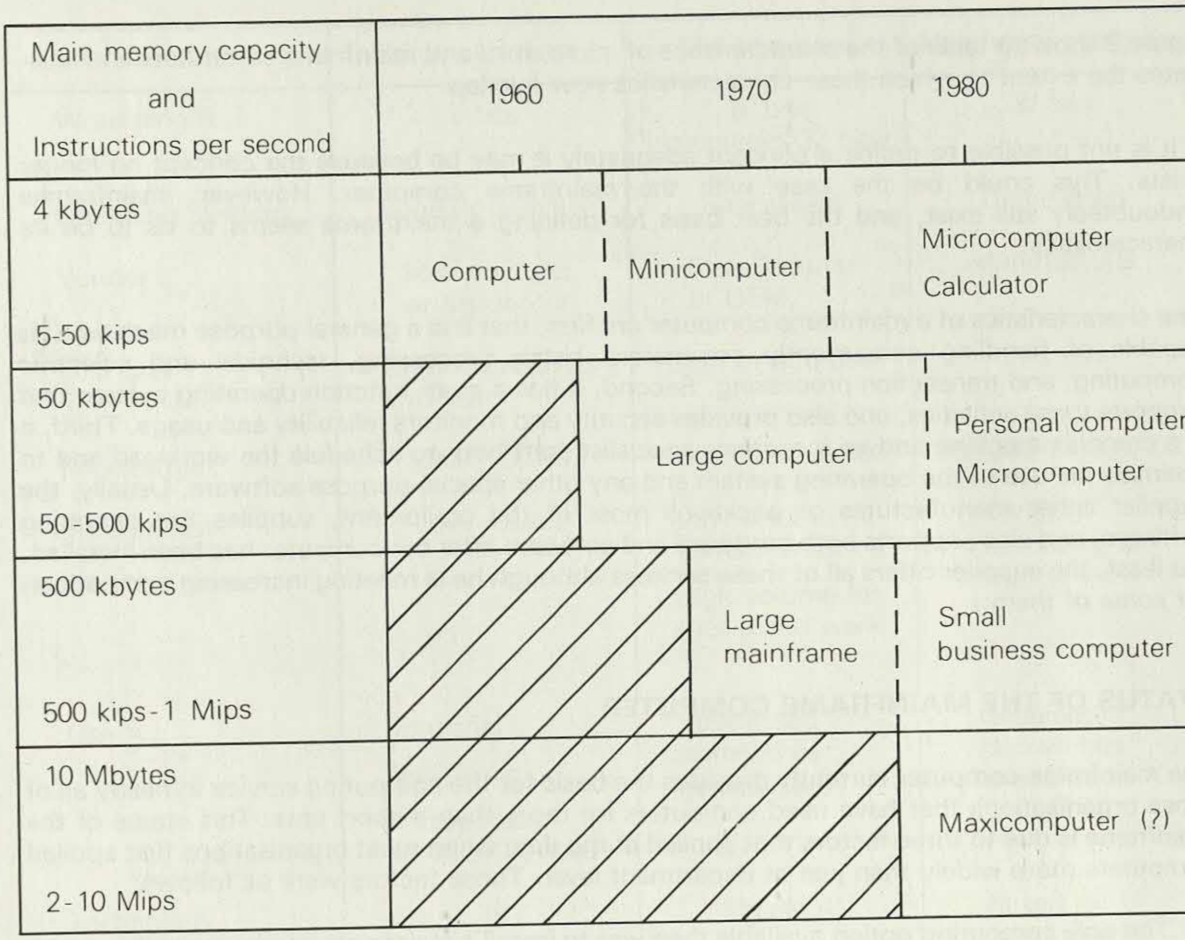

5. Definition based on exclusion
A mainframe computer, defined by what it excludes, is any computer that is neither a micro-computer nor a computer that is specially designed to handle a single application. A cynical commentator even suggested to us that a mainframe was a computer that was not very good at handling transaction processing. This definition assumes that all these categories can be defined, and we record it here mainly to illustrate how difficult it is to produce an acceptable definition.

Clearly, the distinctions upon which the original meaning of the term mainframe was based have now largely disappeared. As the cost of production has fallen, the size and power of computers has increased, so that the differences between small machines and large machines have been gradually eroded. For the price one paid for an IBM 360/65 in 1968 one could buy a 370/168 in 1975, and a 3033 in 1979. Since there is now little cost penalty in providing more power, suppliers offer the same processor under various guises, and they now present their range of products in terms of peripheral devices and software products. Similarly, a supplier who wishes to survive in this competitive market will move away from special-purpose products and towards general-purpose ones. This move will confer benefits in terms of economies of scale, lower marketing costs and lower maintenance costs. This means that the distinctions between micro, mini and mainframe computers will be discernible only at the extremes of the power spectrum.

Figure 2 shows a table of the characteristics of micro, mini and mainframe computers and illustrates the extent to which these characteristics now overlap.

If it is not possible to define a concept adequately it may be because the concept no longer exists. This could be the case with the mainframe computer. However, mainframes undoubtedly still exist, and the best basis for defining a mainframe seems to us to be its characteristics.

The characteristics of a mainframe computer are first, that it is a general-purpose machine. It is capable of handling concurrently commercial batch processing, technical and scientific computing, and transaction processing. Second, it has a multi-function operating system that supports these activities, and also provides security and monitors reliability and usage. Third, it is a complex machine and so it requires specialist staff both to schedule the workload and to optimise the use of the operating system and any other special-purpose software. Usually, the supplier either manufactures or packages most of the equipment, supplies the operating software, and also supports both hardware and software after the computer has been installed. (At least, the supplier offers all of these services although he is meeting increasing competition for some of them.)

STATUS OF THE MAINFRAME COMPUTER

The mainframe computer currently provides the basis for the computing service in nearly all of those organisations that have used computers for more than a short time. This status of the mainframe is due to three factors that applied at the time when most organisations first applied computers more widely than just at department level. Those factors were as follows:

1. The only computing option available then was to install a mainframe computer at the centre of the organisation.

2. Most organisations, at that time, had centralised management structures and, in many of them, computers were considered to be a means of applying central control.

3. Grosch's law (that computing power is proportional to the square of the cost) was still valid.

In those early days, mainframe computers were installed, their high cost demanded high utilisation, they achieved an almost self-perpetuating growth, and they became an essential component in an organisation's administrative systems. Today, even though other computing options are now available, even though the fashion in organisation structures has changed and decentralised profit centres are now in vogue, and even though Grosch's law is now, at best, linear, the status quo is slow to change.

Grosch's law began to break down in 1976, and figure 3 on page 6 shows the relationship between computing power and computing cost, and support costs and complexity, as it existed then. The figure illustrates that computing power rises faster than computing cost, and support costs (i.e. the costs of specialised staff) rise slower than computing cost, both of which factors support centralised mainframe computing. It also shows that complexity (in terms of the requirement for, and the size of, special operating software) rises very much faster than computing cost. It is the rate of growth of complexity that is becoming one of the major arguments against centralised mainframe computing.

Figure 2 Characteristics of typical systems

Characteristics	Microcomputer	Minicomputer	Mainframe computer
Word length	8 bits	16 bits (sometimes 32 bits)	32 bits
Operating systems	None or primitive	Specialised	Generalised
Vendor	Manufacturer or distributor	Manufacturer or OEM	Manufacturer
First user	OEM (sometimes end-user)	OEM or end-user	End-user
Price range	£200-£5,000	£10,000-£100,000	£50,000-£1,000,000
Throughput	Low volume	Low volume for general work High volume for specialised work	Variable according to type of work
Use	Personal Academic Incorporated in other devices	Specialised (sometimes generalised)	Generalised (sometimes specialised)
Technology	IC	Mainly IC	Mixed – becoming IC
Typical main storage capacity	16 kbytes	64 kybtes	512 kybtes
Path width	8 bits	16 bits	32 bits
I/O peripherals	VDU, printer, simple tape, discette	Wide range, but generalised I/O capacity specialised	Wide range, some specialised I/O very large capacity

THE PRESSURE FOR CHANGE

While the central mainframe remained the only economic computing option the user was chiefly concerned with achieving the best possible level of service. The complexity of the central installation, the introverted attitude of the data processing department, the high cost and the

Figure 3 Relationship between computer power and complexity in 1976

Financial view		Practical considerations	
Computer power	Related cost	Support costs	System complexity
1	100	1	1.0
2	175	2	2.5
4	300	3	8.0
8	450	4	20.0

Note: The numbers shown for system complexity are for general-purpose workloads with teleprocessing.

(Source: *Datamation*)

long lead time required for developing and amending systems all were regrettable but all equally unavoidable.

The rapid fall in the cost of hardware has changed that picture, and the change has occurred not only because small computers are much cheaper than they used to be. Equally important is the fact that the advantage that the application of Grosch's law conferred on large central computer installations has now been considerably reduced.

The OEM suppliers of minicomputers have recognised that the mini has a new market and they have developed that market accordingly. With product development cycles of about two years (compared with up to five years for a mainframe product) and with their aggressive pricing policies, the minicomputer suppliers have been able to offer an apparently attractive alternative to mainframe computing. This has not only brought new first-time users into the market, it has also led to the growth of distributed computing in large organisations which already had mainframe systems.

Viewed from the user's point of view, the minicomputer appears to offer the following advantages:

1. Compared with the total recharge arising from centralised computing the hardware is cheap to purchase, and it may even be cheap enough to avoid corporate capital expenditure controls.

2. The hardware is rugged and it can be accommodated in the normal office environment. Often it does not require either much space or special air-conditioning.

3. The simplicity of operation avoids the need for specialised staff. Even where it is necessary to appoint a computer supervisor, he or she can generally carry out all those various tasks (such as job scheduling, equipment security, file security and printer operation) that, in large central installations, frequently require separate specialised staff.

4. The computing system is now within the end-user's control. No longer does his data have to be sent to the central installation, to be scheduled, processed and eventually returned to him.

For any user manager who has suffered the frustrations of using a central installation, but who has never had to manage his own computing, these advantages seem very attractive. From the central DP manager's point of view, there are probably some attractions in restricting the growth of the mainframe workload, such as:

1. The transfer of the time-critical and business-critical jobs from the mainframe to the end-user's own computer relieves the DP manager of a source of trouble and worry.

2. The reduction in the number of jobs that are competing for resources simplifies the scheduling of work.

3. The reduction in the number of modes of processing (especially if it includes eliminating a transaction processing service during the day) reduces the complexity of operating software and increases the flexibility for scheduling batch work (especially rework) during the day.

Although the DP manager might feel that the development of distributed computing might reduce his own responsibility and lead to new problems of control, it might also offer him some relief from his immediate day-to-day difficulties.

What has been said above represents merely a simplified summary of the case for distributed computing. Nevertheless, the advantages given above represent the popular reasons for considering this alternative to the mainframe computer.

In addition, the mainframe now seems to be close to its technological limit. The present type of architecture will soon be limited by the speed of electrical signals. The present type of software is already limited both by the human inability and its own inability to manage complexity. If mainframe computers are to be required either to handle larger workloads or to process larger problems, new hardware and software architectures will be required.

THE INERTIA EFFECT

The significant constraint that limits the move towards distributed computing is the tremendous inertia that is built into the status quo. This inertia has the following three aspects:

1. Systems inertia
   Most of today's applications systems depend at least partly on central mainframe computers. It has been estimated that the replacement value of user code in the US is about $200 billion, and also that the US Federal Government currently spends more than $6 billion each year on software, 50% of which is spent on maintaining existing systems.

   Even if the necessary skills and resources were available, it would take a considerable time to change the basis of that amount of systems investment. Most managers, having spent so much time, trouble and money in achieving the level of computer use that exists today will need positive reasons to take the risk of changing rapidly to any other style of computing.

2. Management inertia
   DP managers are well aware of the very real problems of system development and maintenance. They may feel that the claimed advantages of distributed computing are sometimes based on a lowering of standards, for which the organisation will pay later. It could also be said cynically that DP managers have a vested interest in maintaining the status quo. Their salaries are comparable to those of other senior managers in the organisation, and their level of remuneration is based partly on the large amount of resources that they control and partly upon the general shortage there is of people who are qualified to do their job. Any trend towards either distributed computing or a diminution of the role of the mainframe could

undermine both the level of responsibility and the scarcity value of the DP manager. It would not be naive, therefore, to assume that DP managers may well attempt to control the rate at which distributed computing is introduced.

3. Supplier inertia
Suppliers of computers sell into three main market sectors: new applications, transfers from competitors, and growth of existing applications. Of these, the third is easily the biggest and most important market for the mainframe suppliers.

Inevitably, then, mainframe suppliers will do their utmost to sustain their existing market, even though they may be preparing at the same time against the relative decline of that market in the long term. The systems that are being developed to run on the mainframes that are being sold today will still be running (on these or similar machines) in four or five years' time.

Those who supply alternatives to mainframe computers find it difficult to break into the mainframe supplier's customer base. The high costs of overcoming the defensive tactics adopted by the mainframe suppliers and of overcoming also the systems and management inertias referred to above, threaten alternative suppliers' profit margins. Consequently, suppliers of mini and microcomputers are likely to continue to concentrate on new applications and first-time users.

SUMMARY

The mainframe as we have defined it on page 4 is well entrenched. It is sustained by the interests of buyers and sellers alike. Those organisations that have installed a mainframe will typically have a large investment in applications programs. These programs can be transferred to other suppliers' mainframes only at a considerable risk and at a high cost in time, effort and inconvenience. The risks and costs will be even greater when applications programs are to be transferred to other types of computing equipment altogether.

Mainframe suppliers have an investment in hardware, in systems software, and in their support organisations, and, naturally, they wish to protect that investment. They are under increasing pressure to generate new revenue from their customers as hardware prices (and hence gross profit margins) continue their inexorable, if somewhat jerky, fall. At the very least, mainframe suppliers will wish to control the rate of decline of the mainframe in order to maintain their level of revenues. Only by doing this will they be able to invest in, and then to transfer their customers to, new products which replace the mainframe. Most of them will no doubt hope to achieve something better than that. So far as the mainframe is concerned, the three main possibilities (ranging from the worst to the best) appear to be as follows:

1. Senile decline
The mainframe will not break through its present limits. Its workload will progressively be drawn off onto new products which, without the mainframe inheritance on their backs, will be able to cope with present tasks more effectively and also to take on new ones.

2. Stable maturity
New products will arrive, but these will complement rather than supersede the mainframe. The mainframe will continue to operate as it does now, and will gradually improve on the way it runs its batch programs, drives its database, and handles its terminals.

3. Coming of age
The mainframe will evolve steadily, grappling with and curing its own weaknesses, taking on new applications and providing new facilities as these are required. Through this evolutionary process it will consolidate its position as the hub in an organisation around which data processing, and ultimately information processing, revolves.

The threat to the mainframe comes both from outside and inside itself. The success of the minicomputer in business applications has encouraged the move to distributed computing, and this move is now firmly established. For the mainframe suppliers, this trend can mean loss of revenue from existing customers and ultimately loss of customers as well. Less well established, but no less threatening, is the present slow move to personal computers, which might draw away the mainframe workload.

These two moves represent the outside threats to the mainframe. They are not necessarily threats to the mainframe suppliers, since they may opt to join the movement rather than resist it. Indeed, most of them have already done so, albeit with varying degrees of conviction.

The threat from the inside is perhaps as serious as the threats from outside. It is the increasing complexity of the mainframe computer which, as we suggested on page 4, is one of the major arguments against mainframe computing. It could also be the main reason for the success the minicomputer has had in invading mainframe territory.

This complexity is a natural consequence firstly of the expanding role of the mainframe — from batch processing to timesharing to on-line processing — and, secondly, of attempts to improve its usability — to make it easier to operate, to program, to maintain, and so on. These attempts may have been successful individually, but taken as a whole they have created an environment which is far beyond the comprehension of most of the humble mortals who have to use the system. Of even more significance, as far as the future of the mainframe is concerned, is that complexity alone may be the factor which limits the mainframe's further development. It may become impossible for the supplier himself to develop and support the software, and equally, it may become impossible for the data processing professional or the end-user to make effective use of the equipment.

To determine what is likely to become of the mainframe we look more closely in the next three chapters respectively at those three groups of people that are most likely to influence its development: data processing managers, end-users and suppliers.

CHAPTER 3

THE ROLE OF THE MAINFRAME — THE DP MANAGER'S VIEW

OUR SURVEY OF LARGE COMPUTING INSTALLATIONS

To provide material for this report we carried out in-depth surveys of nineteen large computing installations. Our purpose was to find out what types of service these installations are providing, what opinions people hold about the strengths and weaknesses of the mainframe, and what trends can be discerned for the future. We summarise in this chapter the results of this survey. We give the results in some detail because we believe they will be of general interest as well as contributing to this particular report. For comparison, we have included, whenever they are relevant, the results of the 1979 *Datamation* survey of DP budgets.

1. Type of organisation
   The sample included fifteen manufacturing businesses, one national retail chain, a local government authority, a public utility and a large hospital. All of the organisations had at least ten years' experience of some aspect of computing. Eight of the organisations had a centralised management structure and eleven had a decentralised structure.

2. Type of computer
   Six of the organisations had exclusive use of an IBM mainframe, seven had exclusive use of an ICL mainframe, three had exclusive use of a Univac mainframe, and three had both IBM and ICL machines.

3. Type of service provided
   All of the organisations provided a central batch processing service (i.e. data is input directly to the mainframe, not via a remote terminal). All but two provided a remote batch service. Sixteen organisations provided a transaction processing service using the mainframe. Five provided a timesharing service and eight were using the mainframe for on-line program development.

4. DP department organisation
   Three of the DP departments were organised on a functional basis, with separate sections for systems analysis and programming. All the remaining departments were organised on a project team basis, with each team being responsible for either one business function or one operating unit. In none of the organisations were analysts and programmers located within user departments. (*Datamation* reported the same result.)

5. Staff levels
   The total number of staff in each department varied between 230 and 28. The number of specialist staff (defined for our purpose as staff who have special expertise concerned with supporting the department rather than with developing or maintaining applications systems) varied between 71 and 1. The percentage of specialist staff in the total varied widely. In general, the larger departments had a higher percentage.

   All of the organisations stated that they had fewer specialist staff than they required for their needs.

6. Budgeted cost

Twelve of the organisations provided information on their 1978/79 budget in absolute terms, and thirteen provided percentage breakdowns. The budgets ranged in size from £400k p.a. to £4,000k p.a. The average percentage breakdown was:

	Our survey	*Datamation* survey
Hardware and communications costs	44%	32%
Staff employment costs	42%	53%
Other costs (software, maintenance, consumables, etc.)	14%	15%
	100%	100%

The DP department budget varied between 0.5% and 1.0% of the organisation's annual turnover. In certain cases, however, additional computing expenditure was included in user manager's budgets (e.g. terminal costs and communications costs).

The DP department budget had increased by about 10% over the previous year's budget. (*Datamation* reported 12%.)

7. Software

With the exception of two organisations who used PL/1, all the organisations used COBOL for commercial work. The six organisations that carried out scientific computing generally used FORTRAN for this purpose.

Assembler level languages were used very little and then only by system software specialists.

There was very limited use of the newer, high-level languages, such as APL or CORAL.

Between 60% and 80% of software activity was directed at correcting and amending existing systems.

There was evidence of interest in some of the formal methods for developing systems (e.g. structured programming) but no evidence of their use.

Twelve organisations were already using database management systems. These included four organisations that had unusually complex applications. Interestingly, those four were also amongst the seven organisations that were evaluating proprietary database management systems.

All of the organisations were using the mainframe supplier's communications software, with the exception of one that had written its own. All of the IBM users were using or were intending to use SNA. The arguments they advanced in favour of SNA were:

— It exists now.

— It allows centralised control but with distributed intelligence.

— It may assist with the migration to full distributed computing.

Most of the organisations said they were willing to consider software packages as an alternative to in-house development, the main advantage claimed being the shortening of the system development time. However, apart from payroll packages, personnel packages, and a few simple ledger packages (implemented on local minicomputers), there was little evidence that applications packages were being used. Adverse comments were made concerning implementation difficulties and maintenance problems with applications packages. By contrast, software utility packages (e.g. disc utilities, library systems and program editors) were well regarded.

On average, about 4% of the total budget was spent on external software. (*Datamation* reported 3%.)

8. Distributed computing

Twelve of the organisations were using either micro or minicomputers for independent scientific or process control work.

Five of the organisations were using either micro or minicomputers for independent commercial applications.

Eight of the organisations were using minicomputers within a communications network, the minis being linked with the mainframe computer. In those installations, the mini was generally used as a terminal for data entry, data validation, and data transmission in batch mode. In general, scheduling and job control was done at the central installation, and genuine remote job entry was not common.

All of the organisations used terminals linked to the mainframe, but there was a wide variation in the type of use. The main types of use were:

— On-line enquiry.

— On-line file update.

— Timesharing (generally for either scientific or technical problem solving).

— Program development.

The *Datamation* survey stated that "most of the hardware is still firmly situated at the central site. What is going out to the field seems to be almost exclusively terminals".

9. Mainframe loading

All of the organisations were using their mainframe for more than the standard working day. The shift patterns were:

— Two shifts a day for five days a week (four organisations).

— Two shifts a day for more than five days a week (three organisations).

— Three shifts a day for five days a week (two organisations).

— Three shifts a day for more than five days a week (ten organisations).

Several of the organisations stated that their increased use of interactive systems was requiring service to be available for more than the standard working day in order to match production shift patterns. Two organisations said that they intended to upgrade their machine so as to avoid changing from two to three shifts a day.

Most of the organisations were using machine usage monitors that either had been provided by the mainframe supplier or had been written in-house. The only common factor these organisations had was a difficulty in defining a unit of usage and in interpreting the figures that they collected.

The general pattern of machine switched-on time was:

Production work (including rework)	63%
Program development and testing	25%
Housekeeping and maintenance	12%
	100%

10. Relationship with the user

A common pattern was evident in the involvement of end-users with system development projects. End-users were involved during the early stages (in project identification and feasibility studies) and during the ultimate stages (in file creation and system implementation). They were not, however, involved to any great extent during the stages of system analysis, system design, programming and testing.

The usual form of involvement was either through project steering committees or through working parties. In a few cases, end-user staff were seconded to the DP department.

All of the DP managers said that user involvement in the management of projects was desirable, partly because it compensated for a shortage of skilled DP staff and partly because it contributed to good system design.

In a few specialised systems, end-users were involved with file design or programming.

Ten organisations charged their users for DP services. Charges for system development work were based on staff times plus the cost of materials. Charges for computer operations were based either on CPU usage or on the number of transactions processed. In some organisations, differential cost rates were used to influence the pattern of use (e.g. to discourage batch working during the day).

Interestingly, all the six organisations that made no charge for DP services had centralised management structures.

There was a trend towards including the costs of data preparation and communications in end-users' budgets. However, two organisations that do this said that they experienced problems in allocating charges for multiplexed lines and in coping with uncontrolled growth of the communications network.

ATTITUDES TO THE MAINFRAME COMPUTER

We asked the organisations to rate the performance of their mainframes under twelve headings. Their individual responses are summarised in figure 4 overleaf.

The characteristics of the hardware (i.e. its fitness for the purpose) and the supplier's hardware support were generally rated as either satisfactory or good.

Figure 4 General performance rating of the mainframe in the surveyed user organisations

Characteristic	Number of respondents					Overall rating
	Excellent	Good	Satisfactory	Poor	Dreadful	
Hardware characteristics	1	7	3	–	–	76%
Manufacturer's service (hardware)	–	6	2	1	–	71%
Operating system software provision	–	2	4	4	–	54%
Telecommunications software provision		3	4	4	–	58%
System development support aids	–	–	2	1	–	53%
Manufacturer's software support	–	1	1	3	–	52%
Reliability	5	2	3	–	–	84%
Flexibility	–	4	3	3	–	62%
Ease of use	–	3	5	1	–	64%
Ability to handle batch applications	5	4	1	–	–	88%
Ability to handle real-time applications	1	1	1	6	–	44%
Price/performance characteristics	1	3	3	2	–	67%

The ratings of the operating system, the communications software, the system development software, and the supplier's software support varied between good (6 responses), satisfactory (11 responses), and poor (12 responses).

The reliability of the hardware was generally rated as either good or excellent. This factor was often cited as one of the major strengths of the mainframe. Where the mainframe is controlling a communications network any failure can be disastrous. Two organisations quoted availability figures of 98% and 95% respectively. These figures are not as good as might be hoped, bearing in mind that 95% availability means 25 minutes lost time each day. Flexibility (i.e. the ability to handle different types of work and to change schedules between them) and ease of use were generally rated as either satisfactory or good.

The ability to handle batch applications was generally rated as either good or excellent. By contrast, the ability to handle on-line applications was rated poor. (The one organisation that gave an excellent rating had written its own communications monitor.)

The ratings for price/performance varied, but were generally either satisfactory or good.

These findings, and our conversations with the DP managers, aligned with the commonly-held assessments of the strengths and weaknesses of centralised mainframe computers, which are tabulated in figure 5 overleaf.

TRENDS

Three important trends emerged from our questions regarding the organisations' short-term plans and intentions. They are discussed below:

1. Distributed computing
   All of the organisations were intending to increase the distributing of computing power and to place it within the control of the end-user.

   In general, they were planning to achieve this by creating a network containing a central mainframe, with minicomputers on end-user sites supporting terminals. They were all expecting that the development of the network would increase rather than decrease the load on the central mainframe.

   The major reasons the organisations gave for making this change were:

   — To meet the needs of a decentralised management structure.

   — To reduce the vulnerability of the service either to a failure of the mainframe or, most important, to industrial action.

   These reasons are, of course, based on managerial, rather than technical considerations.

   Other reasons the organisations gave us were:

   — To increase the end-user's control over both the development and the operation of his systems.

   — To increase the cost/effectiveness of the systems (from the end-user's point of view).

2. Distributed, replicated databases
   The trend towards a distributed, replicated database is closely related to the trend towards distributed computing. However, it calls for special comment here.

   The organisations said that one of the main reasons why they had decided to retain a central mainframe within their distributed network was to be able to maintain a central, integrated database. If end-users with local computing power need to have access to the contents of the database a sub-set of the data would be transmitted to the local site for enquiry and update purposes. The central copy of the database would always be maintained as the master copy. We describe this approach as the replicated database approach.

Figure 5 Relative advantages and disadvantages of mainframe computer

Advantages

Is flexible (handles both batch and on-line processing, changing its character over shifts).

Has powerful processors and large capacity store.

- Handles large-scale applications.
- Permits integration of DP systems.

Provides a variety of software facilities.

Provides multi-access to common facilities.

Has ability to do batch processing (which is the most economical form of data processing).

Enables scarce skills to be pooled and optimised.

Allows good career prospects.

Enables control to be centralised and standards to be enforced.

Offers cost-effective DP service (from the central point of view)

Disadvantages

Is remote from end-user (which may result in the view that it is an expensive overhead that does not contribute to company profitability).

Demands expensive skills, which are not readily available.

Has vulnerable central facility.

Involves lengthy timescales for developing applications.

Has high operating system overhead and poor telecommunications software.

Is not fully utilised until all applications have been developed.

Leads to dis-economies of scale (because of control problems and complexity).

Those organisations that have opted for this approach have done so rather than choosing either of the alternatives of providing access only to the central database, or of partitioning the database between the various users.

3. Profit-making applications
The third important trend was the organisations' intentions to start developing computer

systems to support the profit-making activities of the business. The implication behind this move is that they have already developed their cost-reducing applications satisfactorily. Examples of profit-making applications they quoted to us were:

- Financial planning and control, and cash management.
- On-line production planning and control, including real-time order processing and stock control.

SUMMARY OF FINDINGS OF OUR SURVEY

The main findings of our survey of DP installations may be summarised as follows:

1. DP managers approve of the following aspects of mainframe computers:
   - Their hardware and the supplier's support of it.
   - Their reliability.
   - Their flexibility to handle different types of work.
   - Their ability to handle batch work.
   - Their cost/performance.
2. DP managers are critical of the following aspects of mainframe computers:
   - Their operating software and the supplier's support of it.
   - Their limited ability to handle on-line applications.
3. The trend towards distributed computing will continue, and it will be achieved by:
   - Constructing networks which will include a central mainframe computer.
   - Maintaining a central database and replicating sub-sets of the data for local use.
4. The emphasis in applications systems will change from cost saving to profit making.

We conclude from this evidence that the role of the mainframe will change in the following ways:

- The provision of a transaction processing service will decrease, since this work will be transferred to local computers. However, database enquiry work will increase.
- A large batch processing load will remain, but it will be of a different kind. The raw business transactions will be transferred to local computers and will be handled interactively on them. The task of updating the central database will become a major batch job.
- The trend to distributed computing will add to the processing and storage requirements of the central mainframe.

Given the weakness of current operating software and the uncertainty over communications protocols, we believe that the organisations in the survey have made a practical and sensible decision regarding the control of data. The replicated database approach places more emphasis both on good system design and equipment reliability and less emphasis on communications and data management software.

CHAPTER 4

THE ROLE OF THE MAINFRAME – THE END-USER'S VIEW

We realised, when we carried out our survey of DP installations, that the results inevitably would be biased both by the technical and the political attitudes of the DP managers concerned and by their need to justify their own decisions. We are now carrying out a survey of end-users' attitudes to computing, but it is not yet complete. However, we present below a summary of the first twenty responses from this survey of end-users in order to strike a balance with the DP managers' views given in the previous chapter.

SURVEY PROFILE

Our questionnaire was completed by a senior manager in each of the twenty companies. (None of the companies in this survey were included in our survey of DP installations.) The profile of the end-user survey sample is shown in figure 6 overleaf.

The companies' experience with computing systems is shown in figure 7 on page 21. Not surprisingly, there was considerable experience with batch systems. The most popular batch jobs were financial accounting, warehouse stock control and payroll. There was also considerable experience with using on-line terminals, and ten companies had two or more years' experience. The most popular on-line jobs were financial accounting, warehouse stock control and management accounting.

Six companies had two or more years' experience with using a local small computer, the most popular jobs being production planning and control and warehouse stock control.

ATTITUDES TO COMPUTING

Our questionnaire asked the respondents to rate their opinions on various aspects of the computing service that they received. A summary of their responses is set out below:

1. Current performance
   Respondents were asked to rate the current performance of batch processing, on-line terminals, external timesharing, and in-house small computers on a scale of "excellent" to "dreadful". The results are illustrated in figure 8 on page 22, and can be summarised as follows:

   — Opinions on batch processing were dispersed, with the median being "satisfactory".

   — Opinions on on-line terminals and timesharing were either "good" or "satisfactory", but neither "excellent" nor "poor"/"dreadful".

   — Opinions on in-house small computers were either "excellent" or "good".

2. Current disadvantages
   Respondents rated a list of potential disadvantages of the four types of service on a scale of

Figure 6 Profile of end-user sample survey

	Number of respondents in each category
Business sector in which the organisation is engaged	
Finance	1
Distribution	2
Manufacturing	16
Other	1
Size of the organisation	
51 - 200 employees	3
201 - 1000 employees	5
1001 - 5000 employees	9
More than 5000 employees	3
Functional responsibility of the respondent	
Finance	8
Production	1
General management	4
Personnel	1
Divisional management and services	5
Marketing	1
Business environment	
Very stable	3
Changing but predictable	8
Changing and unpredictable	9

"very significant disadvantage" to "not a disadvantage". The number of respondents who rated each disadvantage as either very significant or fairly significant is shown in figure 9 on page 23.

The most serious disadvantages of batch processing, as perceived by the end-users surveyed, are:

— Lengthy development times.

Figure 7 End-users' experience with computing systems

Facility	Period of significant use			
	More than 5 years	2 to 5 years	Less than 2 years	No use
Batch processing systems	15	1	1	1
Terminals on-line to a large central computer	2	8	6	2
Terminals to an external time-sharing service	—	3	6	6
In-house small computer	1	5	4	7

— The difficulty or the risk in adapting systems to changing requirements.

— High development costs.

— High operating costs.

— Poor response to requests for new systems.

The most serious disadvantages attributed to on-line terminal services are:

— Lengthy development times.

— High operating costs.

— Poor response to requests for new systems.

These three disadvantages are all attributed also to batch systems.

3. Current advantages
Respondents described the main advantages of each of the four types of service. The most frequent comments were:

— Batch processing is cheap, reliable, secure and good for large-volume jobs.

— On-line terminals provide instant access, are flexible and are easy to control.

— Timesharing provides similar advantages to on-line terminals.

— An in-house small computer is within the end-user's own control.

Figure 8 End-users' rating of current performance of computing service

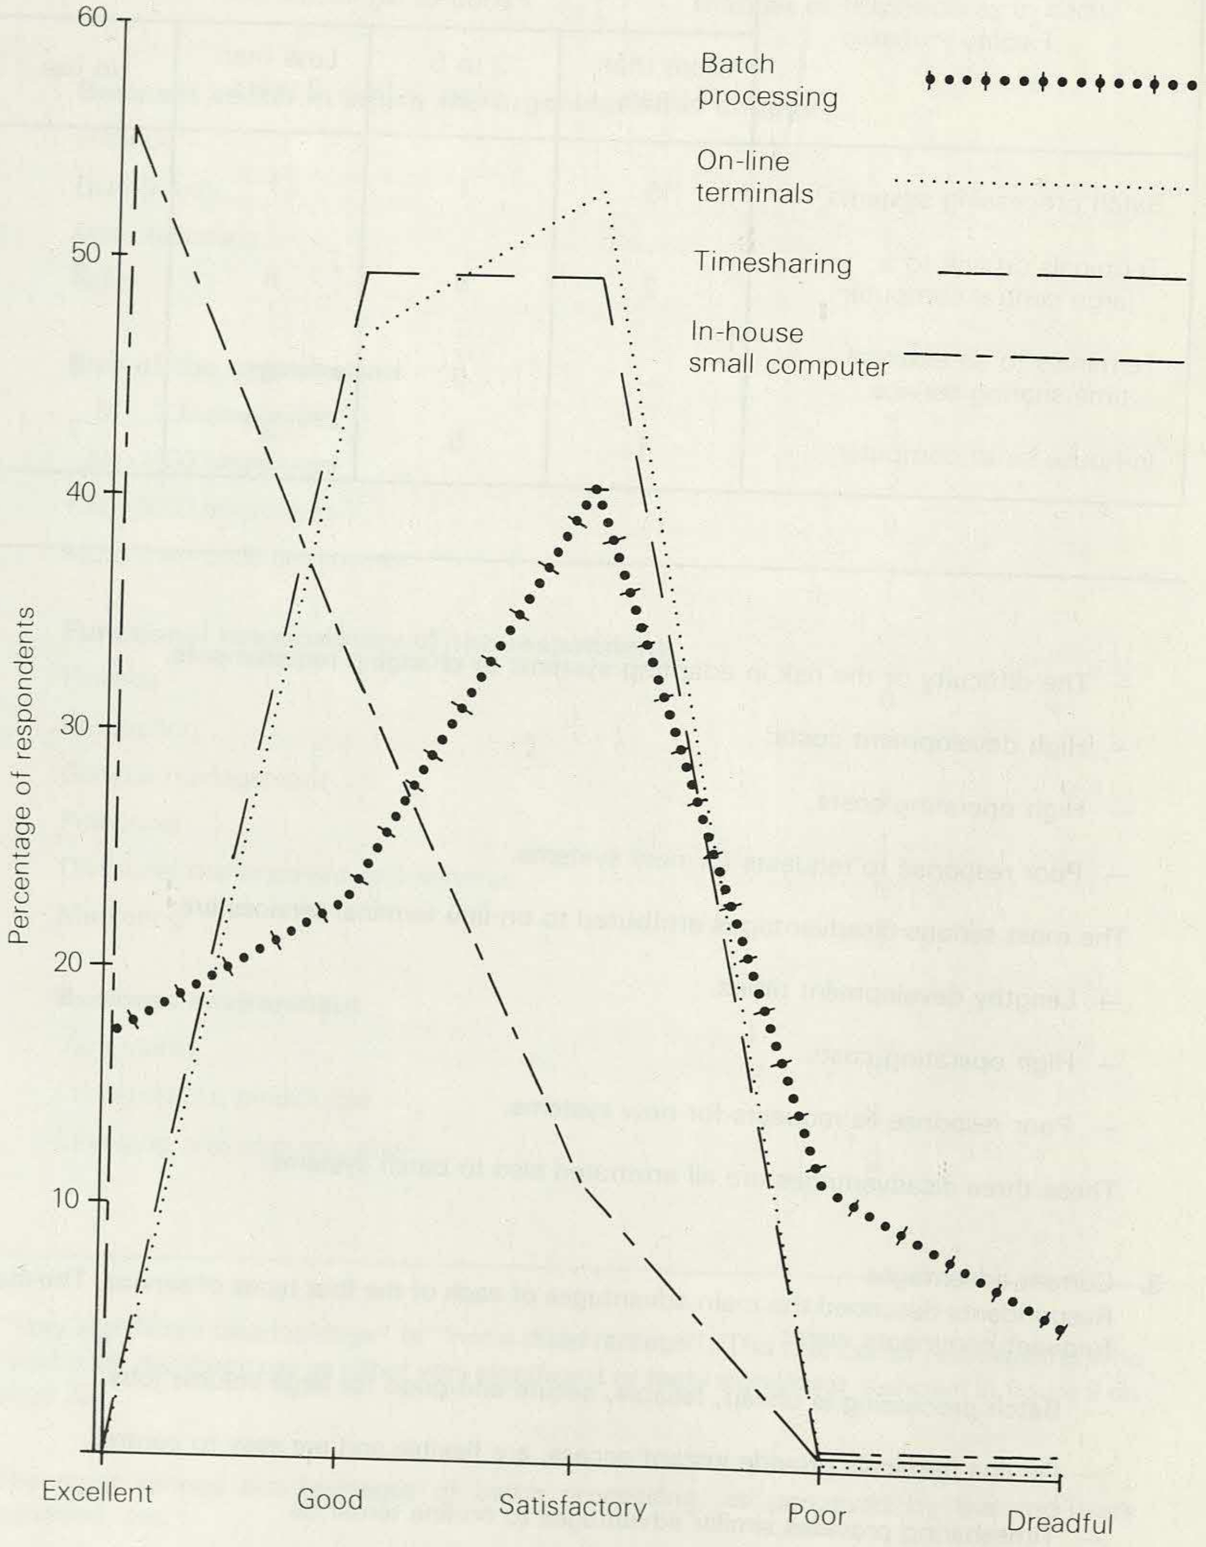

4. Charging policy

Respondents were asked what method is used in charging them for system development and computer operations and in what way charging policy affected their use of the resources. We were able to crosscheck some of their replies with the respective DP managers.

The general picture that emerged was that users either did not know how they were charged or, where they thought they did know, had the wrong impression of the charging method.

The effect of charging was almost always negative. Users either thought that the charges were too high (although it is difficult to see how they would know this), or were cautious in using computing services, or used the charging policy as an argument to gain control of their own computing.

Figure 9 End-users' perceived disadvantages of computer services

Perceived disadvantage	Batch processing	Use of terminals		
		On-line terminal	Time-sharing terminal	In-house small computer
High development costs	8	5	2	–
Lengthy development times	11	7	2	1
Poor response to requests for new systems	7	6	1	–
Poor operational times	4	2	1	–
Difficult or risky to adapt systems to changing requirements	9	4	1	–
Unreliable service	5	3	2	–
Difficult for staff to use	–	2	1	–
High operating costs	8	6	2	–
Lack of control	6	5	1	–
Output produced either late or in error	6	2	2	–
High vulnerability	2	–	2	1

Numbers show totals of disadvantages rated as very significant or fairly significant.

5. Involvement with DP staff

Respondents were asked what contribution they or their staff made to the development and the operation of computer systems. They were also asked about the effectiveness of the methods of communication between themselves and system development staff.

The results correlate closely with the results of our survey of DP installations. End-users are heavily involved in the early stages of a system development project but less so in the later stages. They are involved with project management and with monitoring system performance. Surprisingly, they are not greatly involved either with the selection of the terminal equipment that their own staff will use or with the training of their staff.

Several different communication methods were rated as being either very effective or quite effective.

6. Important needs

Respondents rated a list of suggested end-user needs on a scale from "essential" to "no interest". The results are shown in figure 10.

Figure 10 End-users' perceived needs

Facility required	Degree of requirement				
	Essential need	Strong need	Some need	Important but already adequate	No interest
"User-friendly" terminal devices	–	3	8	2	2
Non-procedural programming languages, to enable end-user staff to develop own systems	4	3	4	2	3
Easy access to data files using query languages	6	4	5	2	1
More direct means of entering data, e.g. by voice or scanner	–	4	7	–	5
Higher reliability and greater availability of computers	3	3	2	8	–
Cheap, stand-alone computers	6	1	4	4	2
Improved communications systems	2	7	2	2	2
Better integration between existing systems, e.g. text and data processing	3	4	6	–	1

The most important needs and the respective "scores" were:

— Easy access to data files, using query languages (31).

— Cheap, stand-alone computers (24).

— Better integration between existing systems (24).

— Improved communications systems (22).

— Non-procedural programming languages, of a kind to enable end-users to develop their own systems (22).

SUMMARY OF PRELIMINARY RESULTS

These preliminary results from our end-user survey correlate in some respects with our survey of DP managers and diverge in other respects.

Leaving aside the perceived need for cheap stand-alone computers (which can presumably be met independently of the mainframe-based services), some of the important needs identified by the end-users could be satisfied by the mainframe. Others could be satisfied by the type of distributed computing system that was being considered by the DP managers. Some of the needs might also be satisfied by stand-alone minicomputers as their software capability develops. We attempt in figure 11 to summarise the match between these end-user needs and the major available equipment options. If all needs exist within an organisation, and if they are all to be met in the optimum way, it seems that no single solution would be appropriate, and a hybrid approach would need to be adopted.

As far as existing services are concerned, the level of end-user satisfaction with batch systems is lower than for any of the other services we surveyed. Overall, however, the level of satisfaction was perhaps higher than many people would have expected.

Figure 11 The matching of equipment options to end-users' needs

Facility required	Suitability of equipment options		
	Mainframe (with DBMS)	Replicated database	Stand-alone mini/micros
Access to files/query languages	Yes	Yes	Possibly
Cheap stand-alone computers	No	No	Yes
Better integration between systems	Yes	Possibly	No
Better communication systems	Unpromising	Possibly	No
Non-procedural languages for users	Yes	Possibly	Possibly

There is also a conflict between end-users' responses to the question on the disadvantages of batch services and some of their other replies, and this applies particularly to their general comments. The latter seem to indicate that computer systems (most of which in our sample are mainframe-based systems) have had a significant and positive impact on those business areas for which the end-user managers concerned are responsible. The adverse reaction seems to result principally from the remoteness (in every sense of the word) of the service and from the inconvenience and frustration that result from this. Dissatisfaction then tends to focus on the obvious targets, such as costs and lengthy development times. Our findings on charging schemes tend to support this interpretation.

It is also noteworthy that despite their complaints about costs, practically all the respondents look forward to extending their use of computer systems in the future. On this subject too, frustration manifests itself. It ranges in intensity from threats of a showdown over the cost of centralised computing to pleas for greater control and more use of terminals.

CHAPTER 5

THE ROLE OF THE MAINFRAME – THE SUPPLIERS' POSITION

In this chapter we look at the mainframe computer from the suppliers' point of view. We describe the pressures that are currently influencing the suppliers, we suggest what their marketing reaction to these pressures is likely to be, we speculate upon the tactics that they may adopt as a defence against competition, and, finally, we reach some conclusions about the sort of mainframe computer that the suppliers will be likely to market during the 1980s.

PRESSURES ON THE SUPPLIERS

To suggest that the market for the mainframe is under pressure may sound surprising against the statements in the press that the queue for new IBM systems represents more computing power than is currently installed in the Western world. A more objective analysis, however, reveals that there are pressures from four sources as we discuss below:

1. End-user requirements

   End-users have recently discovered two basic truths about computing. The first is that there are now alternatives to service from a central mainframe. They may not be able to exploit these alternatives themselves, but the mere existence of the alternatives is cause enough for them to put pressure on their DP managers who, in turn, put pressure on the suppliers. The second basic truth is that computing hardware is now relatively cheap. The DP manager can still point to all the other costs and risks associated with developing and running computer systems, but the end-user is now much more strongly tempted to call the DP manager's bluff. Again, this pressure on the DP manager rebounds on to the supplier.

   The sorts of products and services that users and DP managers are now demanding have already been described in chapters 3 and 4. The minicomputer suppliers (who have shorter product development cycles and no inheritance of obsolete products such as burden the mainframe suppliers) are able to respond more quickly (if not more adequately) to these demands. The mainframe suppliers, therefore, cannot afford to ignore the demands.

2. Technology

   The technology of computing equipment is developing rapidly, particularly with regard to storage devices and manufacturing methods. The changes in technology make existing products obsolete, and although user organisations do not normally want the latest technology *per se*, they do want the improved performance, the greater availability and the increased reliability that the technology makes possible. New technology also brings new competitors into the market, and these intensify the pressure on the established suppliers.

3. Finance

   The mainframe market is still growing by at least 10% p.a. There is no doubt, however, that this rate of growth is lower than it used to be and is also lower than the rates of growth of other sectors of the computing market. In addition, the reduction in the prices of equipment has reduced the absolute profit per unit sold.

   These two factors combine to put financial pressure on the suppliers. The historical source of their growth in profits is now drying up, so they must seek growth in profitable sales from other sectors of the market.

The present trend in costs is not a new phenomenon. Figure 12 shows the trend in various computing costs from 1960 to 1985. The trend is also reflected in the changing pattern in DP budgets. The cost of the CPU and the main memory has fallen from about 80% of the hardware budget to about 35% to 45%. Terminals and peripheral equipment have increased in cost during the same period from about 20% to about 30% to 45%, with the balance being made up with data entry equipment.

Figure 12 Trend in computing costs

	1960 costs	1975 costs	1985 costs
Hardware			
Processors and internal storage	100	.5	1
Fast access main storage	100	2	2
Communication lines	100	61	32
Software			
Developed in-house	100	28	13
Purchased externally	100	6	2

4. Competition
   The mainframe suppliers' existing customer base is also being attacked from two directions by competitive suppliers. First, plug-compatible manufacturers (PCMs) are offering not only compatible hardware but also compatible operating software, at a better price and/or on better terms than the mainframe suppliers. This means that the DP manager can continue to offer the same systems service, in the same way, but more cheaply.

   Second, suppliers of small business systems are appealing directly to end-users. They are offering alternative solutions to user problems that can be implemented quickly and that appear to be cheaper than the mainframe service.

Although it is still difficult to forecast the long-term impact these pressures will have on the profitability of the mainframe suppliers, they are clearly causing the suppliers to adjust their marketing strategy, which we discuss next.

MARKETING STRATEGY

The likely reactions of the mainframe suppliers to the pressures described above have been widely discussed, and were discussed in Foundation Report No. 14. Naturally the mainframe suppliers will persevere with their attempts to smooth over their existing problems. In particular, they need to attempt to deal with the two most unsatisfactory aspects of mainframe computing, namely the long lead times and the high cost of developing systems, and the difficulty of maintaining systems. No doubt they will offer new software packages and new system development tools and techniques (e.g. IBM's IPTs and ADF, and ICL's Data Dictionary System). They will also probably concentrate on improving the reliability and the ease of use of their systems.

Most DP managers will approve of developments of this kind. Other aspects of the mainframe suppliers' approach, however, will be more difficult to evaluate and may sometimes be positively unwelcome to DP managers. We outline below four such aspects of mainframe suppliers' policy. They are not new developments, but, as competitive pressures increase, the mainframe suppliers may begin to apply these aspects of their policy more ruthlessly than they have done in the past. In fairness, it must be pointed out that it is not the mainframe suppliers alone who apply these policies. So also do all those suppliers who offer a systems solution rather than merely hardware. In all cases, their motivation is the same — to protect their investment in both a systems capability and in a support organisation (which is what the user pays most for) against those suppliers who offer hardware alone (which is a small and decreasing proportion of the total cost) and against those suppliers who (like the PCM suppliers) offer something in between.

1. Locking in the customer
   Suppliers will be strongly tempted to keep their new products different from competing products, in order to deter their customers from changing to alternatives. This will be particularly evident with network architectures. As we see it, the sequence of events will be:

   — The supplier will offer attractive new products that are supported only by his network architecture.

   — The customer will become committed to the network architecture.

   — The supplier will then plan to control the rate at which he introduces additional products, irrespective of the competition, so as to maximise profitability.

2. Maintaining obscurity
   The distinctions between hardware and software are being confused. Depending mainly on the economics involved, functions can be implemented either in silicon logic, or in microcode, or in stored software. The intermixing of hardware and software helps the supplier to make his designs obscure, and in this way handicaps the activities of the PCM suppliers.

   If the interfaces between the components of the computer do eventually become standardised, suppliers, as a strategy to maintain their obscurity, may resort to the encryption of the control signals that pass across the interfaces.

3. Providing total solutions
   To counteract the fall in profit per sales unit, suppliers will seek ways of increasing their sales volume. One obvious way to achieve this will be to evolve from the numerical data processing business by offering the ability to handle all types of information — numbers, text, images and voice.

   Users will see this trend, rightly, as a move to increase either managerial or office productivity or both. Suppliers will see it as a new outlet for processing power, but they will choose their products and services carefully. The development of a specialised logic chip is very expensive, and so suppliers will seek high-volume applications through which they can recover their development investment. (The interest in retail point-of-sale equipment provides a good example.) For the same reason, suppliers will avoid specialised, low-volume applications and will leave them to systems houses.

The biggest potential market is in office multifunction terminals. The rising cost of people will soon cross the falling cost of equipment and an explosion in the availability of products of this kind can then be expected.

The emphasis in the future will be on integrated, total solutions. Users must weigh up whether the limitation of choice that this brings with it is acceptable to them.

4. Changing the pricing policy

Suppliers will almost certainly change their pricing structures. As we discussed in Foundation Report No. 14, IBM's price umbrella has shrunk, or perhaps has collapsed altogether. Inevitably, therefore, mainframe suppliers will no longer be able to rely on generous profit margins on hardware sales and will need to look elsewhere for revenue.

The obvious source for this new revenue is software, although it is doubtful whether enough revenue can be raised from software to compensate fully for the declining revenue from hardware. The collecting of software revenues presents considerable difficulties, as the experience of the record companies with the collecting of copyright dues on cassette tapes demonstrates.

The consequences of this new attitude to software could be that the level of bundled support available from mainframe suppliers will decline, that the cost of bespoke features will climb rapidly, and that the level of generality in general-purpose software products will increase in order that they can be offered to the widest possible market. It could also mean that the ability of the mainframe suppliers (including the market leader, IBM) to develop major new software products will be severely constrained. If this happened, the present generation of operating systems might be with us for a very long time.

SUMMARY – A DESIGN BRIEF FOR THE MAINFRAME

We have discussed above those pressures that currently influence the mainframe suppliers, and we have also discussed the mainframe suppliers' likely marketing strategy. Finally, by way of a summary, we set out below a design brief for a possible mainframe for the 1980s. We deal first with the users' requirements and second with the suppliers' requirements.

To meet the users' requirements, a mainframe of the 1980s must:

1. Be fully compatible with today's equipment and today's applications software. Most users cannot even contemplate the possibility of undertaking a large conversion exercise.

2. Have improved reliability, with self-diagnosis capability and, if possible, self-correction capability.

3. Be capable of handling information of all types, in very large volumes, and in all processing modes.

4. Be easy to operate and use.

5. Be supported by effective system development tools and methods.

6. Be compatible with standard communications protocols and public network interfaces.

7. Be easily expandable to accommodate additional functions.

To meet the suppliers' requirements, a mainframe of the 1980s must:

1. Have a flexible architecture that will support new types of components that have not yet been designed.

2. Make good use of standard components and manufacturing methods.

3. Satisfy standard interface requirements, whilst still retaining scope for obscurity.

4. Have hardware and software that are easy and cheap to maintain.

5. Offer a degree of integration of functions that limits the user's choice of products (and particularly the choice of peripheral and terminal products) from outside the supplier's own range.

6. Be competitive both in terms of functions and price/performance.

CHAPTER 6

MAINFRAME DESIGN IN THE 1980s

Before we draw our final conclusions, we consider, in this chapter, those possible developments in technology that might improve the mainframe's capabilities. Those developments might enable the mainframe to overcome the forces of inertia that we discussed on pages 7 and 8, or might enable it to begin to cope more successfully with its own complexity.

In particular, we discuss the directions that the designers of large general-purpose computers may take during the 1980s. We base our discussion partly on the evidence we have set out earlier and partly on what can be derived from those machines that are already available or whose designs have been published.

DESIGN INFLUENCES

In seeking to extend the capabilities of today's mainframes, designers are influenced by the four external factors that are discussed below:

1. Limits on performance
   Until recently, mainframe computers have generally had a monolithic processor architecture. The single box that one saw when looking at the CPU contained a single electronic module — the processor — which was closely integrated with the main memory.

   The performance of the monolithic processor is rapidly reaching its limit. That limit is imposed by the speed at which electrical signals can be transmitted (i.e. the speed of light) and by problems of heat dissipation. Improvements in performance may still be possible, but they are increasingly expensive to achieve. The price/performance graph is illustrated in figure 13, and designers are already at the point at which the gradient rises rapidly.

   An alternative to the monolithic processor must be found if computers are to achieve greater power. That alternative is likely to take the form of a complex of closely-linked modules.

2. Functional trends
   A mainframe computer may be regarded as having three main components which are linked by channels. The three main components are:

   — A communications module, which handles input from and output to local and remote terminals, readers and printers.

   — A processor module, in which programs are run and data is transformed.

   — A data module, in which programs and data are stored.

   This macro-architecture is illustrated in figure 14.

   Traditionally, the processor module has been regarded as the centre of the design.

Figure 13 Price/performance graph for monolithic processors

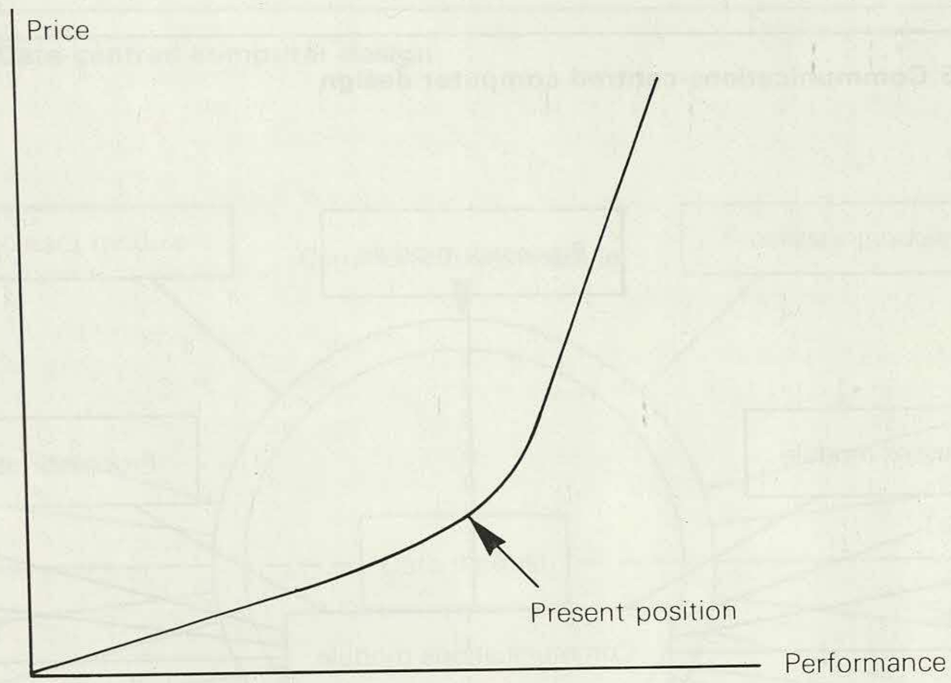

Figure 14 Traditional computer design

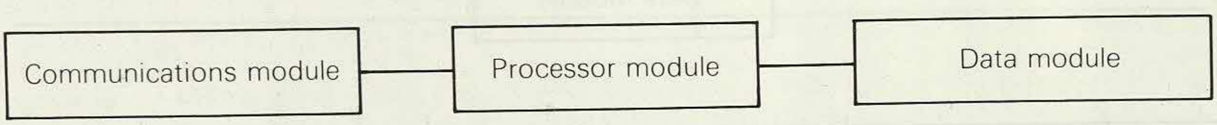

As the price of processing logic has fallen this concept has become rather inappropriate. The reason for this is that processing power can now be sited at various places in the architecture and a large modern mainframe may easily have a dozen or more processors within it.

If, as we have suggested, the rationale for the mainframe is its ability to integrate DP systems and to provide common access to a corporate database, then future mainframe architectures may be expected to be either communications-centred or data-centred instead of processor-based.

The communications-centred design is illustrated in figure 15. It shows a communications module (which might be either a linear bus, or a ring bus or a very fast switch) which links out to various kinds of terminal and links in with various kinds of processor modules and data modules.

Figure 15 Communications-centred computer design

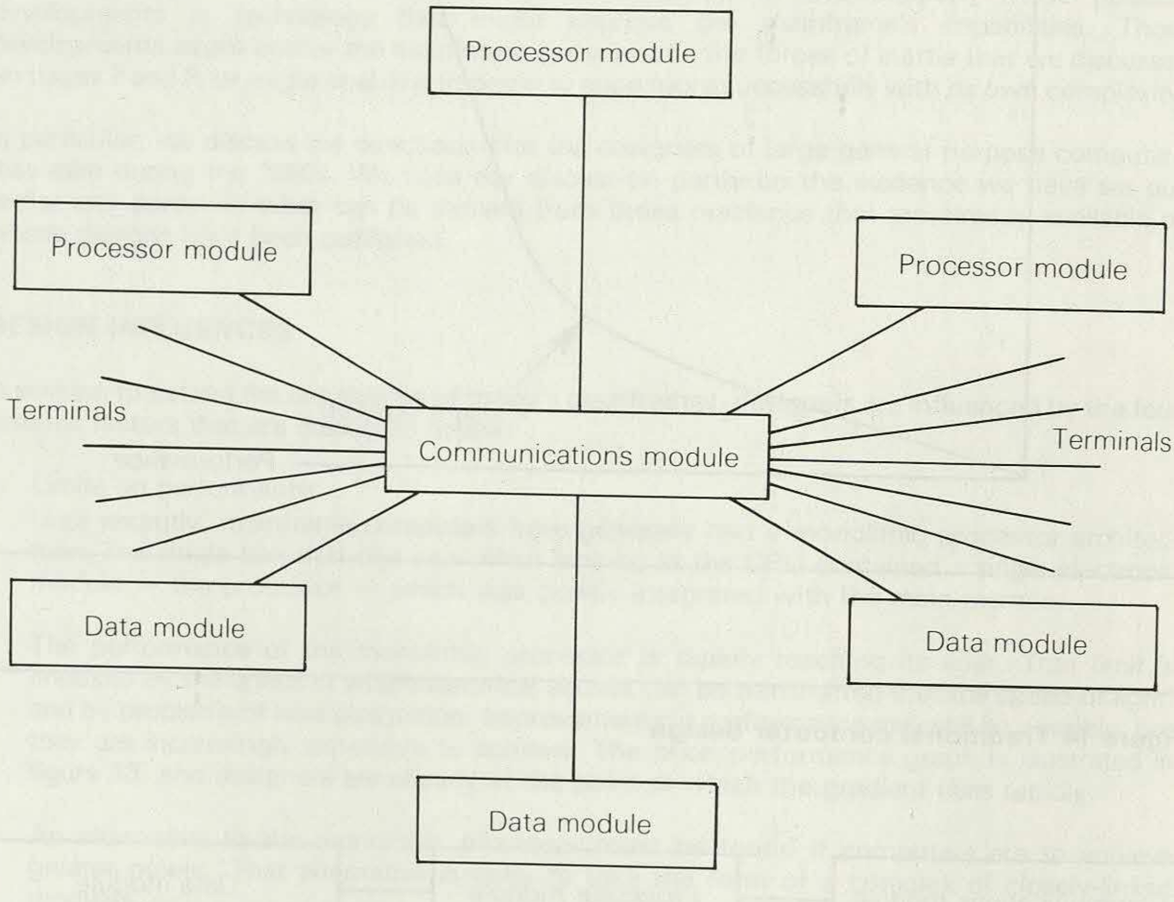

The data-centred design is illustrated in figure 16. The data module would be a very large virtual-storage system, comprising various types of storage device. In the figure, the communications module is shown as a ring, but it might equally well be either a linear bus or a very fast switch. Various kinds of processor module would be linked to the communications module.

In engineering terms, the two alternative designs illustrated in figures 15 and 16 might not be very different. But conceptually and functionally they would be different. The communications-centred device would be designed primarily to act as a switch. It would provide connection between all the end-users in an organisation, plus various types of service for handling all types of information.

The data-centred device would be designed primarily for information storage and retrieval, but it would include powerful information processing facilities. This design is close to the proven capabilities of today's mainframes.

Figure 16 Data-centred computer design

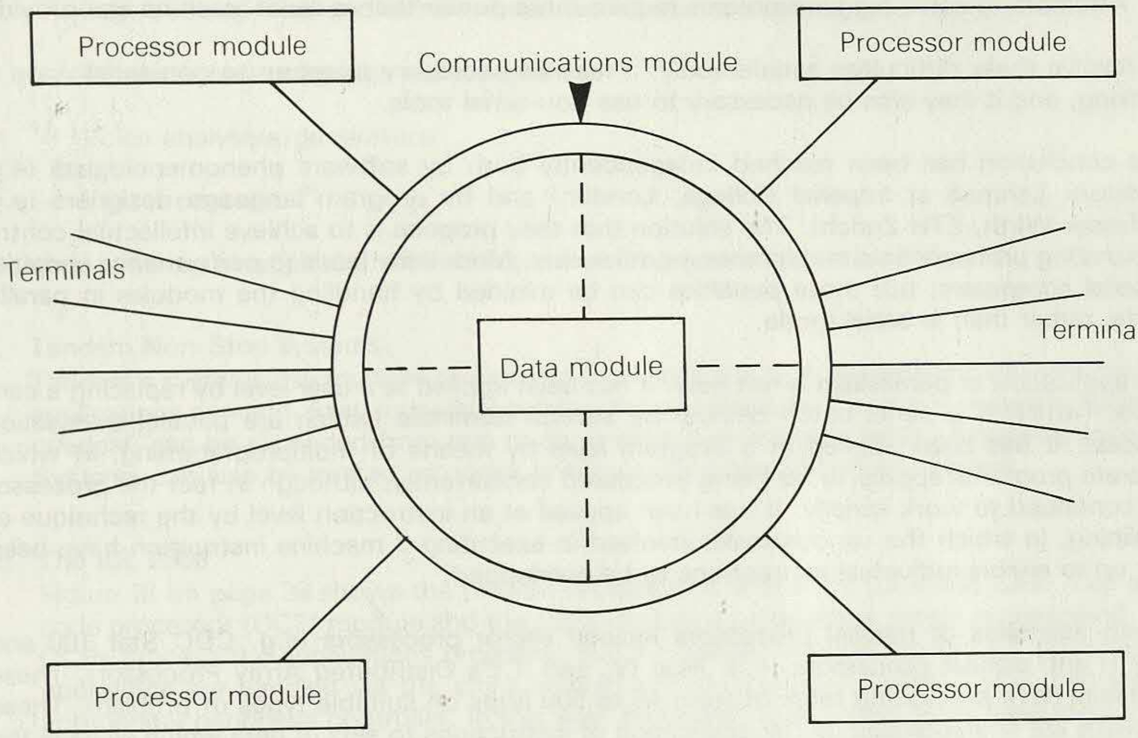

---

3. Flexibility
   As we suggested in chapter 5, progress from the suppliers' point of view will come not from great leaps forward, but rather from a steady process of evolution across an increasingly wide front. This means that designs must be capable of accommodating new devices. Again, this suggests that future designs will be based on a complex of closely-linked modules, the composition of which would be easy to change.

4. Reliability
   With the declining cost of integrated circuits, increased reliability will be achieved by duplicating or even triplicating components. It is still not clear whether this redundancy can best be provided within the chip (as in the new IBM 64k store chips), or on the board, or by replicating boards or whole processors and stores. One of the processor modules is likely to act as a reliability controller, and, in this role, it will monitor performance and switch in alternative components if any components fail.

SERIALISM AND PARALLELISM

The traditional computer is very firmly based upon the von Neumann principle of the sequential application of processing instructions to data. This serial approach has led to the following two difficulties:

1. Attempts to consider and to analyse problems in a serial way are producing confusion rather than solutions. Analysts and programmers have lived with serial processing for so long that they have forgotten that many applications are inherently parallel in nature. This is true not just of weather forecasting and nuclear physics, but also of modelling and stock control.

2. Attempts to solve bigger problems require more power than a serial machine can provide.

To resolve these difficulties satisfactorily, it may be necessary to return to non-serial ways of thinking, and it may also be necessary to use non-serial tools.

This conclusion has been reached independently both by software phenomenologists (e.g. Professor Lehman at Imperial College, London) and by program language designers (e.g. Professor Wirth, ETH Zurich). The solution that they propose is to achieve intellectual control by handling problems in a much more modular way. Modularity leads to performance penalties in serial computers, but these penalties can be avoided by handling the modules in parallel mode, rather than in serial mode.

The application of parallelism is not new. It has been applied at a user level by replacing a card reader (which is a serial batch device) by several terminals (which are parallel transaction devices). It has been applied at a program level by means of multiprogramming, in which separate programs appear to be being processed concurrently, although in fact the processor has continued to work serially. It has been applied at an instruction level by the technique of pipelining, in which the various tasks involved in executing a machine instruction have been split up to enable individual instructions to be overlapped.

Recent examples of parallel processors include vector processors (e.g. CDC Star 100 and Cray 1) and cellular processors (e.g. Illiac IV, and ICL's Distributed Array Processor). These machines have processing rates of from 40 to 500 Mips on suitable types of problem. These problems are characterised by the application of instructions to sets of data which all have the same format (e.g. vectors and matrices). Ways now need to be found of applying this approach to problems of a more general nature.

Conceptually, there seems to be no reason why effective processing power could not be radically increased by applying parallel processing techniques. In practice, there have been three historical barriers to parallel processing. The first is the cost of providing a worthwhile number of processors, with their associated memory, in which to run the parallel modules of the job. The second is the difficulty of breaking down the problem into independent modules that can be processed in parallel. The third is the difficulty of providing an operating system to control the parallel processing.

THE MULTI-MODULE MAINFRAME

We have suggested in this chapter so far that the limits of serial processing, the advantages of parallel processing, and the need for flexibility and reliability, all indicate that the mainframe computer in the 1980s will have a multi-module basic architecture. This trend is already evident. Large Burroughs and Univac computers have been designed on this principle for some years. The larger members of the ICL 2900 series and the IBM 303X series are also multi-module machines.

The sorts of modules that could be included in the mainframe complex might include:

- Fast serial processors, probably incorporating pipeline techniques.
- Special-purpose processors, e.g. vector processors or cellular processors.
- Many smaller medium-speed processors, to handle parallel processing.
- Communications modules, to provide both internal and external switching.
- Virtual-storage modules, with a storage controller and various kinds of storage device (e.g. bubble, fixed and moving head disc drives and optical memory devices).
- Database processors (as discussed below).
- Voice analysers/generators.
- Text processors.

Some interesting examples of multi-module mainframes are discussed below:

1. Tandem Non-Stop systems
Tandem's systems use multiple processors to provide the high availability that many on-line applications demand. Each individual system, the architecture of which is shown in figure 17 overleaf, can be expanded from two to up to sixteen processors, without re-programming. Systems can also be treated as nodes in a geographically dispersed network.

2. The ICL 2900
Figure 18 on page 39 shows the present architecture of the ICL 2970 and 2980. The order code processor (OCP) module and the store module may be either single or replicated. The OCPs process the applications programs that are loaded into the stores. All input/output operations are handled via the store module access controller (SMAC) and then by the appropriate peripheral controller. In this way, input and output operations have been transferred from the OCP to the SMAC. This architecture does not correspond with either the data-centred or the communications-centred design that we described, but it does illustrate the disintegration of the monolithic processor.

3. The MU5 and MU6 systems
Staff at Manchester University, having built the Atlas computer, started designing its successor, the MU5, in about 1968. This computer has been in use, as part of a multiprocessor complex, since about 1973.

   Figure 19 on page 40 shows the architecture of the MU5. The main components are:

   - The MU5 computer, a very fast, multi-programmed computer with a power rating of about twenty times Atlas (i.e. about 4 Mips).
   - An ICL 1905, which is used as a batch input/output controller.
   - Several PDP-11 machines, which are used as interactive terminal controllers.
   - The exchange, which is a very fast switch which operates at about 80 Mbytes per second, and which links any of the processors with any of the stores, or with any of the other processors.

Figure 17 Schematic diagram of Tandem Non-Stop architecture

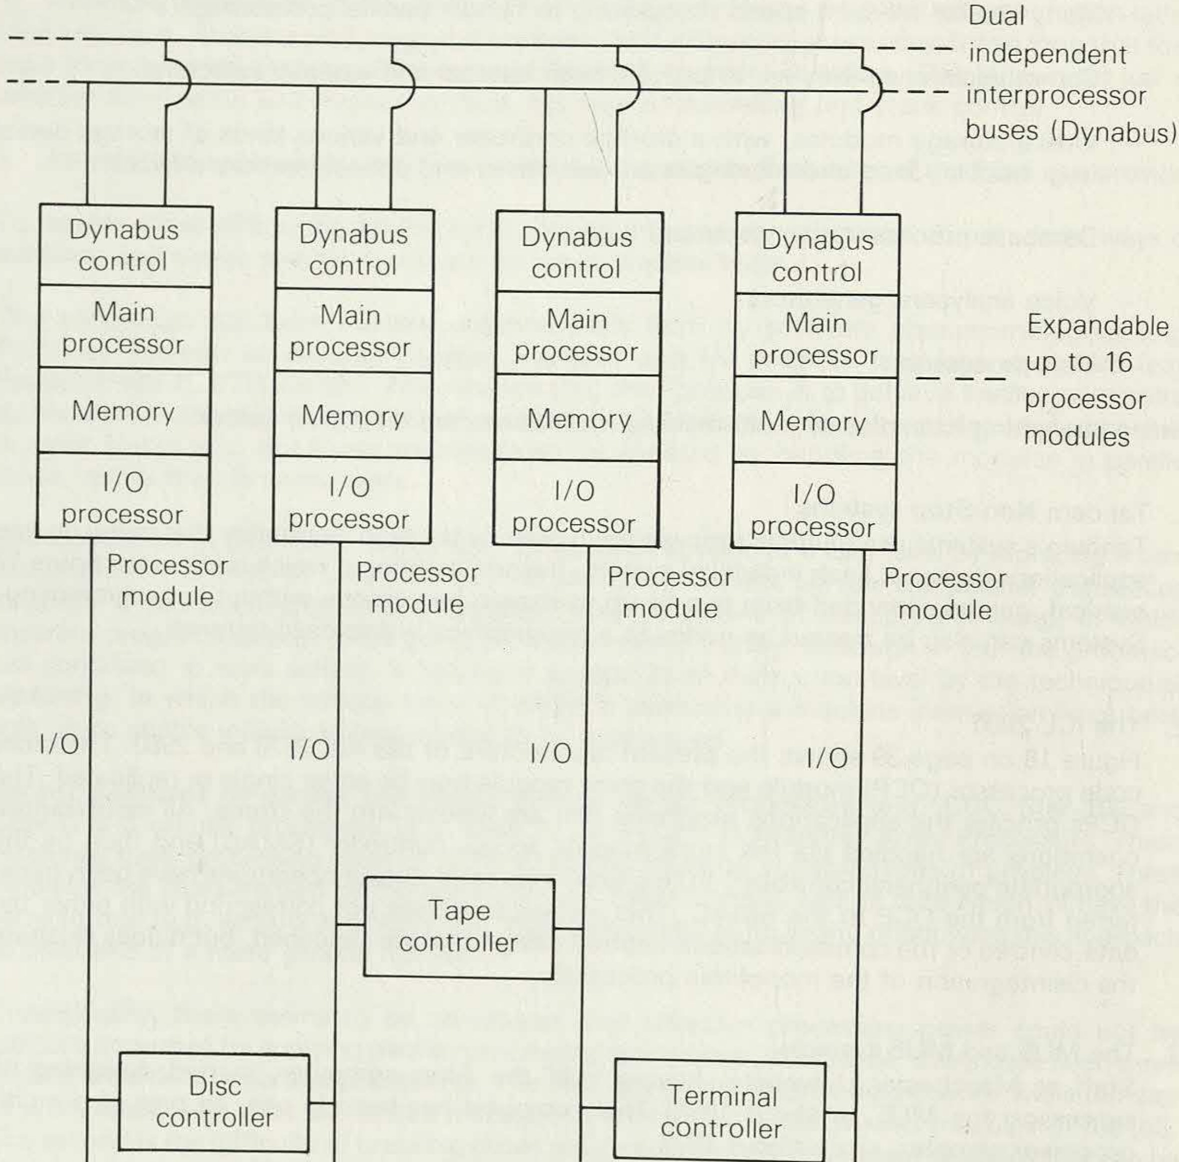

Figure 18 Architecture of the ICL 2970 and 2980

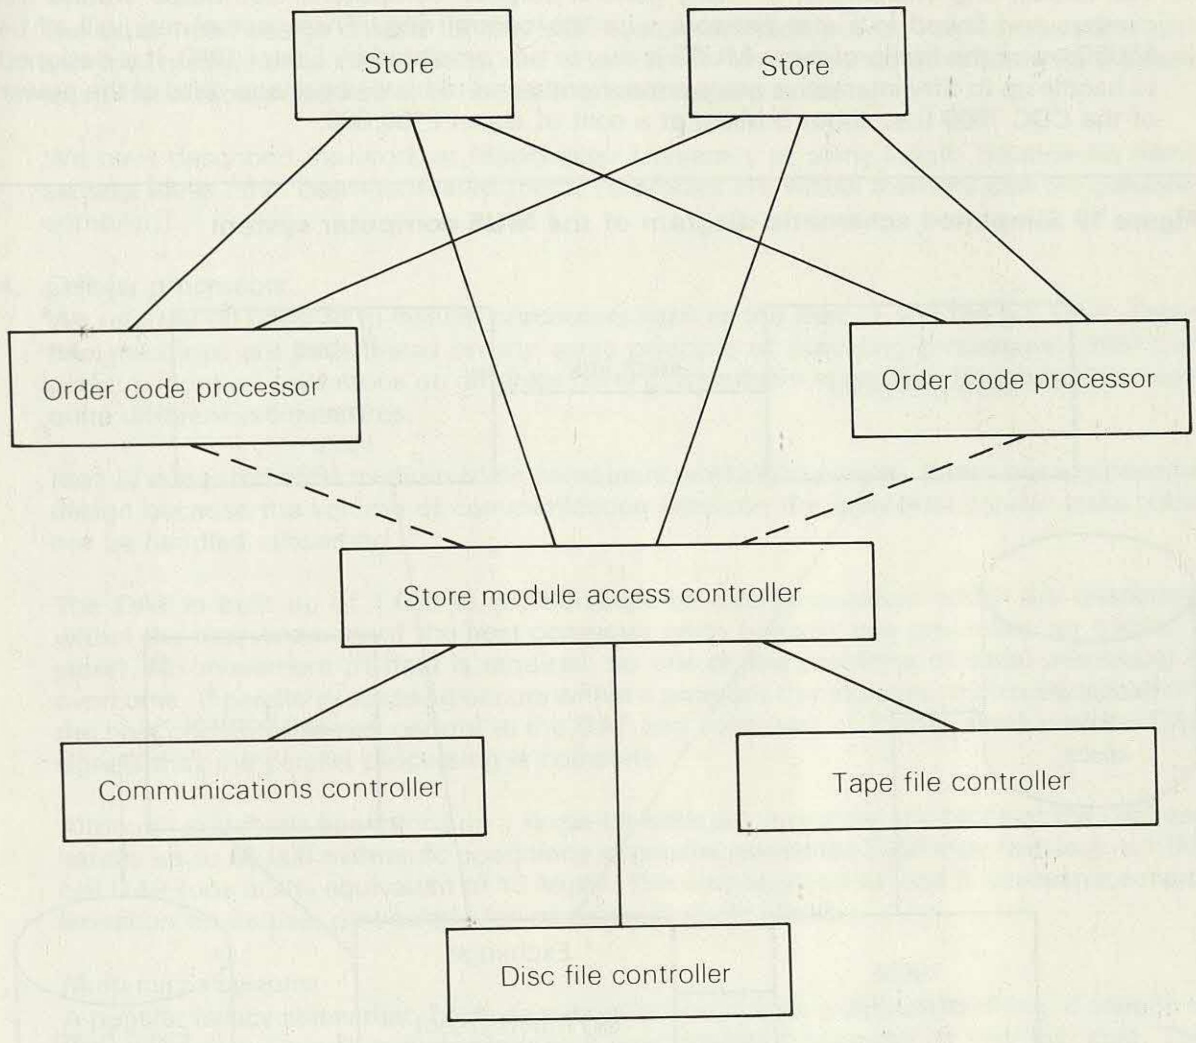

The whole system is controlled by a purpose-built operating system, MUSS, which supports multi-programming and the transfer of control of a program between processors, but does not support the partitioning of one program to enable the parts to be run in parallel on separate processors.

The university department is currently working on the successor to the MU5 system, which will be called the MU6. The eventual MU6 system will have three concentric rings of processors.

The central ring will be a multi-module computer, with a data-centred design. It will include a specially designed database processor, with a very high degree of parallelism, probably similar to the ICL DAP. Many different types of specialised processor will be included in the ring (e.g. dataflow processors, vector processors, cellular processors and specialised peripheral devices). All the processors will be interlinked with each other and also with the database processor via an exchange similar to the one in MU5. The function of the central ring will be to provide specialised processing power and to act as the principal filestore.

The middle ring will consist of many general-purpose computers, distributed around the campus and linked in a star network with the central ring. These computers will all be MU6Gs, and the first prototype MU6G is due to be completed by Easter 1980. It is designed to handle up to fifty interactive users concurrently and to have about one-third of the power of the CDC 7600 (i.e. about 3 Mips) at a cost of about £150,000.

Figure 19 Simplified schematic diagram of the MU5 computer system

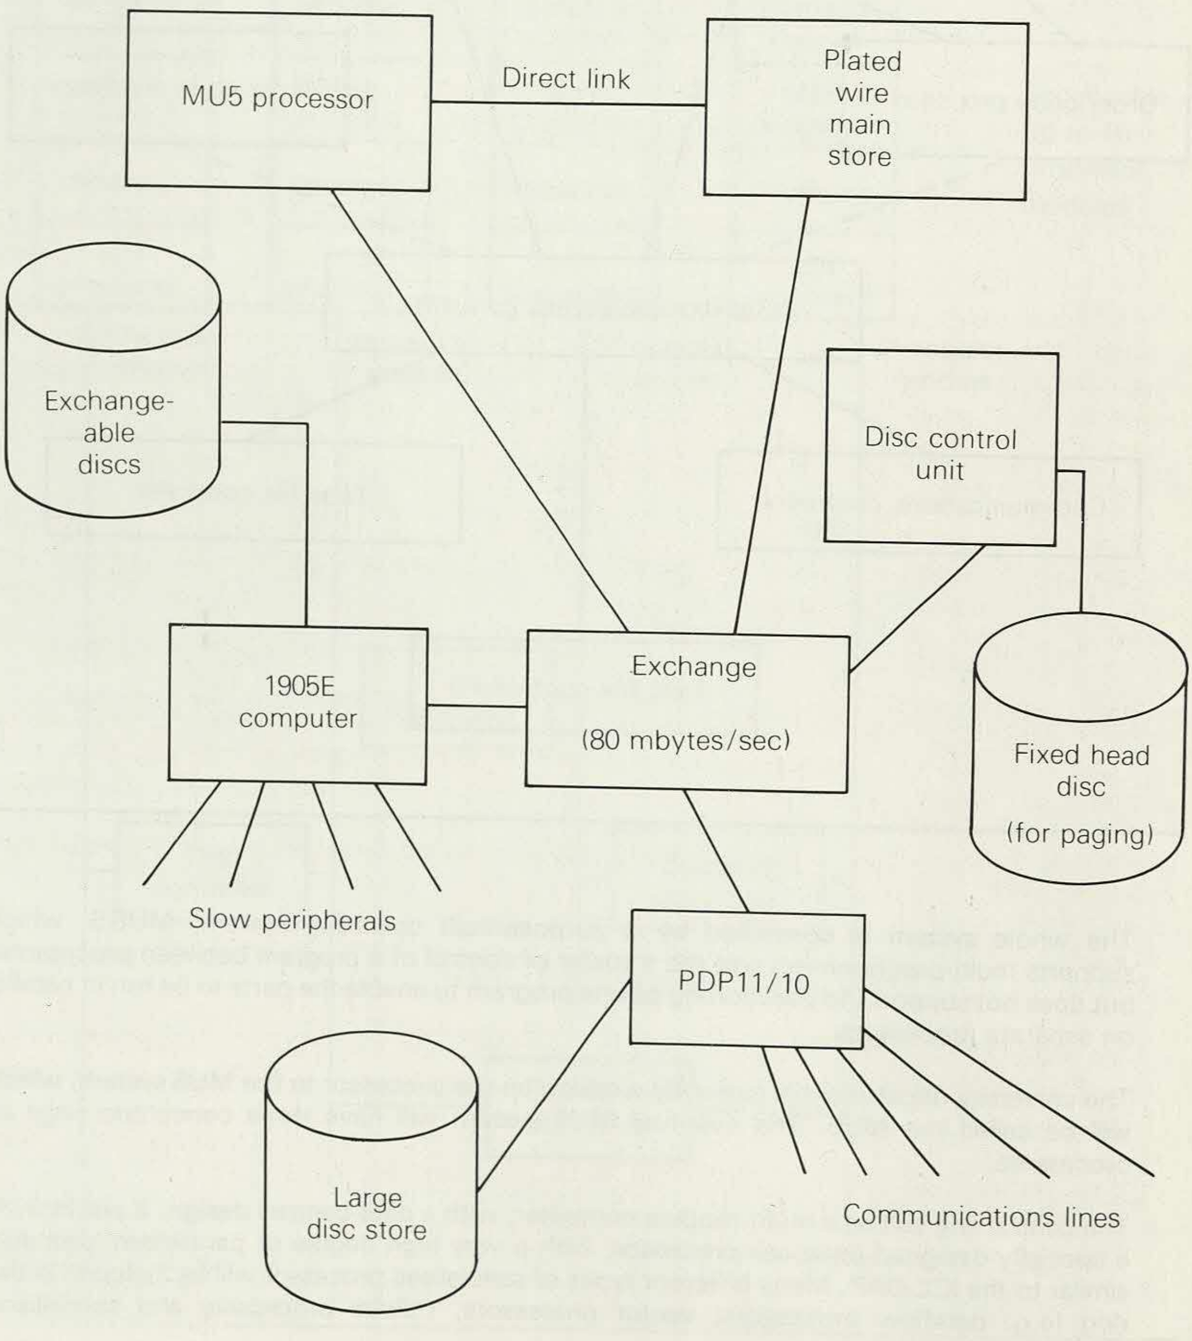

The outer ring of the MU6 will consist of a medley of small computers (such as minis, micros, personal computers, intelligent terminals and special-purpose microprocessors), in fact almost anything that is likely to be found in an academic environment. These small computers will link, whenever necessary, into the MU6Gs in the middle ring.

The three main design criteria for the MU6 system are high power, low construction cost and low maintenance cost. These are important in an academic environment, whereas reliability, although desirable, is not an essential requirement.

We have described the work at Manchester University at some length because so many seminal ideas have been generated there. (Examples are virtual memory and the compiler compiler.)

4. Cellular processors

   We referred on page 36 to cellular processors such as the Illiac IV and the ICL DAP. These two machines are both based on the same principle of providing processors (cells) that perform identical operations on different parts of a problem in parallel. However, they have quite different architectures.

   Illiac IV consisted of 64 medium-scale computers working in parallel. It was not a successful design because the volume of communication between the individual parallel tasks could not be handled efficiently.

   The DAP is built up of 1,000 to 65,000 single-bit microprocessors which are distributed within the main memory of the host computer (with typically one processor per 4 kbits of store). No movement of data is required, so one of the problems of serial processing is overcome. If parallel processing occurs within a program (for example, matrix manipulation) the host computer passes control to the DAP and continues with other work until the DAP signals that the parallel processing is complete.

   Although arithmetic operations on a single-bit basis are very slow, the fact that the DAP can handle up to 65,000 arithmetic operations in parallel makes the DAP very fast (e.g. a 1,000 cell DAP runs at the equivalent of 10 Mips). The time required to load it represents a major limitation on its use, particularly for commercial applications.

5. Multi-micro systems

   A popular fallacy states that, because microprocessors cost practically nothing, if enough of them are strung together they produce a very powerful computer at very low cost. This statement may be true in itself, but it falls far short of the whole truth. It omits the difficulty of expressing problems in a parallel way, and it omits also the difficulty of providing a suitable operating system. Nonetheless, several research projects have been carried out with multi-micro systems.

   Figure 20 overleaf shows the architecture of the PULSAR computer that Digital Equipment Corporation has developed using its own LSI-11 microprocessors.

   PULSAR has been designed to be software-compatible with the PDP-11 range. Access by the processors to the main memory is via a "P-bus" and is controlled by a bus controller. The processors run independently, except when they are using certain critical sections of the operating system. There can be contention for the operating system, and so the number of processors that can be added to the P-bus is limited in practice to 16.

   Figure 21 on page 43 shows the architecture of the DEMOS computer that is being developed jointly by the UK National Physical Laboratory and Scicon Limited. Each processor in the system has its own store and its own version of a simple operating system (called a

Figure 20 Architecture of the PULSAR computer

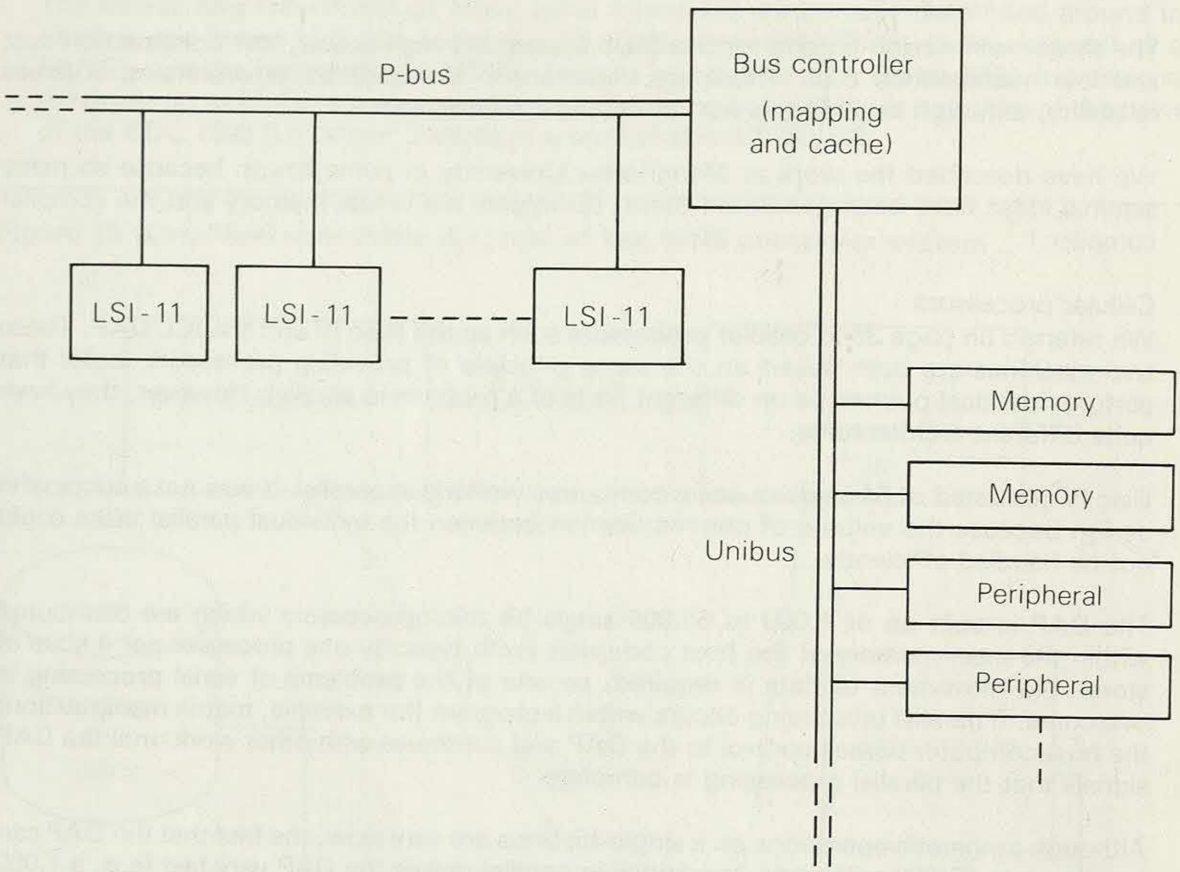

kernel), which handles communications between the processors. The processors are linked to standard ring ports, which are themselves linked into a ring bus. The ring bus has a transmission rate of about 5 million words per second and can handle up to 250 ports.

6. Database processor

   Content-addressable filestores are claimed to be particularly suitable for handling databases. In a large database, the data indices can occupy between two and four times the space of the raw data. A single data retrieval may require the execution of more than 1,000 instructions. However, a more complex retrieval, which requires searches of secondary indices or chaining, may increase the number of instructions ten-fold. One way of solving this problem is to partition the data into separate files, each containing related data, then to search through the files in parallel, and, finally, to correlate the retrieved data. This approach fits in well with the relational data model.

   ICL's Content Addressable Filestore (CAFS) is based on a conventional disc system. Data is read simultaneously from all heads (and thus from all surfaces), into an equal number of fast registers. The CAFS controller then searches the transferred data for values of interest. At present, there may be up to ten searches on data from ten heads, giving a maximum theoretical advantage over serial processing of 100. CAFS shows great promise for certain kinds of information retrieval work.

Figure 21 Architecture of the DEMOS computer

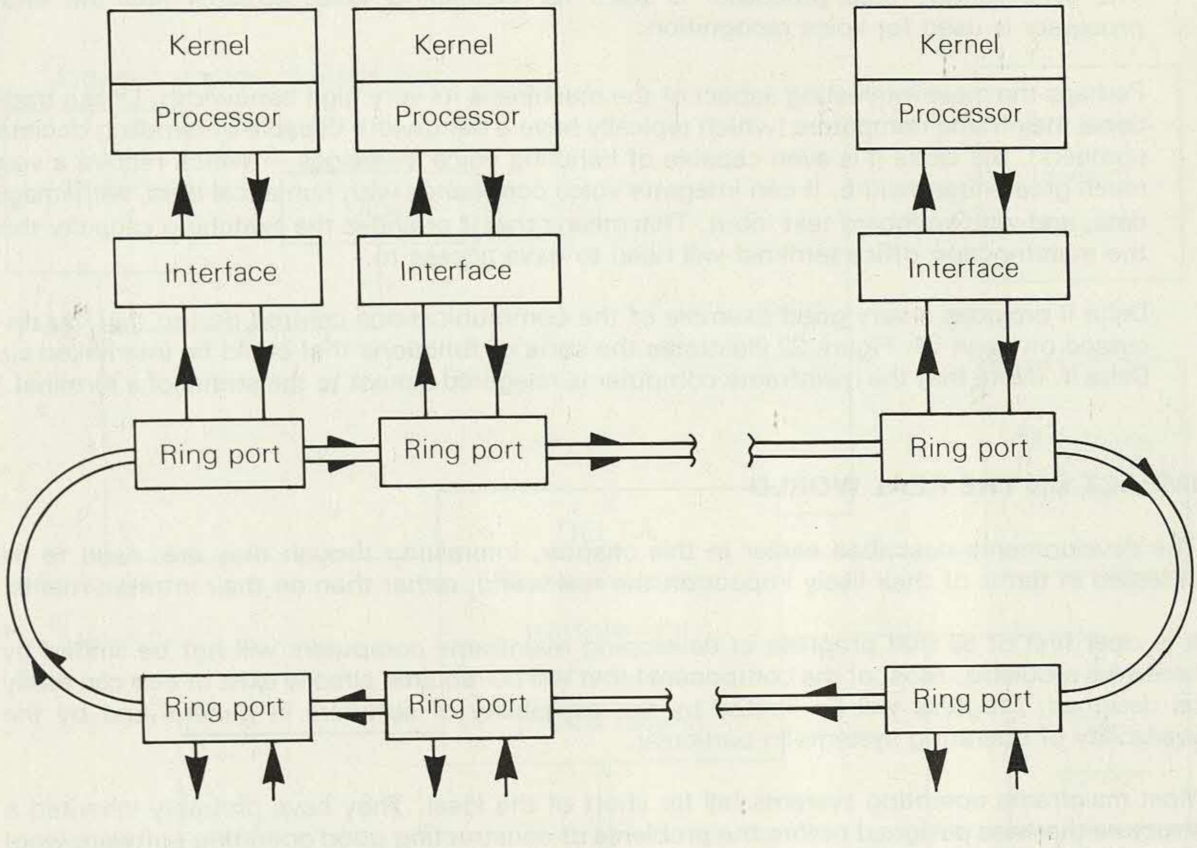

7. Delta II
   Considerable interest was aroused in the UK recently by the decision of the government-owned office products company, Nexos Limited, to buy the European marketing and manufacturing rights to the Delta II computer, developed by the Delphi Communications Corporation in the USA.

   The Delta II computer is designed for high reliability and great flexibility. The duplicated databus runs at 120 Mbits per second and is designed to handle up to thirty-two processors of sixteen different types, although only the following six types have been designed so far:

   - General-purpose processor.
   - Peripheral bus processor.
   - Synchronous bus processor.
   - Disc data controller (two types).
   - Synchronous data processor.
   - Array processor unit.

There can be between two and twenty-six general-purpose processors, which can be programmed in PASCAL, FORTRAN or BASIC. Up to four pairs of peripheral bus processors provide support for up to 4,000 asynchronous peripheral devices, and similar support for synchronous devices. The disc data controllers provide for up to nine 300 Mbyte disc drives. The synchronous data processor is used for examining voice streams and the array processor is used for voice recognition.

Perhaps the most interesting aspect of the machine is its very high bandwidth. Unlike traditional mainframe computers (which typically have a bandwidth capable of handling decimal numbers), the Delta II is even capable of handling voice messages — which require a very much greater bandwidth. It can integrate voice commands with numerical data, with image data, and with keyboard text input. This means that it provides the switching capacity that the multifunction office terminal will need to have access to.

Delta II provides a very good example of the communications-centred design that we discussed on page 34. Figure 22 illustrates the sorts of functions that could be interlinked via Delta II. (Note that the mainframe computer is relegated almost to the status of a terminal.)

IMPACT ON THE REAL WORLD

The developments described earlier in this chapter, interesting though they are, need to be assessed in terms of their likely impact on the real world, rather than on their intrinsic merits.

It is clear first of all that progress in developing mainframe computers will not be limited by hardware problems. Most of the components that will be required already exist or else can easily be designed. Progress will be limited by the availability of software in general and by the availability of operating systems in particular.

Most mainframe operating systems fall far short of the ideal. They have probably inherited a structure that was designed before the problems of constructing good operating software were properly understood. They may have been implemented by large teams of programmers who were inadequately managed, and consequently they generally have poor standards of quality and poor interface control. They are very large and typically contain two to three million lines of code. Perhaps worst of all, their original structure will have degenerated during their life, so that they are now over-complex, error-prone, and difficult to amend further. As an example, figure 23 on page 46 illustrates the increasing complexity of IBM's OS/360 structure.

The end result of all the above is that the mainframe designers now face a dilemma of choosing between two alternative courses of action. The first course is to develop their existing operating systems to control their newly-designed mainframes. This course would please their existing customers but it would leave the suppliers with even more complex operating systems. Their second choice is to develop new operating systems at considerable cost and face the wrath of the users who pioneer the new software.

This dilemma is not unique to mainframe suppliers. It faces all suppliers who wish to add a new product to their range. But the particular problem that the mainframe suppliers face is even more acute. They need to design an operating system that is capable of running on a multi-module mainframe, that provides forward compatibility for existing applications systems, and that is able to control the modules of a single program running in parallel on different processors.

The mainframe supplier has already made progress towards meeting the first two of these three requirements. Simple multi-module computers are already in use, and they do offer forwards compatibility. However, they are either multi-programming machines or machines that allow

Figure 22 Schematic diagram of the interconnecting of functions with the Delta II computer

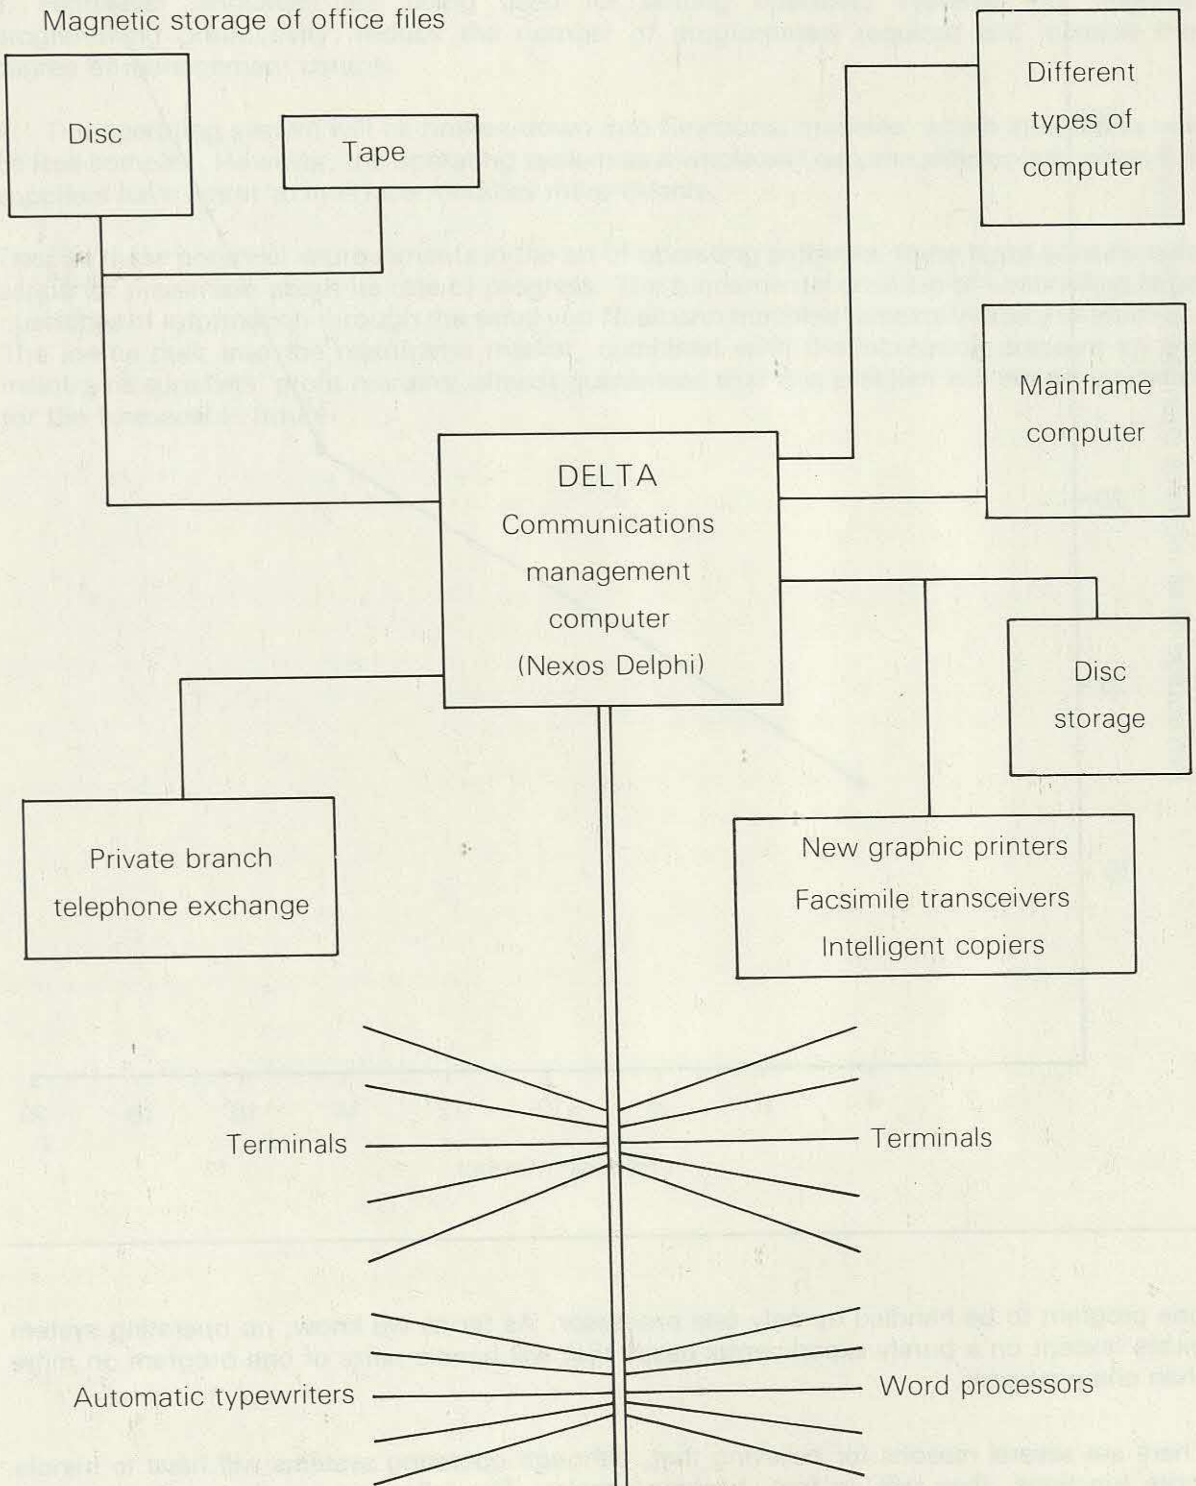

Figure 23 Increasing complexity of the IBM OS/360 structure

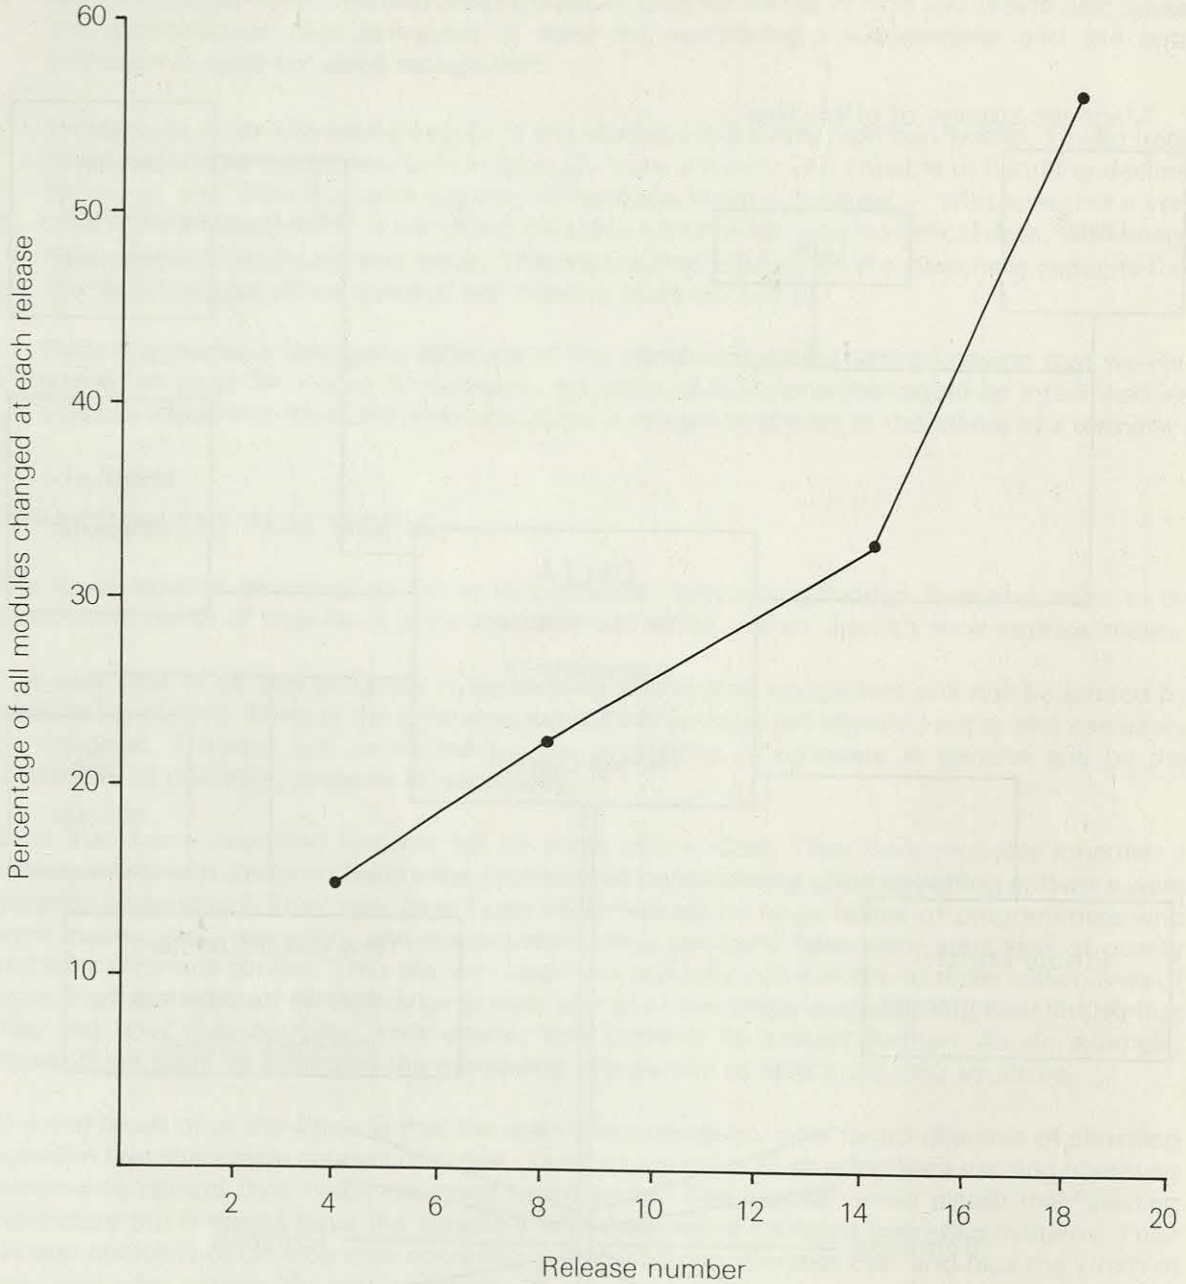

one program to be handled by only one processor. As far as we know, no operating system exists (except on a purely experimental basis) that will handle parts of one program on more than one processor.

There are several reasons for believing that, although operating systems will have to handle more functions, they will, in fact, become simpler. The following are the most important reasons:

1. The design of operating systems is now quite well understood. The only severe unsolved conceptual problem is knowing how to control parallel processing.

2. Some of the traditional functions, like peripheral device control, will be removed from the operating system and will be reprogrammed (probably in microcode) on the device controllers.

3. High-level languages are being used for writing operating systems that improve programming productivity, reduce the number of programmers required and increase the degree of management control.

4. The operating system will be broken down into functional modules, which individually will be less complex. However, the operating system as a whole will become simpler only when the suppliers have learnt to interface modules more cleanly.

Despite these potential improvements in the art of operating software, there is still considerable scope for pessimism about its rate of progress. The fundamental problem of channelling large quantities of information through the serial von Neumann machine remains virtually untouched. The inertia built into the mainframe market, combined with the increasing pressure on the mainframe suppliers' profit margins, almost guarantees that this problem will remain unsolved for the foreseeable future.

CHAPTER 7

CONCLUSIONS

DEVELOPMENT OF THE MAINFRAME

On pages 30 and 31 we put forward a design brief for the mainframe of the 1980s, which we summarise again in figure 24. As we pointed out then, this brief contains incompatibilities. Firstly, and not surprisingly, the interests of users and suppliers conflict. Secondly, the suppliers have to make some difficult choices between the respective priorities of manufacturing, design and marketing. Undoubtedly, these conflicts will handicap the development of the mainframe and its ability to meet users' needs in the future.

Other factors that constrain the advance of the mainframe are:

1. The complexity of the software
   Although software technology is advancing, progress is slow. There is no evidence that the usability of software is improving fast enough to make a real impact on the shortage of skilled staff, which threatens to become chronic.

   This deficiency of software technology handicaps the suppliers as well. Advanced operating software should be cheaper to develop, to maintain and to enhance. This may well be true of the latest software products, but, if it is, it is not yet feeding through in the form of any obvious benefits to the users. It is legitimate therefore to ask whether the complexity of software systems is either approaching or has already reached the limits beyond which any enhancement becomes self-defeating.

2. The suppliers' inheritance
   The suppliers' need to preserve their customers' massive investment in applications software presents them with a serious problem. If users are to be enabled to make significantly more effective use of the large quantities of processing power that are now available to them a radical new approach to operating system design will be required. It is difficult to see how suppliers could make such a leap forward without leaving many of their customers behind in the process.

   In addition, because all the suppliers (including IBM) are now experiencing reduced revenue from hardware sales, it is difficult for them to justify the massive investment that is necessary in order to design, build and launch a radically new product. In other words, the suppliers are likely to continue to build on top of their existing software base and thus perpetuate many of its present inadequacies.

We conclude from this that the 1980s will see a steady evolution of the mainframe computer with no spectacular changes in effectiveness or capabilities. The mainframe will, then, remain broadly what it is now — a numerical data processing computer that has powerful file handling capabilities and can support a varied workload, but makes heavy weather of communications.

This is not necessarily bad news for those organisations whose operations rely so heavily on the mainframe today. There is convincing evidence of a continuing demand for the batch processing at which the mainframe excels, with an increasing trend towards accessing on a remote job entry basis. Distributed computing is expected to add to, rather than to reduce, this batch workload. The database applications that are being implemented now can be expected to generate an increasing volume of enquiries in the future, particularly as query languages improve and also as users learn about their capabilities. Some applications (such as reservations systems) are by nature centralised, and so they are logical candidates for the mainframe approach. Very large mass storage will also tend to be associated with the mainframe, which will have the power and the software necessary to exploit it fully.

Figure 24 Design brief for the mainframe of the 1980s

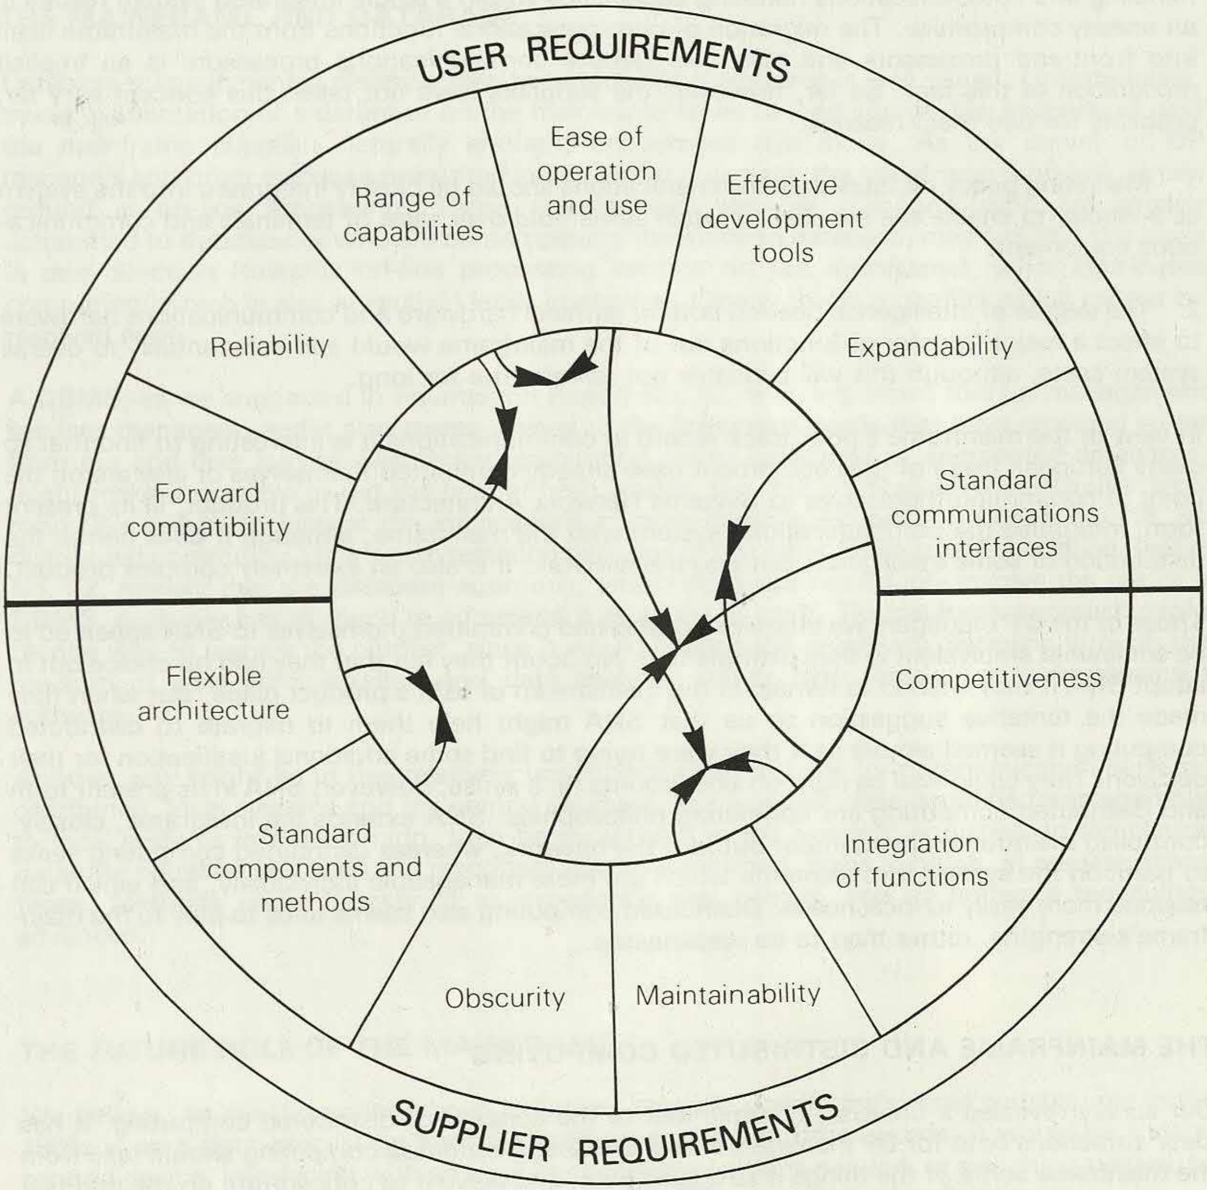

Opposing arrows illustrate some of the potential conflicts in the set of user and supplier requirements.

THE MAINFRAME AND COMMUNICATIONS

However, the issue which will concern many organisations is not so much the mainframe's potential workload (which they can gauge for themselves from existing trends) but the mainframe's correct relationship with other systems. A particular concern at present is the mainframe's relationship with the communications system. In this, we believe that the distinction hardware designers draw between data-centred and communications-centred designs is a significant one, and that these two different approaches to design do result in radically different pieces of equipment. This difference is demonstrated by the design of the Delta II discussed on pages 43 and 44. The relatively poor performance of the mainframe in the communications environment provides supporting evidence. It shows that the attempt to incorporate both data handling and communications handling capabilities within a single integrated system results in an uneasy compromise. The migration of communications functions from the mainframe itself into front-end processors and also into remote communications processors is an implicit recognition of this fact. So far, however, the suppliers have not taken this concept very far, probably for two main reasons:

1. Marketing policy dictates that communications should be closely integrated into the system as a whole, to enable the supplier to retain some hold over sales of terminals and communications equipment.

2. The degree of intelligence needed both in terminal hardware and communications hardware to effect a major transfer of functions out of the mainframe would add substantially to overall system costs, although this will probably not remain true for long.

In view of the mainframe's poor track record in communications it is interesting to find that so many European users of IBM equipment have already committed themselves or else are on the point of committing themselves to Systems Network Architecture. This product, in its present form, integrates the communications system with the mainframe, although it does permit the distribution of some intelligence out into the network. It is also an extremely complex product.

Those of the DP managers we interviewed who had committed themselves to SNA appeared to be somewhat ambivalent in their attitude to it. No doubt they felt that they had no choice but to adopt SNA if they wished to remain in the mainstream of IBM's product plans. But when they made the tentative suggestion to us that SNA might help them to migrate to distributed computing it seemed almost as if they were trying to find some additional justification for their decision. They could well be right on both counts. In a sense, however, SNA in its present form and distributed computing are competing philosophies. SNA extends the integrated, closely-controlled mainframe environment out into the network, whereas distributed computing seeks to partition the system into elements which are more manageable individually, and which can respond more easily to local needs. Distributed computing also seems to us to play to the mainframe's strengths, rather than to its weaknesses.

THE MAINFRAME AND DISTRIBUTED COMPUTING

Our survey revealed a universal commitment to the concept of distributed computing. It has clear attractions both for DP managers and end-users. Distributed computing should take from the mainframe some of the things it struggles with, and leave it to concentrate on the things it does well.

There are of course, some potential problems associated with distributed computing. In particular, there are potential problems of communication between the periphery and the centre, and there is little experience available to help solve those problems. Also, the technology of distributed database, which complements distributed computing, is, as we indicated in

Foundation Report No. 12, some way from realisation. What SNA, and indeed the network architectures of all the mainframe suppliers promise to do, is to solve these problems by allowing the necessary level of overall control to be exercised via the communications system, while also permitting computing power to be placed where it is needed.

The concept is an admirable one. The danger is that these networking products so increase the complexity of the operational environment that it becomes impossible to realise the benefits of distributed computing. In other words, the objective of a software solution intended to provide added flexibility could be defeated by the complexity of that solution.

THE MAINFRAME AND DATABASE

Database management is clearly a task for which the mainframe is well suited. Unfortunately, the implementation of a database on the mainframe tends to lead into on-line processing, and the mainframe supplier, naturally enough, encourages this move. As our survey of DP managers and other evidence presented in this report suggests, the mainframe is on less certain ground in on-line processing. Those management services managers who are already committed to database or who are contemplating the move to database, may feel that this leads in one direction (towards on-line processing centred on the mainframe), while distributed computing (which is also attractive) leads in another. Clearly this is a conflict which cannot be resolved easily.

A DBMS, as we suggested in Foundation Report No. 12, is an important tool for management services managers, and it also meets several of the important needs that have emerged so far from our end-user survey. Distributed computing, using replicated or segmented databases, might make it possible to get the best of both worlds, but this is still relatively unfamiliar territory. (We will be looking at the implications in more detail in Foundation Report No. 18, on distributed computing.) It is worth restating here one of the other findings of Foundation Report No. 12, namely that the database approach, which need not necessarily involve the use of a DBMS, probably has as much to commend it as a DBMS itself. The database approach might be one way to resolve the conflict, since it would impose the necessary discipline through the medium of data administration and data analysis, rather than through data management software.

Another way might be to use relatively unsophisticated methods for communication between distributed, local systems and the central database. For example, files could be transferred out to local systems in the morning, then be recaptured in the evening; enquiries on the central database could be batched, and responses could be returned, again in batch, at predetermined times. Methods could be refined as experience was gained, and as software technology advanced.

THE FUTURE ROLE OF THE MAINFRAME

We believe, as the discussion above indicates, that the mainframe's most suitable role in the 1980s is as a data processing and file handling system. In other words, it would be a data-centred device, interfacing with an intelligent communications network or switching system, to which would be attached the user organisation's other communicating devices.

We would not exclude the possibility that the mainframe might evolve into a system that was capable of taking on the role of a data- and communications-handling system. However, for the reasons we put forward earlier in this report, we doubt whether this dual role would provide the most satisfactory solution to the information processing needs of any but small organisations.

At present, most of those applications that occupy data communications systems are data processing based, and so it is natural to look first to the mainframe supplier for data communications products. When one looks further ahead, it is legitimate to question whether the integrated approach, which (with variations) the mainframe suppliers prefer, will remain appropriate when communicating devices within organisations multiply, and when new applications (such as electronic mail) develop and produce a more complex pattern of traffic.

As the market for communicating terminals grows, so additional and better-supported communications products will become available, and this development will enable users to build their own networks without a major commitment of resources and without serious risk. The new market that this development will create can be expected to bring in specialist suppliers such as Codex and Comten (marketed by ITT in Europe), the telecommunications companies such as Plessey and Siemens, and office products companies such as Xerox, Olivetti, Philips, Nexos and, indeed, the office products arm of IBM (in competition with IBM's data processing division). These companies will not have the same need to discourage their customers from attaching other companies' products to their networks as the mainframe suppliers do at present.

As this market develops, communications standards will become a diminishing problem, as the European PTTs, in particular, begin to exert the influence on the market that their investment and monopoly power gives them.

SUGGESTIONS FOR A STRATEGY

We recommend that data processing systems and communications systems should be separated functionally. Our reasons for making this recommendation are:

- To mitigate the deficiencies the mainframe has in a communications environment.
- To take advantage of the wide range of communications products which will become available over the next few years.
- To reduce the complexity of the systems environment.

We do not put this approach of functional separation forward either as a universal rule, or as an approach that can be achieved overnight. We put it forward rather as general guidance to be set alongside business requirements and existing commitments. We give more detailed suggestions below:

1. Those organisations that are already heavily committed to on-line applications will naturally place high priority on the need to preserve their investment in the applications concerned. Business requirements will therefore dictate the future course of action for those organisations.

This means that there is probably little scope for them to deviate from their present approach — a data communications system based either on the mainframe supplier's software or (for some major users like airlines and banks) on purpose-built networks, some of which are already independent of the mainframe.

2. For those organisations that have no such commitment we suggest that the level of integration between data processing and communications systems should be kept as low as the applications permit. Distributed computing systems using replicated files are a good example of this approach, in that they enable batch-oriented techniques to be used for communication.

3. We have expressed our reservations about mainframe-based network architectures as they

stand at present. These products will evolve over the next few years. As the evolution takes place, users will be better able to judge how much freedom those products permit in the locating of equipment and how well they are able to cope with those more-flexible traffic patterns that are likely to be characteristic of office systems. If they do not measure up to requirements in this respect, organisations may need to look elsewhere for communications products that suit their requirements.

4. DP managers may be concerned that they will be left with unused mainframe capacity if they change their plans for on-line applications, and they may, of course, not wish to change their plans for a number of valid reasons. We believe that few DP managers are likely to be troubled with problems of over-capacity for long, but, even so, they might consider whether cheap mainframe power could be better used either to improve fallback provisions for existing systems, or to give programmers better system development facilities, or to reduce dependence on shift working at the central site.

5. Batch systems are not necessarily dead, or even moribund. We found clear evidence of their value to users, and, in some cases, this value might be enhanced by on-line facilities. In many cases, it is probable that changes in procedure or even simple public relations work might improve their perceived value to users.

In summary, we feel that the future of the mainframe is at least as much in the hands of management services and DP managers as it is in the hands of suppliers. Clearly, the mainframe does have a role to play in applications of a certain type. Clearly, it also has limitations. We believe that those limitations are not likely to be removed easily or quickly. The mainframe is no longer the only choice for data processing applications, let alone for new office applications.

Management services managers should be aware of the choices open to them. They will need to look backwards at their existing investment in applications software, as well as forwards to new opportunities to exploit computing and communications technology. The mainframe dominates the backwards perspective and this means that it will also figure in the forward picture. It will be up to management services managers to ensure that the mainframe takes its rightful place amongst the expanding range of alternatives.

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

The Butler Cox Foundation

Report Series No 16 The role of the mainframe computer in the 1980s

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

Completed by ....

Job title ....

Organisation ....

Address ....

....

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

The Butler Cox Foundation

Report Series No 16 The role of the mainframe computer in the 1980s

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

In order to take account of your comments we need to receive them from you by 17 January 1980.

Completed by ....

Job title ....

Organisation ....

Address ....

....

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

The Butler Cox Foundation

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

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

Akzo Systems B.V.
Velperweg 76, Arnhem, The Netherlands
Tel 85-662629

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

Printed in England