Presentation Summaries

BUTLER COX FOUNDATION

Coast to Coast

Study Tour 1989

BUTLER COX
FOUNDATION

1989 Study Tour of the United States

Introduction

This document has been produced as an aide-memoire for the participants in the Butler Cox Foundation 1989 Study Tour. Delegates visited supplier and user organisations in Massachusetts and California. Each section records the details of the presentations made during the tour.

The information presented here is essentially that provided by the host organisations at the time of the visits. Care has been taken to reflect this information as faithfully as possible, although working from spoken presentations, and without a full transcript, neither completeness nor total accuracy can be guaranteed. Many of the host organisations provided Butler Cox with copies of the visual aids used during the presentations, and a selection of these has been included where appropriate.

We should again like to record our thanks to all of the organisations and individual people we visited. We believe that these notes will help the delegates to derive the fullest benefit from a very informative and highly intensive tour.

Published by Butler Cox plc
Butler Cox House
12 Bloomsbury Square
London WC1A 2LL
England

Photoset and printed in Great Britain by Flexiprint Ltd., Lancing, Sussex.

BUTLER COX
FOUNDATION

1989 Study Tour of the United States

Contents

Wednesday 10 May	Wang	1
	Open Software Foundation	5
Thursday 11 May	Babson College	8
Friday 12 May	Bachman Information Systems	22
	Thinking Machines	26
Monday 15 May	Northern Telecom	29
	Xerox Palo Alto Research Center	33
Tuesday 16 May	Hewlett-Packard	35
	Sun Microsystems	39
Wednesday 17 May	McDonnell Douglas	43
	Intel	46
Thursday 18 May	Speech Systems	49
	Pixar	53
Friday 19 May	Relational Technology	55
	Natural Language	58
	Sybase	61
List of delegates		63

Wednesday 10 May

Wang

Tim Sloane, Director, Consultant Relations, welcomed Study Tour delegates and introduced the speakers and the program for the visit. The speakers were:

- Ken Olisa, Senior Vice President, Worldwide Marketing.
- Roger Sullivan, Director, Wang Integrated Office Systems Marketing.
- Steve Lavine, Director, User Interfaces and Artificial Intelligence.
- Tom Haywood, Planning Manager.
- Ken Russo, Operations Manager, MIS.

Wang's vision of the future

Mr Olisa described the background to Wang's vision of the future by demonstrating how computer technology has been introduced into organisations to date. He identified five stages:

- The isolated mainframe era.
- The era of online terminals linked to mainframes and the integration of mainframes.
- The introduction of minicomputers around mainframes in order to satisfy the departmental needs of businesses.
- The introduction of microcomputers around the minicomputer network. (IBM has sold approximately 26 million PCs, and yet systems managers are still likely to ignore the technology.)
- The era of local area networks providing gateways to mainframes.

Mr Olisa described the present situation, where many systems managers have lost control and where technical infrastructures are increasingly fragile because they have not been well designed. The case used to cost-justify the introduction of technology in many organisations is often shallow, and the opportunity to exploit technology is often lost. The job of the systems manager should be to manage *all* of an organisation's information. This information may be in the form of data, text, voice, or image. Wang's vision of the future is one in which all this information coexists, is integrated, and is actively managed and exploited.

At present, use of facsimile equipment is increasing rapidly. It is being introduced into organisations in much the same ways as PCs were introduced. Wang sees facsimile technology as an integral part of image processing and, as such, has integrated it into the Wang Integrated Image System (WIIS). Problems with the use of computer technology are perceived as a potential limitation, and consequently, Wang is actively developing user interfaces to WIIS that will overcome this problem. To this end, Freestyle (described later) has recently been introduced.

The Wang Integrated Image System (WIIS)

In 1988, Wang carried out a survey of a large US-based insurance organisation. Mr Sullivan said that the survey showed that only 10 per cent of the company's active information was being processed and managed with any form of electronic technology. The other 90 per cent was managed by the use of paper, or was held on microfilm, as shown in Figure 1, overleaf. The fact that, in many organisations, the whole systems budget is being spent managing just 10 per cent of the total information, shows the potential for systems that can include paper-based information. Because much of the information recorded on paper or microfiche is handwritten, existing computer systems are not able to deal with it. Image processing systems are being designed to fill this gap.

Figure 1 The impact of IT on information management

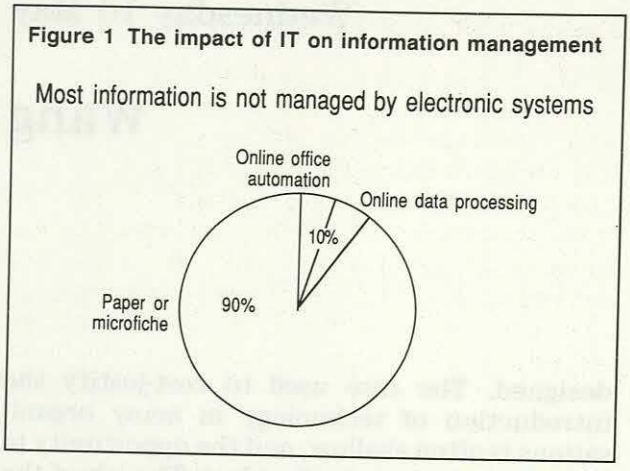

Currently, packaged image processing systems account for only 5 per cent of the market. Most image applications are being developed in-house. This is expected to change in much the same way as the market for payroll packages changed. The main purchasers of image systems are end users; systems departments often simply react to events, and try to integrate image systems with other corporate systems. Mr Sullivan believes that the majority of image processing systems will ultimately need to be integrated with corporate IT facilities, and that systems managers therefore need to become more active in managing the technology.

Potential areas of application

WIIS was introduced in 1986 and, to date, 250 systems have been sold, 50 per cent of which have been installed. The breakdown of sales by industry sector is shown in Figure 2. Mr Sullivan listed some areas in which image systems might potentially be applied:

Figure 2 Sales of WIIS, by sector

Sector	*Per cent*
Government	49
Banking	8
Insurance	8
Manufacturing	13
Services	14
Other	8

- The effectiveness of many *government* bodies depends on their ability to manage and make available large volumes of documents. A particularly good example of this is the management of records such as birth, marriage, and death certificates.
- *Banks* increasingly compete on a global basis, but have little opportunity to differentiate their products except through customer-service levels. The ability to track correspondence and to route it through the organisation improves responsiveness to customer needs and to customer enquiries. Image-based processing is becoming a powerful weapon in the search for competitive advantage in this area.
- Both in claims processing and in underwriting, the *insurance* industry is dependent on the efficient processing of large volumes of paperwork. Traditionally, this has meant that clerical staff effectively control the business, retrieving and processing documents. Image systems are making it possible for insurance companies to redefine the flow of work through the organisation, and to improve efficiency and customer service.
- The advent of just-in-time systems has meant that *manufacturing* companies are able to process a greater volume of paperwork in support of smaller batch sizes. Much of the transaction-based paperwork associated with purchasing, receiving goods, manufacturing, and invoicing can be administered by image systems.

The components of WIIS

An image processing system consists of various combinations of the following components:

- Scanner: It is possible to scan up to 200 pages per minute, on double-sided paper, if required.
- OCR indexing: Documents can be indexed using the information from a fixed position on the page, such as a purchase-order number. Similarly, different forms can be associated by using common identification data from the forms, as in the case of multiple-page tax returns.
- Optional disc drives: WORM technology is used. Image-compression algorithms allow

20,000 images per side to be stored. Current discs are 12″, but are soon likely to be reduced to 5.5″.

- Facsimile: Facsimile can be used to input or output images to or from WIIS.
- Jukebox: A jukebox allows access to multiple optical discs and drives, with each jukebox providing for 3.5 million images. By adding multiple jukeboxes, storage capacity can be increased to over 100 million images.
- Wang VS products: The heart of WIIS is the Wang VS processor and software. Multiple VS machines can be integrated using local area networks, and integration with mainframes can be provided.
- Screens: Wang provides high-resolution monitors, but alternatively, PC ATs or clones equipped with Hercules boards can be used. Colour is not yet supported.
- Applications development: Applications can be developed by using standard languages such as Cobol, Fortran, Oracle, and Focus.

Reliability and performance

Reliability levels have been very high, although the scanner works on the same principle as a photocopier, and therefore needs to be operated with care. Performance levels currently achievable provide access to an image in under 30 seconds. If the technology is to be used to support telephone enquiries, careful attention will need to be paid to staff training to ensure that the wait period is used constructively and does not lead to customer dissatisfaction. Response times are likely to improve as repository technology develops and as communications technology improves.

Payback

Experience to date has shown that savings in floor space resulting from the reduction in paper storage facilities, and productivity improvements arising from the redesign of workflows, lead to payback periods of about 12 or 13 months. Other unquantified, but probably more significant, advantages arise from improved levels of customer service. The applications with the highest returns are those that involve responding to random queries, or that require access to random information in the organisation.

The development of Freestyle

Mr Lavine described the Freestyle product, which he had been responsible for developing as a front end to WIIS. He drew a comparison between using a pencil and using a computer. The pencil is the ideal communications device. Once a person learns to write, the facility is easy to use. This contrasts with the computer keyboard, which is easy to learn, but complex to use because of the need to interact, via software, with the machine. For many computer users, the ad hoc nature of computer usage causes further difficulties.

Freestyle, which consists of a digitiser pad and an electronic pencil, has been developed to overcome these problems and to make WIIS easy to use. By using the pencil, it is possible to 'write' an electronic document. The image of whatever is 'written' on the digitiser pad is immediately available on the screen. The document can also be erased, using the 'eraser' end of the electronic pencil.

Voice messages can also be associated with, and stored with, an image. Any image can be recalled or routed elsewhere on the system, along with handwritten notes (entered via the digitiser), and verbal notes (entered via a telephone). Via a window/icon facility, and using the pencil to point to the required facility, a user can operate the system without a keyboard, although a small keyboard device is available.

Freestyle was demonstrated running on a 286-based PC equipped with 640k of memory and a Hercules graphics card. Response was immediate via the digitiser, and voice playback was of very high quality.

The use of WIIS in Wang

The contracts maintenance organisation (CMO) within Wang represents a business of $650 million turnover, and its administration has traditionally been based on the batching and clerical processing of paperwork. It is responsible for contracts maintenance and billing. In 1986, a decision was taken to use image processing technology in order to improve productivity and increase customer satisfaction. The implementation process has been phased, and during this process, more users have been brought onto the system.

The largest part of the current application is the call-charge system. This involves the processing of engineers' activity reports, arising from maintenance work that is to be charged on a time-and-material basis. Activity reports show the details of the work, which is used as the basis for billing. Some 125,000 reports per month are processed. The major benefit of using an image-based system is that CMO staff are now able to respond to customer queries rapidly and efficiently, and are also in a better position to manage the outstanding debt. The payback period on the original investment was nine months. Previously, many disputes were unresolved because of the difficulties associated with retrieving information and confirming signatories on work-completion documents.

Image processing is also used for archiving invoices. This is an integral part of the contracts-management process. CMO has provided wide access to this information, and made it possible to produce duplicate copies, or to fax a copy to a customer, automatically.

Demonstration

During the lunch period, delegates were able to see demonstrations of Freestyle and to watch a video illustrating users' experiences of implementing image processing systems. The demonstration showed how easy Freestyle is to use. The technology is new (it was introduced early in 1989) but Wang is committed to the principle of making it easier to access WIIS.

User experience confirmed that the pioneering days with image processing are over. Several significant applications were described, including those of British Airways in the United Kingdom, AAA in the United States, and government applications in France and the Netherlands.

Open Software Foundation

Presentations on the Open Software Foundation (OSF) were given by David Tory, President and Chief Executive Officer, Alex Morrow, Vice-President, Strategy, and Donal O'Shea, Vice President, Operations.

OSF was founded in May 1988 by a group of IT suppliers who were concerned about barriers to growth in the Unix market. They realised that they could no longer afford to support their own proprietary versions of the Unix operating system. This was too expensive, and there was a need for standard industry products. At that time, the Unix market was very fragmented, with around 200 versions of Unix available. They were all basically similar, but they were incompatible in terms of their functionality. The result was that independent software vendors were not prepared to develop application software for the Unix environment, and this was thought to be limiting growth in the market. AT&T, the Unix licence owner, had effectively divided the market in order to rule over it, so the OSF sponsors decided to combine in order to lower their costs of supporting the Unix operating system, and thereby enlarge the total Unix market. OSF would also provide a very good market research capability for them.

The sponsors are Apollo, Bull, Digital Equipment, Hewlett-Packard, Hitachi, IBM, Nixdorf, Philips, and Siemens. They agreed to provide funding of $122 million on the basis of $4.5 million each, for three years. OSF has three kinds of organisation as partners. These are end users of operating software, independent software vendors who wish to develop application software for the market, and the hardware vendors themselves.

Commercial organisations pay an annual subscription of $25,000, not-for-profit organisations pay $5,000 per annum, and university departments pay $2,000 per annum. There are currently 94 members. OSF believes that this wide membership gives it a balanced view of the industry and of market needs. The OSF management team is made up of a chairman and chief executive, five functional vice-presidents, and a general counsel. It aims to employ around 220 staff by the end of 1989 — 60 per cent in software engineering, and the remainder in research and development. Currently, it employs 125 people. Its operations are located in the east and west of the United States, in Europe, in Japan, and shortly in China.

OSF is a not-for-profit research and development organisation, and its charter is the development and licensing of open software for the benefit of the whole information-processing industry. It aims to be self-funding by 1992. Its seven guiding principles are:

- Its offerings are based on relevant industry standards.
- It uses an open process to solicit inputs and technology.
- It applies a timely, vendor-neutral decision process.
- It provides early and equal access to specifications, and continuing development for all its members.
- Its implementations are hardware-independent.
- It offers fast, reasonable, and stable licensing terms.
- It maintains awareness through research and university participation.

It intends to provide offerings of three types:

- Software products consisting of specifications, architecture-neutral code, and validation suites.

- General services, such as publications, education, and consultancy.
- Member services, such as meetings and 'porting' laboratories.

OSF's product strategy is to provide open software that will act as a buffer between three groups. Between end users and the independent software vendors, it will provide an independent user interface, and between the independent software vendors and the hardware vendors, it will provide an applications environment — in particular, a standard version of the Unix operating system.

OSF will develop its product offerings, first, by issuing an open invitation to any organisation to submit sample products. It will then carry out an open review-and-development process, which will be controlled by the OSF members. It will make its product offerings available to anyone, on payment of a standard licence fee.

Its first offering is OSF/Motif, which is a graphical user interface. The request for technology (RFT) was published in July 1988, and it expects the first sample product to be available in July 1989. It chose the graphical user interface as its first product partly because there is a pressing need for a good standard, and partly because the user interface typically accounts for 30 per cent of operating-system code. It is an important competitive area and OSF expects to issue hundred of thousands of OSF/Motif licences. IBM itself has welcomed the development of Motif and intends to use it as its strategic user interface for its Unix products. The main inputs from which Motif is being developed came from Digital Equipment, Microsoft, and Hewlett-Packard. OSF expects that not only will Motif be used in a Unix environment, but that it will also be transferred to non-Unix environments, such as OS/2 and VMS.

OSF's second product will be a standard kernel for the Unix operating system, to be called OSF/1. This will be based on IBM's AIX/3, which is a completely new version developed by IBM and is very different from its earlier AIX/2. However, OSF intends to augment the AIX/3 specification. The key difference between OSF/1 and AIX/3 is that it will be more portable — in fact, OSF has incorporated some ideas drawn from the Carnegie Mellon University MACH operating system. It will also be more modular, efficient, extensible, and maintainable.

OSF has proposed a development partnership programme with its members, with the objective of reducing the time required to bring a new product to market, by providing members with early access to the technology in the form of a vendor kit, an applications kit, an untested university version, and a fully-tested commercial version. OSF will also set up a 'portability' laboratory in which vendor members can locate their own software development team to work closely with OSF staff to develop the portability features of any new software product.

Another OSF initiative is the 'architecturally neutral distribution facility' (ANDF). OSF is trying to develop a scheme under which an application can be developed once, but be distributed across many different architectures. In other words, it wants a single format for software distribution that is hardware-independent — what it describes as 'shrink-wrapped' software. Instead of compiling its applications right down to operating machine code, OSF would have an intermediate step that would be distributed, possibly to many different environments, and then be finally compiled by a local translator on the target hardware.

The next RFT that OSF intends to issue will be in the field of inter-operability. At present, there are many proprietary approaches to this problem, but there is no single industry standard. OSF has set up three special-interest groups. One is in basic communications, one is in the distributed-computing environment, and one is in network management and systems administration. These working groups have been looking at what needs to be done to achieve inter-operability, and at how a single software environment can be created. OSF expects to issue the very complicated RFT in July 1989.

In reply to a question about security, OSF said that OSF/1 would have a much higher level of security than any current versions of Unix, but that it would certainly not satisfy central government or banking requirements. These would need to be provided through hardware rather than software.

In reply to a second question, about prominent Unix suppliers who are not members of OSF, such as Unisys, NCR, and Apple, OSF said that these companies had not joined originally, either because they were suspicious of IBM's motives in sponsoring OSF, or in Apple's case, because they believed that they already had a significant advantage in the market. However, OSF did say that it was in very close contact with these companies and expected them to join at some convenient later time.

The Foundation delegates were impressed by the objectives and methods of OSF, but some expressed scepticism about the motives of the OSF sponsors.

Thursday 11 May

Babson College

Study Tour delegates attended a one-day conference at the Center for Information Management Studies (CIMS) at Babson College, entitled 'Information Systems in Transition'. CIMS is a consortium of major companies in the New England area that have joined with Babson College to find new ways of improving the use of information systems. The conference was also attended by US systems executives from organisations that are sponsors of the CIMS programme.

The conference agenda was designed to provide insights into the 'winning technologies' of the 1990s and into ways of managing them. Summaries of the presentations made at the conference are set out below.

Megatrends for the 1990s

The scene-setting session was given by Jerry Kanter, Executive Director of CIMS. He suggested that there were ten megatrends that were important for business people who wish to maintain effective systems operations. He briefly described these as follows:

More power in less 'real estate'
Mr Kanter began by saying that technological developments, like the Intel 486 chip with 1.2 million transistors, was continuing the trend towards more power in less space. Other powerful trends are CASE tools, optical-disc storage, and optical fibre. Although, to some extent, these were changes in degree, the time was approaching when such changes in degree would become changes in kind.

The medium influences the message
The means by which communications are transmitted are increasingly affecting the message itself. Mr Kanter quoted optical fibre and EDI as examples.

The office-less office
Although the need for paper is unlikely ever to disappear, technology offers the opportunity to remove much paperwork, and even offices, by allowing people to work away from city centres, in their homes. The technology must, however, become more user-friendly in new ways before such a vision can be realised.

'The shoemaker's children shall not go barefoot'
Mr Kanter pointed out the growing importance of CASE. There is continuing pressure to cut costs in the systems environment, and since about 60 per cent of a systems department's budget is spent on people, CASE can have a tremendous impact. Systems people are, however, reluctant to use their own tools.

Data: an inexhaustible resource
The United States stores vast quantities of data in both public and private research centres. Data, or at least references to it, can be captured and used as part of a public or private network to provide a powerful resource.

Cluster management and 'groupware'
The cluster-management concept is different from the classic hierarchically structured organisational approach. Essentially, it means that problems are solved by small groups of people, selected on the basis of their ability to contribute, and not on the basis of where they feature in the organisational hierarchy. These people do not necessarily have to share the same geographic location. 'Groupware' is a generic name for the technology (such as local area networks, wide area networks, and video-conferencing) that links members of the cluster.

Information literacy
The emphasis in the future must be on information literacy, even though, in the past,

it has been easier to obtain funding for hardware expenditure than for training the people who are to use it.

Parallel processing, neutral networks, and expert systems

Powerful parallel processors, like those installed at the Dow Jones news retrieval service, will herald a new era of information access and processing. Neutral networks are designed to emulate the human thinking process. Expert-system applications are likely to be commonplace in the 1990s.

Strategic vision means losing sight of the forest and the trees— enter pragmatic planning

A more pragmatic view of planning is emerging as management is becoming weary of long-range strategic planning. The emphasis will be on a more pragmatic approach to systems planning, even though this may mean that it will have to be more short-term.

The redoubtable human factor

Mr Kanter said that although we talk a lot about the human factor in technology, we have not really understood it. Executives will not, as a rule, wish to spend much time learning to use the technology.

The strategic use of information systems

Chris Demos, Senior Business Advisor to the IT division of Federal Express, began by emphasising that business competition is now global. Companies from a multitude of nations are all in the process of deciding what factors will give them competitive advantage.

He said that, unquestionably, IT provides Federal Express with its competitive edge; it is what sets Federal Express apart from its competitors. Federal Express is part of a $10 billion air-express delivery industry. Its own share is over $4 billion. The company also knows, however, that it must continue to innovate and enhance its systems to maintain its competitive position. To meet the challenge, Federal Express has created an extensive worldwide air/ground network of over 200 aircraft, over 22,000 vehicles, and more than 63,000 people. These numbers demand skilled management, both in human and technological terms.

The ability to manage information is critical for productivity in Federal Express. The company considers it to be so critical that it devotes almost five per cent of its revenue (that is, about $200 million) to information and telecommunications systems. Pouring money into technology however, is not enough. Before deciding to invest, Federal Express asks itself whether the technology will support its employees, and whether it will serve its customers.

Mr Demos emphasised the importance of corporate culture in achieving Federal Express's objective. The basic corporate philosophy is 'people, service, and profit'. The company believes that people are the basic ingredient for innovation, and that it is people who make information systems strategic. It is people who serve customers better and increase productivity.

He then went on to describe the software environments in which Federal Express operates, and the supporting computer operation and telecommunications network. Federal Express's information systems fall into four categories:

- Software environments.
- Commercial.
- Air/ground operations.
- Customer service.

Of the commercial systems, the human resources information system, PRISM, is worth noting. Through PRISM, each manager can access accurate, timely information for each employee, from any Federal Express terminal.

The COMETS system supports the operational side of the business. It is an airline software system. COMETS gives operating units near realtime updates of traffic demands, weather, equipment status, and any operational data necessary to meet Federal Express's service commitments. The company's service goal is 100 per cent on-time deliveries, which it meets between 98 and 99 per cent of the time. It is a huge operation, with over 1,300 facilities worldwide.

The SUPERSYSTEM is a ground system that incorporates all information about every flight. Information includes the mechanical status of

the plane, the gate for each flight, the fuel requirement, load data, the names of crew, and so on. The plan is to extend the system to incorporate graphic displays, tracking each flight's progress on route maps, and determining weather information. Modelling techniques and artificial intelligence should improve night-time operation.

COSMOS is the customer-service system, and all deliveries are tied in, one way or another, to the COSMOS network. With over 65,000 online terminals, it is one of the largest interactive systems in the world. It contains all the basic customer information, including names, account numbers, addresses, and pick-up locations. It also interacts with several other systems and devices to maintain complete records of each shipment. When a courier picks up a package, he uses a hand-held computer, called Super-tracker. The location of each package is continually updated, and it should never change hands without the system being updated. The result is that any Federal Express agent around the world can ask COSMOS about a package or document and tell the customer the exact time it was delivered and who signed for it.

To keep up with the growing demand for the movement of information, Federal Express's voice and data networks have been migrated to a digital network. A satellite-based system has also been developed and, together with other communication developments, an integrated network, called I-NET, has been created.

All this, Mr Demos reiterated, has been achieved by creating an organisational environment that fosters autonomy and risk-taking and tolerates failure — what he called a 'power environment', where people have more power and are more willing to take risks, and where technology therefore thrives.

The future for strategic or competitive-advantage systems

John Cunningham, of Competitive Technologies, Inc., began by quoting the results of a survey in which CIOs and CEOs were questioned about the significance of the strategic use of IT. For both types of respondent, the pattern had changed between 1985 and 1989 (see Figures 1, 2, and 3).

Figure 1 A CIO priority

Strategic use of IT is very significant

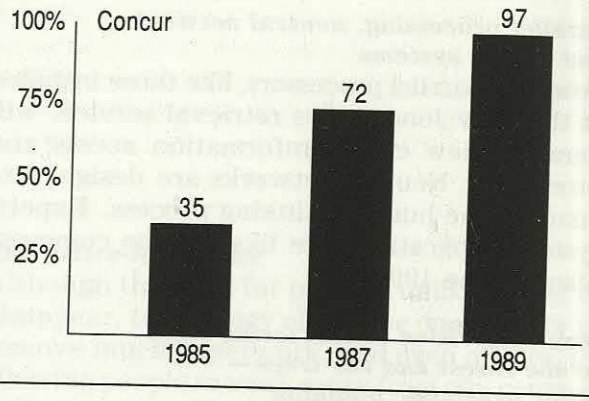

Figure 2 A growing CEO awareness

Strategic use of IT is very significant

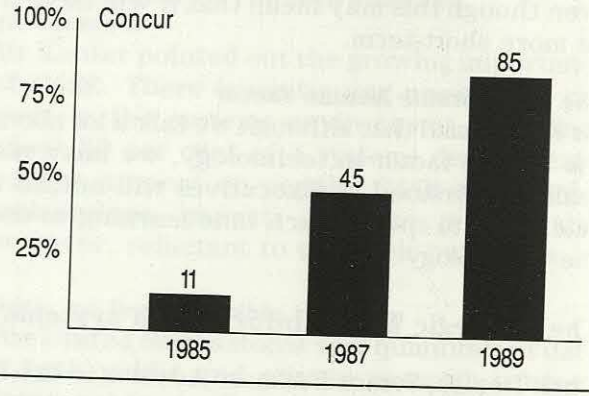

Figure 3 A business-unit priority

Strategic use of IT is very significant

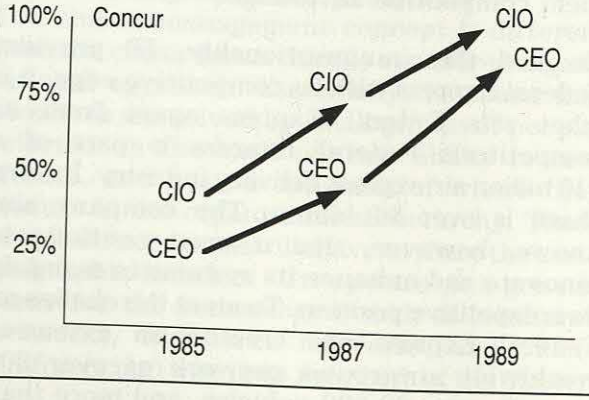

He then went on to describe the role of service (as opposed to products) in many markets. Eighty per cent of purchase decisions are based

on a combination of the quality of the product and the service (as opposed to either without the other).

Traditionally, people use technology to support operations and improve productivity; by using technology to support product *and* service developments, a company can differentiate itself and improve market access. The benefits of technology are thus multiplied, and the organisation becomes a winner rather than just an 'also ran' in the race to be ahead of the competition.

He also supported the view that Europe was likely to be at a disadvantage in the competitive race in the future because of high wage rates, and he showed figures to justify this view (see Figure 4).

Figure 4 Wage rates in advanced technology industries

	1980	1985	1988	Average 1983-1988	Average 1988-1993
Germany	$12.33	9.56	18.16	12.93	24.55
Japan	5.61	6.47	12.94	8.75	17.03
United States	9.84	12.96	13.98	13.03	15.75

Issues for the 1990s — a survey of IS executives

Jerry Kanter, Executive Director of CIMS, introduced the discussion of issues that arise in managing the information systems function by describing the results in a survey of 55 US systems managers and the 29 non-US systems managers who were on the Butler Cox Study Tour. A copy of the questionnaire is included on pages 19 to 21. The results of the survey are shown in Figure 5. Some of the interesting points to emerge from the survey are:

- The five most highly ranked issues are the same for US and non-US executives, although the order varies.
- Two issues are rated significantly higher by US executives — the importance of effective application development procedures, and cost containment.
- Two issues are rated significantly higher by non-US executives — the availability of a tested security/back-up plan, and company-wide education programmes for end users and information systems personnel.

During a panel discussion, the results of the survey were discussed and the following points were raised.

Availability of a tested security/back-up plan

It was noted that during the morning's presentation, Federal Express had admitted to giving low priority to a tested security plan in the past. This seemed to reflect a general attitude among US information systems executives, who rated this issue 12th out of 18 listed in the questionnaire. Non-US executives rated it sixth. This was probably due to the fact that the issue had received wide media coverage throughout Europe. Both US and non-US executives agreed that it was particularly difficult to gain approval for investment in security and back-up arrangements.

The skills mix and motivation of information systems personnel

Both groups rated this an important issue although, during the discussion, it was pointed out that the importance attached to education programmes for end users and information systems personnel was very low. US executives ranked it 17th, and non-US executives ranked it 13th. One possible explanation for this was

Figure 5 Information systems executives resident outside the US accord different levels of importance to particular issues

Issue	Ranking by: US	Non-US
Rapport and credibility with senior management	1	2
Knowledge of the business	2	3
Ability to recognise and exploit strategic systems opportunities	3	1
Long-range vision and plan	4	5
Skills mix and motivation of I/S personnel	5	4
More effective use of telecommunications	6	9
The establishment of the overall corporate I/S architecture	7	7
Employment of effective application development procedures	8	12
Improved level of system connectivity	9	8
Cost control and containment	10	14
Reporting level of the I/S organisation	11	10
Availability of a tested security/back-up plan	12	6
Knowledge of state-of-the-art technology	13	15
Development and implementation of I/S standards and policies	14	11
Selection and use of vendor/third-party systems	15	18
Managing use of personal computers (set standards, training, and so on)	16	17
Company-wide education program for end-users and I/S personnel	17	13
Exploitation of new technologies as they become available	18	16

that the wording of the question on education implied that such education programmes would be coordinated centrally within an organisation. Often, this is difficult to achieve. There was widespread agreement during the discussion on the need for improved levels of education and training, particularly for information systems staff.

Managing the use of personal computers
It was noted that both groups accorded this issue little importance, reflecting the fact that responsibility for personal computer systems had been clearly delegated to end users.

Focusing on the right issues
As a general conclusion to the discussion on the survey, the question of whether or not information systems executives were tackling the right issues was discussed. It was pointed out that those issues that had been rated very low were, in fact, the issues that were likely to give credibility to the systems function within the organisation. By contrast, those issues that had been rated most highly were generally associated with long-term paybacks. Several speakers from the floor suggested that information systems executives would be better off pursuing some of the more short-term issues and increasing their standing within their organisations.

The impact of globalisation on information systems management

Tony Brewer, UK Director of the Butler Cox Foundation, suggested that there are seven factors critical for the success of the systems department in supporting organisations with international operations. These are:

- Supporting the global business strategy.
- Identifying and reconciling application requirements.
- Standardising data across countries, functions, and maybe, businesses.
- Establishing a coherent technical architecture.
- Implementing and supporting *multiculturally*.
- Staying flexible.
- Learning and improving.

Traditional international companies
Until the early 1980s, academics distinguished three types of international company. They called them global, multinational, and international. Global firms produced standard commodity products that were distributed worldwide. Their strength was economies of scale in R&D, production, and marketing. Examples included consumer electronics and oil companies. Multinational firms produced differentiated products, designed for national markets. Their strength was market responsiveness. Examples included food and fast-moving consumer-goods companies. International firms produced a family of similar products at different stages of product evolution and maturity, which they were able to distribute according to the maturity of worldwide markets. Their strength was the ability to acquire technological skills and knowledge, and to manage their product life cycles. Examples included telecommunications switching and ethical pharmaceuticals companies.

These traditional distinctions are now breaking down. Consumer needs, advertising, access to product information, and wide product distribution are becoming universal. Costs and risks are large and rising, so there is a need to share fixed costs. Competition is intense, and competitive advantage is very difficult to sustain, so organisations need rapid access to worldwide markets.

Global companies
A global company now needs the strengths of all three types of international company simultaneously. Its competitive position in one national market is affected (for better or worse) by its competitive position in other national markets. It is therefore different from a domestic company that exports via overseas outposts. Examples of global companies include Ford, Matsushita, Procter & Gamble, and IBM.

There are three organisational dimensions of global management (see Figure 6). These are *business*, with a revenue and profit orientation, *function*, with an activity orientation, and *geography*, with a customer and product orientation. As firms develop from domestic, via international, to global, the relative strengths of these dimensions can change (see Figure 7).

Figure 6 Three organisational dimensions

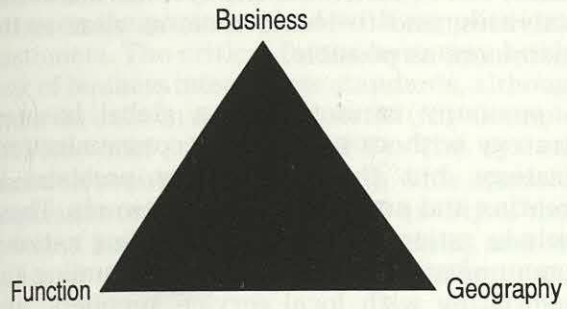

Figure 7 Examples of dimensional dominance

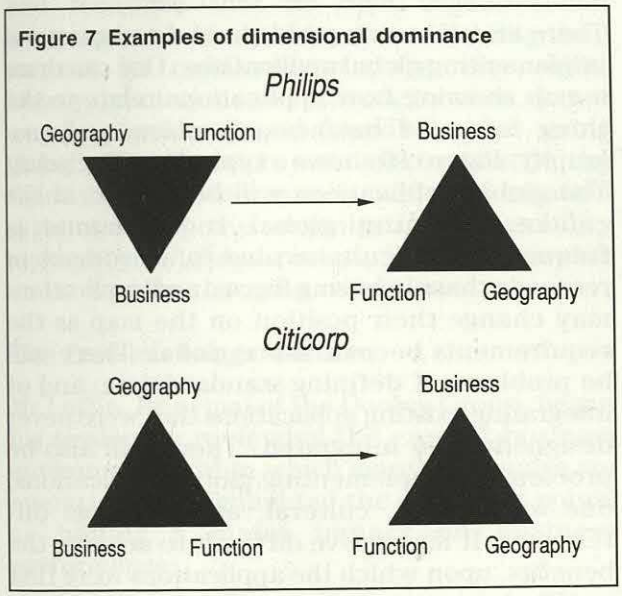

There are three sources of competitive advantage available to global firms. The first is their ability to exploit geographical differences, such as the availability of staff, or the cost of materials, or differences in market demand, or price flexibility. The second is their ability to generate economies of scale, to spread fixed costs, to speed up learning benefits, and to maximise buying and selling power. The third is their ability to achieve economies of scope, such as exploiting access to customers, or a global brand name, or technical skills and knowledge.

Global management

Successful global firms have found that they need to replace three traditional managerial assumptions by three new rules, in order to maximise the effectiveness of the whole firm. Instead of treating all their operational units as uniform, they have found that units should be differentiated, each doing what it is best at, in its own way. Instead of making all their operational units independent, they have found that they should deliberately make them interdependent. Instead of assuming that leadership and control will come from the top, they have found that objectives, values, and control should be shared throughout the firm.

There are three main management tasks in a global firm. The first is *configuration* — determining which units should do what, and where. This depends on local market needs, costs, and access to resources, and can be depicted by a global investment matrix (see Figure 8). The second is *coordination* — establishing flows within functions to transfer skills and information. The third is *linkage* — establishing flows between functions, and with suppliers and customers, thereby creating economies of scale and scope. It is now possible to design a global firm so as to maximise its effectiveness, with the use of IT providing the 'glue' and the 'oil'.

A crucial characteristic of a global firm is its organisational environment. This is the unique combination of structure, roles, skills, styles, attitudes, aspirations, and history. It is stable, durable, not easily changed, and can be both an asset and a liability. It should be an asset during stable conditions, creating internal stability and efficiency, but it may be a liability during unstable conditions, resisting change and inhibiting learning. The organisational environment determines the ability of the global firm to devise and implement a global strategy. Limited organisational capability, rather than poor business or marketing strategy, is often the

Figure 8 The global investment matrix

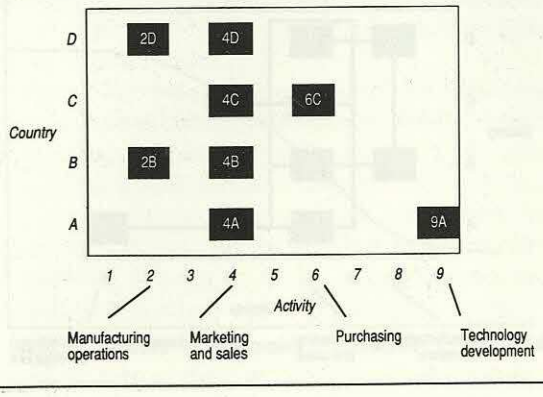

greatest constraint on effective global management.

The role of information systems management in global companies

The role of information systems management in global companies is to enable the firm to maximise the three sources of competitive advantage, and to support the three main management tasks. To support configuration, the firm needs systems for standard costing, for budgeting, and for monitoring and comparing performance. To support coordination, the firm needs team technologies, such as electronic mail, data storage and retrieval, conferencing, business TV, and video. To support linking, the firm needs systems for market analysis and forecasting, for purchasing and production scheduling, and for logistics management. It may also need inter-organisational systems to link with suppliers and customers. The firm needs to use IT to provide communication and linkage across the global resources matrix (see Figure 9), and this requires the existence of data standards and good communications.

Implications for information systems management

To provide appropriate systems support in a global firm, the systems manager must apply the three new rules, and carry out the three global management tasks, within the systems function. He or she must configure, coordinate, and link, to determine who does what, and where, and to establish what flows where, and to whom. There are no standard models for developing an effective structure for the systems function in a global organisation. However, Mr Brewer's advice was to concentrate those activities that are dependent on technical skill and knowledge, and to locate them at the most appropriate place, but to distribute the systems marketing activities, and to locate them as close to the customers as possible.

A company cannot have a global business strategy without a global telecommunications strategy, but there are many problems in creating and managing global networks. These include rationalising and reconciling national communications standards, understanding and negotiating with local service suppliers, and managing a network internationally.

There are also many problems in developing and implementing global applications. One can draw a map showing how applications relate to the three forces of business, function, and geography. Figure 10 shows a typical starting point. The global applications will be located at the centre. Defining global requirements is frequently difficult because of the need to reconcile these opposing forces, and applications may change their position on the map as the requirements become more global. There will be problems of defining standard data, and of integrating existing applications that were never designed to be integrated. There will also be problems of implementing global applications, due to national, cultural, and language differences. It may prove difficult to achieve the benefits, upon which the applications were first justified, in every country. Finally, there will be problems of applications support, release management, and customer service in many different countries.

Alliances and acquisitions are frequently the means by which global firms develop rapidly, and establish themselves in worldwide markets. The systems function should be able to

Figure 9 Use of IT to provide coordination and linkage

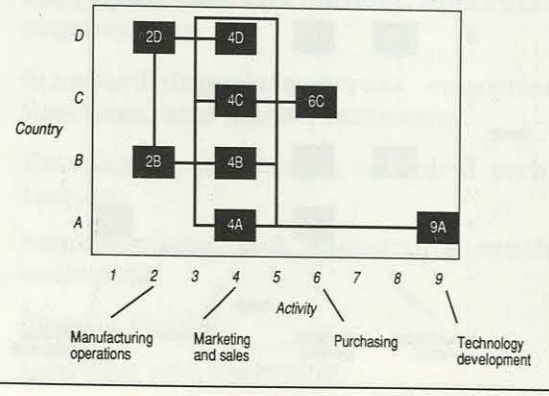

Figure 10 Application map showing opposing forces

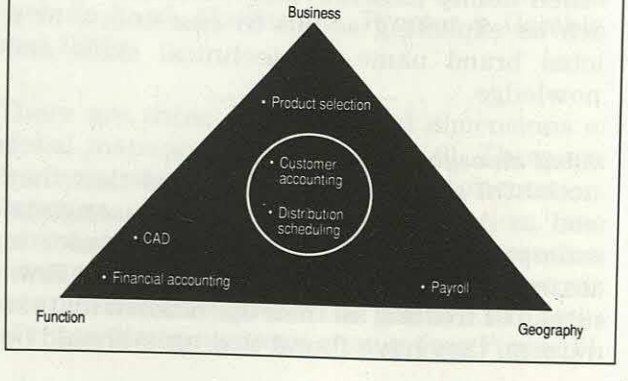

contribute to their success. Inter-organisational systems may be needed to establish links between the partners, and with suppliers and customers. The critical factor here may be the lack of business interchange standards, although these do exist in some industries (for example, SITPRO, MAP, X.400, SWIFT). The systems function should be able to help with assessing acquisition targets by judging the quality of their systems management and their data, and the value of their installed base of applications.

Finally, the systems function must be realistic and recognise that its own organisational environment is likely to be more a liability than an asset. Its traditional structure, roles, skills, history, reputation, and attitude to users, do not usually provide a good foundation for 'going global'. The function must also reconcile the rigidity that often comes with information systems, with the organisational flexibility needed by its global parent, if it is to survive under rapidly changing conditions.

The free-exchange culture

Mr Loftin, Principal of the Dooley Group, began his session by reviewing the current business environment within which many companies are operating. He highlighted the following points as having a major impact on business performance:

- The need to compete in a world market for customers, capital, and resources.
- The effect of rationalisation, which usually means mergers or acquisitions.
- Increasing market segmentation, which means that companies are having to recognise the demands being made on them by more and more classes of customers.
- The effects of regulation and deregulation.
- The growing pressure to use and exploit technology.
- Changing corporate values, often with a much more short-term orientation than before, and structural changes, of a longer-term nature, caused by national and international economic situations.

Mr Loftin referred to these pressures as the seeds of chaos. In bad times, organisations are likely to try to reduce costs, to discontinue operations, and to narrow the range of their investments. Conversely, in healthy economic conditions, companies are likely to try to diversify rapidly, to expand by acquisition and organic growth. For many companies, the reaction to the need to switch between good times and bad times very rapidly leads to an increasingly short-term focus; nowhere is this more apparent than in the area of information technology. In organisations that are going through a period of transition, there is often a lack of clarity about the roles of employees, an increased rate of organisational change, and some instability in product and market development. This is likely to result in a feeling of crisis among management, a lower quality of product, and shifting and unclear priorities. These are clear indications of an organisation under stress.

Mr Loftin suggested that organisations that are innovative are better equipped to handle the negative effects of transition. It follows that, as the need for transition is likely to increase dramatically as a result of business and competitive pressures, it would be beneficial for organisations to adopt an innovative approach. This is illustrated in Figure 11.

The characteristics of an innovative company

Innovative companies can be considered in terms of their structure and their communications channels. Structurally, they tend to be very flexible, and to function on the basis of multiple linkages or task-force activity rather than being hierarchical. Staff have broad roles.

Figure 11 The relationship between management style and the effect of organisational transition

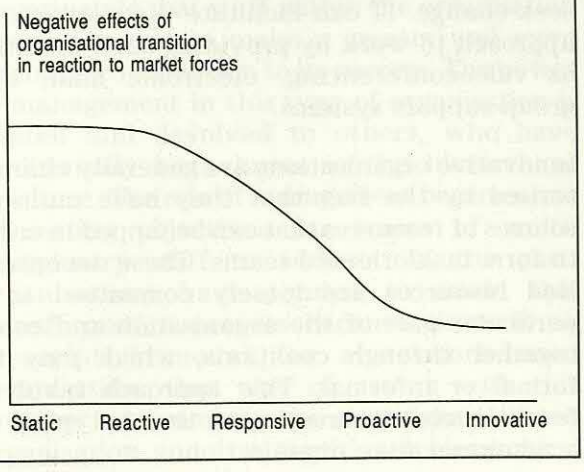

The ownership of resources is very transient, and certainly not a political issue.

In terms of communication channels, innovative companies have very open forms of communication. Members of the organisation share a common purpose and common motives; communication, both horizontally and vertically within the organisation, is facilitated. Mr Loftin referred to '*The Change Masters: Corporate Entrepreneurs at Work*', by R M Kanter, which suggests an approach to creating innovative organisations. This is based on three principles:

— Open communications.

— The use of personal networking.

— The decentralisation of resources.

The openness of communications within an organisation is significantly influenced by the physical arrangements of factories and offices. Mr Loftin advocated the open-office concept as being a particularly important tool in this respect. The free-exchange culture encourages the dissemination of information, but this also places a responsibility on management to ensure that information is not underused, or passed to people merely for the sake of it.

Personal networking describes the team-based approach to work. Teams will be formed on the basis of the expertise of individual people rather than on their status within an organisation. Team working can be facilitated by encouraging mobility within an organisation, allowing people to establish relationships and to explore new ideas together. This will generally not succeed unless employees feel secure in their employment and are prepared to take risks and seek change. IT can facilitate the networking approach to work by providing such facilities as videoconferencing, electronic mail, and group-support systems.

Innovative organisations are generally characterised by the fact that they have multiple sources of resources that can be tapped in order to form task-oriented teams. These decentralised resources are loosely committed to a particular part of the organisation and come together through coalitions, which may be formal or informal. This approach involves fewer levels of management than is typical of a non-innovative organisation.

Information systems functions are generally very reluctant to decentralise; there is a feeling that economies of scale can be obtained by centralisation. However, as IT represents only a very small proportion of the total expenditure of most organisations, it is perhaps more rational to concentrate on business needs and to use IT to support the business by decentralisation. Information systems should not be treated any differently from any other resource.

In addition to the three basic principles, managers will, of course, also need to be equipped with what have been called 'power tools' to enable the innovation process to be effected. These power tools consist of information, resources, and support from senior management. Information is essential for the various parts of an organisation to work together successfully. Resources are money, space, and the time to pursue innovative ideas. Support is in the form of managerial legitimacy and encouragement to innovate.

Conclusion

In conclusion, Mr Loftin suggested that transition within organisations was now a permanent state of affairs. Those organisations that had adopted innovative cultures were less likely to suffer in terms of their ability to react to the transition. Innovative organisations were generally very open and very successful at mobilising their human resources to achieve business success. All businesses and all sections within a business have a 'transition tolerance'. It is important to understand this, and thereby, to exploit transition.

Transformation: what companies must do to exploit IT investments in the 1990s

Ms Wilson, Director, Transformation Consortium, Nolan Norton & Co, and Principal Researcher, Sloan School of Management, MIT, reviewed some of the earlier presentations of the day, and concluded that before organisations were fully able to realise the benefits of implementing the networking approach or the benefits of global operation, some pragmatic changes and some changes in attitude were needed. By and large, IT managers and their CEOs did not have a common understanding of the important issues affecting the exploitation and the use of IT. A recent survey in the United

States showed that CEOs were generally unhappy with their ability to quantify and evaluate IT investments.

In particular, they were concerned to know whether they were doing the right things, doing the right things well, and more importantly, whether or not they were well placed to face future competition with their IT facilities. Despite the Utopian views about the future use of IT, held by many systems managers, some more fundamental issues need to be understood with respect to business organisation. Businesses need to be radically transformed to acquire a network structure, and IT can facilitate and support this change. Organisations with a network structure have the following features:

- Tasks are allocated to teams, not to positions, in a bureaucracy.
- Departmental boundaries are permeable, allowing the organisation to be more responsive.
- There are fewer levels in networked organisations. Typically, IT can be used to enrich the role of the levels that remain.
- Responsibilities become shared, and decisions become group decisions.
- The form of the organisation is organic, rather than rigid and inflexible.

Organisations are failing to exploit the opportunities offered by IT because of their inability to transform the traditional organisational structures, based on bureaucracy, to the newer and more relevant organisation structures, based on networking.

There are three critical dimensions to handling the transformation from a hierarchy- to a network-based organisation. These are shown in Figure 12 and were explained by Ms Wilson in terms of a case study based on GE in Canada. At the beginning of the 1980s, GE recognised the need to make radical changes to its organisation in order to retain its competitive position. It therefore embarked upon a process of organisational change through multi-skilled, self-managed team work, increased information-sharing, and commitment to quality and productivity throughout the organisation. The process of transformation that GE Canada

Figure 12 Critical dimensions in transforming an organisation's style from hierarchical to networking

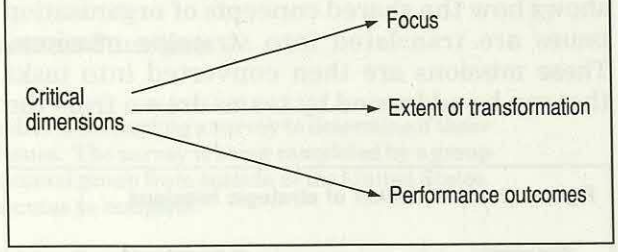

underwent involved the consolidation of corporate functions and businesses into one facility, the appointment of a steering committee, and the setting up of cross-functional task forces for the analysis and recommendation of new organisational design.

The results of the transformation were monitored; significant head-count reductions were achieved at the same time as motivation levels were increased within the organisation. One unforeseen consequence of the organisational change was a slow-down in the organisation's decision-making process. This was largely a result of the change in role of senior managers. Work descriptions and remuneration packages were not consistent with the policy of assigning work to teams rather than to individuals, who consequently found it difficult to adapt to the new organisational structure. This highlights one of the major difficulties of moving from traditional hierarchical organisations to network-based organisations — the fact that the management style and management role within organisations must change.

Network-based organisations are founded on the principle that staff within the organisation should be able to make a greater and more effective contribution to its success. The power of management in this type of organisation is diluted and devolved to others, who have traditionally been lower in the hierarchical structure. The role of management becomes one of setting objectives for teams of multi-disciplinary staff to address. In the past, organisations have tended to avoid this transformation process because there is no prescriptive model for a network-based organisation. One way of coping with this problem is to combine the objectives-oriented organisation and the traditional hierarchical

organisation and to form teams to undertake strategically important missions.

This approach is illustrated in Figure 13, which shows how the shared concepts of organisation issues are translated into strategic missions. These missions are then converted into tasks that can be addressed by teams drawn from the various parts of the organisation, as shown in Figure 14. Several organisations have found that this approach to introducing the networking organisational style is appropriate, and that it provides a vehicle for a high degree of executive involvement in both the visionary process and the process of organisational change.

Figure 13 Identification of strategic missions

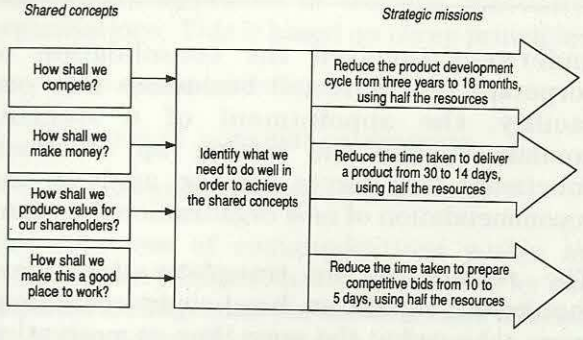

Figure 14 Satisfying strategic missions by task-oriented teams: empowering the people

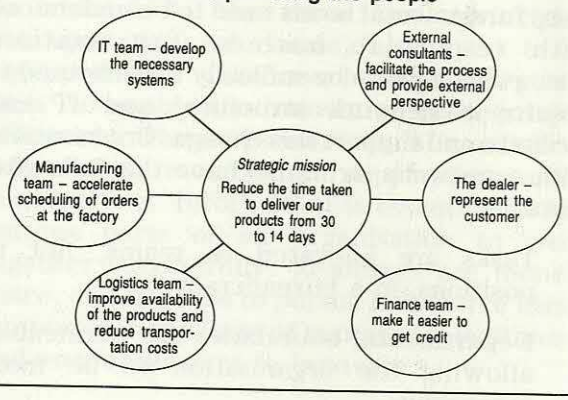

Survey questionnaire used by Babson Center for Information Management Studies

Issues in managing the information systems function

The Babson Center for Information Management Studies is conducting a survey to determine if there are any cultural differences in the importance of IS issues. The survey is being completed by a group of IS executives in the United States, and an international group from outside of the United States. The questionnaire should take no longer than 10 minutes to complete.

Please indicate below how important you think each issue is to being an effective manager of the IS function. Circle the appropriate number under each issue.

We appreciate your cooperation, and will send you a report when it is completed.

1. **Knowledge of state-of-the-art technology**

1	2	3	4	5	6	7
Irrelevant		Possibly useful		Important		Very critical

2. **Ability to recognise and exploit strategic systems opportunities**

1	2	3	4	5	6	7
Irrelevant		Possibly useful		Important		Very critical

3. **Knowledge of the business**

1	2	3	4	5	6	7
Irrelevant		Possibly useful		Important		Very critical

4. **Rapport and credibility with senior management**

1	2	3	4	5	6	7
Irrelevant		Possibly useful		Important		Very critical

5. **Cost control and containment**

1	2	3	4	5	6	7
Irrelevant		Possibly useful		Important		Very critical

6. **Reporting level of the information systems organisation**

1	2	3	4	5	6	7
Irrelevant		Possibly useful		Important		Very critical

Continued/...

7. Long-range vision and plan

1 — 2 — 3 — 4 — 5 — 6 — 7

Irrelevant; Possibly useful; Important; Very critical

8. Availability of a tested security/back-up plan

1 — 2 — 3 — 4 — 5 — 6 — 7

Irrelevant; Possibly useful; Important; Very critical

9. The establishment of the overall corporate information systems architecture

1 — 2 — 3 — 4 — 5 — 6 — 7

Irrelevant; Possibly useful; Important; Very critical

10. Exploitation of new technologies as they become available

1 — 2 — 3 — 4 — 5 — 6 — 7

Irrelevant; Possibly useful; Important; Very critical

11. Development and implementation of information systems standards and policies

1 — 2 — 3 — 4 — 5 — 6 — 7

Irrelevant; Possibly useful; Important; Very critical

12. Establishment of a company-wide information systems education programme for both end users and systems personnel

1 — 2 — 3 — 4 — 5 — 6 — 7

Irrelevant; Possibly useful; Important; Very critical

13. Skills mix and motivation of information systems personnel

1 — 2 — 3 — 4 — 5 — 6 — 7

Irrelevant; Possibly useful; Important; Very critical

Continued/...

14. Employment of effective application development procedures

1 — 2 — 3 — 4 — 5 — 6 — 7

Irrelevant — Possibly useful — Important — Very critical

15. Selection and use of vendor/third-party hardware and software systems

1 — 2 — 3 — 4 — 5 — 6 — 7

Irrelevant — Possibly useful — Important — Very critical

16. Improved level of system connectivity

1 — 2 — 3 — 4 — 5 — 6 — 7

Irrelevant — Possibly useful — Important — Very critical

17. More effective use of telecommunications

1 — 2 — 3 — 4 — 5 — 6 — 7

Irrelevant — Possibly useful — Important — Very critical

18. Managing the use of personal computers (setting standards, providing training, consulting , et al)

1 — 2 — 3 — 4 — 5 — 6 — 7

Irrelevant — Possibly useful — Important — Very critical

Please review the list and mark (1, 2, 3) the top three issues in importance.

Friday 12 May

Bachman Information Systems

Presentations were given by Arnold Kraft, President and CEO, Dick Manasseri, Product Marketing Manager, and Mike Fritz, Product Manager. A customer viewpoint was presented by Dan Novick-Timmons of Shearson Lehman Hutton, and there was a short video demonstration.

Bachman Information Systems was founded in 1983 by Charlie Bachman, who was an early worker in the field of data management and the inventor of the Bachman diagram. He set up the firm to develop and market a new type of CASE system to help maintain, enhance, and migrate existing information systems, and to develop new systems. Bachman's products are designed to address the dramatic increase in the cost of developing application software, the fact that maintenance and improvement of existing systems absorbs up to 80 per cent of development resources, and the need of systems departments to adapt their applications to new business requirements. The company now employs 80 staff in seven sales offices in the United States, with six distributors in Europe.

The characteristics of today's CASE market are:

- Many vendors/competing claims.
- Incomplete product lines.
- Individual products attempting to interface with each other.
- Products addressing only those applications developed within their own environment, and ignoring the enhancement and maintenance of other existing applications.
- Products imposing the adoption of a strict methodology and associated tool set.
- Lack of concurrent support for both data and process activities.
- Products failing to address database access or performance optimisation.
- High initial investment in training and deployment.
- High risk.

Bachman is developing an integrated set of CASE tools that will address the needs of application development and maintenance in the four quadrants shown in Figure 1. On the data side of the diagram, the basic concept is of data re-engineering. The six steps involved here are:

- Capturing the existing database design from the data-description language.
- Reverse-engineering this physical design to a logical design.
- Validating and enhancing the logical design.
- Forward-engineering to a new physical design.
- Optimising the physical design.
- Generating a new database description.

Figure 1 Bachman re-engineering product set

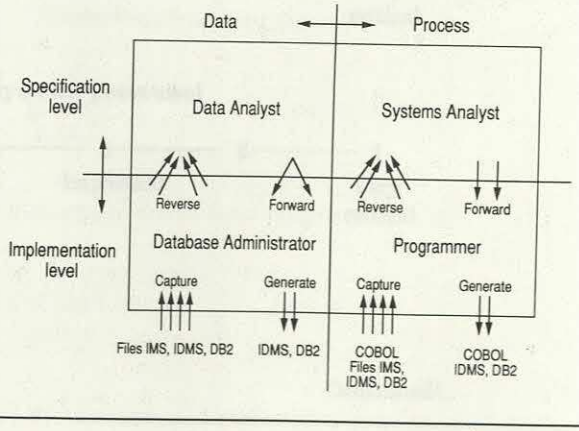

The intentions are to facilitate the re-use of previously written databases, to extend the life of existing applications, to reduce the development resources tied up in maintenance, and to make migration to new databases easier. The products will also help to automate routine tasks, give timely advice, and allow users to capture design decisions so that essential knowledge is not lost when key staff leave. When completed, the re-engineering product set will support these activities for the process side of the diagram as well as the data side.

The products are deliberately aimed at the IBM/MVS mainframe environment. They allow data descriptions to be captured from conventional VSAM files, from IMS databases, from IDMS databases, and from DB2 databases. Currently, the tools will generate data descriptions for IDMS and DB2 databases.

The products are being named to reflect the jobs that they assist. Existing products are called Database Administrator and Data Analyst, and future products will be called Systems Analyst and Programmer. The roles of these tools in the re-engineering cycle are shown in Figure 2. The market coverage of the Bachman products, compared with today's competing products, is shown in Figure 3. The current Bachman re-engineering product set, Release 2.1, available since the first quarter of 1989, is illustrated in Figure 4. The benefits of the Bachman products are:

- They focus on a real MIS problem — maintenance.
- The product benefits are understood immediately.

Figure 2 Bachman re-engineering cycle

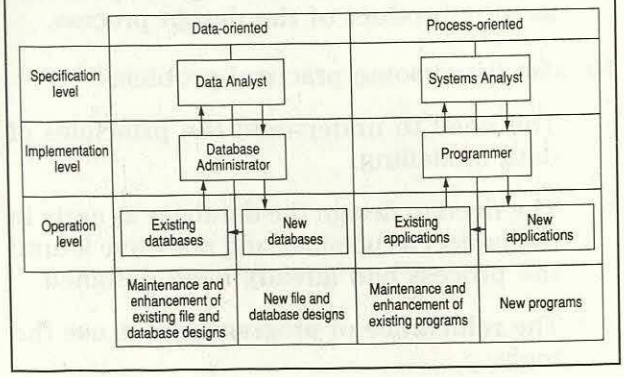

Figure 3 Bachman competitive positioning

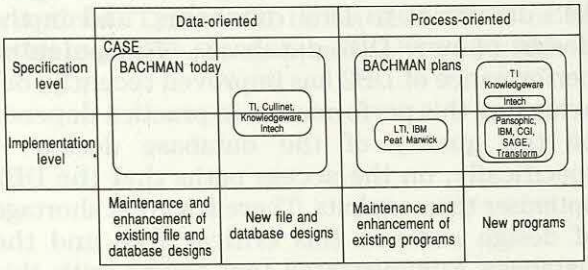

Figure 4 New products

BACHMAN/Re-engineering product set release 2.1
Available first quarter 1989

- BACHMAN/Data Analyst
- BACHMAN/Database Administrator (DB2) (Supports DB2 Versions 1 and 2)
- BACHMAN/DBA (DB2) Catalog extract
- BACHMAN/DA capture (IMS)
- BACHMAN/DA capture (files)

- Product components are individually valuable now, and do not depend on completion of the entire product line.
- They support a critical need to capture and understand data.
- Initial investment is low (all components run on PC workstations).
- Products serve as consultants and teachers to their users.
- The advisor log captures situations, recommendations, and the decision.
- The risk is low (does not impose methodology and improves individual performance immediately).
- Payback is immediate (addresses the currently installed base, thereby protecting investment).
- Shows data relationships graphically on workstations.

The most useful application of the Database Administrator tool will be in the conversion of IMS databases to DB2 databases, and in the design of new DB2 databases. The potential performance of DB2 has improved recently, but achieving this performance in practice depends on the quality of the database design — specifically, on the access paths that the DB2 optimiser then exploits. There is a great shortage of design skills in this critical area and the Database Administrator tool assists with this task and helps to optimise DB2 performance. Additional versions of Database Administrator may eventually be developed to interface with other relational database management systems.

Database Administrator will be enhanced in line with developments in DB2. For example, it is already capable of supporting the first level of database distribution that is now a feature of DB2.2. Notably absent from the Bachman product set is anything to help the development of a data-dictionary system. Bachman assumes that IBM's Repository will satisfy this need.

Bachman has developed the data tools before the process tools because it is possible to understand data structures without process information, but not vice versa. A new product, Systems Analyst, is being developed for the analysis and specification of applications. The company expects that it will be available within 18 months. It will have four components:

- An information model that will link with the data tools.
- An information-flow model.
- A functional-decomposition facility.
- A graphical-process description language.

This product is based on the functional-decomposition method developed by Chris Gane, who now works for Bachman, and who has been heavily involved in its development. It will provide expert-systems facilities that are designed to prevent invalid activities and to suggest correct action. The tool can either be user-driven, with all options available, or used in structured-dialogue mode, when the tool will guide the user through the process, thereby avoiding the problems of getting started, knowing to what do next, and knowing when the process is complete.

Another product, Systems Designer, will generate code under CICS, MVS, VM, and DB2. Bachman claims that these process tools will be the first products that will lead to optimal database design, and will automatically include up-to-date information in the design dictionary. There will also be 'seamless integration' with its other data tools. These tools will also facilitate the consolidation of individual data or process models to form a total corporate model, and will highlight any omissions or inconsistencies.

Dan Novik-Timmons provided a customer's viewpoint of two projects in which the Bachman tools had been used. The first was a mailing-list-management system that used a massive IDMS database that had been poorly designed, and was to be rewritten using DB2. The second was a trade-history system that needed to be rewritten. It was based on at least five conventional VSAM files, together with an IDMS database, all of which were holding similar information. In both cases, the Database Administrator product was used to extract a logical data design from the existing physical design, to analyse and optimise it, and then to create a new DB2 design. On the trade-history project, a pilot DB2 design was developed in eight hours, compared with several days using conventional methods. Mr Novik-Timmons listed the benefits of using the Bachman products as:

- The greater speed with which the database could be designed and reviewed.
- The ability to use graphic representations.
- The ability to make immediate changes.
- The value of the expert advisor in checking proposed designs and giving helpful advice.
- The fact that documentation was produced as a by-product of the design process.

He also listed some practical problems:

- The need to understand the principles of data modelling.
- The need to design the database as early in the project as possible, and not leave it until the process had already been designed.
- The reluctance of programmers to use the tools.

– The difficulty of implementing this new technology.

A typical price range for one copy of the software product is between $25,000 and $50,000 in the United States. In addition, a large-screen, high-resolution workstation is required. This could be an IBM PS/70 or 80, which should not cost more than $20,000.

Friday 12 May

Thinking Machines

The Thinking Machines presentations and demonstrations were given by Mr Paul Rosenblum, Mr Ralf Fiebrich, Director of Advanced System Development, Mr Robert Millstein, Manager of Applications Development, and Mr John Mucci, Vice-President for Marketing and Sales.

Thinking Machines' business is to make and sell computer systems for what Mr Rosenblum described as data-intensive applications, as opposed to processing-intensive applications. With the latter, the objective is to get a better answer by carrying out more processing on the available data (for example, computing the most accurate route for a flight to the moon). With data-intensive applications, the objective is to get a better answer by considering more data — for example, simulation, document-retrieval, or image-processing applications.

Thinking Machines was founded in 1983, and is still completely privately financed. Its first Connection Machine was introduced in April 1986, and the next generation machine (CM2) followed a year later. Prices range from $1 million to $10 million, depending on configuration and applications — a typical price is $2 million, which is relatively inexpensive for the amount of processing power provided. The machines are built out of the same parts as those used for PCs, which means that they benefit both from the mass production of PCs and from the fact that improvements in PC technology lead to improvements and possible cost savings for Thinking Machines. The power and relatively low cost of Thinking Machine's computers makes it possible to solve problems that, hitherto, were either insoluble or too expensive to solve.

The Connection Machine architecture is different from that used with other supercomputers, which divide a program up so that different processing elements work in parallel on different parts of the same application. When additional power is required, additional processing elements are added, which means that the methods for controlling and synchronising the subdivision of tasks has to be amended. The Connection Machine's 'data-parallel' architecture overcomes this problem by assigning one processing element to each piece of data. Adding further processing elements does not change the fundamental structure of the computational process.

The total amount of silicon in a Connection Machine is not significantly different from that in any large computer; by rearranging the architecture, the performance available from a given amount of silicon has been improved markedly. The architecture is also fault-tolerant, because a Connection Machine comprises many replicated, identical components.

Figure 1 shows where, in the view of Thinking Machines, parallel-data architectures fit in the evolution of computer technology and how they can offer outstanding cost/size ratios compared with conventional systems.

Elements of the architecture

The architecture of the Connection Machine is illustrated in Figure 2. A 64,000-processor machine offers 2,500 mips in performance — which is an order of magnitude or two greater than that offered by a typical mainframe. The concept underlying the architecture allows thousands of processing elements to cooperate in parallel (for example, in an image-processing application, each pixel may be assigned to one processor). The communication subsystem allows all the processing elements to

Figure 1 Data parallel architectures provide superior cost/size ratios

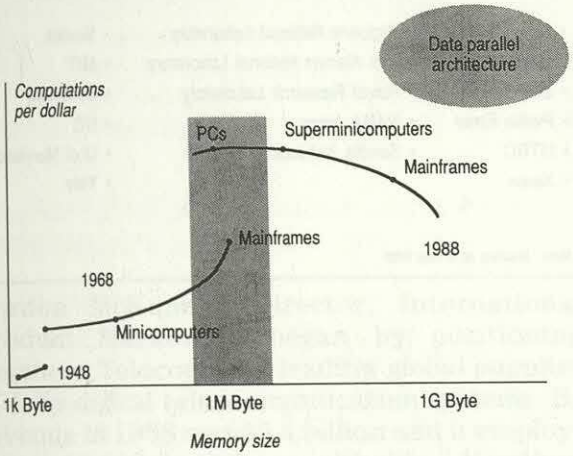

Figure 2 The Connection Machine system

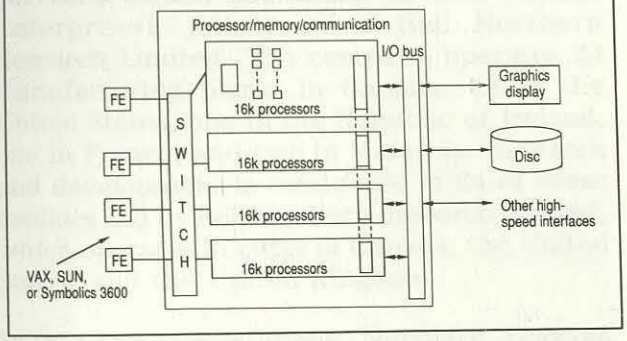

communicate with each other so that each processor can access the entire memory of the machine; it has also been constructed so that all the processing elements can communicate at the same time. The I/O subsystem is also constructed in a parallel fashion, as is the disc subsystem. The Connection Machine also has a graphics subsystem capable of very high I/O rates. In addition to the inter-processor communication subsystem, there is another communication system (based on front-end processors), which controls the main processors.

The total memory capacity for a 64,000-processor machine is 2 gigabytes; 32- or 64-bit floating point operation systems are available as options. The architecture can be expanded if more performance is needed, either by adding more processors or by improving the performance of the circuitry.

In terms of software, standard operating systems can be used with relatively minor changes, because a traditional computer system is used. In terms of programming systems, the parallel version of Fortran is mostly used, but other high-level languages, like C and Lisp, are also used. At a lower level, PARIS is used, which is roughly equivalent to an Assembler language. Programming parallel-architecture machines is a key issue, however.

Competition

All of the traditional vector-systems architectures, such as those used by Cray, NEC, IBM, and Hitachi, are pushing the limits of technology. Since little can be done to improve the performance of the processing elements, these systems will also have to use parallel architectures to achieve improved performance. However, it will not be easy to add parallel processing capabilities to these existing architectures. By comparison, the Connection Machine has been designed so that additional parallel-processing elements can be added easily.

Thinking Machines also competes with other companies like AMT, Intel, NCUBE, and Meiko, each of which provides computers with highly parallel architectures. Machines such as the Apollo Domain do not really compete with the Connection Machine. They provide limited parallelism and much lower processing performance (in the 10 to 50 mips range). They are more suited for working in parallel on relatively unconnected tasks; their communication speed is too slow to make them suitable for subdividing one program into several tasks that can be processed in parallel.

Applications and users

The Connection Machine is more powerful than other supercomputers by orders of magnitude and can be used for a whole variety of data-intensive applications (see Figure 3, overleaf). So far, the Connection Machine is being used largely, but not exclusively, by the scientific community; Thinking Machines' major clients are listed in Figure 4, overleaf. All of them are based in the United States, but a French organisation was expected to announce an order for a Connection Machine within a week or two of our visit.

Figure 3 Computer systems for data-intensive problems

- Image processing (machine vision, optical character recognition, graphics, medical imaging)
- Simulation (VLSI logic simulation, VSLI electrical simulation, fluid flow, neural networks, finite element analysis)
- Numeric processing
- Information management (free text searching, banking and investment management, statistical analysis of very large databases)
- Artificial intelligence (machine vision, memory-based reasoning, neural networks, intelligent access to huge databases, optical character recognition)
- Seismic processing (interactive 2D and 3D migration and modelling, whole earth 3D elastic wave modelling, reservoir modelling
- Computational biology

Figure 4 Thinking Machines' customers

Commercial	*Government*	*Universities*
• Dow Jones	• Argonne National Laboratory	• Boston
• GE/RCA	• Los Alamos National Laboratory	• MIT
• Lockheed	• Naval Research Laboratory	• Syracuse
• Perkin Elmer	• NASA Ames	• UC
• UTRC	• Sandia National Laboratory	• U of Maryland
• Xerox		• Yale

Note: Statistics as at May 1989

Northern Telecom

Gordon McKenzie, Director, International Product Marketing, began by positioning Northern Telecom as a leading global supplier of fully digital telecommunication systems. Its revenue in 1988 was $5.4 billion and it employs some 50,000 people worldwide. Northern Telecom is 52 per cent owned by Bell Canada Enterprises and, along with Bell Canada (a 100 per cent owned subsidiary of Bell Canada Enterprises), jointly owns Bell Northern Research Limited. The company operates 24 manufacturing plants in Canada, 13 in the United States, one in the Republic of Ireland, one in France, and two in Malaysia. Research and development is conducted in 24 of these facilities and by Bell Northern Research Limited, which operates in cities in Canada, the United States, and the United Kingdom.

During the presentations, Northern Telecom described its products and strategies in the area of ISDN links with digital PABXs, voice processing, and multivendor connectivity via digital PABXs.

ISDN links with digital PABXs

Northern Telecom has implemented ISDN in its Meridian SL-1 switch. This provides for communications between a central switch and a PABX or mainframe computer. Northern Telecom has implemented the lower three levels of the OSI reference model on the Meridian SL-1, thereby providing basic connectivity. The upper four levels, which provide the functionality, have not been completely described by CCITT, and standardisation between vendors at this level has not therefore been achieved. Northern Telecom has developed its own proprietary standards for the application-oriented layers of the OSI model. It is unlikely that standards will emerge for this level of the OSI model for some time, although Northern Telecom participates in standard-setting bodies worldwide. The Meridian SL-1 also supports the Common Channel Signalling System No 7 (CCS7) standard for digital links to the PTT network. Figure 1 shows the configuration options for the Meridian SL-1, and illustrates the following options:

- A link to a fully digital PTT link.
- Private digital links between Meridian SL-1 systems.
- Links from the Meridian SL-1 to a mainframe or other processor.

Mr McKenzie described some of the benefits of using ISDN services:

- The CLID (call line identification) facility, which allows the called party to identify the name and address of the caller. This information is obtained via the link to the computer processor, and the information is available on the recipient's computer screen simultaneously with the call. This provides call screening.
- High-speed communication for use in videoconferencing, high-speed facsimile, and high-speed data communications applications.

Figure 1 Configuration options for the Meridian SL-1

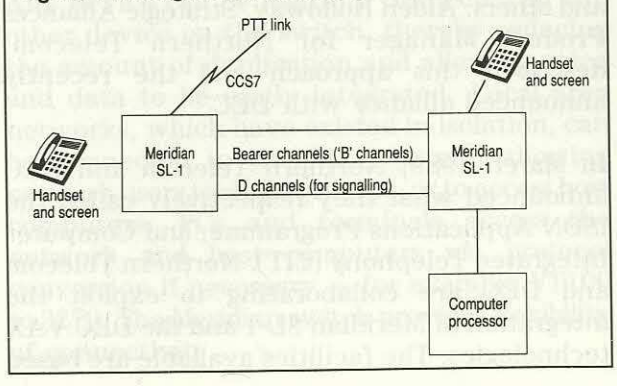

- Call routeing, which allows the recipient to re-route calls as the result of information about the caller, obtained from the processor link. This results in improved customer service.

- Faster dial time and elimination of the need for internal printed telephone directories. The user scans a database of employees on the computer processor and selects a name, and the computer instructs the PABX to make the connection. This has the additional benefit of coping with environments where staff frequently move their location. A central database provides the information for calling. This results in improved productivity.

- Integration of data. This facility allows information pertinent to the caller to be displayed on the recipient's computer screen. It is becoming widely used in the mail-order industry, and results in improved customer service.

- Simultaneous call and file transfer. This allows a call to be transferred to another extension, together with the caller-specific information displayed on the screen. This facility is used in customer-service applications, where customer enquiries need to be re-routed.

- Call segregation. The origin of a call can be identified, so that the call can be automatically routed to the most appropriate recipient.

Northern Telecom has been developing non-exclusive alliances with computer vendors in order to exploit fully the integration of voice and data. Alliances have been formed with Apple, NCR, DEC, Bull, IBM, Hewlett-Packard, and others. Alden Holloway, Strategic Alliances Product Manager for Northern Telecom, described this approach and the recently announced alliance with DEC.

In March 1989, Northern Telecom and DEC announced what they respectively called the ISDN Applications Programme, and Computer Integrated Telephony (CIT). Northern Telecom and DEC are collaborating to exploit the integration of Meridian SL-1 and the DEC VAX technologies. The facilities available are based on software developments, some on the SL-1, and some on the VAX. A developers' tool kit is now available for the VAX. This can be used to develop customised applications that integrate voice and data technology at the desktop.

Some of the potential applications for ISDN were described as follows:

- Healthcare: medical crisis calls in which the recipient will be able to utilise the CLID facility to identify the caller.

- Retail: credit authorisation, using links to customer databases.

- Car rental: booking via links to reservations and national driving-licence systems.

- Telemarketing: agents will be able to use customer information on computer screens, re-route calls and data simultaneously, and have their outgoing-call schedule automated by the voice/data integration offered by ISDN. Where multi-line telephones are used, stacked calls can be displayed, allowing the agent to decide on the priorities.

All of these applications offer opportunities for improved productivity, better customer service, and cost savings.

Voice processing

Telephones are used inefficiently by modern businesses. Calls to an individual may not be answered (perhaps because of time zone differences). From an organisation's point of view, problems manifest themselves in terms of manning of switchboards and the dissemination of information to people throughout the organisation. Northern Telecom has developed voice-messaging and voice-processing systems to overcome these problems, and this technology is available within the SL-1 product range.

The 1989 market for voice-processing systems is estimated to be some $800 million, and is forecast to grow to some $1,800 million by 1993. The growth in services on customer premises is continuing, with voice-mail the norm in the United States. Significant growth is forecast in the two other market sectors identified by Northern Telecom. Service providers — that is, PTTs — are beginning to differentiate their

services by providing voice-messaging services. Growth is forecast as a result of the development of interactive voice response (IVR) services. These services provide menu facilities via the telephone, resulting in routeing of calls to the most appropriate recipient. During the presentation, the following facilities were described:

- Voice menus.
- Meridian mail access.
- Meridian mail PC.
- Networking.

These services are in addition to the now widely accepted facility of providing the caller with an 'answering machine' type of service.

Voice menus

The 'automated attendant' feature allows callers to leave messages at any time, regardless of whether a telephonist or the recipient is present. Alternatively, callers can be routed to specific extensions, specific mail boxes, or a menu of further choices. Information mail boxes allow the recipient to pre-specify the messages and announcements that are presented to callers upon successful connection. One of the most powerful facilities available is the voice-menus application, which uses the responses keyed in by the caller via the touch-tone-dialling keys to work through menus up to 20 deep. For example, the caller could be given a voice message inviting him to key '1' for information on one product, or '2' for information on a different product. Depending upon the response he gives, the caller is led through the appropriate hierarchy of menus, which means that connection to an actual person occurs only when it is necessary. In future, this type of functionality may be provided by voice recognition.

Meridian mail access

This facility provides new business opportunities by allowing callers to access the database of information that is available via the PABX. Messages can also be called up by the caller. These can be a combination of electronic mail and voice-mail messages that have been stored in the PABX awaiting a connection from the caller. This facility is being used by hotels as an added-value service to its clients. The electronic mail facility is currently available on IBM PCs or compatibles.

Meridian mail PC

Both electronic mail and voice-mail messages can be indexed and managed by a PC-interface to the Meridian SL-1. This interface enables a mailbox owner to receive both types of messages by using the PC function keys. Messages stored on the Meridian SL-1 can be deleted, composed, sent, and so on, by this method.

Networking

Meridian SL-1 can support up to 500 nodes, connected either by analogue or digital links. Meridian mail services can be provided locally at each node, or, for those locations that do not justify a voice-processing facility, they can be provided centrally, and made available over a network. Networking standards for voice processing do not currently support multivendor connectivity, although the Northern Telecom product supports X.400. It is likely that voice-processing standards will emerge later in 1989, with implementations following in 1990.

Developments in voice processing are likely to concentrate on industry-specific applications, where the equivalent of mail-access facilities will be provided for voice messages. Also planned are increased message-handling capacity for the PABX handling, and the provision of facsimile store-and-forward facilities.

Multivendor connectivity

The Meridian SL-1 switch is based on 'open' principles, and during recent years, one of the most significant growth areas has been in configurations where the SL-1 acts as a gateway for data services (see Figure 2, overleaf). The SL-1 incorporates data cards that handle the connections of multiple devices to the switch. Any device can be logically connected to any other device on the switch, thereby reducing the amount of duplication and allowing voice and data to be easily integrated. Local area networks, which have existed in isolation, can be connected to the SL-1 system, allowing network users to share printers, or to access host computers. PCs and terminals access the network and host computers via protocol conversion if necessary — for example VT100 to 3270. The Meridian switch provides flexibility of connectivity.

The connectivity illustrated in Figure 2 is available at the desktop through a single twisted-pair cable — usually, no additional wiring is necessary. Consequently, Northern Telecom is able to offer connectivity and flexibility in linking processors and peripherals that are currently isolated within the organisation. The PABX market is growing at some 29 per cent per annum and a significant part of this growth is due to the need to provide connectivity.

The application of Meridian SL-1 data services by Martin Marietta Astronautics Group was described. Local and wide-area networks, local processors, PCs, workstations, and remote hosts are all connected via the Meridian SL-1. Any user on the network has access to any facility on the network, subject to security checks. Should any user move to a different location on the network, service access is maintained, and all outbound communication can be redirected at one central point on the Meridian SL-1. All services are available regardless of the position on the network, with the Meridian software providing the necessary intelligence.

Figure 2 Meridian SL-1 data services

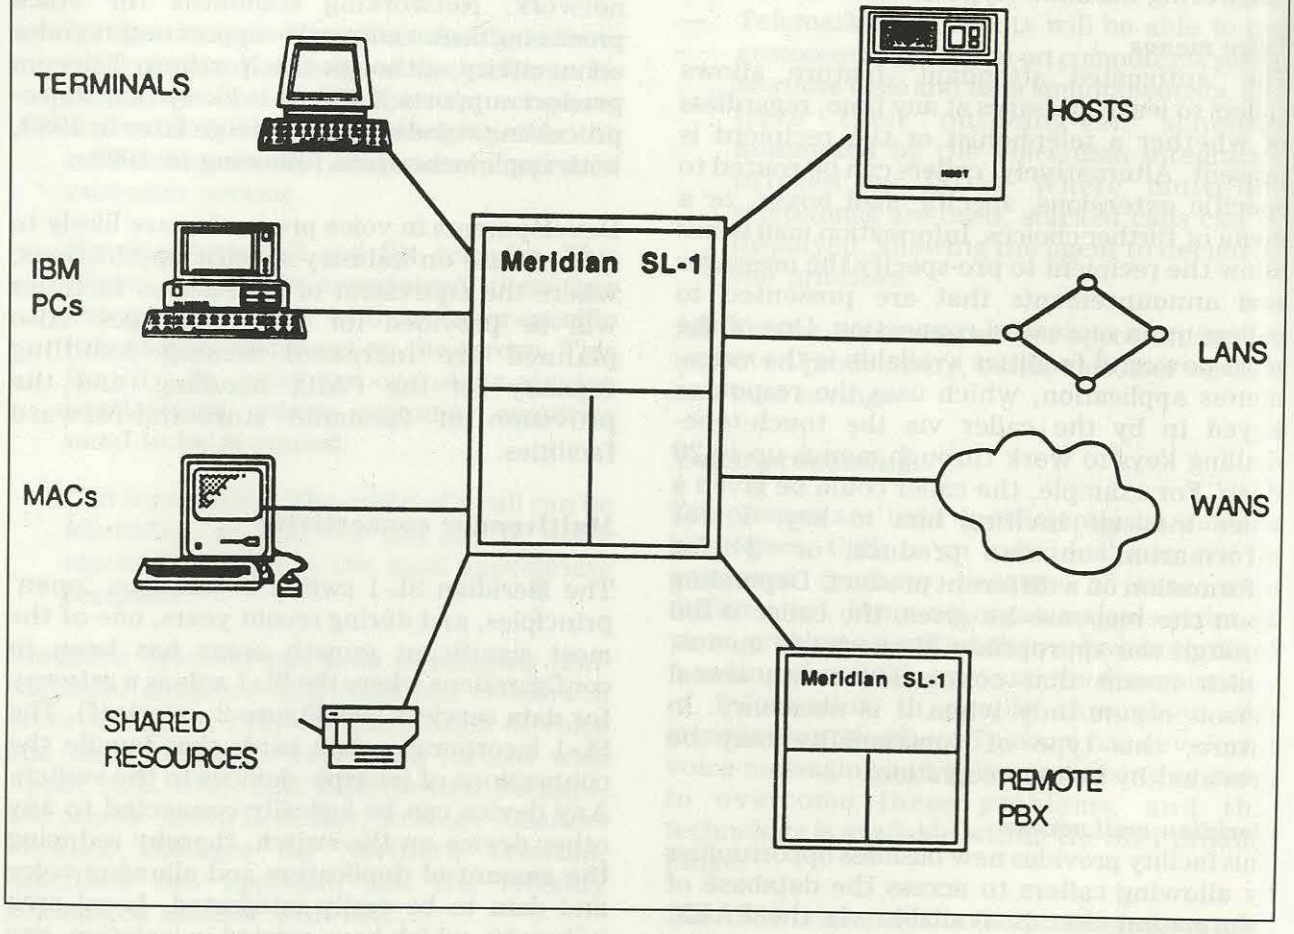

Xerox Palo Alto Research Center

Frank Squires, Vice-President, Research Operations, gave an introduction describing the history of the Xerox Corporation and of the Palo Alto Research Center (PARC). Xerox had total revenues in 1988 of $15 billion. Systems and business products contributed nearly $11 billion, the remainder being contributed by financial services. The company spent about 7 per cent of revenue on research and development, of which $100 million was shared between the three research centres. Research at PARC concentrates on information systems, computing, and artificial intelligence, and is carried out by six research laboratories. These cover computer science, electronic documents, electronics and imaging, electronic materials, and system science. Europarc, in Cambridge, United Kingdom, concentrates on the human/computer interface.

There are 350 permanent staff, of whom 265 are highly qualified scientists. It also employs around 100 visiting scientists. Examples of PARC's product successes include laser printing, personal workstations, Ethernet, programming environments (MESA, Smalltalk, Interlisp-D), solid-state lasers, and several new businesses and joint venture companies. Xerox's research strategy is to encourage inter-disciplinary activity to exploit discontinuities in technology, to test and demonstrate prototype products, and to enhance its research investment with very strong university relationships and up to 10 per cent external funding. Its current research themes include document processing, intelligent access and retrieval methods, collaborative technologies, knowledge-based productivity tools, document recognition, colour imaging, electronic materials, multifunction devices, and novel I/O devices.

Document recognition

Dr Gary Kopec then gave a presentation on document recognition. This is a relatively new activity at PARC, which has increased greatly in importance over the last two years. Xerox uses the term 'document' to mean any encoding of information for human communication in a medium. A document is therefore any kind of medium for man/machine communication, and includes recorded audio and video as well as conventional information on paper. There are four aspects of document recognition:

- Editing scanned documents.
- Document 'dry cleaning'.
- Paper as a user interface.
- The use of embedded data.

The process of editing a scanned document is complicated by the existence of different encoding methods — for example, text, diagrams, and half-tone pictures. There are therefore four parallel aspects of editing. These are:

- Using faster techniques to read the graphics.
- Using character-recognition techniques to read the text.
- Using half-tone de-screening techniques for photographs.
- Analysing the overall structure of the document.

The edited content can then be passed on to the document translator.

Xerox uses the term document 'dry cleaning' to describe all of the activities involved in removing unwanted annotations from an existing document. These could include highlighting, underlining, unwanted comments, and even coffee stains. Dry cleaning might also include maintenance of the history of revisions to a document and the provision of keyword storage and retrieval.

The use of paper as a user interface covers such activities as the automatic creation of cover sheet instructions for copiers, facsimile machines, and electronic catalogues.

Data can be embedded within a document for such things as giving formatting instructions (such as type fonts, and layouts), for maintaining a history of the use of the document, for providing pointers to other related documents, for providing instructions on the language in which the document is to be reproduced, and even for security purposes, such as prohibiting copying of the document. Dr Kopec ended by emphasising that document recognition was not a product or service on its own, but only an enabling technology, which could be used in many other systems. He said that there was no single market for document-recognition technology, but that there were many different but narrow market niches.

Electronic colour printing

Dr Richard Beach then gave a presentation on electronic colour printing. He said that PARC's interest lay in the problem of maintaining the fidelity of coloured graphical computer images when these images are reproduced outside the computer, either as photographs or colour prints. The problem arises because a given colour on the screen looks different when photographed or printed because the camera and the printer weight the colour data differently from the computer screen. In an attempt to overcome the problems, PARC has developed a set of colour-encoding standards, which define the full range of colours. However, two items that have colours that are spectroscopically identical may seem to have different colours when viewed in different kinds of light — for example, natural daylight, tungsten light, halogen light, and neon light. He then showed a humorous film with a serious purpose, which described some of PARC's research on the colour of flamingo feathers.

At the end of the session, a panel of five senior research managers answered questions from Study Tour delegates. Asked what would be the prevailing operating software-environment in the coming years, they said that they are converting their workstations to run under Unix, but they admitted that Unix would need to be considerably improved. Asked what developments were likely in the field of user interfaces, they said that they expected these to change significantly and they demonstrated a stylus-driven, hand-held input tablet which had been developed by a company in Japan. Asked what they saw as the likely developments in artificial intelligence, they said that in the field of knowledge-based systems, the problem is no longer technical, but more a problem of project management. This includes recognising potential applications, finding champions to promote them, and training staff to use the existing tools.

Hewlett-Packard

Michael Mahon, General Manager of the Client Server Architecture Laboratory, explained that cooperative processing is a concept that has evolved from two different strands — timesharing systems, where users share one machine, and personal computing systems. Timesharing systems had the flexibility of communications and the advantages of economy of operation, but they were slow and unresponsive to the needs of individuals, and when PCs appeared, there was a major move to the PC environment. However, the PC environment had the disadvantage of using peripherals and mips inefficiently, and difficulties in communications. Hence, the model of cooperative processing was developed, based on local area networks that provide access to shared resources, such as printers, plotters, and high-performance workstations.

Cooperative processing is defined as multiple machines cooperating through logical and physical connections. It is the *logical* connections that are difficult to achieve — that is, to make machines behave as a single, integrated, information system.

Hewlett-Packard believes that it is important to stop looking at a network as a means of interconnecting different equipment. It should, instead, be looked at as a utility. Hewlett-Packard sees the overall environment consisting of a backbone system with gateways to sub-networks covering separate workgroup areas — for example, engineering or manufacturing. The objective is to provide a set of facilities — for example, graphics, file storage, and database access — all provided transparently by the network. It views this process as a client/user concept, with object-oriented services provided by a broker, as shown in Figure 1.

A distinction must be drawn between the use of distributed resources and distributed

Figure 1 The client/broker concept

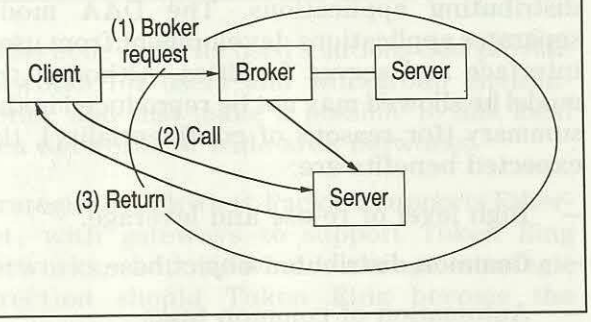

applications. The latter is only now beginning to be implemented, as shown in Figure 2. New tools are needed to enable distributed applications to be developed fully, but Hewlett-Packard expects them to grow quite rapidly because resource sharing is already well established.

The issues facing Hewlett-Packard now are:

— Standards for inter-operability. This is seen largely as a political problem.

— Hiding the complexity of the system. This is the most difficult challenge.

Figure 2 Distributed resources/distributed applications

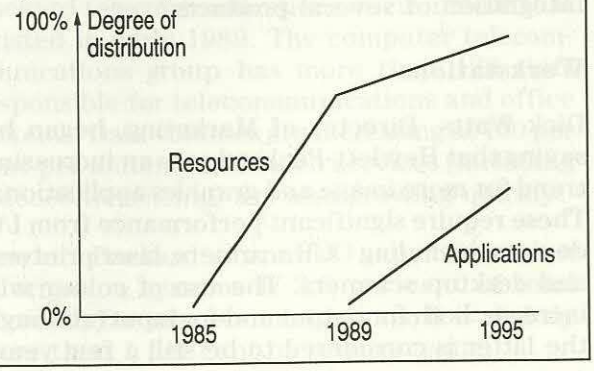

— High-performance interconnection. This is less difficult because it is generally known what needs to be done.

— Security: This may not, however, be perceived as an important problem by the market.

Distributed Applications Architecture

Phil Sakakihara, General Manager, Distributed Applications Architecture Laboratory, explained that the Distributed Applications Architecture (DAA) offered a framework for distributing applications. The DAA model separates applications development from user-interface and server facilities. Although the model he showed may not be reproduced in this summary (for reasons of confidentiality), the expected benefits are:

— High level of re-use and leverage.

— Common distributed object base.

— Automation of common tasks.

— Consistent user interface.

— Consistent set of API services.

— High level of portability.

— Extensions and scalability standards, workstations, and hosts.

— Encapsulation for existing applications.

Some of these benefits are still a vision today, but Hewlett-Packard expects that the underlying structure will be in place by the early 1990s. Mr Sakakihara said that NewWave was the first manifestation of the DAA concepts. A demonstration of NewWave was given, showing the various products. These included spreadsheet, voice mail, electronic mail, and production of regular reports involving integration of several products.

Workstations

Dick Watts, Director of Marketing, began by saying that Hewlett-Packard sees an increasing trend for more image and graphics applications. These require significant performance from I/O devices, including OCR scanners, laser printers, and desktop scanners. The use of colour will increase, both for output and for input (although the latter is considered to be still a few years away). The demand for colour copiers, colour facsimile, and so on, will grow.

There will also be a growing demand for mass storage due to the need to store graphics and images. The use of optical-storage technology will grow rapidly, as will large file servers (stacking discs), which among other functions, can provide automatic back-up.

The three main types of chip architecture — Intel's, Motorola's, and RISC — are expected to continue to provide increased performance, offering 5 to 10 mips (for PCs), 10 to 25 mips (for workstations), and 20 to 50 mips (for RISC-based 'superstations') (see Figure 3). The current base of applications supports the view that all three architectures will continue to exist. For cooperative processing, this means that identical interfaces are necessary for each architecture.

Another trend is the ability to integrate Unix and DOS/OS2 environments via the use of products such as LAN Manager. This will mean that the workstation will increasingly be viewed as a 'window' to the networks, and less emphasis will be placed on the particular working environment.

Mr Watts concluded his session by describing the market shares shown in Figure 4. At the time of our visit, Hewlett-Packard was in the process of acquiring Apollo. With this acquisition, Hewlett-Packard will have 30 per cent of the worldwide workstation market and 45 per cent of the European market.

Networking trends

Ed Muns, General Manager, Information Networks Division, described the current business environment as being characterised by:

Figure 3 Computing technology

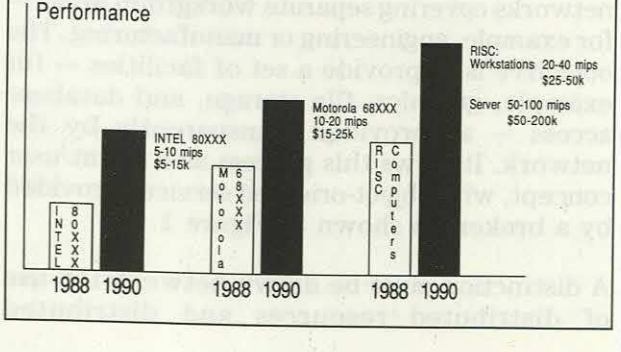

Figure 4 Worldwide workstation market

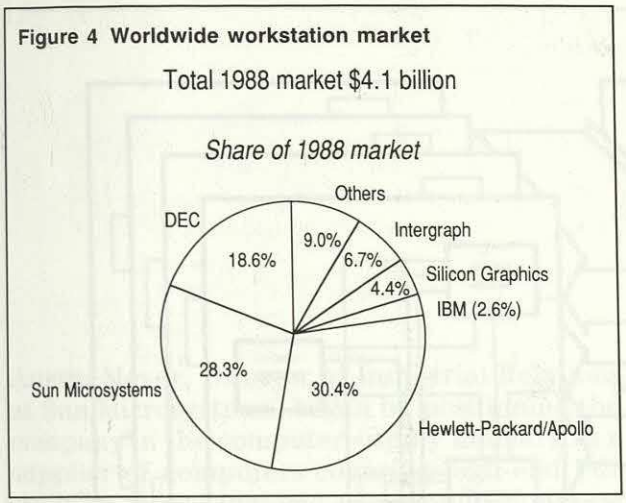

- Global competition.
- Inter-enterprise partnerships and dependence (such as those between Hewlett-Packard and Microsoft, and between Hewlett-Packard and Apollo).
- Worldwide distribution of operations.
- New technologies and industry standards.

Hewlett-Packard conceives the network as a multi-user system expanding into wide-area networks (see Figure 5).

In Hewlett-Packard, the networking approach is marketed under the name Hewlett-Packard AdvanceNet. This is not a proprietary product, but represents a strategy. It comprises multivendor communications based on OSI. Recognising that IBM would be the last supplier to adopt OSI standards, Hewlett-Packard has also implemented SNA, and it has a strong TCP/IP offering.

Basically, Hewlett-Packard's view on standards is as follows:

- If a standard exists, the company will implement it.
- In response to customer requirements, and where there are no standards, it will innovate.
- It will also take steps to promote the use of the standard (for example, through joint ventures).
- If, after some time (say, for example, five years), the standard is not adopted, Hewlett-Packard will accept the alternative.

Figure 5 Cooperative computing services

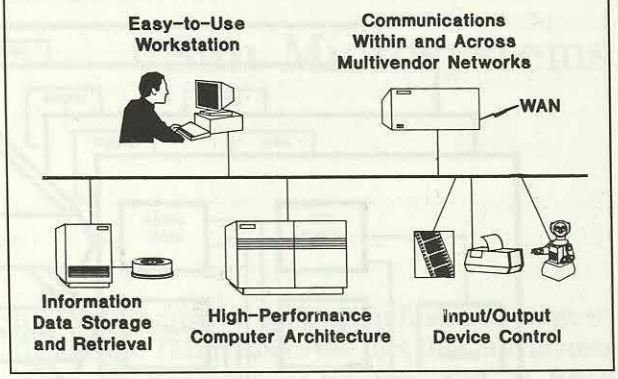

It also believes in focused solutions that provide networks for users and workgroup environments, and that make it possible to link local area networks to wide-area networks.

Strategically, Hewlett-Packard supports Ethernet, with gateways to support Token Ring networks, but it would change strategic direction should Token Ring become the dominant standard.

Hewlett-Packard's corporate information systems

Hank Taylor, Director Corporate Telecommunications, explained that the company's strategy is to evolve to a private, worldwide, digital, high-speed network to link all Hewlett-Packard sites. This network will support a mixture of voice, data, and video communications and is to use a combination of private facilities, linked, where appropriate, to public networks.

He described the Hewlett-Packard Information Architecture (Figure 6, overleaf). Hewlett-Packard has separate IT departments functionally aligned to manufacturing, marketing, and so on. Figure 7, overleaf, describes the Hewlett-Packard telecommunications environment that existed in early 1989. The computer telecommunications group has more than 120 staff responsible for telecommunications and office systems. Data volumes are increasing at 100 per cent per annum, and video services (including videoconferencing) are also growing quickly.

Hewlett-Packard's annual expenditure on information systems is 3.5 per cent of total revenue; for telecommunications, it is slightly more than 1 per cent of total revenue.

Figure 6 Hewlett-Packard Information Architecture

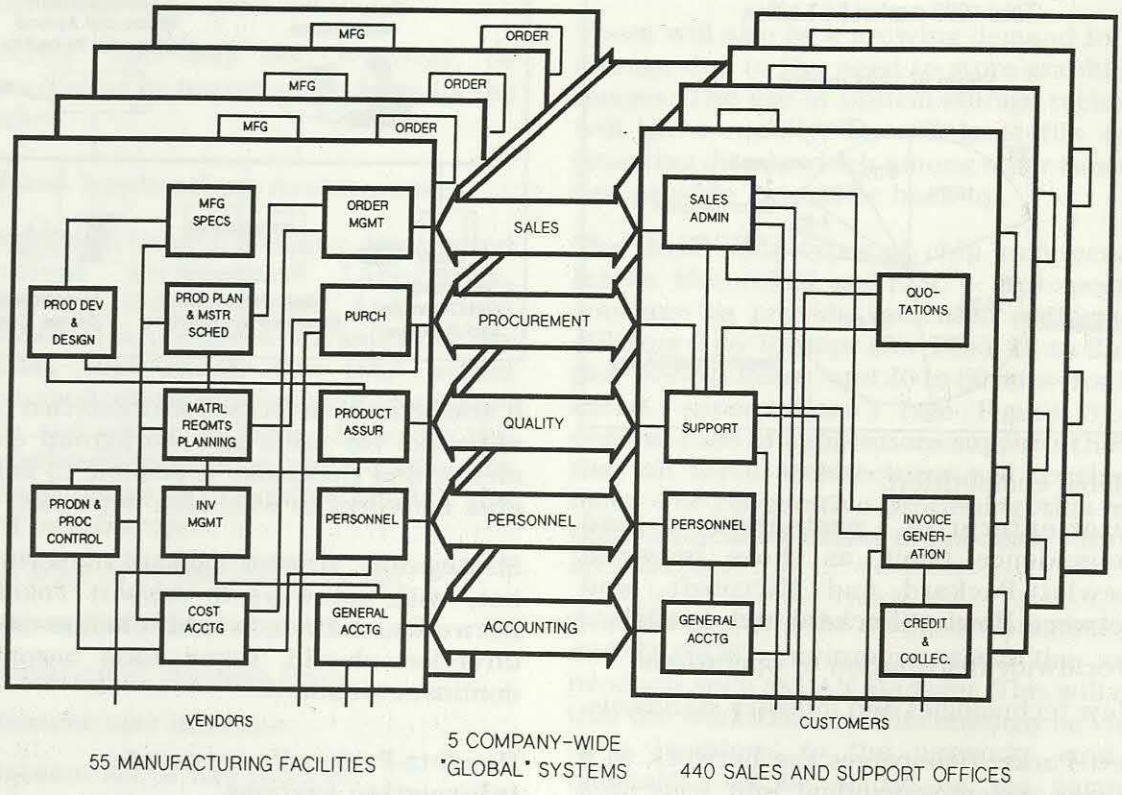

Figure 7 Hewlett-Packard telecommunications environment, January 1989

Transmission

- 1,800 Private Circuits, 95% Digital
- 2,650 Virtual Private Circuits, 99% Digital

Data Networks

- 300 WW Sites Served
- Services Include: Bil/Char/Mo:
  - Store & Forward (Batch) 55
  - Interactive (X.25/PPN) 45
  - Engineering (Internet) 100
- Speeds from 4.8 to 1,544 KB/sec

Voice Network

- 275 WW PBX's, 60% On-Net
- 70,000 Stations, 90% On-Net
- 9 Million Calls/Mo, 65% On-Net

Video Networks

- 10 Video Conference Rooms
- 1 Video Broadcast Net with Encrypted Satellite Uplink & 110 Receive Only Earth Stations in 11 countries

Networking Services

- EDI, 20% of Contract Purchase $ < 1% of Customer Orders
- 75 World Wide Public Net Connections
- 260 External Messaging Links

Site Networks

- 2,500 Host CPU's
- 7,500 Portable PC's
- 4,500 UNIX Work Stations
- 36,000 MS DOS PC's
- 30,000 Terminals
- 90 Data Switches
- 350 LAN's

Office Messaging/Conferencing

- 79,000 Electronic Mail Subscribers 200,000 Messages / Day
- 9,500 Electronic Conferencing Users
- 40 Voice Messaging Sites, 16,500 Users
- 600 Grp III FAX Units, 400,000 Pages / Month

Tuesday 16 May

Sun Microsystems

Austin Meyer, Director of Industrial Relations at Sun Microsystems, began by positioning the company in the computer-supply industry as a supplier of computers covering high-end PCs through to the low-end superminicomputers. This sector, known as the workstation sector, is based on the networking of distributed computers and the use of Unix as an open systems architecture. During the period 1987 to 1989, this market has grown from $14 billion to $29 billion worldwide, and during the same period, Sun's market share has grown from some 6 per cent to 9 per cent.

Sun competes mainly with DEC, Hewlett-Packard, Apollo, IBM, and SGI. In the technical workstation market, it is the market leader, with about a 28 per cent market share in 1988 (see Figure 1). Sun's strategy is to provide the highest-performance open systems workstations, servers, and software, based on industry standards and designed to operate in a distributed computing environment. Its objectives are to achieve price/performance leadership, continually to improve quality, and to reduce the time required to bring product developments to market. The time currently taken to bring products to market is 18 months for Sun, compared with an industry average of five years. This reflects the fact that Sun invests some 13 per cent of its turnover in R&D, compared with an industry average of 7 to 8 per cent.

Sun's revenues for the first nine months of 1989 were $1.3 billion compared with $0.68 billion in the same period of 1988, and as it continues to expand, the geographical proportion of business derived from different parts of the world is changing. For 1989, it is projected that the United States will account for only 50 per cent, compared with 62 per cent in 1988. The Pacific basin and Europe will account for the other 50 per cent. Hardware revenues, by application area, are shown in Figure 2.

Sun actively encourages 'cloning' of its products, believing that market expansion is in its interest because this provides an opportunity to add value. Recent product announcements include two Motorola-based products — the Sun 380, a 3-mip processor, priced at around $6,000, and the Sun 3470, a 7-mip processor, priced at $41,000. Three SPARC-based processors, using RISC technology, have also been announced. These are the SPARC 1 processor, offering

Figure 1 Technical workstation market share (1988)

Sun	28.3%
DEC	18.6%
Hewlett-Packard	16.9%
Apollo	13.5%
Other	22.7%

Figure 2 Sun's hardware revenues by application area

Design automation and manufacturing	29%
Software engineering	27%
Manufacturing	4%
Finance	4%
Electronic publishing	5%
Emerging markets	16%
Imaging	5%
Scientific research	10%

12.5 mips at $9,000, the SPARC 330 processor, offering 16 mips at $30,000, and the SPARC 370 processor, offering 16 mips at $41,000. Sun's intention is to continue doubling the performance of its workstations every year, and on this basis, the 200-mip chip is projected for the period 1991/92.

Sun distributes its products through OEM channels, but also sells to government bodies, academic institutions, and technical users, whom it serves directly. Sun does not provide application software but relies on value-added resalers to develop it. To date, over 1,500 applications have been developed or transferred to Sun Unix-based workstations, and while a large proportion of these are in Sun's traditional markets of CASE, CAD, artificial intelligence, and education, application software is also becoming widely available in areas such as finance and manufacturing.

Sun sees its main competition as DEC and Hewlett-Packard, although Apple and Compaq are also strong threats for the 1990s, owing to their manufacturing efficiency and cash-management control. Sun's response to this is to enhance its sales opportunities by concentrating on technology, the manufacturing process, and marketing. In the area of technology, only focused R&D is undertaken, and fewer, more focused products are brought to market. In the manufacturing process, Sun is concentrating on automation and quality improvement, and in the marketplace, is being very selective and seeking global opportunities.

In conclusion, Austin Meyer pointed to the fact that Sun's penetration of the total worldwide computer market was only 1 per cent. A large part of this market would never be available to Sun; it was still a very small company, but given its success in developing the Unix workstation market, it was likely that its business would continue to expand considerably. It would, however, continue to maintain performance on the key operating ratios of gross margin, R & D expenditure, operating margin and net margin, a historical perspective of which is shown in Figure 3.

Figure 3 Historical operating data

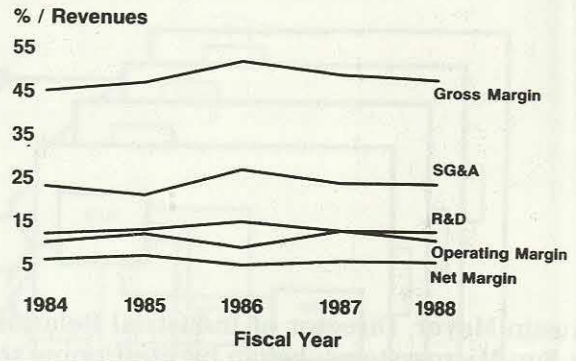

Unix unification and strategy

Sri Rajeev, Product Manager, Unix System V, Release 4, began by describing the rationale behind unifying Unix. Open systems have become widely accepted during the last few years, largely in response to the fact that standards have emerged. Open systems features were described as:

- Continuity: software investments can be preserved.
- Availability: large body of software available.
- Interoperability: different vendors' systems work together.
- Portability: application software can be used on many vendors' systems.
- Scalability: same software environment on PCs to supercomputers.

By 1987, there were two major versions of Unix offering the potential to realise the benefits of open systems These versions were Berkeley 4.2 from Sun, and Unix System V from AT&T. There were also many other versions of Unix, with the result that the critical mass necessary to attract applications developers was not being reached.

AT&T and Sun therefore formed an alliance late in 1987, with the objective of developing a unified version of Unix. The result of this will be Unix System V, Release 4, which will be available in the summer of 1989. It consists of elements of the earlier AT&T and Sun versions of Unix, as well as features of Microsoft's Xenix. Figure 4 illustrates the convergence, and Figure 5 illustrates the sources of the major components of this unified version of Unix.

Program compatibility is at two levels — the application program interface (API), and the

Figure 4 The heritage of enhanced Unix System V

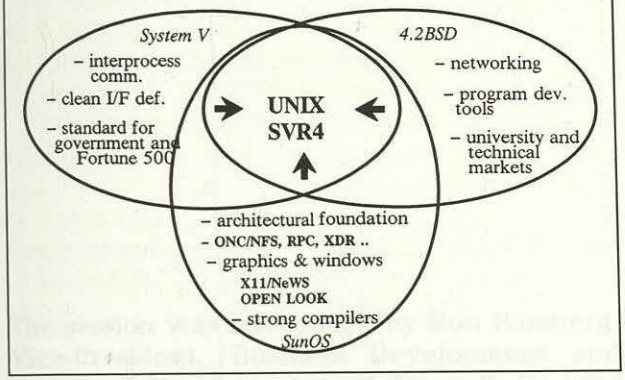

Figure 5 Unix systems consolidation

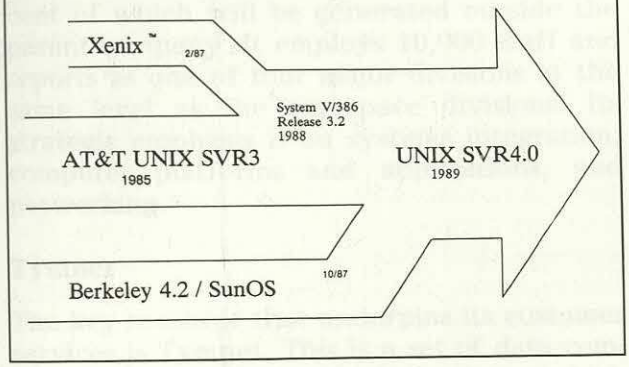

applications binary interface (ABI). There are two components of the ABI — generic and architecture-dependent. The generic component is the same across all architectures, whereas the architecture-dependent portion defines specific operating-system interfaces for each architecture. ABI provides the ability to develop 'shrink-wrapped' software for the Unix market. ABIs are under development for most chip vendors, giving the possibility of transferring software from 386 PCs to 680n0 and SPARC workstations, through S/370 mainframes, to Cray supercomputers. APIs provide a set of functionality which is common at the source-code level across different hardware ranges. This allows source code to be moved from machine to machine, and to be recompiled for the environment of the recipient machine. At the user-interface level, Open Look provides a tool kit for the specification of the user interface.

Unix will continue to develop around the 'kernel' specified in System V, Release 4, and conformance testing and compliance testing will be offered by Unix International. Vendors will be able to add to the functionality, but not to detract from the functionality of the kernel, as shown in Figure 6. As Unix System V develops, the following enhancements are planned:

- Concurrency and multiprocessing.
- Security architecture.
- Realtime capabilities.
- Commercial extensions for transaction processing and file mirroring.
- Fault tolerance.
- Support for very large multi-user systems, which implies high performance and very large databases.

The Unix suppliers now form two groups. The OSF group is developing its product from the original AIX base, while the Unix International Group is developing Unix System V. Both groups are supported by vendors and major users, and consequently, multiple versions of Unix are likely to persist.

Mr Rajeev contrasted the features of Unix System V, Release 4 with those of AIX/3. This comparison was, of course, favourable to System V, particularly as it was made in the absence of details about the eventual OSF product. The comparison is shown in detail in Figure 7, overleaf.

In conclusion, Mr Rajeev made the point that Unix has become clearly established as the open systems standard. The opportunities offered may, however, be tarnished by the battle between OSF and Unix International, and in the short term, this battle will be over user-interface

Figure 6 Unix releases and standards conformance: a schematic

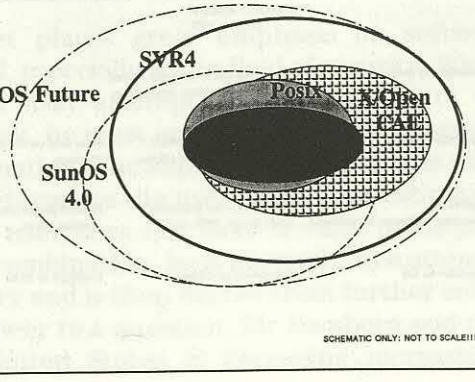

standards — OSF has announced Motif, and Open Look has been selected by Unix International. Asked about the future for Apple, if and when its competitive edge is eroded by these developments, Mr Rajeev predicted that it might join both OSF and Unix International.

Figure 7 A feature comparison of SVR4 and AIX/3

Feature	SVR4	AIX 3
Virtual Memory	Page Shared	Segment Shared: best for segmented architectures with 33 or more bits
	100s of lines to port	10,000 lines to port
	Best technology, can share pages and parts of files	Older technology, limited to sharing only segments or whole files
Dynamic linking	Application program, available to user programs	Part of kernel, not available to user programs
Shared libraries	VM from SunOS	Available
Kernel data structures	Dynamic allocation	Preallocated: less flexible
Multiprocessor support	Many preemption points in kernel; also SVR3 based multiprocessor systems available from AT&T, others	Arbitrary pre-emption in kernel, but this does not provide multiprocessing by itself; AIX has not run on multiprocessor systems yet
File System	Virtual File System supports multiple file systems, eg. NFS, RFS, UFS, S5FS	Only AIX and NFS file systems provided
File sys. integrity	FSCK for checking	Partial journaling: no more reliable than FSCK
Disk Mirroring	Not available	Available
Media Compatibility	System V and BSD	BSD only
Real Time	Hundreds of preemption points No locking necessary No bound on interrupt response Scheduler supports processes that can run to completion uninterrupted	Preemptible kernel reduces latency Locking code may cause delays No bound on interrupt response
Compatibility	Binary, object compatibility with Xenix, BSD, System V apps. Actual merge, only one mode	Only compatible with older AIX releases Needs compatibility libraries and multiple modes
Networking	Streams, Sockets interfaces can support DARPA, OSI	Sockets only: not powerful enough for OSI, SNA
Security features	B1 certifiable SVR4.1 B2 capable	C2 certifiable B2 capable??
Swap space management	Swaps to file system; dynamically allocates swap space	Swap to preallocated disk partitions: vulnerable to overallocation and underallocation bugs
Modularity	Modular, orthogonal functions eg. Streams-based networking Virtual Memory, Virtual File Systems	Modular?
Binary standardization	ABIs for many architectures	No binary standards
Quality	Origin from BSD, SVR3 and Xenix, all branching off the same base of AT&T code Well defined interfaces: Streams, VFS	Origin from several incompatible AIX versions and 4.3BSD Brand new, untested IBM code
Evolutionary features	File system types, dynamic linking Protocol independent networking	??
Hardware support	Runs on 3B2, 386, 68000, and SPARC before release	Runs only on RT PC before release
Migration/ preserving software investment	Compatible, easy migration from Xenix, BSD, SysV	Difficult migration, need larger porting effort
Expertise in OS	Widespread in industry	Confined to IBM

McDonnell Douglas

The session was introduced by Ron Bamberg, Vice-President, Business Development and Planning. He said that McDonnell Douglas Information Systems Company (MDISC) expects to have revenues in 1989 of $1.3 billion, 90 per cent of which will be generated outside the parent company. It employs 10,000 staff and reports as one of four major divisions at the same level as the aerospace divisions. Its strategic emphasis is on systems integration, computer platforms and applications, and networking.

Tymnet

The key resource that underpins its customer services is Tymnet. This is a set of data communications products and services, which allows any kind of terminal computer or local area network to communicate with any other kind, economically, accurately, and securely. Tymnet was established 20 years ago by the Tymshare company to support its own internal activities. The components of the Tymnet network are illustrated in Figure 1. Tymnet is now a telecommunications service, used by organisations, where flexibility, functionality, and changing requirements are important, where telecommunications has a high strategic value, where customer needs are geographically dispersed, and where multivendor technology is involved.

Figure 1 Tymnet network components

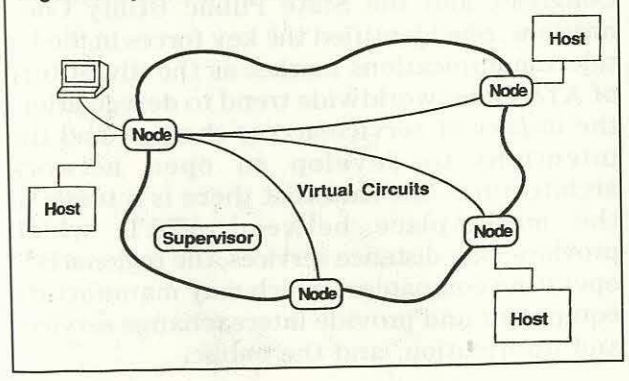

Tymnet has two primary lines of business — public data network services and private data network systems. The market for public data network services in the United States is worth $450 million, of which Tymnet has 40 per cent and Telenet a further 40 per cent. The market is currently growing by between 10 and 15 per cent per annum. There are now 2,000 customers, with over 3,000 host computers, and millions of end users. The market for private network services in the United States is worth $400 million, of which Tymnet has about 15 per cent.

Tymnet is a centrally controlled network, with many access nodes. Currently, there are five separate supervisor nodes, of which only one is in control. The active supervisor node regularly sends 'sleeping pills' to the other supervisor nodes. If a 'sleeping pill' is not received, the other supervisor nodes initiate a controlled-takeover procedure. This takes one minute, during which no new calls to the network are accepted. Control responsibility is passed to a new supervisor node in a controlled way. There has been no total failure of the network in 19 years.

Tymnet places great emphasis on software control, especially in the field of security. Whenever a user attempts to gain access to the network, he must enter a user name/password combination. The supervisor node checks these and determines the user's allowed destination. If the user does not have a valid name/password combination, he is allowed three attempts at entry and is then barred from further entry. In answer to a question, Mr Bamberg said that the United States is becoming increasingly

security-conscious, especially as there is a trend away from private to public communications services.

Financial transaction services

The second presentation, on financial-transaction services, was given by Jace Cuni, Vice-President, Sales Network Applications. Mr Cuni described the three financial-transaction services that are handled by McDonnell Douglas Payment Systems Company (MDPSC). The first service is point-of-sale (POS) authorisation for payment cards. There are 150,000 POS terminals connected via Tymnet to an MDPSC switch that distributes authorisation requests to the various card services. Sixty-five different types of POS terminal are supported, and the service supports 27 of the top 50 banks in the United States.

The second service is electronic draft capture. This replaces paper payment vouchers with electronic vouchers, and is an automated method for capturing and balancing non-cash transactions at POS terminals. This service has grown very rapidly since the end of 1986, and it now accounts for over 60 per cent of all new business.

The third service is called TeleCheck. This is a cheque-guarantee and payment service for retail merchants. In 1988, TeleCheck guaranteed 74.4 million cheque payments, worth $7.7 billion, and there are now 94,000 merchants linked to this system. Currently, this service generates a bad debt of approximately 0.4 per cent.

In the longer term, MDPSC intends to capture not only payment and customer information, but also information on the goods sold. Clearly, it is positioning itself to control the flow of information across the interfaces between traditional industries such as retail and banking.

Electronic messaging

The third presentation, on electronic messaging, was given by Nancy Browning, Marketing Manager, Applied Communications Services. She said that there is a continuum of messaging services from text messaging, via file transfer, corporate applications interchange, corporate information access, and proprietary EDI, to standard EDI. She suggested that the reasons why a third party should use a store-and-forward service are:

- Insufficient in-house capabilities.
- Limited in-house network connectivity.
- Security.
- Limited staff experience and expertise, especially in the field of data communications and messaging.

Text messaging and file transfer are used for sales-force communication, financial reporting, and communications with suppliers and customers. Corporate-applications interchange is used for sales reporting from POS devices, remote order entry, and invoicing and inventory management for distributors. Corporate-information access is used for software updates, inventory status, and pricing updates. EDI is used for exchange of business documents between computers in a standard format. She said that the MDISC service, EDI*NET, will be based, from July 1989, solely on Tandem computers. The keys to integration of all these messaging services are the X.400 standard for electronic mail and the X.12 standard for business documents.

US communications regulatory environment

Jo Ann Couche, Director of Regulatory Affairs, then gave an extremely witty presentation on the somewhat dry subject of the regulatory environment. She identified the key policy makers in the United States as the Federal Communications Commission (FCC), Judge Green, responsible for the modified final judgement on AT&T divestiture, the US Congress, and the State Public Utility Commissions. She identified the key forces in the US telecommunications market as the divestiture of AT&T, the worldwide trend to deregulation, the impact of service-access charges, and the intentions to develop an open network architecture. She said that there is a tussle in the marketplace between AT&T, which provides long-distance services, the regional Bell operating companies, which may manufacture equipment and provide interexchange services and information, and the public.

She concluded with the following observations. First, the regional Bell operating companies will gradually be deregulated and their entry into new network services will continue. Second, the global trend towards deregulation of basic services will affect the United States — local exchange rates will increase and long-distance rates will decline and stabilise. Third, the US Congress will attempt to pass legislation to limit Judge Green's authority, and his response will be to relax restrictions but to retain authority. Fourth, non-structural safeguards will replace structural separations. Fifth, users will continue to be confused and will pay more for their services.

Global networking

The final presentation on global networking was given by Mike Rude, Director, International Business Development. He described Tymnet services outside the United States.

Wednesday 17 May

Intel

Intel was founded in 1968, with the mission of exploiting the potential of large-scale integration of transistors into silicon chips. It was originally a supplier of semiconductor memory for mainframe computers and minicomputers, but has become one of the leading suppliers of microcomputers. Intel's growth in revenues and net income faltered in 1985 and 1986, but recovered in 1987 and 1988, owing largely to the launching of the 286 and 386 processors.

Other noteworthy statistics on Intel are:

- Over half (57 per cent) of its 1988 sales were in North America, 24 per cent in Europe, 7 per cent in the Asia/Pacific area, and 11 per cent in Japan.
- Sales are to OEM companies who incorporate them into their products. Seventy-five per cent of sales are made to computer companies, 10 per cent to industrial companies, and 8 per cent to communication companies, and the remaining seven per cent is split between automotive, military, and consumer applications.
- Expenditure on R&D was $318 million in 1988.
- The company's strategy is to maintain staff at about the current level of 20,000, using subcontractors to cope with additional demands, where necessary.

Product offerings

Intel's products fall into three groups: superchip sets, embedded controllers, and processors used as the basis for microcomputers. The presentation concentrated on the 32-bit 386 chip family, now being shipped, and the powerful 486, for which systems should begin to be shipped by 1990. The 286 is seen as being suitable for the low end of the PC market. Figure 1 shows how the density of chips has grown.

Intel's view is that the market is moving to 32-bit — hence the success of the 386. Every major software vendor is writing 32-bit software. The price/performance ratio of Intel's 32-bit processors is shown in Figure 2. The 386 also offers the capability to run multiple applications from different operating systems concurrently, such as MS-DOS and OS/2, or MS-DOS and Unix. (Intel has joined both Unix International and the Open Software Foundation, as it believes that although one Unix standard would be preferable, two is better than the multitude that have grown into existence.) Intel's ultimate aim is to dominate the Unix market, with DOS and OS/2 seen principally as a commercial stopgap. Figure 3 shows the choice of software on different Intel chip generations, and the superiority, in this respect, of the 386.

Essentially, the 386 will be priced to replace the high end of the 286 market. Intel's expectation is that, by 1992, more than 50 per cent of all PC shipments will be based on the 386, with the 486 accounting for around 20 per cent. The 486 is designed for the high end of the market — for

Figure 1 Growing logic chip density

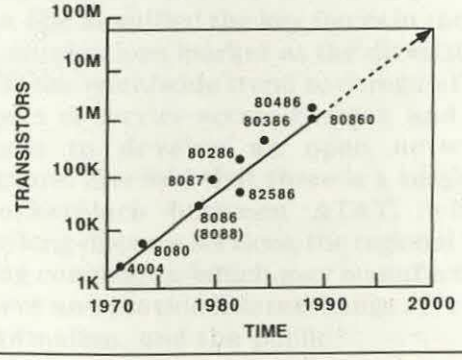

Figure 2 32-bit performance for everyone

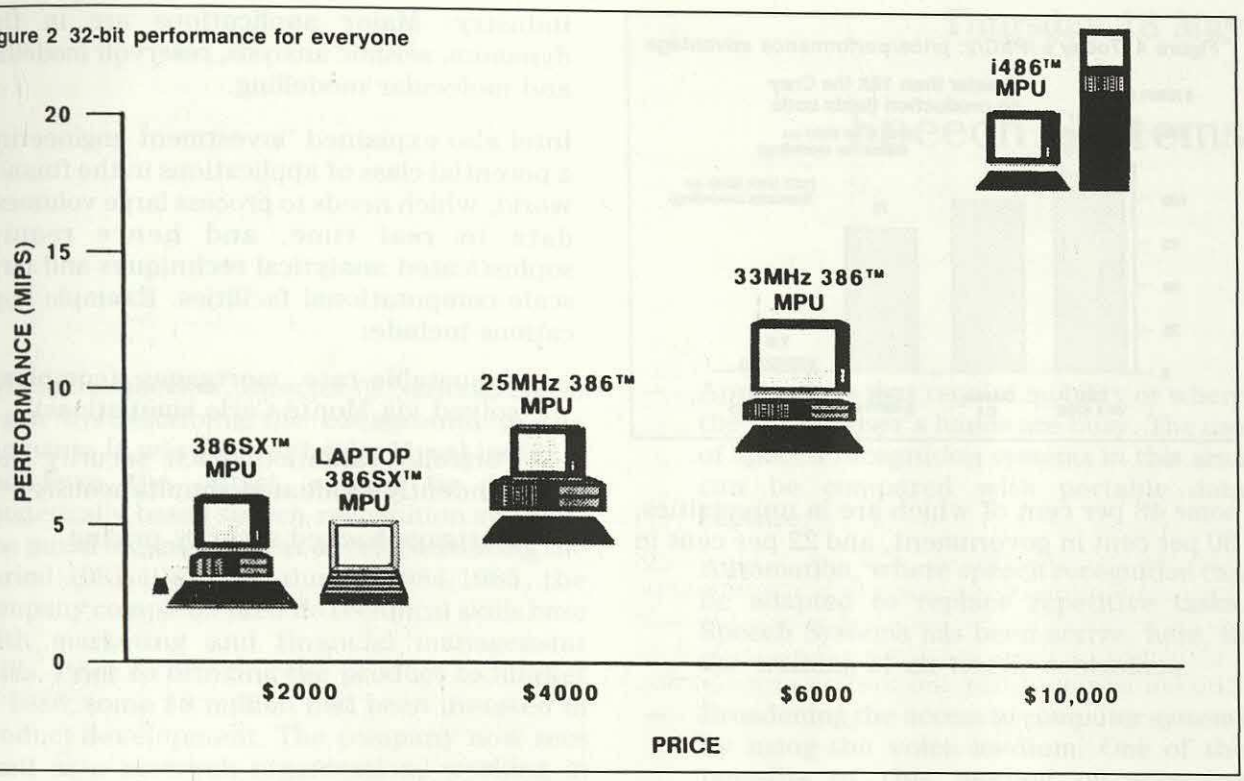

Figure 3 Not all software runs on all sytems

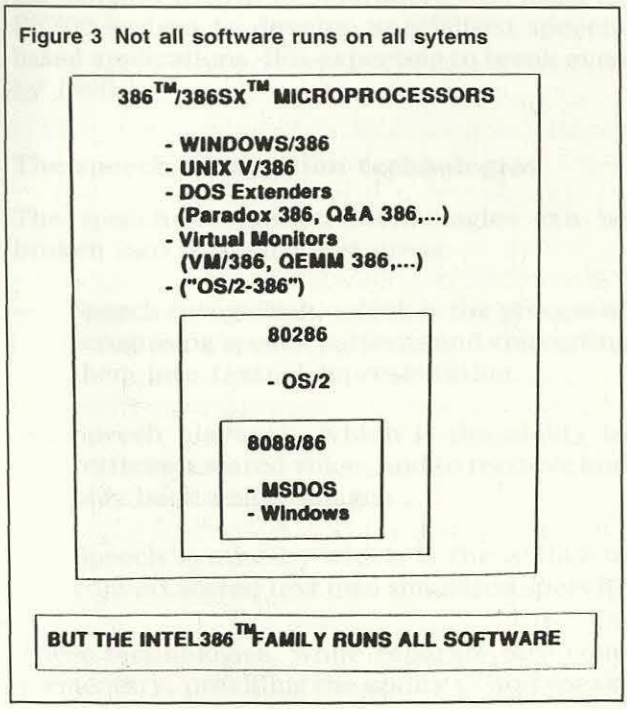

example, workstations, minicomputers, and high-end PCs. With more than one million transistors, the 486 microprocessor provides 50 times the performance of the original PCs. The objective is to achieve 100 per cent compatibility from laptop to mainframe. There are 25 million PC users and a $15 billion software base, so compatibility protects the users' investments and offers them a choice of vendors. The 386 can multi-task all MS-DOS applications, enabling the user to move rapidly from task to task.

Intel Scientific Computers

Intel Scientific Computers was formed in 1987 to sell large-scale parallel computer systems. In 1985, the iPSC/1 was launched, followed, at the end of 1987, by the iPSC/2. In October 1988, concurrent I/O facilities for the iPSC/2 were introduced, which provide mass-storage options, giving the system the potential of 500 gigabytes total disc capacity. Typical characteristics of the iPSC/2 are:

- Number of processors: 16 to 128.
- Peak processing power: 64 to 512 mips; 4 to 1,280 megaflops.
- Memory size: 16 to 2,048 megabytes.
- Storage: 1.2 to 512 gigabytes.

The price ranges from $185,000 to $2,000,000, although customers typically pay about $500,000. Figure 4, overleaf, shows price/ performance compared with other machines. There are 170 iSC systems installed worldwide,

Figure 4 Today's iPSC/2: price/performance advantage

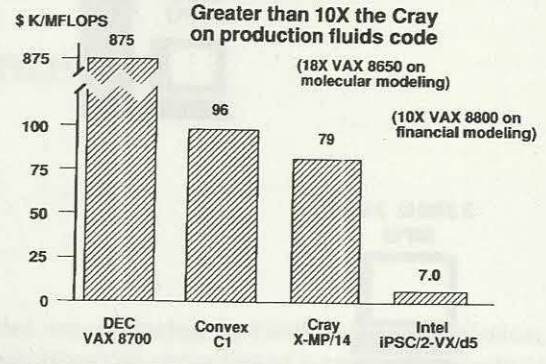

some 48 per cent of which are in universities, 30 per cent in government, and 22 per cent in industry. Major applications are in fluid dynamics, seismic analysis, reservoir modelling, and molecular modelling.

Intel also explained 'investment engineering', a potential class of applications in the financial world, which needs to process large volumes of data in real time, and hence requires sophisticated analytical techniques and large-scale computational facilities. Example applications include:

- Adjustable-rate mortgages (complexity solved via Monte Carlo simulations).
- Portfolio valuation (each security independently evaluated simultaneously).
- Mortgage-backed security pricing.

Speech Systems

Kyle Sturzenbecker, Director of National Sales, began by describing the background to the company. It was formed by Dr Mysel in 1981, and from the outset, sought to develop phonetically based speech-recognition systems. The initial technology was developed during the period 1981/1983, and during 1984/1985, the company complemented its technical skills base with marketing and financial management skills. Prior to bringing the product to market in 1986, some $8 million had been invested in product development. The company now sees itself as a research organisation, working in partnership with system builders, and using its PE200 system to develop specialised speech-based applications. It is expecting to break even by 1990.

The speech-recognition technologies

The speech-recognition technologies can be broken into three distinct areas:

- Speech recognition, which is the process of recognising speech patterns and converting them into textual representation.
- Speech playback, which is the ability to retrieve a stored voice, and to retrieve and play back voice 'images'.
- Speech synthesis, which is the ability to convert stored text into simulated speech.

These technologies, while separate, are complementary, providing the ability to add speaking and listening functions to computer systems.

Application potential

The applications of speech recognition are beginning to develop rapidly, and three main areas were described:

- Applications that require mobility or where the system user's hands are busy. The use of speech-recognition systems in this area can be compared with portable data capture.
- Automation, where speech recognition can be adapted to replace repetitive tasks. Speech Systems has been active, here, in the training of air-traffic controllers.
- Broadening the access to computer systems by using the voice medium. One of the benefits of this method of accessing computer systems is ease of use. However, this is still a very costly facility, and is unlikely to be available in the foreseeable future.

While speech-recognition systems have been available for some time, they have generally been prohibitively expensive. However, this is changing, and it is predicted that there will be a huge growth in the market for speech systems in the 1990s. As a consequence of developments in chip technology, the price/performance ratio for chips is improving at a rate of 40 per cent per annum, and by 1992, it is predicted that 15 mips will be available for about $2,000, making the provision of speech-based systems economically viable for a wide range of applications. The basic technology is now stabilising, and it is likely that workstation vendors will include speech-processing capabilities in order to add value to what will increasingly become commodity hardware products.

Technology comparison and competition

One of the advantages claimed by Speech Systems Inc for its technology is that it is both speaker-independent and vocabulary-independent. Speaker independence means that the machine does not need to be trained to

understand the user's voice in advance, and vocabulary independence means that the machine has been taught to recognise words. The products are therefore general-purpose, although most of the current applications are specialised and are based on knowledge of the application being embedded into the vocabulary.

The technologies used by speech-based systems, and some of the vendors, are shown in Figure 1. Speech Systems Inc believes that it has little competition in the marketplace, because the only product of any significance is from IBM, and this product is based on probablistic interpretation. Mr Sturzenbecker illustrated the deficiencies of the technique by asking the audience which number was likely to follow three, in a speech-based system that was based on probability. The answer is not four, as one might imagine, but 60, because of the fact that the IBM 360 is so often quoted that probablistic interpretation results in 60 being chosen by the system as the number that is most likely to follow three.

Uses of speech systems

Five examples of the use of speech systems were quoted:

- Speech-to-text processing. This technology is used in radiology, quality assurance inspection, annotation, and order entry.
- Conversational user input. In this environment, speech is understood by the machine, and this leads to the machine's responding in the required manner. It is applied in the command-and-control systems used in air-traffic control, and also in CAD, where it is possible to dispense with menus and mouse control.

Figure 1 Speech-based systems: technologies and vendors

	Technology base		
	Phonetic model	*Word model*	*State transition model (probabilistic interpretation)*
Speaker-independent	X		X
Vocabulary-independent	X		
Speaker-dependent		X	
Vocabulary-dependent		X	X
Vendors	Speech Systems Inc		IBM, BBN, CMIS

- Vocabulary retrieval. In this application area, speech-based systems can be used for database retrieval, and Speech Systems Inc is currently working with Natural Language Inc. Applications are likely to be in publishing, media, and general-purpose, public, database-access systems.
- Speech to word to speech. Because of the technology used by Speech Systems Inc, it is possible to code speech represented in text format by using data-compression techniques. As a result, the bandwidth requirement for speech transmission can be greatly reduced. This is likely to find applications in the telecommunications marketplace. The same technology can be used for language translation.
- Speech to speech. This application of the technology would be used for voice-store-and-forward systems. Because voice would be translated to compressed textual representation, there is a significant reduction in the level of disc storage needed.

Current projects

Many of the projects that Speech Systems Inc is working on are currently covered by non-disclosure agreements, although three projects were described:

- Federal Aviation Authority, which is developing training systems for air-traffic controllers.
- NASA, which is developing a training and simulation environment.
- Department of Education, which is developing a computerised language-competency assessment centre.

The US Defense Advanced Research Projects Agency (DARPA) is also funding a great deal of general language-based research, in conjunction with Speech Systems Inc.

Mr Sturzenbecker described the technology components of the company's product. These are shown in Figure 2, and discussed below:

- Continual speech is transmitted, via a headset containing a microphone, to a phonetic engine (or speech encoder). This device takes out word breaks and represents speech in phonetic codes, which

Figure 2 Components of the Speech Systems Inc product

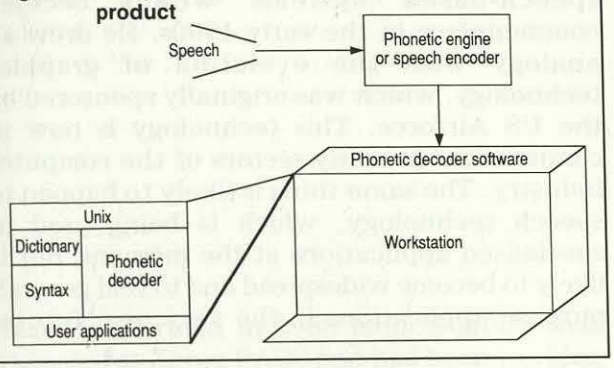

are passed to a workstation. During the presentation, Study Tour delegates saw a demonstration, using a 3-mips Sun 360 workstation running Unix. The software on this machine consists of a phonetic decoder, a dictionary, and a syntax 'code book', as well as the user application. During the demonstration, a number of delegates used the speech-recognition system, and despite their different accents, were able to communicate successfully with the machine.

- A dictionary on the machine consists of some 30,000 words, against which the phonetics arriving at the workstation via the speech encoder are compared and understood. For specific applications, a subset of additional words is installed in the machine.

- The syntax-checking facility is a knowledge-based system that works by word association to interpret homonyms correctly (words that sound the same but that have different meanings) — for instance, 'two', 'too', and 'to'. By word-association methods, the correct version can be derived. Similarly, variations in pronunciation are dealt with by this software. For instance, the word 'government' is often said in America as 'govment'. The system has the capability of reconciling this. By the same process, words such as 'um' and 'ah', often embedded in normal speech, can be omitted.

During questions, delegates asked whether background noise affected the ability of the system to recognise speech. It does, but the technology is being developed to enhance the quality of the speech that reaches the phonetic engine rather than to remove background noise, and there are likely to be technological developments in this area.

Demonstrations

The applications that were demonstrated were:

- Air-traffic control. Aircraft movements were controlled by speech, which was interpreted by a computer system. This is a highly specialised application that is based on a large number of pre-defined speech rules that apply in the area of air-traffic control. However, the method of interface with the system was through continuous speech throughout.

- As a front end to the Intergraph CAD product, where speech commands were used instead of the functions that would normally be controlled by the mouse, such as setting a template, changing line representations and line widths, and zooming in and zooming out on detail.

- Information retrieval. This application showed how voice could be used to select from a menu, and thus retrieve information from a database. The system demonstrated its ability to ask for the input to be repeated if it had little confidence in being able to understand the original input. Another application for information retrieval would be to position speech-based systems in shopping malls, where shoppers would be able to enquire about the availability of products and receive information on special offers available.

It was pointed out, however, that the demonstration system operated at only one-third of the speed that would normally be employed in a commercial application.

In summary, Mr Sturzenbecker commented on the cost of speech-based systems, which are currently about $8,500 for the voice-recognition hardware and software to 'front-end' an existing workstation. To this would need to be added the cost of the application-software development. The hardware costs are high at present because of the small volumes that are currently being built. They are likely to fall to between $5,000 and $6,000 very shortly, and as low as $3,000 when volume manufacturing commences.

It was again stressed that the application depended more on software than on hardware, and as such, it is likely to be marketed as an added-value product by workstation vendors in the future. In response to a question about the use of parallel-processing techniques, it was pointed out that chip-technology developments, and in particular, the recent release of the 860 chip from Intel offering 33 mips, were likely to offer a much better price/performance than was available from using parallel-processing techniques.

In conclusion, Mr Sturzenbecker predicted that speech-based systems would become commonplace in the early 1990s. He drew an analogy with the evolution of graphics technology, which was originally sponsored by the US Airforce. This technology is now in common use in many sectors of the computer industry. The same thing is likely to happen to speech technology, which is being used in specialised applications at the moment, but is likely to become widespread and to find general-purpose applications in the very near future.

Pixar

Pixar was formed in 1986 from a group that originated at Xerox PARC and had been working together for 15 years. The prime investor in Pixar is Steve Jobs, and Pixar is a sister company of NeXT.

Alvy Ray Smith of Pixar said that he had two objectives in giving his presentation. The first was to explain the essential difference between the two sides of image processing — geometrical modelling, and imaging. The second was to 'calibrate' our minds, in terms of the difficulty of creating realistic images and the amount of processing power required.

The two sides of image processing are represented in Figure 1. The geometrical modelling domain is concerned with representing objects by means of lines, polygons, surfaces, and three-dimensional drawings. Surface colour and texture, light source, and atmospheric haze can all be represented in the model. The imaging domain is concerned with representing pictures in a realistic way. The components of the picture have no individual identity; they are represented simply by thousands of picture elements (pixels), which are merely coloured dots filling a two-dimensional or three-dimensional space. The process of making geometrical models look realistic is called rendering, and the reverse process of extracting a discrete object from an amorphous mass of pixels is called feature extract or pattern recognition. The dotted line in the figure separates the world of computer graphics from the world of image processing. The former includes most of the activities of geometrical modelling and a little of imaging; the latter contains most of imaging and a little of modelling.

Pixar has proposed a set of standards to define the interface between modelling and rendering.

Figure 1 The two sides of image processing

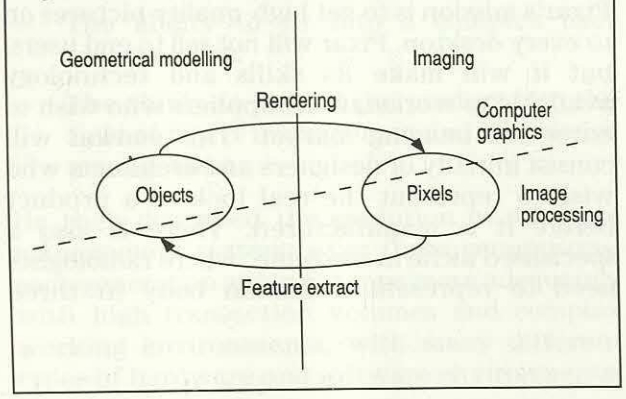

These are the RenderMan standards, and they are increasingly being accepted by the image-processing industry. Mr Smith drew the analogy between RenderMan and Postscript, which is now the accepted standard for defining what a page looks like, and has led to the development of the desktop publishing industry.

Mr Smith projected a series of beautiful photographs to illustrate various aspects of geometrical modelling and image rendering. Whereas geometrical models are created by mathematical techniques, pixel images may be created by scanning real objects in three dimensions. He explained that, given sufficient mathematical skill and processing power, a picture built up from geometrical models can look as realistic as a photograph of a real object. A picture of bicycles in a shop, for example, was made up entirely of geometrical models, and consisted of 40,000,000 polygons. A three-minute animated film required the computational power of six Cray and 15 VAX supercomputers. To represent realtime images at TV standard resolution requires between 1,000 and 3,000 mips.

Until very recently, this kind of processing power has not been available to the image

industry, nor has there been the capability to handle pixels in bulk. The Pixar image computer has been designed to do this. The standard 24-megabyte model can store 4,000,000 colour or 16,000,000 monochrome pixels. This would be required to represent a 256 x 256 x 256 block of three-dimensional space. Its architecture has been specially designed to process pixels and so provide realistic rendering. The standard model costs $29,500, and is used as a front end to a graphics workstation, which might be a Sun or MicroVAX machine.

Pixar's mission is to get high-quality pictures on to every desktop. Pixar will not sell to end users, but it will make its skills and technology available to workstation suppliers who wish to enter the imaging market. This market will consist initially of designers and architects who wish to represent the real look of a product before it is manufactured. There is also a specialised niche in medicine, where radiologists need to represent a human body in three dimensions so as to locate injuries or tumours without invasive surgery. The image computer has also found applications in areas like fault detection in turbine blades and in rocket engines.

Mr Smith ended his presentation by predicting that the 1990s will be dominated by graphics. Realistic rendering and visualisation will start to have an impact within five years, and within 10 years, will be commonplace. The workstations displaying the best pictures will be the most successful products. Pixar intends to supply the technology that will create this new industry. Mr Smith ended by showing three short animation films, the most recent of which won an Oscar. They were fully computed, in the sense that every image was created within the computer, rather than via external photography. Very little of Pixar's business is in the film industry, but it makes such films periodically to put its skills and technology on display.

Friday 19 May

Relational Technology

The opening presentation was given by Paul Newton, President and Chief Executive Officer, who introduced Relational Technology Inc (RTI) and the Ingres product line. RTI provides a family of software products for database management and application development. The company was founded in 1980 and is now one of the world's top 20 software companies. It is a leading supplier of software tools to users of Digital VAX computers, and minicomputers running the Unix operating system. Its products are also available on mainframes, other minicomputers, workstations, and personal computers from many suppliers.

Mr Newton said that relational data management is now the fastest-growing sector of the market for data management tools, out-selling conventional systems by five to one. He said that the market for relational database management systems breaks down into three parts. The first contains management systems that concentrate on data enquiry and analysis for decision support. The principal suppliers in this part are Informix and Ashton Tate. The second part contains management systems for mainline production applications. The main suppliers are IBM, Digital Equipment, RTI, and Oracle. The third part contains products for high-performance, online, teleprocessing applications. The main suppliers here are Tandem and Sybase.

The market for mainline production applications is the largest and fastest-growing part of the market. RTI has to compete with IBM and Digital in this sector, but it sees Oracle as its main competitor. Mr Newton said that the main product requirements to compete in this sector are:

- A fast, ruggedised, industrial-strength, processing engine.
- Integrated application development tools.
- The ability to fit into a complex user environment.
- The ability to extend the product into the field of object management and knowledge management.

He then described the evolution of database management systems over three generations, each generation aiming to cope more adequately with high transaction volumes and complex working environments, with many different types of hardware and software environments.

Mr Newton then briefly outlined the family of Ingres application-development tools. This comprises CASE design tools, prototyping tools, fourth-generation production languages, and high-performance languages for performance tuning. He said that the Ingres family forms an integrated set, all of which have the same look and feel, and all of which interface with the same data dictionary system.

The second presentation was given by Dr Michael Stonebreaker, who was one of the founders of RTI. He said that, in the future, data management tools will need to be able to handle not only data management, but also object management and knowledge management.

He also showed how the capability of Version 6 of Ingres, which has only recently been released, will be developed over the next five years. He described Version 6 as 'the Sybase killer'. It already has significant performance advantages and it has been written to include object-management capability. Dr Stonebreaker said that providing this capability had required a complete re-write of the product, and he believed that this gave RTI a two-year advantage over its competitors who do not yet include this capability.

All the current relational database management systems are designed to handle business data processing applications, and do not have the capability to handle objects such as geometrical models in CAD applications, or text, or spatial data in geographical information systems. Dr Stonebreaker illustrated this by showing how difficult it would be to handle a relatively simple geometrical problem using conventional SQL, and how easy it would be with object-management capabilities. He said that the main benefits of object management are that it facilitates efficient management of more general kinds of information and it provides reusable code that can be shared throughout many applications.

Dr Stonebreaker then illustrated some of the concepts of knowledge management. He included allerters, figures, forward chaining, backward chaining, super views and exceptions. Again, he said that conventional relational database management systems are not able to handle these concepts easily, whereas systems designed for knowledge management would be able to do so. The new version of Ingres, available later this year, will have the ability to handle all three dimensions — data management, object management, and knowledge management. He predicted that there would be a spate of announcements from other suppliers as they attempt to provide these facilities using various supersets of SQL that will be incompatible.

In the final part of his presentation, Dr Stonebreaker illustrated how it is possible to save two orders of magnitude in hardware cost. First, main memory is falling in cost to around $500 per megabyte. This implies that we should be thinking in terms of gigabytes of main-memory storage, so that databases with high transaction volumes should be resident in main memory. However, this implies that it would be desirable to have uninterruptible power supplies. It also implies that the current optimiser routines in database management systems will need to evolve, from handling storage on auxiliary discs to handling storage in main memory. Second, he said that the cost of central processing units is falling to below $5,000 per mip, and even as low as $1,000 per mip. This trend is illustrated by Joy's Law, established by Bill Joy of Sun Microsystems, who said that chip performance can be predicted by the equation:

Single chip performance = 2^n, where n equals current year minus 1984.

Thus, today, in 1989, chip performance should be 2^5 = 32 mips.

Chips with this sort of performance will be available this year. This high level of performance is being achieved with Reduced Instruction Set Chip (RISC) architecture rather than conventional chip architecture. Processor vendors are therefore dividing into two groups — those using proprietary conventional chip architectures designed before 1984 (including Intel 80XX, Motorola 68XXX, VAX, 370, AS400, PS/2), and those using RISC. In general, RISC-based processors run three times as fast as non-RISC-based processors. Because RISC chips are simpler and easier to design, RISC processors based on them will develop faster than their non-RISC equivalents.

Figure 1 shows the price per mip, across various speed ranges, of some of the popular processor models. It illustrates vividly that RISC-based processors are 10 to 50 times cheaper than non-RISC processors with the same speed. He said that several vendors are now building shared-memory multiprocessors, using RISC chips, that sell at around $2,000 per mip. In answer to a question, Dr Stonebreaker agreed that his figure was slightly over-simplified. The key factor is transactions per second per dollar. However, the overall trend is quite clear — all users should be planning how to transfer from proprietary hardware architectures to RISC-based architectures, because RISC chips will become universal commodity products.

Figure 1 RISC chips provide much better price/performance than other architectures

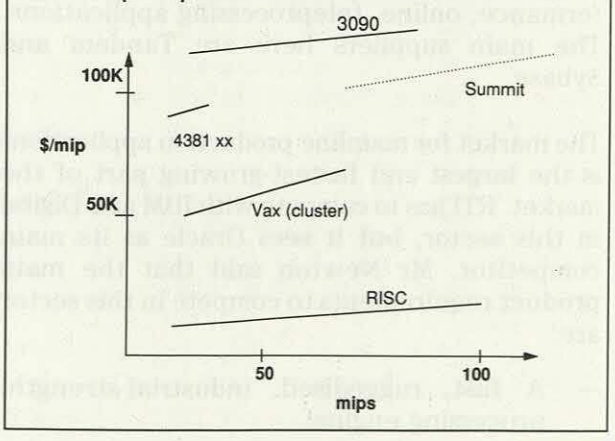

Dr Stonebreaker said that all the vendors with proprietary architectures now have strategies for basing their future products on RISC, while maximising their revenues from existing products. Intel will move to the 860 family, and Motorola to the 88000 family. Digital is already offering RISC-based workstations. He expects an aggressive announcement from IBM, probably a PC/RT workstation running AIX, shortly. Ultimately, OS/2 will have to be modified so that it can run on RISC architectures, and this will add to its uncertainty as a mainline IBM product.

In answer to a question, Dr Stonebreaker agreed that users with simple needs will be well supported by 386/486-based products, but mips are addictive, and once users get used to high-quality graphical interfaces, to multi-user multi-tasking, and to high-performance workstations, their demand for power will be insatiable.

He said that a universal, server-based operating environment is more functional and practical than a proprietary workstation-based operating environment, such as OS/2. Unix is the only universal environment on offer, and RTI is betting that Unix will win. He said that many database vendors, including RTI, will provide high-performance transaction-processing capability for Unix. He commented that the apparent differences between OSF and Unix International are more political than real, and that, ultimately, the market will demand a single version of Unix, capable of running on the cheapest RISC chip available.

Dr Stonebreaker then said that the third way to save cost and improve performance was by using arrays of multiple 3.5″ disc drives. All sizes of disc drives now have the same cost per byte, reliability, latency, and data-transfer rate. However, to store, for example, 7.5 gigabytes would require either one 14″ conventional drive or 75 PC-type 3.5″ drives. The advantages of using an array of small drives are that they provide many more read heads (and can therefore access data faster), they require less space, and they consume much less power. Their disadvantage is that the total array is less reliable than a single drive. This problem may be overcome by storing data 'horizontally', with the bits for one byte being stored on five separate drives including a parity drive, rather than 'vertically', with all the bits stored on a single drive. This change would greatly increase data resilience, allowing data to be reconstructed using the parity drive in the event of a head crash, and would also avoid the need for data mirroring.

Dr Stonebreaker summarised by saying that the way to 'run blindingly fast and cheap' was:

- To move to RISC architecture.
- To use gigabytes of main memory.
- To use arrays of 3.5″ drives.

Delegates commented afterwards that this presentation perhaps provided an answer to their questions about the motives of the OSF sponsors. They wondered if OSF was part of IBM's strategy for transferring to RISC.

Friday 19 May

Natural Language

Tania Anochaev, Chairman, President, and CEO, gave an introduction to Natural Language Inc (NLI). She said that the opportunity for the creation of NLI came about as the result of the convergence of a number of market forces — in particular, the growth of departmental processing, the growth of powerful workstations (particularly those running Unix), the emergence of SQL as an industry standard database language, the growth in the market for relational database management systems, and the storage of strategic information in relational databases. This convergence is creating a demand for easy access to information from an increasing number of non-computer professionals. NLI provides a family of information-management products that allow computer users simultaneous access to a heterogeneous mix of databases on either remote or local computers.

NLI products

NLI's principal product is Natural Language. This is an intelligent interface, providing access to relational databases by translating conversational English into SQL. NLI claims that it is the only commercially available information-access product that requires no end-user training, since queries and responses are all posed in conversational English. Natural Language resides on a workstation, is able to access all the major relational database management systems using their own version of SQL, can be easily transferred to different operating systems and hardware, and allows concurrent and transparent access to multiple, heterogeneous, relational databases.

Other associated products are:

- NLI Connector, which is an interactive knowledge-based tool that is used to define and model the relationships in a specific database. It is used to provide Natural Language with the additional vocabulary and concepts required to handle specific types of database.
- NLI Gateway, which is an interface that allows Natural Language to run on a variety of relational databases.
- NLI Database Interface, which generates the specific SQL commands for each type of relational database.

The product set costs $17,000 upwards, depending on the hardware used and the number of connected workstations.

The future market

The need for these products is demonstrated by a recent survey of Fortune 500 companies in the United States. Eighty-one per cent will have applications that access data from more than one computer. Sixty-two per cent will require access to data on different vendors' computers. Eighty-four per cent will have applications accessing data from different databases and file systems.

The market for relational database management systems will have grown to nearly $6 billion by 1992 (see Figure 1). The market for database tools will have grown to over $13 billion by 1991 (see Figure 2). Of this, NLI believes that it will achieve a share worth $300 million. However, in addition to the market for natural-language access to relational databases, there will also be the larger and complementary market for information services, advisory systems, and speech interfaces (see Figure 3). NLI sees market opportunities especially in manufacturing, pharmaceuticals, insurance, financial services, telecommunications, consumer goods, health care, and government.

Figure 1 Forecast market for database management systems, 1987-1992

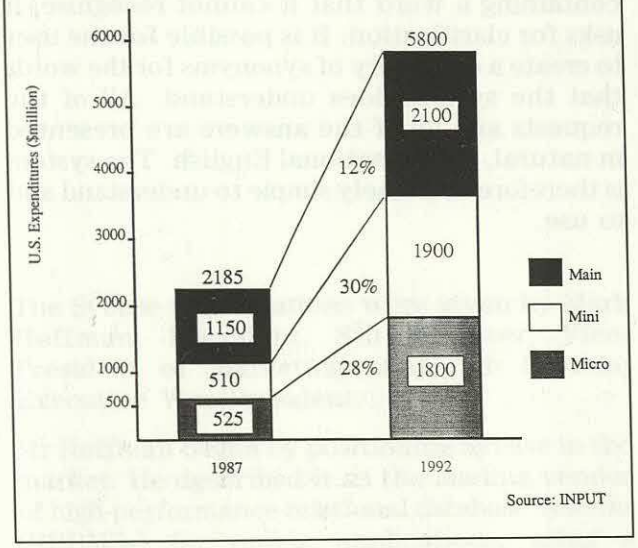

Figure 2 Forecast market for database tools (includes all database-related products)

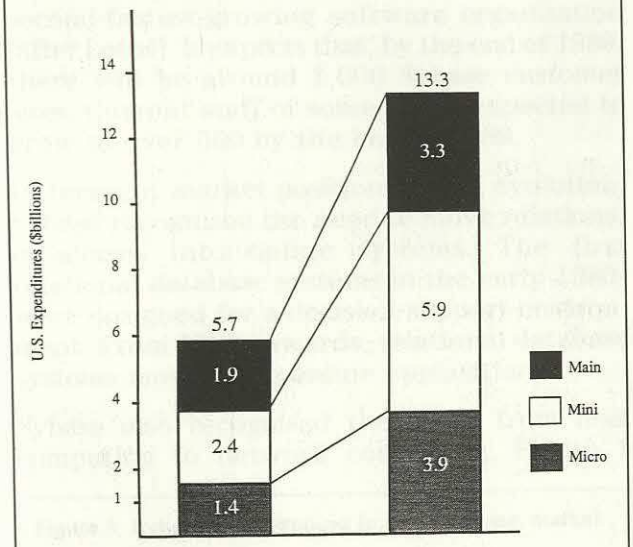

Relationships with other companies and equipment

NLI has formed strategic alliances with major software vendors (including Oracle, Relational Technology, Sybase, Informix, and Microsoft), and with hardware vendors, including Sun, Digital Equipment, IBM, and Hewlett-Packard.

Currently, NLI products support 11 operating-system environments, including MVS, VM, VMS, OS/2, and various versions of Unix. They support nine relational database management systems, including Digital's RDB, Ingres, Informix, IBM's DB2, Oracle, and Sybase. They support five communication network protocols, including asynchronous, SNA LU6.2, TCP/IP, and DecNet. They also support ten different ranges of hardware, including Apple Macintosh, Digital VAX, Hewlett-Packard 9000, IBM 370, IBM PS2, and Sun.

The company

NLI is a private company, funded with venture capital. It was founded five years ago, and now has a staff of 47 (expected to increase to 65 by the end of 1989). It started selling products during 1988, and currently has 50 'blue chip' customers, mainly in North America. NLI has considered the possibility of creating foreign-language versions of its products, but it would need to rewrite at least one-third of the program logic, and would consider doing this only as a joint venture with foreign partners. It has also considered the possibility of linking its Natural Language query tools with speech-recognition products (such as those demonstrated during the Study Tour by Speech Systems Inc). NLI has five offices in the United States, and has also appointed a distributor for Europe.

Figure 3 1992 potential market

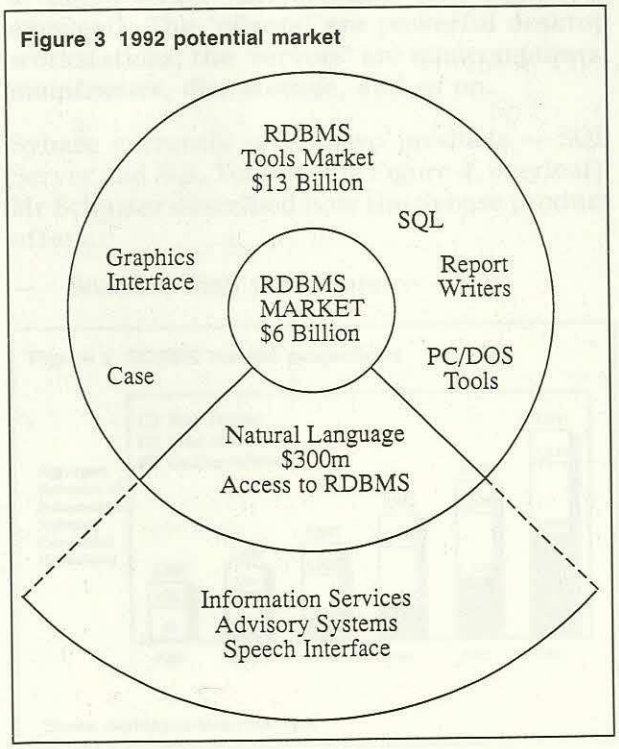

Demonstration of Natural Language

NLI demonstrated its Natural Language product, using a simple database of personnel data on a Compaq PC. When a simple request is typed in, the system translates that request into terms that it can understand. The translated request is displayed so that the user can confirm that this is exactly what he wants. The system will then execute the request. If the system receives a request that it does not understand, it requests more information. If it receives a request containing a word that it cannot recognise, it asks for clarification. It is possible for the user to create a dictionary of synonyms for the words that the system does understand. All of the requests and all of the answers are presented in natural, conversational English. The system is therefore extremely simple to understand and to use.

Sybase

The Sybase presentations were given by Mark Hoffman, President, Stu Schuster, Vice-President of Marketing, and Bob Epstein, Executive Vice-President.

Mr Hoffman began by positioning Sybase in the market. He described it as the leading vendor of high-performance relational database systems (RDBMSs) for online applications, using a client/server architecture. Sybase was founded in 1984 with venture-capital funding. Production shipments began in May 1987. It is the second-fastest-growing software organisation (after Lotus). It expects that, by the end of 1989, there will be around 1,000 Sybase customer sites. Current staff of some 300 is expected to grow to over 550 by the end of 1989.

In terms of market positioning and evolution, Sybase recognised the need to move relational databases into online systems. The first relational database systems in the early 1980s were designed for a decision-support environment. From 1986 onwards, relational database systems moved into online applications.

Sybase also recognised the trend from host computing to network computing. Figure 1 shows a schematic of Sybase's positioning in the database market, and Figure 2 shows the growth and relative positioning of different DBMSs to 1993. Sybase's objective is to achieve one-third of the 1993 market for online RDBMSs, forecast to reach $1,900 million by that time. Its distribution strategy is based mostly on direct sales, which account for some 70 per cent of total sales. OEMs, international distributors, and value-added retailers account for 10 per cent each.

Figure 1 Sybase's positioning in the database market

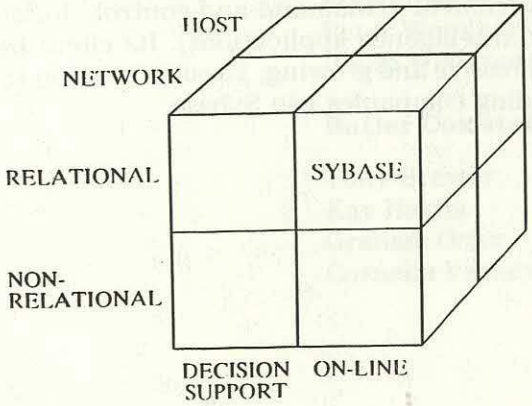

Sybase has strategic alliances with companies like Microsoft, Ashton Tate, Apple, Sun, Status, NeXT, Pyramid Technologies, and TRW. These vendors are building interfaces to the Sybase products.

Sybase's products are based on the concept of a 'client/server' architecture (see Figure 3, overleaf). The 'clients' are powerful desktop workstations; the 'servers' are minicomputers, mainframes, disc storage, and so on.

Sybase currently offers two products — SQL Server and SQL Toolset (see Figure 4, overleaf). Mr Schuster described how the Sybase product offers:

— Scalable high performance.

Figure 2 RDBMS market projections

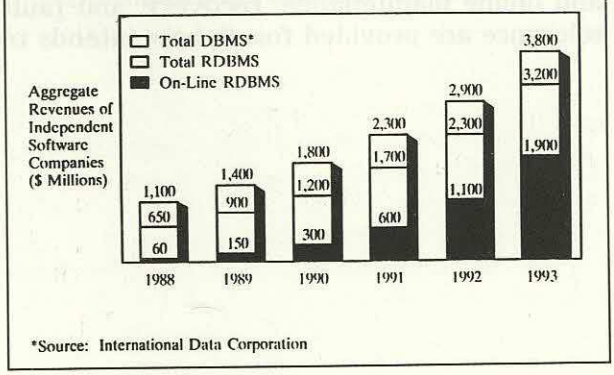

*Source: International Data Corporation

Figure 3 Sybase client/server architecture

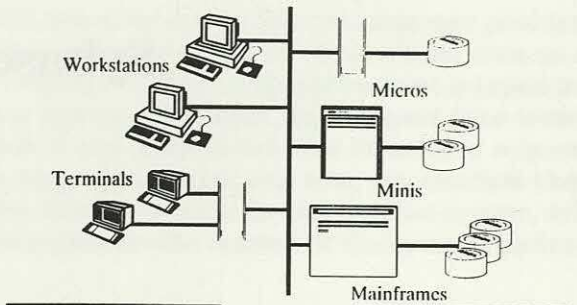

Figure 4 Sybase client/server products

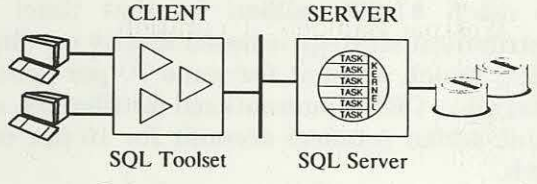

- Window-Based
- Decision Support
- Application Development
- Open Tools Interface
- Performance
- Integrity/Security
- Availability
- Open Distributed

— Server-enforced integrity.

— Multi-level security.

— High availability.

— Open distribution.

— Window-based tools.

Sybase products can perform in a host computer as well as in a networked-computing environment.

Current performance is five transactions per second per mips. Integrity is server-enforced, and online maintenance, recovery, and fault tolerance are provided for. Sybase intends to provide the capability to integrate its own workstation presentation with the workstation interface being used (for example, DEC Window, OS2/Presentation Manager, OSF/Motif).

Sybase's vision of the future software world is based on:

— Unix (both OSF and Unix International).

— OS2.

— VMS.

— IBM (SAA).

The move towards 'client/server' computing that Sybase foresees will mean many changes. It will not simply be a matter of adding a network to a system, but will involve separating the applications into their logical components. Instead of system programmers and application programmers, there will be client programmers and server programmers. It also means that RDBMSs must keep the application level simple, and hide the complexity in the network. Transaction control must be moved out of the application. Client computing will be driven by users who are unlikely to be persuaded to adopt a standard they do not want, and are likely to wish to continue to use what they know. Server computing will be driven by the more quantifiable issues of performance, reliability, and connectability.

Sybase's target markets are financial services (asset-management, security-trading, and customer-service applications), manufacturing and distribution (CIM, inventory-control, and point-of-sale applications), telecommunications (network-management, problem-tracking, and installation-scheduling applications), and government (command-and-control, logistics, and intelligence applications). Its client list is impressive and growing; 13 out of the top 15 US trading companies use Sybase.

1989 Study Tour of the United States

List of delegates

David Batts	RHM Computing
Jo Connell	FI Group
Bernard Gautier	Banque National de Paris
Raffaele Giordano	Senato della Repubblica
Alex Harris	Westpac Banking Corporation
Ian Harris	Qantas Airways
Paul Johnston	Inland Revenue
Roger Jones	ISTEL
Jean-Paul Kennis	Régie Renault
Ian Lawrenson	Department of Trade and Industry
Jean-Louis Lunel	Total — CFP
Warner Manning	Hong Kong Bank
Reen van Marion	Koninklijke Ahold NV
Christopher Robinson	ISTEL
Patrick Roussel	EDF/GDF
Jean-Pierre de Ruddere	Digital Equipment NV/SA
Henk van Schie	Nederlandse Dagbladunie BV
Martin Shelton	British Coal Corporation
Uday Shetgiri	Citicorp Overseas Software
Derek Shute	Royal Insurance
Janick Taillandier	RATP
Colin Talbot	Lloyd's of London
Tony Tomlinson	Citibank NA
Geoge Vogel	Swiss Aluminium
Fred de Waard	RAET NV Software & Computer Services
Iain Watson	Ministry of Defence
Stewart Wilson	Rolls-Royce
Ronald Winkler	Dalgety
Peter Woodhart	Southern Electricity
Philip Woodnutt	Burton Group Financial Services

Butler Cox staff

Tony Brewer
Kay Harris
Graham Otter
Cornelia Varney

BUTLER COX
FOUNDATION

Butler Cox

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

Objectives of the Foundation

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

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

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

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

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

Membership of the Foundation

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

The Foundation research programme

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

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

The report series

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

Butler Cox plc
Butler Cox House, 12 Bloomsbury Square,
London WC1A 2LL, England
☎ (01) 831 0101, Telex 8813717 BUTCOX G
Fax (01) 831 6250

Belgium and the Netherlands
Butler Cox BV
Burg Hogguerstraat 791,
1064 EB Amsterdam, the Netherlands
☎ (020) 139955, Fax (020) 131157

France
Butler Cox SARL
Tour Akzo, 164 Rue Ambroise Croizat,
93204 St Denis-Cédex 1, France
☎ (1) 48.20.61.64, Télécopieur (1) 48.20.72.58

Germany (FR)
Butler Cox GmbH
Richard-Wagner-Str. 13, 8000 München 2, West Germany
☎ (089) 5 23 40 01, Fax (089) 5 23 35 15

United States of America
Butler Cox Inc.
150 East 58th Street, New York, NY 10155, USA
☎ (212) 891 8188

Australia and New Zealand
Mr J Cooper
Butler Cox Foundation
Level 10, 70 Pitt Street, Sydney, NSW 2000, Australia
☎ (02) 223 6922, Fax (02) 223 6997

Finland
TT-Innovation Oy
Meritullinkatu 33, SF-00170 Helsinki, Finland
☎ (90) 135 1533, Fax (90) 135 1091

Ireland
SD Consulting
72 Merrion Square, Dublin 2, Ireland
☎ (01) 766088/762501, Telex 31077 EI,
Fax (01) 767945

Italy
RSO Futura Srl
Via Leopardi 1, 20123 Milano, Italy
☎ (02) 720 00 583, Fax (02) 806 800

The Nordic Region
Statskonsult AB
Stora Varvsgatan 1, 21120 Malmo, Sweden
☎ (040) 1030 40, Telex 12754 SINTABS

Spain
Associated Management Consultants Spain SA
Rosalía de Castro, 84-2°D, 28035 Madrid, Spain
☎ (91) 723 0995