Management Conference
Session Summaries

BUTLER COX
FOUNDATION

Managing High-Technology Projects
Bournemouth, 7-9 June 1987

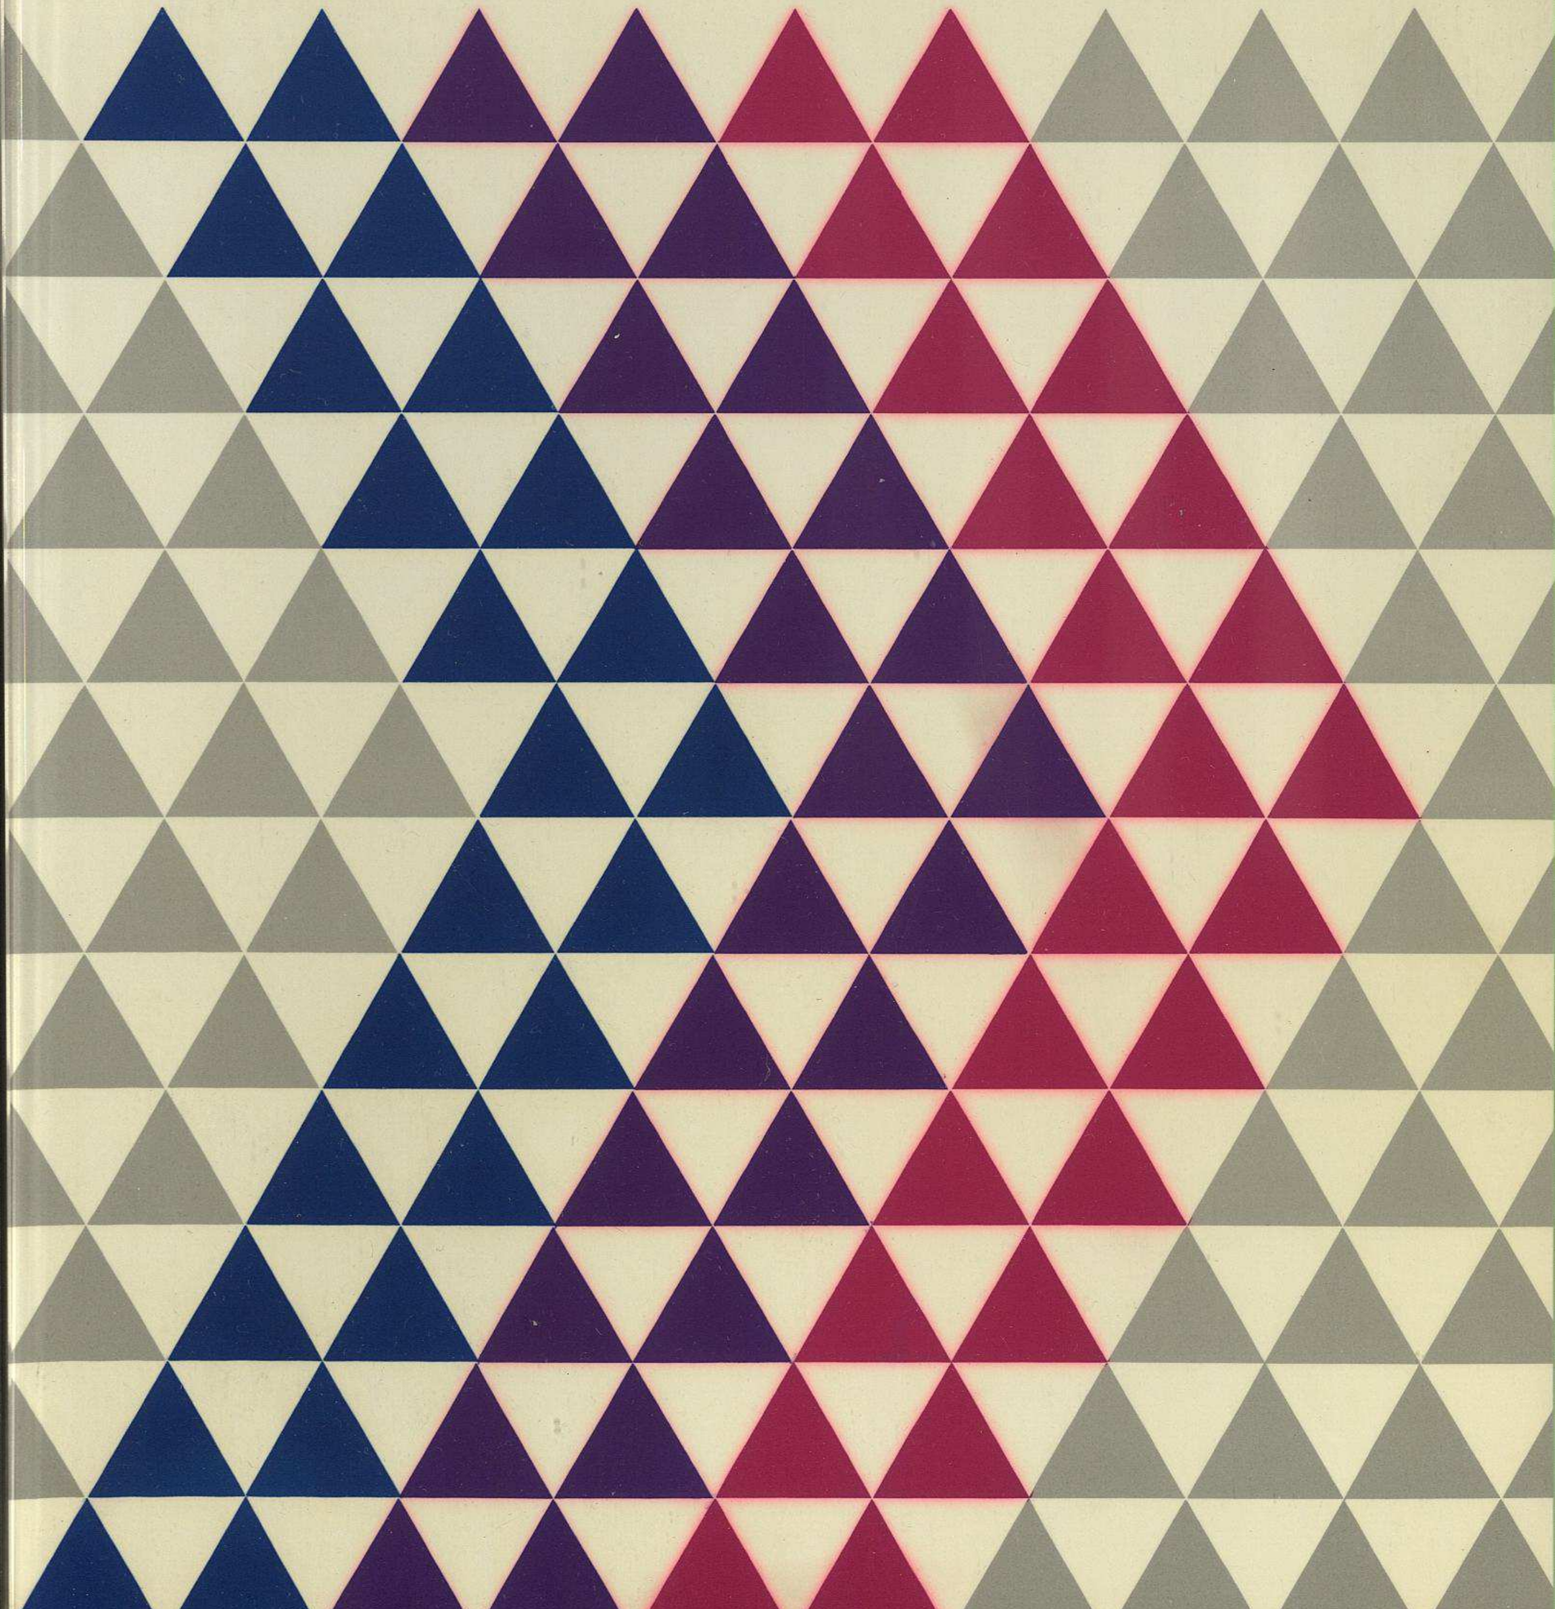

BUTLER COX
FOUNDATION

Managing High-Technology Projects

Management Conference
The Dormy Hotel, Bournemouth, 7-9 June 1987

INTRODUCTION

The 1987 UK Foundation Conference was held at the Dormy Hotel, Bournemouth, between 7 and 9 June. The conference focused on the issues associated with managing high-technology projects. This document contains summaries of the presentations made at the conference.

The summaries were prepared by Butler Cox consultants during the conference and are intended as an aide-memoire. They are not a verbatim transcript but present, as faithfully as possible, an interpretation of the main points made by each speaker. For the sake of brevity, some points have necessarily been condensed or omitted.

Where appropriate, the summaries include a selection of the visual aids used by the speakers.

CONFERENCE SUMMARY

High-technology projects use technology that is new and untried. They are therefore much more risky than conventional data processing projects so they should only be considered if the potential payoff is much higher than usual. The factors that have to be considered if they are to be managed successfully are also very different from those in conventional projects, so the project has to be designed much more carefully than usual. An inappropriate project design will guarantee failure but a suitable design can greatly increase the chance of success. Conventional project-control methods are not appropriate and different methods and management styles must be developed to suit the particular requirements of the project.

Four aspects of information technology that are developing rapidly, and which hold the promise of high payoffs, were examined. These were parallel computing architectures, expert systems, speech processing, and image processing. Most notably, Professor Feigenbaum gave several examples of expert systems that have been developed very quickly and inexpensively but have delivered tremendous benefits. Three interesting case histories were also described, each illustrating various aspects of high-technology project management.

This conference revealed the possibilities associated with harnessing emerging technology, indicated some of the difficulties of managing such projects successfully, but also gave practical advice on how to overcome them.

BUTLER COX
FOUNDATION

Managing High-Technology Projects

Management Conference
The Dormy Hotel, Bournemouth, 7-9 June 1987

Contents

Printed in Great Britain by Flexiprint Ltd., Lancing, Sussex.

Keynote Address

Roger Woolfe, Butler Cox

Roger Woolfe is director of group consultancy at Butler Cox. He began by defining high-technology projects and contrasting them with their more conventional counterparts, and then went on to examine how businesses can identify and exploit high technologies to their advantage.

High-technology projects are innovative and risky. They involve technology that is relatively new and untried, but that promises a significant business advantage. New technologies of this sort abound. They include supercomputers, expert systems, optical storage, mobile data communications, and cordless business telephones. None of these technologies are blue sky, in the sense of being years away from normal service. Instead, they are relatively close to hand. Yet all are immature. To explain the concept of maturity, Roger Woolfe described the well-known product life-cycle curve (see Figure 1). Here, the curve passes through four stages of maturity, from embryonic to base. The maturity of a technology is not absolute, stressed Roger Woolfe. Instead, it should be seen in the context of a particular industry sector. A technology that is pacing in one sector may be key in another. Therefore, in examining and predicting the impact of a technology on a specific business, it is important to do so within the context of the sector.

Figure 1

TECHNOLOGICAL MATURITY

As they advance along the product life cycle curve, generic information technologies pass through stages:

1 embryonic : still in the lab
2 pacing : field trials
3 key : commercial service amongst pioneers
4 base : widespread commercial service

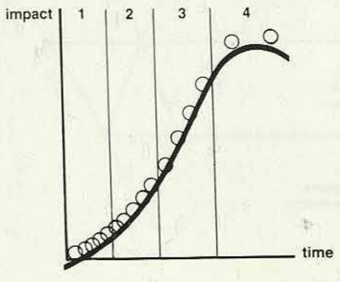

Leading-edge technologies are in one of the earlier stages in the product life cycle in a given sector. There are many examples of successful applications of leading-edge technologies. As an illustration, Roger Woolfe quoted several, taken from different industry sectors: fashion clothing, engineering tooling, catalogue shoppers, wholesale groceries, and letter distribution.

A case in point was Benetton, the Italian fashion clothes-store business, whose success is in large measure due to close coupling between demand and production. Benetton uses in-store POS terminals linked to computer-controlled knitting and cloth-cutting machines in the central factory, and robots for item handling in the warehouse, which is claimed to be one of the world's most automated.

High-technology projects are risky, however. An example was Federal Express's Zapmail system, which has recently been withdrawn. Zapmail was an advanced letter-distribution system spanning the United States, using digital facsimile tranceivers and Federal Express's own high-speed data network. Although the technological problems were resolved, Federal Express's market predictions turned out to be unrealistically optimistic.

Reviewing the experience of the MIS population as a whole, Roger Woolfe noted that instances of successful exploitation of high technology were in fact quite rare. ''Most companies'', he noted, ''relegate IT to providing a reactive service for middle management'' (see Figure 2). The opportunities for exploiting IT were permeating the whole of the business value chain, however. Explaining Porter's value-chain analysis (Figure 3) he noted how information technologies could be mapped on to all nine cells in the chain (Figure 4), a situation that did not apply even a few years ago.

How does a business spot technologies when they are still embryonic, yet are likely to have an important influence on the business sector later on? One approach is the advanced technology group (ATG). The ATG is a small, dedicated department within the MIS department that specialises in looking over the horizon at emerging technologies.

Figure 2

SUCCESS STORIES ARE RARE

"Most companies relegate IT to providing a reactive service for middle management"

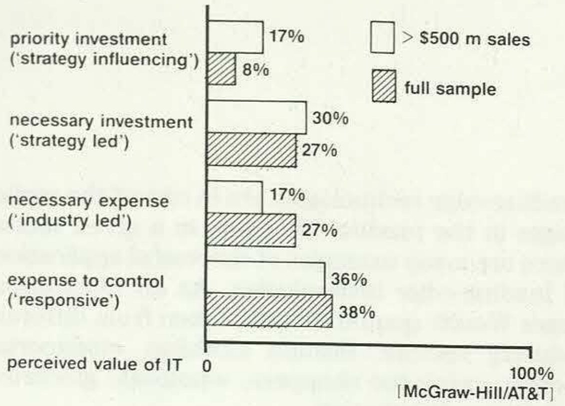

[McGraw-Hill/AT&T]

Figure 3

APPLICATION OPPORTUNITIES ARE BROADENING – 1

IT opportunities permeate the business value chain:

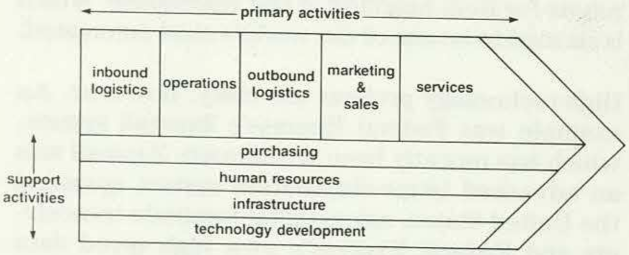

Figure 4

APPLICATION OPPORTUNITIES ARE BROADENING – 2

IT opportunities permeate the business value chain:

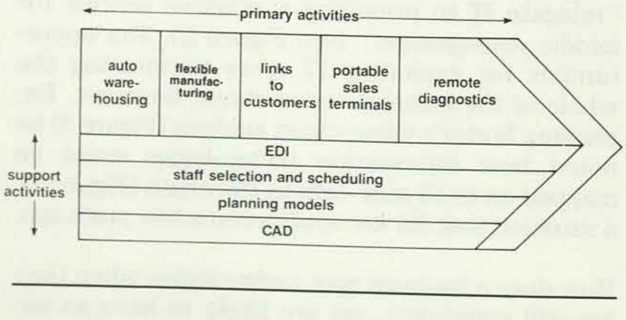

Managed with a loose, entrepreneurial style, ATGs have time to look beyond end users' immediate needs. They forecast and experiment with new technologies, often taking them to the point at which a business case can be made for their exploitation. A problem with ATGs, however, is the guru syndrome: it is hard in practice to avoid too much emphasis on technology spotting and too little on business requirements. What is needed is a balance of the two. Ideally, each business should follow the simple steps set out in figure 5. Technology spotting is one step, but only one.

An alternative to the ATG is the use of outside consultants. As an illustration Roger Woolfe explained the ways in which Butler Cox's own consultants have assisted companies to cover the steps outlined in Figure 5, using an approach bringing together business requirements and information technologies at an intensive management-decision workshop (see Figure 6). There is a preworkshop stage at which the consultants filter the technologies likely to be most pertinent, and familiarise themselves with the business sector and competitive position, as well as the specific strengths and weaknesses, of the company in question.

Figure 5

STEPS IN THE PROCESS

"Every pound spent when you have competitive advantage puts you further ahead, while others have to spend just to stand still."

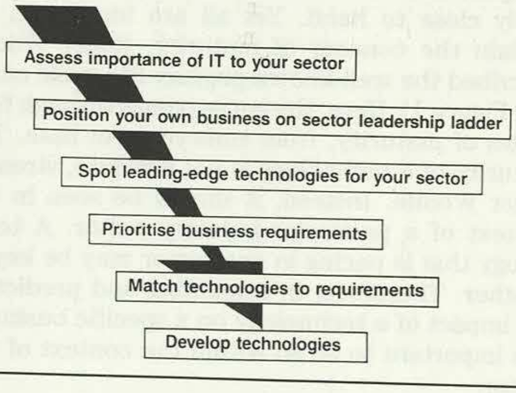

Figure 6

SPOTTING LEADING-EDGE TECHNOLOGIES

An alternative to the ATG: Butler Cox's management decision workshop

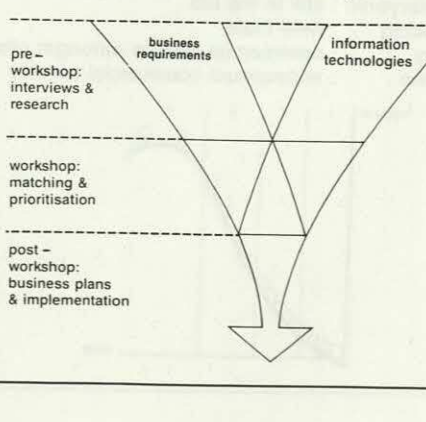

The decision workshop projects the impact of new technologies on the sector, prioritises business requirements, and matches the two together. An important component is that of technology forecasting for the sector as a whole. This is an exercise in technology intercept, and the outcome of a typical exercise is illustrated by the technology-take-up chart in Figure 7. What matters for an individual business is for it to recognise the technologies that will become pacing and key for their sector in the next few years, and then to position itself on the take-up chart. After the workshop, action plans are hammered out, leading first to the generation of a business case and finally to technology implementation. Butler Cox has conducted management-decision workshops of this sort over a wide range of sectors including retailing, wholesaling, pharmaceuticals, distribution, manufacturing, and banking. "It is", said Roger Woolfe, "most gratifying to see the results of these efforts, where companies we have worked with exploit new technologies that have been examined and matched in workshops we were involved with".

Figure 7

TECHNOLOGY TAKE UP

The chart predicts technology take up within a sector

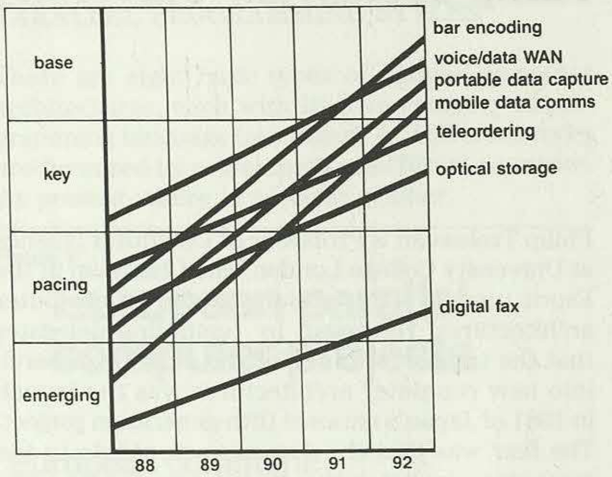

Roger Woolfe summarised by making four points: New technologies are proliferating — the rate at which they are emerging is accelerating; businesses are becoming more IT intensive; technology spotting is therefore of greater importance than in the past; and, finally, formal means do exist for organising technology spotting and exploiting new technologies to advantage.

A number of management lessons could be drawn, and they were summarised in separate slides. The risk element could not be escaped, of course. But, in the words of Peter Drucker "In business, taking no risks is often the biggest risk of all".

Advanced Computer Architectures

Philip Treleaven, University College, London

Philip Treleaven is Professor of Computer Science at University College London, and Chairman of the Esprit (project 415) working group on computer architectures. He began by reminding delegates that the trigger for much of the current research into new computer architectures was the launch in 1981 of Japan's national fifth-generation project. The fear was that the Japanese would do to the computer market what they have done to the traditional watch-making and the motor-cycle industries. Other countries responded by setting up their own government-funded programmes (see Figure 1). Many of these programmes are aimed at using artificial-intelligence (AI) technologies in new generations of 'intelligent' electronic consumer products. Research is concentrated in four areas — new types of computers, expert systems (which are seen as the dominant application area of the 1990s), the human-computer interface (natural-language input/output, pictures, etc), and the design of small, powerful integrated circuits.

Figure 1

International research programmes

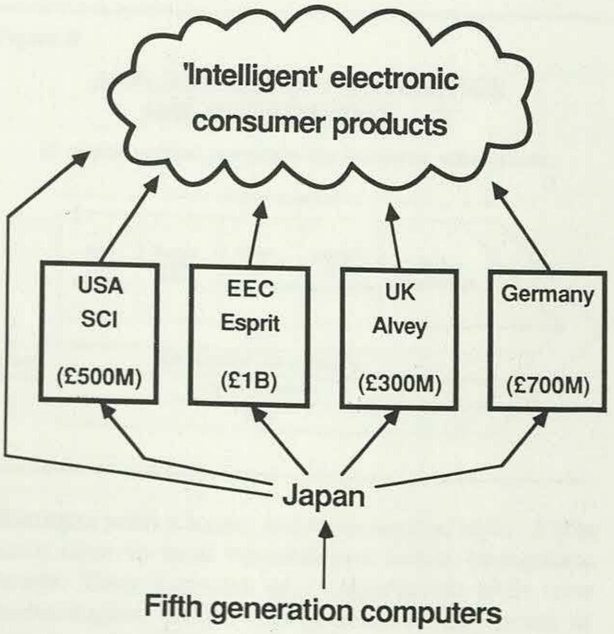

Figure 2 shows the five levels in a fifth-generation architecture. At the top is the applications level. One example might be an intelligent television that searches for a channel currently transmitting a news programme, or a television that can translate a foreign-language programme in realtime. Below the applications level will be the software components, typically expert systems providing a high-level of sophistication. In turn these will require new programming languages (level 3) such as Lisp

Figure 2

Fifth generation computing

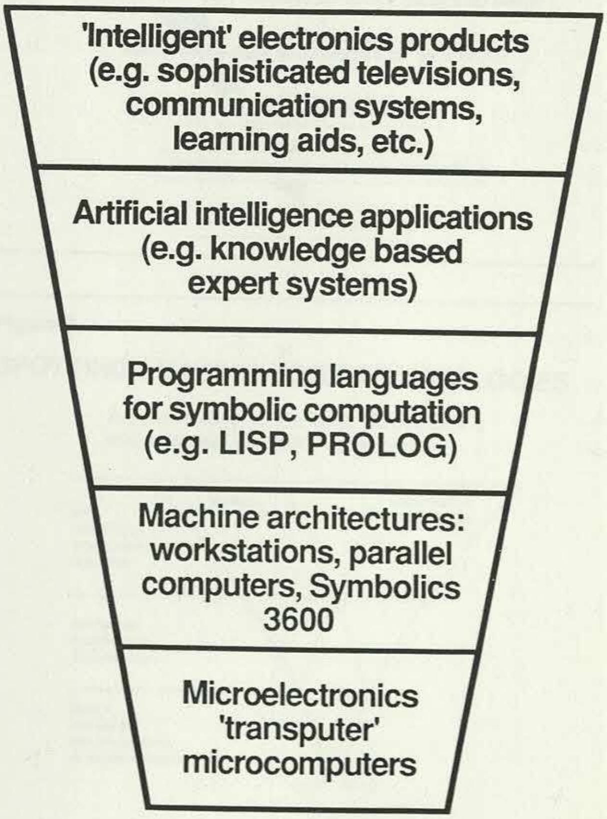

and Prolog, which will allow parallelism to be expressed, and will allow networks of microcomputers to be programmed.

The fourth level is the machine architecture. AI technologies require considerable computing power and they require parallel architectures. In turn this requires new basic microelectronics structures such as the Transputer. The structure of the chips must match the structure of the applications.

The current research interest in new computer architectures is being driven by the emerging needs to process symbols, speech, and pictures, and the trend away from conventional sequential processing to parallel and distributed processing, where several processors are all working on parts of the same task. The conventional machine architecture is now 40 years old, and most of the complexity today exists in the software that has to overcome the constraints of the architecture. The research community believes that the time has come to replace conventional hardware with something better that can support today's software. The vision is that parallel processors will consist of thousands of processors. The processors will be able to communicate with all other computers in the organisation and, eventually, in the world. The worldwide computing network will then become like the worldwide telephone network today.

Figure 3 lists the major European parallel computer projects. The Esprit project 415 is the largest, with $32 million funding and involves six industrial companies (Philips, AEG, GEC, Bull, CSECT, and Nixdorf) and support from various universities. Each company is investigating a particular type of parallel processing architecture (object-orientated, functional, logic/Prolog, logic/functional, and logic/connection-method/data flow). In the United Kingdom, the emphasis is on the so-called 'reduction' machines with the Flagship and Grip projects, which involve ICL. In Germany, the Suprenum project is aimed at building an inexpensive supercomputer out of thousands of inexpensive Motorola chips.

PARALLEL PROGRAMMING STYLES

There are eight basic types of parallel computer architectures, each with its own category of programming language (see Figure 4). Different styles are favoured by researchers in different countries. At present, there is no clear winner.

Figure 3

European parallel computer projects

European community

- **P415 – Parallel architectures and languages**
- **P1085 – Supernode**
- **P1588 – Span**

Britain

- **Flagship**
- **Grip**

Germany

- **Suprenum**

Figure 4 Parallel programming styles

Computer architectures							
Control-flow machines	Data-flow machines	Reduction machines	Actor machines	Logic machines	Rule-based machines	Connectionist machines	Input/output machines
Transputer	LAV	Alice	Apiary	ICOT PSI	Dado	Boltzmann	Clip 4
Programming languages							
Procedural	Single-assignment	Applicative	Object-oriented	Predicate logic	Production system	Semantic nets	Natural language, speech, etc.
ADA, OCCAM	ID, VAL	Pure Lisp	Smalltalk	Prolog	OPS5	ETL	English

CONTROL-FLOW MACHINES

Control-flow machines are associated with procedural languages such as ADA and OCCAM. In a control-flow computer (for example, Sequent Balance and Inmos Transputer) explicit flow(s) of control cause the execution of instructions. In a procedural language the basic concepts are a global memory of cells, assignment as the basic action, and (sequential) control structures for the execution of statements.

DATA-FLOW MACHINES

Much of the early work on data-flow machines was done at Manchester University, although there is now declining interest in this type of architecture. Data-flow machines are associated with single-assignment languages such as ID, Lucid, VAL, and Valid. In a data-flow computer the availability of input operands triggers the execution of the instruction that consumes the inputs. In a single-assignment language the basic concepts are that data 'flows' from one statement to another, execution of statements is data-driven, and identifiers obey the single-assignment rule.

REDUCTION MACHINES

Reduction machines are associated with applicative languages such as Pure Lisp, SASL, and FP. In reduction computers (for example, Alice and Grip) the requirement for a result triggers the execution of the instruction that will generate the value. In an applicative language the basic concepts are application of functions to structures, and all structures are expressions in the mathematical sense. There is a lot of interest in this type of machine in the United Kingdom. The Alvey Programme has "bet its shirt" on reduction machines.

ACTOR MACHINES

Actor machines are associated with object-oriented languages such as Smalltalk. In an actor computer (for example, Apiary) the arrival of a message for an instruction causes the instruction to execute. In an object-oriented language the basic concepts are that objects are viewed as active, they may contain state, and objects communicate by sending messages.

LOGIC MACHINES

Logic machines are associated with predicate-logic languages such as Prolog. In a logic computer (for example, Icot's PIM) an instruction is executed when it matches a target instruction. In a predicate-logic language the basic concepts are that statements are relations of a restricted form, and execution is a suitably controlled logical deduction from the statement. The Japanese are pursuing this path in their fifth-generation project.

RULE-BASED MACHINES

Rule-based machines are associated with production-systems languages such as OPS5. In rule-based computers (for example, Non-Von and Dado) an instruction is executed when its conditions match the contents of the working memory. In a production-system language the basic concepts are that statements are If . . . Then . . . rules and they are repeatedly executed until none of the If conditions are true. This is the path favoured by US researchers. Some very large machines are being built (64,000 processors).

CONNECTIONIST MACHINES

Connectionist machines are associated with semantic-net languages such as NETL. A connectionist computer (for example, Connection Machine) is based on the modelling of interneural connections in the brain. In a semantic-net language networks are used to define the connections between concepts, represented as structured objects.

INPUT/OUTPUT MACHINES

Input/output machines (such as Clip4) are associated with natural-language and speech processing.

PARALLEL COMPUTING APPLICATIONS

Researchers say that current research will lead to commercial parallel computing products by about 1995. However, there are already many commercial applications that require parallelism, and innovative start-up companies are already beginning to bring parallel-processing products to the marketplace. Commercial parallel-processing applications can be classified into four main groups:

- Parallel Unix systems (Such as the Sequence Balance 8000, ELXSI System 6400, and Encore Multimax).
- Fault-tolerant systems (Tandem NonStop TXP, AT&T 3B20D, Apache, Stratus, etc).
- Database systems (Enmass Computer Corporation, Teradata database machine, etc).
- Communications systems (BBN's Butterfly, for example).

The characteristics of some commercially available parallel-processing systems are shown in Figure 5. However, all of the suppliers readily admit that they do not yet know how to program parallel systems effectively. The dominant group of products is parallel Unix systems, where many start-up companies have realised that they can bring their parallel-processing machines to the market 'underneath' Unix, exploiting the inherent parallel-

Figure 5

Parallel transaction processing systems

Company	Price	Parallelism	Memory	Processor
Alliant FX18	$270,000	2-20	Global (64 Mb max)	64 bit CMOS Gate array
ELXSI 6400	$600,000	2-12	Global + local (800 Mb max)	64 bit ECL Gate array
Encore Multi Max	$114,000	2-20	Global + local (32 Mb max)	32 bit NS 32032
Flexible Flex 32	$115,000	2-2480 (20 per box)	Global + local (64 Mb/cab)	32 bit NS 32032
Sequent Balance 8000	$60,000	2-12	Global + local (28 Mb max)	32 bit NS 32032
Sequoia Sequoia	$200,000	2-64	Global + local (252 Mb max)	16 bit MC 68010

Figure 6

Sequent (UNIX-based) balance 8000

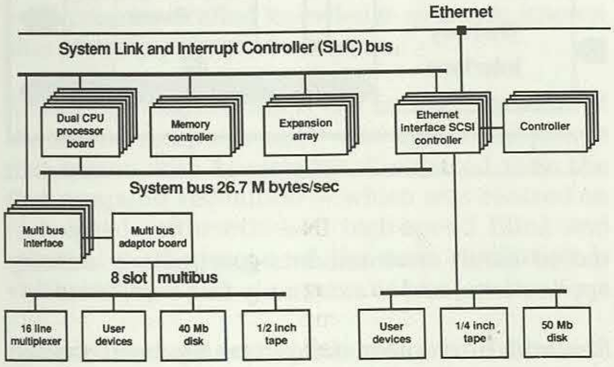

Figure 7

Sequent balance 8000 - processor

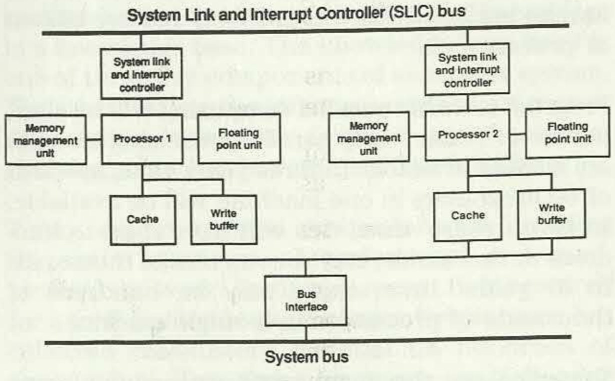

ism of Unix. The market leader is Sequent; the Balance 8000 architecture is shown in Figure 6, and the structure of the Balance 8000 processor is shown in Figure 7.

Professor Treleaven believes that, in general, parallel computers will be used in three main areas:

- To provide high-performance computing (supercomputers).
- To provide a 'graceful' growth path, by allowing additional processors to be plugged-in as required. He believes that this will become a more important area than high-performance.
- To provide greater reliability.

The traditional application area has been in high-performance numerical supercomputers, such as the Cray vector computer and array processors. However, there have been some spectacular failures, despite the large amounts of money that have been spent. An emerging area is that of MIMD (multi-instruction, multi-data stream) computers, where hundreds of thousands of high-performance chips will be plugged together to form a fully parallel architecture. Early versions of such machines are available today.

FIFTH-GENERATION MACHINES

Researchers in different parts of the world are pursuing different architectures depending on the particular type of language they favour. In the United Kingdom and Europe, the emphasis is on functional computers where hundreds of microcomputers are connected to support Lisp. Examples include the GMD reduction machine, Imperial College's Alice, and University College's Grip. In Japan, the emphasis is on logic machines and Prolog. Most Japanese universities are building such machines. Examples include Tokyo University's PIE (Parallel Information Engineering), ICOT's PIM-Rand and PIM-D, and Fujitsu's Kabu-Wake. In Europe, Bull's MultiSchuss project is using relational-database machine technologies to support the execution of Prolog.

In the United States the emphasis is on knowledge-base machines (Columbia University's Dado and Non-Von, Fairchild's FAIM-l, and MIT's connection machine). Other researchers are developing object-oriented computers (for example, Philips' Doom, Keio University's Zoom, and Berkeley's Soar).

Many of these developments share common attributes, particularly in the way they solve the problems of memory management (see Figure 8).

Professor Treleaven concluded this part of his presentation by working through an example of the way Prolog statements are interpreted, demonstrating that it is very difficult to do this with sequential processing. Hence the need for new parallel architectures.

Figure 8

Computer architecture mechanisms for symbolic computation

- **Structured memory:**

  Dynamic list-structured memory (eg Lisp nodes + garbage collection)

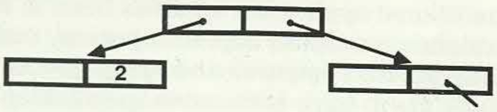

- **Tagging:**

  Associating a type 'tag' with each item of data

- **Control stacks:**

  Stacks for holding partial results and control information during execution of a program

Figure 9

Transputer chip

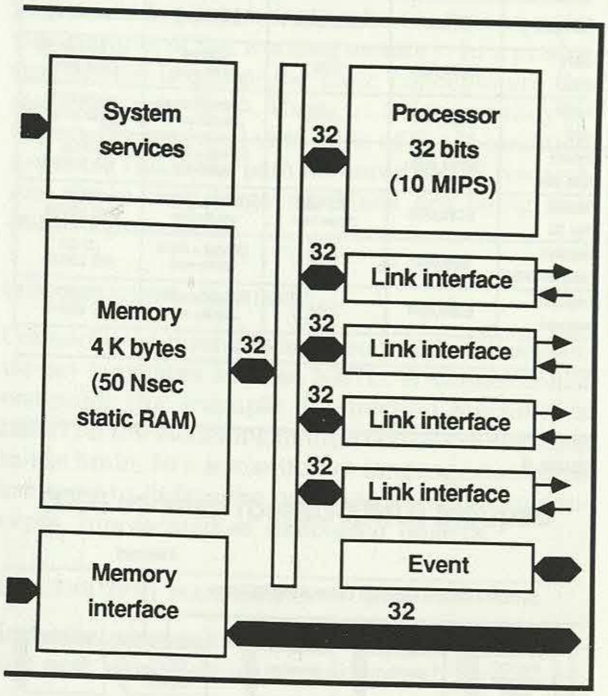

PARALLEL MICROCOMPUTERS

Most semiconductor manufacturers are working towards placing many simple interconnected microcomputers onto a single chip because they see parallel processing being the next big market opportunity. Moreover, making ever-more powerful single-processor chips is self-defeating; the larger chip size slows down the total time required to process an instruction because of the longer communications paths. A typical example is the Inmos Transputer (see Figure 9), which consists of primitive 32-bit processors operating at 10 mips, and links to the outside world.

NEW TECHNOLOGIES

Professor Treleaven concluded by identifying three technologies that will be at the front of computing developments:

— Gallium arsenide.

— Optical computers.

— Biological and molecular computers.

Gallium-arsenide chips are available today, but they will not supersede silicon for commercial products for the foreseeable future. Their use will be restricted to harsh environments (particularly military applications) and to extremely fast computers.

Research into optical computers has been boosted by the 'star-wars' programme. One advantage of optical computers is that photons do not interfere with each other, which means that it is easier to create very good parallel performance.

Research into biological and molecular computers is aimed at mimicking the performance of the human brain, which will probably require billions of processors.

Progress towards parallel computing will be slow, however. Today, early parallel-processing products are already available. In three years' time, upwards of 60 processors in one machine will be available. In seven years' time, this will have risen to hundreds or thousands, maybe even tens of thousands. In 15 years' time, there may be hundreds of thousands of processors in a single machine.

Nevertheless, the novel architectures now being researched will not have much of an impact on the types of parallel machines that will be available. The biggest impact will come through software developments. Optical computers and biological/molecular computers will not have an impact on commercial products for at least 20 years.

Expert Systems

Edward A Feigenbaum, Stanford University

Edward Feigenbaum is Professor of Computer Science at Stanford University, where he directs the Heuristic Programming Project, a leading laboratory for work on knowledge engineering and expert systems.

The principles and methods for building machines that reason spring from the science of artificial intelligence and its applications discipline known as knowledge engineering. Knowledge engineers build programs called knowledge systems, known alternatively as expert systems.

Knowledge engineering is the technology base of the second computer age — the age of computers that reason with knowledge. Compared with the first computer revolution — which was centred on high-speed arithmetic and high-speed filing and retrieval — the scope of the new revolution is vastly wider.

Knowledge engineers collaborate with experts in the various fields of professional work to represent for computer use two forms of knowledge. The first is factual knowledge, which is widely available, rigorous and (most often) undisputed. The second is heuristic knowledge, which is judgemental, nonrigorous, based on experience, and sometimes referred to as 'the art of good guessing'. Expert systems hold both of these two forms of knowledge in a knowledge base. The knowledge base itself is one of three key components of an expert system. The other two are the inference engine (which uses Aristotelian logic with forward and backward chaining), and the human/machine interface.

Today, expert systems applications are invading the entire spectrum of professional and semiprofessional work. During his recent field research for a forthcoming book, Professor Feigenbaum has collected case-history material on hundreds of applications. Ten times as many are following behind them, for the moment still at the prototype or field-testing stage.

Northwest Airlines' Seat Adviser expert system is a good illustration of a high payoff application now in regular use. The Seat Adviser effectively clones the expertise of the airline's seat-inventory controllers. Their job is constantly to change the allocation of discount seats on a given flight, in order to maximise revenue. The need for this has come about since deregulation; airlines want to sell as many seats as possible at full fare, but prefer the marginal revenue from late seat discounts to no revenue at all from empty seats. To perform the allocation, the inventory controller has to watch parameters such as airliner capacity, booking patterns, time to departure, and what the competition is doing. The benefit from optimum allocation is huge: just one seat improvement per flight can yield millions of dollars per year on the bottom line. With hundreds of daily flight departures, Northwest has insufficient inventory controllers to cover all flights.

Seat Adviser, developed as a collaborative effort by Northwest with Sperry (now Unisys) and Intellicorp, took just four weeks to design and implement. Now it is connected into the central seat-reservation system and used constantly as an aid by the inventory controllers.

Professor Feigenbaum went on to cite other instances where expert systems have successfully been put to commercial use. One was American Express's credit authorisation adviser. First installed in the company's Fort Lauderdale systems facility, it is now being implemented at other major sites across the United States to clone the expertise of human authorisation advisors. Every American Express card transaction over $50 (in the United States) has to be cleared before acceptance. Most clearance requests are filtered by mainframe computers using conventional techniques. The remainder go through to human credit authorisers who have in the past used a complex series of information displays (showing past useage history, payment records, late payments, and so forth) to assist with the decision process. Today, the expert system is able to emulate the human reasoning process, so helping to lift the burden from human shoulders.

A second illustration of expert systems at work comes from Northrop, the military aircraft builders. This company has developed an expert system for process planning — working out the shop-floor processes for piece-part manufacture. What it has done is to capture the expertise of a human master process-planner. The consequence is a reduction from 12 hours to 15 minutes in the time taken to produce a new process plan, and with no deterioration in the resulting process efficiency. Northrop is now extending the process-planning expert system to link with a computer-simulated model of the shop floor, which eventually should enable it to plan parts, and load and schedule the shop, in realtime.

Other applications are as diverse as: the diagnosis of machine faults, as well as human and plant disease; manufacturing planning and control; financial services' decision making; interpreting complex engineering data; military operations planning; and constructing high-rise buildings.

Professor Feigenbaum expanded on the benefits of expert systems. Benefits spring from a variety of sources. The most frequently observed is that of internal cost saving. Du Pont, for instance, has achieved a payoff of between $100,000 and $200,000 annually from just one month's investment. A second source is a radical increase in speed of working: improvements in the range 10:1 to 100:1 are commonplace. Other sources of benefit are improved quality of work, improved consistency of decision making, and new sources of revenue.

Professor Feigenbaum went on to provide advice on achieving success in planning and building corporate expert systems. He encapsulated his advice as follows:

1. Begin early. Do not miss the window of opportunity. Though the payoffs are large, there is a learning curve, making it self-defeating to "wait a while" before entering.
2. The "many pebbles in the pond" approach contrasts with the "one large rock makes bigger waves" approach. Both are used; both have been successful. So far, the United Kingdom has adopted the first approach.
3. Select initial problems carefully; concentrate on those that are knowledge-intensive and that push the limits of human ability to perform.
4. Problems must be chosen where significant bodies of human expertise exist, the experts are available, and the problem area is narrowly defined. A leading construction company, for instance, was extremely dependant on one, highly experienced estimator. He was able to estimate both the time and the cost of major new building projects with great accuracy while they were still at the concept stage, thus strengthening the sales position of the company.

   To reduce its dependence on the key staff member, the company agreed to his working as an expert with a team charged with responsibility for capturing his expertise in a knowledge system.
5. The problem must have a substantial and definable payoff, estimated in advance if possible. Management must see the payoff before authorising the project.
6. Bring the ultimate end users into the process from the first moment. End users must feel that they 'own' the system. Convince end-user managers to pay for at least part of the development costs. People value what they have to pay for.
7. Study the system-integration issues of end users at an early stage. Should the application be based on mainframes, minicomputers, workstations, PCs? What databases will be required, and how will they be accessed?
8. Plan to deliver a working prototype (even of limited capability) early in the life of the project. Knowledge-engineering software-development environments are excellent for rapid prototyping. Buy this system-development software; do not build your own.
9. Do not buy cheap software; the more expensive packages that include training and substantial hand-holding will pay off better. You can always economise later at the implementation stage.

Finally, Professor Feigenbaum stressed the importance of managing people when building successful expert systems. He made three telling points:

- The most successful companies have had tireless, enthusiastic, persuasive, almost charismatic champions of the new technology. Find a champion, and give him or her some organisational freedom and resources.
- Your highly competent experts and your most competent computer specialists can do the work. You do not need trained specialists in artificial intelligence, although one or two do help the process. The others need training, which can be bought if necessary.
- Company executives near the top, including the CEO, are usually enthusiastic about expert systems. Young people at the lower levels are also enthusiastic. The inertia comes from middle managers, who are afraid to innovate. The most inertia is seen amongst DP/MIS managers. Maybe this will change now that IBM is active in expert systems.

Image Processing

John Stonham, Brunel University

Dr John Stonham is reader in electrical engineering at Brunel University where he is head of the twenty-strong pattern-recognition and image-processing research group.

Dr Stonham defined the scope of his presentation to include both 'image processing' and 'pattern recognition'. Image processing is a transformation that endeavours to normalise or convert an image into a standard format (see Figure 1 for an example). The input and output are both representations of an image. Other examples of image processing are contrast enhancement or noise reduction to reveal hidden detail.

Pattern recognition is a labelling operation whereby a set of images that contain some common, or recognisable information is identified uniquely in a meaningful way. For example, a handwritten figure '3' is identified as the number three which may then be used in an arithmetic operation.

A complete image-recognition system will consist of an image processor and pattern-recognition element linked together (see Figure 2). There is a trade-off between the relative performance of these two elements. More sophisticated image processing, resulting in a more normalised representation of the original image, simplifies the pattern-recognition process.

The main theme of Dr Stonham's talk was that image (and speech) recognition will be an important component of artificial intelligence (AI) systems. Earlier workers in the field, using the computing architectures and algorithmic approaches of the 1960s and 1970s were over-optimistic in expecting hearing and seeing machines to evolve naturally from their research. There is a more realistic view at the moment that image recognition is far from being a trivial task, and that some more revolutionary approach, or novel technology, will be required to solve the problem. But he was optimistic that commercial solutions would emerge in the future — without being too precise as to when.

He supported his optimism by describing the direction his own research has taken and the results he has achieved. He has used natural systems, such as the human brain, to stimulate ideas for a more novel approach. He illustrated his talk by showing a video of how the system he has developed performs in practice. The recording showed how the system could be trained quickly to recognise and distinguish between the faces of six of his research students — despite the fact that their facial expressions were changing, and their heads were moving in front of the video camera used as input to the system. The training of the system, and, recognition of the various students only took a matter of seconds for each student. His research system is now being developed commercially by Computer Recognition Systems of Wokingham.

IMAGE RECOGNITION IS AN IMPORTANT COMPONENT OF AI SYSTEMS

Dr Stonham contrasted three areas of technological evolution:

Figure 1

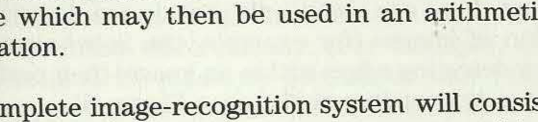

Figure 2 An image-recognition system

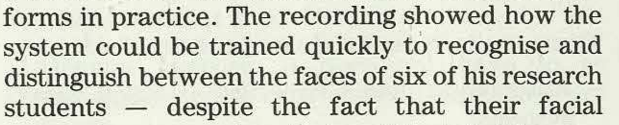

Figure 3

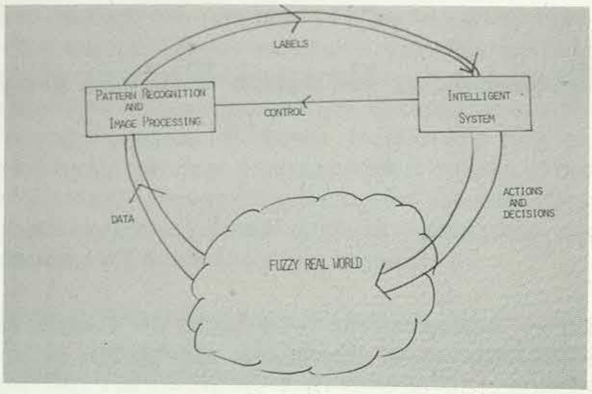

- The industrial era, in which manual tasks were mechanised.
- The computer era, in which routine clerical and calculation tasks were mechanised.
- The AI era, in which man's intellectual ability will be mechanised.

In both the industrial and computer eras, the advantages from mechanisation were immediately evident. The machine could out-perform the human. This is not so for AI — at least yet.

There are two essential components of AI:

- Automating man's expertise — using expert systems.
- Providing an interface to the real world for such systems (see Figure 3).

The problems of developing the interface systems for AI are certainly equal to, and probably greater than, those of developing expert systems. Nevertheless, it is important to solve such problems, as humans are increasingly the largest cost in making products and providing services.

THE PROBLEMS OF IMAGE RECOGNITION

Dr Stonham demonstrated the difficulties of image recognition by reference to the enormous numbers of options possible in any bit-mapped image, and the problems of identifying the key features of images by algorithms or formulae.

For example, if an image is divided into a 3 x 3 matrix of bits, there are 2^9 (or 512) possible patterns. If we increase the image to a 4 x 4 matrix, there are 2^{16} (or 64,000) possible patterns. The number of possible patterns increases exponentially as the size of the matrix gets bigger (see Figure 4). A monochrome television picture requires a matrix of a quarter-of-a-million cells to represent it, with an eight-bit word to represent the shade of grey in each cell. There are an unimaginable number of different patterns possible — 1 with 600,000 noughts after it. There is no way in which patterns can be recognised by matching them against templates, even using the fastest computers imaginable.Pattern recognition cannot be treated as a deterministic problem but rather one of matching a fuzzy representation of the real world to the images known by the system.

Although the text books may give the impression that there are algorithms that do aid the recognition of images (for example, the Sobel algorithm for detecting edges within an image) their performance is poor in practice (see Figure 5).

Today, image recognition is in a position similar to that of the automation of calculation prior to the arrival of the electronic computer. The theory of how to compute mechanically had been known from the time of Babbage. But it required the

Figure 4

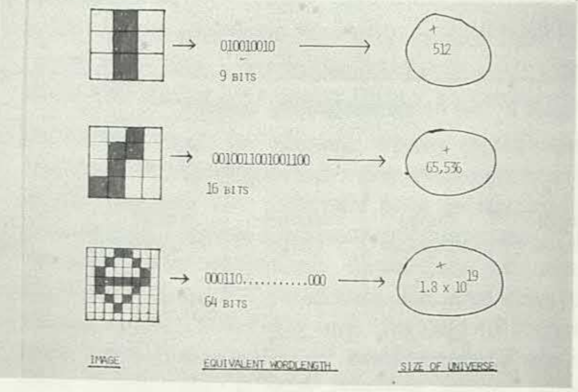

Figure 5

9. Experimental Example

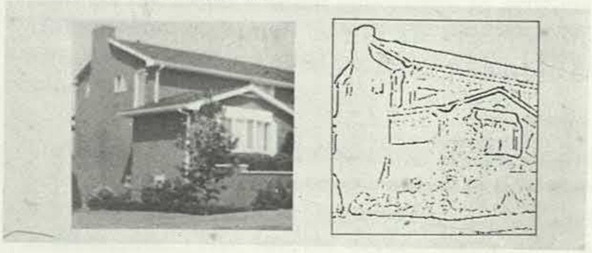

(a) Original image (256 levels)

(b) Processed result

development of electronics to enable the theory to be applied economically in practice. A new technology, more appropriate than traditional computing, is required to provide an economic and practical solution to image recognition.

WHAT IS THE SOLUTION TO THE IMAGE-RECOGNITION PROBLEMS?

Dr Stonham explained that his own attempts to solve the image-recognition problem had been stimulated by looking at the structure of natural systems such as the human brain. We have extremely good image-recognition systems in our brains. Our brains differ from computers in that they have a completely different architecture:

- They have a vast number of parallel processors (billions of neurons).
- They have very high interconnectivity, each neuron being connected to many other neurons and each having up to 5,000 different inputs.
- They do not have explicit storage of information comparable to a computer memory file.
- They do not have a predefined program. We learn to recognise images, rather than having a set of instructions on how to do so fed into our brain.

It would be absurd to try to replicate the brain electronically. But over the last twenty years or so, the elements of the brain structure (neurons) have been modelled electronically. (For example, the earliest models, made by McCulloch and Pitts, replicated the behaviour of small numbers of neurons (see Figure 6). Models of neuron networks can be 'trained' to respond to certain patterns of input bits and not others.

They do not have to be explicitly programmed with the features that characterise those patterns. They can memorise the patterns of interest and discriminate between them although not perfectly. The bit patterns they recognise can, of course, be generated from image scanners or video cameras.

Figure 6

THE MCCULLOCH & PITTS NEURON MODEL

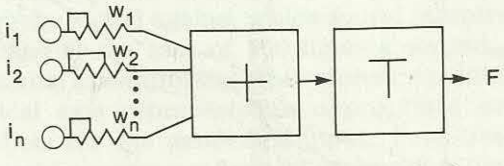

F = 1 if and only if

$i_1 = w_1 + i_2 = w_2 + \ldots \geq T$

After the early work on such models in the 1960s, interest in them declined as the component technology available did not enable the scale of the models to be increased easily. However, we can now build more complex networks — Dr Stonham's own experimental system can deal with a 512 x 512 element picture from a TV screen. It can recognise images, such as people's faces, once it has been trained by exposing the images to be recognised to it. The image need not be precisely the same for it to be recognised. The system will recognise a face that is sufficiently close to the one on which it has been trained — even if the image is moving (by small amounts) or has a different expression, or is distorted by (electronic) noise signals. Both the training and recognition processes can be accomplished in a matter of seconds for a human face.

Compared with the McCulloch and Pitts model of 15 years ago, which had just five neurons modelled on one printed-circuit board, the laboratory machine now has 64,000 functions on one board. The commercial version will be 16-times more powerful. Training and recognition should be accomplished in realtime.

The advantages of the system developed by Dr Stonham are that it is:

- General purpose; it does not have to be reprogrammed to deal with different kinds of image.
- Self-evolving.
- Capable of realtime operation, responding within 40 milliseconds, which is the time available to process normal TV images.

But its performance is not as good as a human being. In this, it differs from the advent of the machines and computers that heralded the industrial and computer ages. They were clearly seen to be more powerful than the people they replaced. But for many applications, better performance in image recognition will not be required. Less-than-perfect image recognition can still be used effectively . It can be applied to tackle the 80 per cent of problems that are easy rather than the 20 per cent that are difficult — leaving the human to concentrate on them. For example, if a system can be set up to inspect the components manufactured on a production line automatically, so that it can select the 90 per cent that are definitely perfect, then a manual inspection of the remaining 10 per cent will be more effective than inspecting a 10 per cent sample selected at random from the complete output. In the latter case, only some of the faulty components would be identified.

Speech Processing

Ken Davies, IBM Research Center

Ken Davies is manager of speech systems design at the IBM T J Watson Research Center, Yorktown Heights, where he is developing realtime dictation systems based on speech recognition.

ISOLATED WORD VERSUS CONTINUOUS SPEECH

Mr Davies started his talk by defining the four dimensions of speech recognition. The first is that of isolated words versus continuous speech. Figure 1 was used to illustrate the differences between the two recognition problems. The diagram represents two or three seconds of speech for the phrase ''I shall invite''. The top two patterns represent continuous speech, with the shaded (upper) position showing the pattern resulting from a digitised spectral analysis. The pulse (lower) representation is the speech pattern in the spoken sentence and is characterised by no gaps, or gaps which do not differentiate for example the pause in ''shall'' . The bottom two patterns represent the isolated words — the words can be distinguished by big gaps.

Figure 1

Isolated vs Continuous

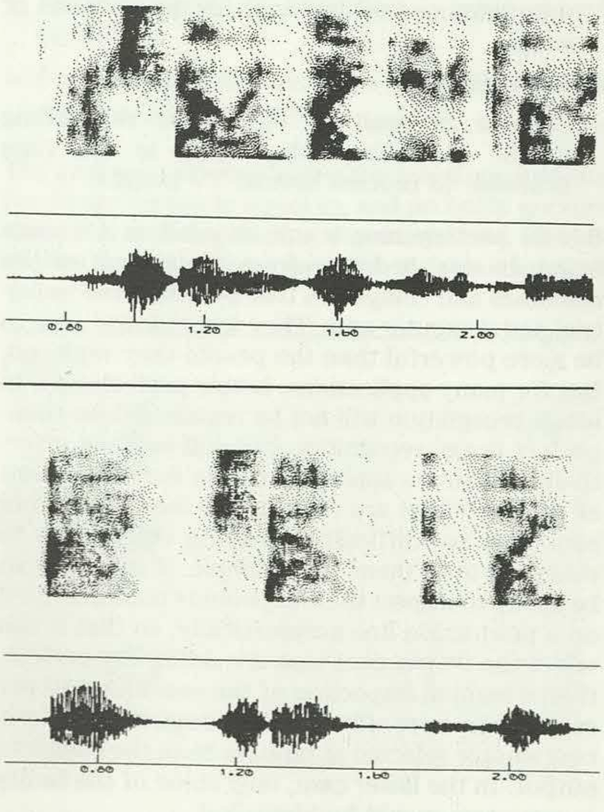

Mr Davies demonstrated some of the problems arising from loss of signal in sampling procedures and sound differentiation, and arising from voice and nasal interference.

VOCABULARY SIZE

The second dimension is vocabulary size. Small vocabularies are suitable for constrained tasks — for example, digits, commands, alphabet. Training can be achieved by word repetition. Large vocabu-

Figure 2

Dynamic Time Warping (DTW)

Isolated

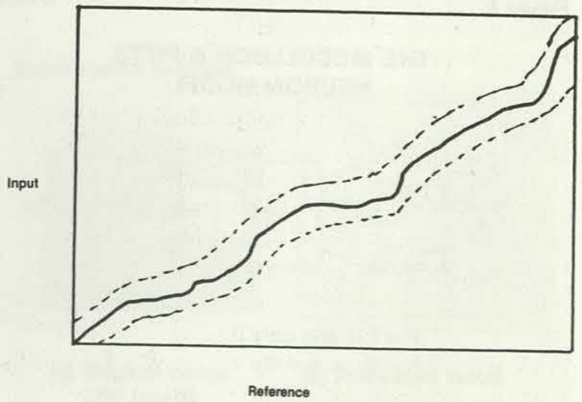

Figure 3

Hidden Markov Models

Isolated + Continuous

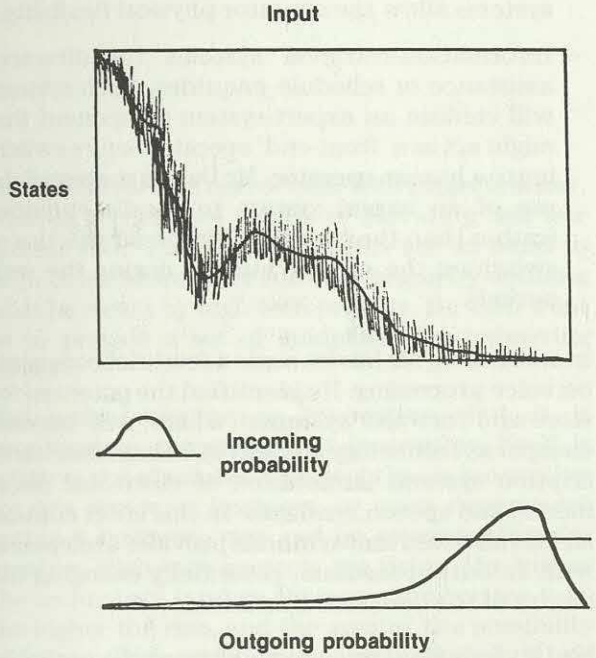

laries operate with no known constraints in applications such as the creation of documents. It is easy to exclude words but more difficult to determine what should be included. For business documents, phrases such as 'sulphur carbohydrate' can be excluded, but there is no clear definition of what is relevant.

RECOGNITION OF PHRASES

The third dimension is that of sound recognition using phrases rather than words. The two basic techniques — Dynamic Time Warping (DTW) and Hidden Markov Models (HMM) — are illustrated in Figures 2 and 3. Time warping is used for isolated word recognition and is the technique used in all commercial products today. Basically it uses a word reference model against which sound samples are matched at the rate of 100 times a second. The horizontal axis represents the reference points, the vertical axis represents an appropriate set of numbers for the generated input. Penalties are defined for variances from a line (running diagonally along the reference matrix). This 'cost function' is prescribed at the time the system is defined and is limited to one or two steps. There is no well-defined procedure for starting and finishing this process.

With Markov Models the expected word states are represented on the vertical axis and the input on the horizontal axis. The transition between states defines the word. A probability function is applied to determine the likelihood of a sound being in a succeeding state. There is a high probability for the region in the downward, diagonal direction. The outgoing probability model utilises a threshold for matching of a best score — it looks at the next sequence of sound. Probability is defined for a range of time and a range of values. Mr Davies contended that this was a better approach for continuous speech recognition.

He then discussed the measure of accuracy, arguing that most systems achieve approximately the same level. The TI database is used as a standard measure for small systems. This is applicable for small vocabularies, but can result in product 'skews', because a product vocabulary can be designed to pass the test. For large-vocabulary systems there is no widely agreed measure. Mr Davies referred to ways in which systems are tested. For test scripts of office correspondence, using six speakers, error rates were less than 5 per cent. Spontaneous creation of documents in this context did not appreciably worsen the error rate. However, with tests using text from different areas and modes of discourse, such as the New York Times, error rates of up to 24 per cent were experienced. In Mr Davies' view these results indicate that acoustic recognition alone is not sufficient and that a linguistic-context analysis is also required. For this purpose, IBM uses data derived from a statistical analysis of office documents.

SPEAKER ENVIRONMENT

The fourth dimension is that of the 'speaker' environment. All systems perform better in a quiet environment and Mr Davies speculated that speech-recognition systems will only perform well in single-person offices. Mr Davies contrasted the different performance capabilities of the human ear and speech-recognition techniques, a difference that is accentuated when speech is transmitted across telephone lines. The narrow frequency range (less than 3.5 KHz) on a telephone line hides many sounds. While people can infer sounds and meaning from the context, recognition techniques are obviously much more constrained. Sounds have to be inferred from bursts of energy, which is a highly error-prone process.

VOICE-RECOGNITION APPLICATIONS

Mr Davies cited electronic mail as an application with high potential for voice-recognition systems.

Before such applications can come into widespread use, application-engineering progress is required so that corrections can be made (to mistakes in the original recognition). Mr Davies described a system where corrections had to be spelled. Even with this potentially onerous approach, successful results were achieved.

To illustrate some working systems Mr Davies outlined briefly the history of his research group. From a small-vocabulary continuous system developed in 1978 the group has advanced to a 20,000 word system based on a PC-AT supported by five DSPs (array processors). He illustrated the system by use of a video demonstration in which a speaker input English sentences from an office context. Response was almost in realtime. However, for a Shakespearean sentence, computation and resolution took several seconds.

Mr Davies moved into the final phase of his talk by discussing the capability of current products. Without being vendor-specific, he noted that there is a wide range of products with vocabularies of between 100 and 200 words. He was more sceptical about the performance of products with more than 200 words in their vocabularies. All such systems are currently constrained to isolated word recognition for dependent speakers using a lip microphone.

The current emphasis for research and development is on continuous speech. In particular, the emphasis is to increase the accuracy and independence of systems in the context of telephone speech. Mr Davies mentioned the increase in activity in military areas, such as the development of speech systems for helicopters.

The areas with most immediate potential for applications were identified as:

- Text/data entry, where today's restricted alphabet, command, and numeric systems will be replaced by (human-oriented) words and phrases.
- Shop-floor systems where a head-microphone can be utilised. Examples are in the areas of quality control and materials handling, where speech systems allow the operator physical flexibility.
- Information-retrieval systems for directory assistance or schedule enquiries. Such systems will contain an expert-system component that might act as a 'front-end' operator before switching to a human operator. Mr Davies suggested the use of an expert system to handle enquiries (rather than the current widespread practice of switching the call to 'musak' during the wait period).

In conclusion, Mr Davies made a few brief comments on voice processing. He identified the potential for store-and-forward systems (which will become cheaper as technology advances), recognition/transcription systems (annotation of electronic documents) and speech synthesis. In this latter context Mr Davies noted that synthesis provides a telephone with an output medium, potentially changing the nature of the device to a terminal.

Mr Davies's final remark underlined the potential use of speech systems with other technologies especially expert systems.

In the commentary and discussion period that followed, several delegates noted their interest in a system that might handle 80 per cent of a data-entry application by speech input. This obviously allows people to concentrate on the 20 per cent of exceptions. Finally, Mr Davies noted that shouting does not obviate the need for a high-quality lip microphone. The noise context associated with shouting makes word recognition even more difficult.

Principles of High-Tech Project Management

Robert M Alloway, Alloway Incorporated

The typical success rate of high-tech projects is low. What is needed is a means of elevating this low success rate. Fortunately, means are at hand to help. In his address, Dr Alloway began by defining what he meant by high-tech projects. He then went on to provide a set of guidelines for improving success.

High-tech projects are innovative, but it is important to recognise that innovation itself is relative, not absolute. A project that is innovative for one sector or business, or even department within a business, may not be so innovative for another. High-tech projects are risky. The higher the technology (and so the more innovative it is) the higher the risk, and the greater the possibility of failure. The payoff should justify the risk. If not, it would be better to avoid the risk and focus instead on alternative projects that are less innovative, and less demanding, but more assured of success. The reality, however, is that business competition, together with the accelerating rate at which new technologies are emerging, are putting more and more pressure on MIS departments to tackle high-tech projects.

The first step in managing high-tech projects is to recognise them for what they are. Managers should do this in a way that is conscious and explicit. The reason is that they are not the same as more conventional projects: they demand management principles that are different and distinct.

Next, be prepared to throw away the conventional project management rule book. Virtually every company has its rule book, often set in tablets of stone, that lays down the procedures for designing and implementing systems projects. In most instances the emphasis is on control. The procedures usually include breaking down the work programme into defined steps, sizing each step, setting budgets, controlling performance to plan, and so forth. With high-tech projects, such traditional approaches should be avoided. There are too many unknowns. Breaking down the work programme into steps no longer makes sense when the steps themselves and their associated resource and work requirement are unknown. High-tech projects are not conducive to the wisdom and procedures enshrined in the conventional project-management standards book. Indeed, using the conventional procedure is likely to be counterproductive.

For high-tech projects, there is no one right project design. Research work undertaken at Harvard in recent years has led to the development of an idea known as the 'basic contingency concept'. The message of this idea is that, to be successful, project design should match local conditions (see Figure 1). Local conditions include factors such as the extent to which the project objectives are firm and clear, the organisational context of the project, the range of departments impacted, the levels of staff involved, the management resources available, and so forth.

High-tech projects require a balance between project design, the nature of the risk, and the organisational context (see Figure 2 overleaf). For instance, the more complex the nature of the task, the more innovative the project design needs to be, and the wider and more complex the organisational context, the more the project design has to cope with. In a recent study conducted at MIT, 20 high-tech projects were surveyed. The purpose was to measure the relationship between success and project design. The 20 projects in question were selected to be representative of a wide range of sectors. Despite the difficulties inherent in measuring the parameters of success, the result of the study was convincing : 72 per cent of the

Figure 1

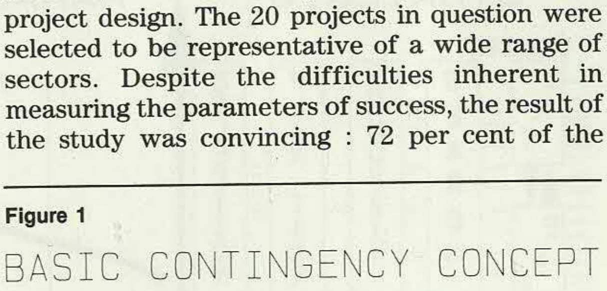

Figure 2

TEMPORARY MANAGEMENT SYSTEM

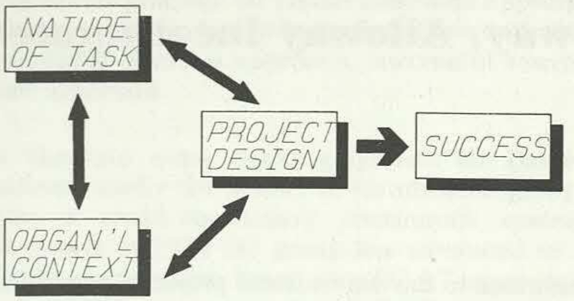

variation in success was due to project design (the remaining 28 per cent was due to other factors such as hassle, luck, hard work, and dedication).

Project design is, therefore, crucial. Inappropriate design will guarantee failure of the project, while an appropriate design can greatly increase the chance of success. The next question is how to manage a high-tech project so as to achieve the right project design. The key is to recognise that there are categories of project design, each with a corresponding management principle.

Figure 3 shows nine categories of project design. In the nine-cell matrix, each cell corresponds to a baseline project design and each one demands a different management approach, and different controls. Consider, for instance, the top right-hand cell. Projects corresponding to this cell are best avoided: they should be split up.

Figure 3

BASELINE PROJECT DESIGN

NATURE OF TASK	LOW	MED	HIGH
HIGH	DECISION SUPPORT	HIGH-TECH CONVERGE	FAIL/AVOID SPLIT UP
MED	TRANSACTION PROCESSING	L.A.G. DATABASE	HIGH-TECH COORDINATE
LOW	MAINTENANCE	MULTI-DEPT TRANSACTION	COMMON SYSTEMS

ORGANIZATIONAL CONTEXT

Finally, Figure 4 shows the classic operational project life-cycle curve, characterising a project as it passes through feasibility, design, programming, implementation, and maintenance. This operational life-cycle curve is supplemented by a second one, depicting managerial effort. This second curve goes through stages as well: diagnosis, planning,

Figure 4

COMPLETE CBIS LIFE CYCLE

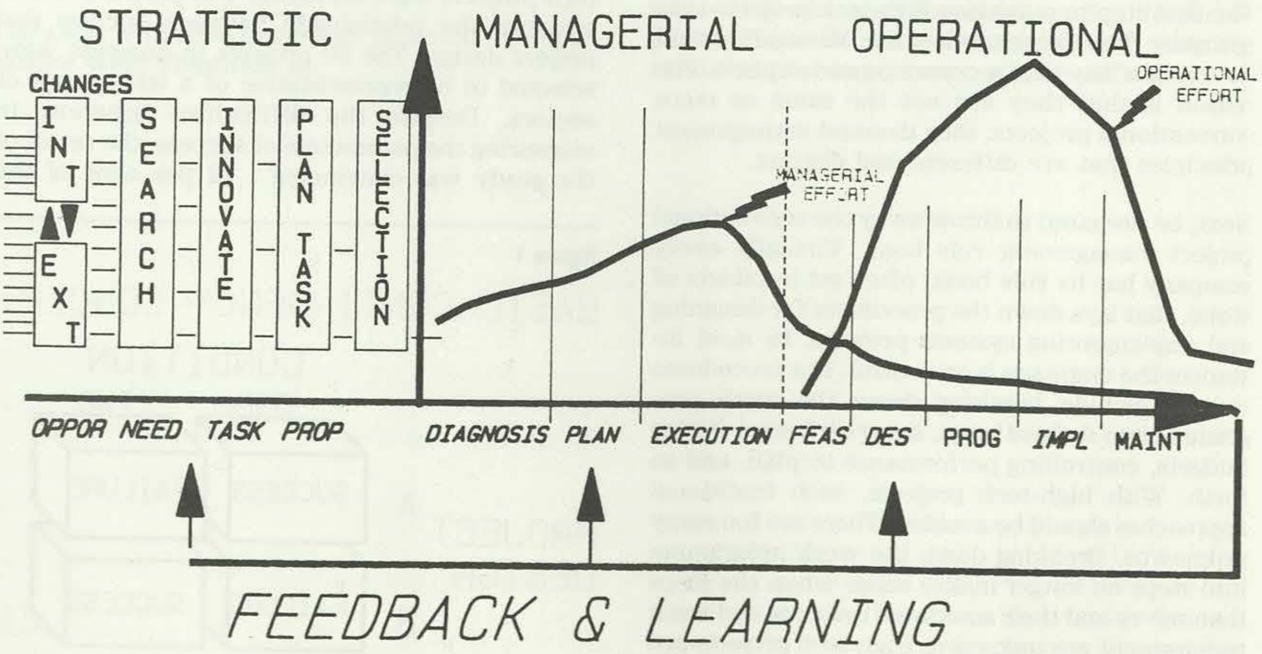

execution. It is these activities, preceding the operational effort, that are crucial in high-tech projects. The scale of the two curves is by no means equivalent. The diagnosis stage of the management curve, for instance, which defines the project design, typically takes no more than a week or so.

Adopting the principles encapsulated in Figure 4 brings two significant benefits. The first is that of improving the likely success of the project. The second, often equally important, benefit is that of ensuring further successes in the future. Success breeds success, and failure breeds unwillingness to take risks again in the future. The successful implementation of a high-tech, high-risk project encourages management to pursue other high-tech projects — leading to further significant business advantages.

Case History: An Integrated Technical Strategy Applied to a Complex Manufacturing Industry

Roger D Butler, Austin Rover Group

Dr Roger Butler is Manager, Forward Planning and Control, within the Austin Rover Group manufacturing operations. He is one of the architects of the group's integrated technical strategy and he has been deeply involved in restructuring the car manufacturing operations, to integrate what were previously fragmented elements using different systems and technologies.

He spoke as a user manager, not as a systems expert. His role is to provide the environment in which novel technologies can flourish but, at the same time, not one in which everyone can develop in incompatible directions.

THE NEED FOR AN INTEGRATED TECHNICAL STRATEGY

During the 1970s, various application islands of different technologies had been developed to improve the productivity of Austin Rover's design and manufacturing operations. The process was evolutionary, but in the late 1970s the company recognised it was likely to finish up with a muddle of incompatible systems that would be very difficult to link together. The problem of establishing standards for systems to avoid incompatibility is recognised throughout the motor industry — hence the MAP initiatives in the United States. At the end of the 1970s, Austin Rover decided that it needed to rationalise its various uses of technology. But it could neither wait for nor afford MAP. Instead, Austin Rover decided to develop its own standards. The result was an 'integrated technical strategy' linking together, and transforming, the ways in which the company designed, engineered, and manufactured products.

Figure 1

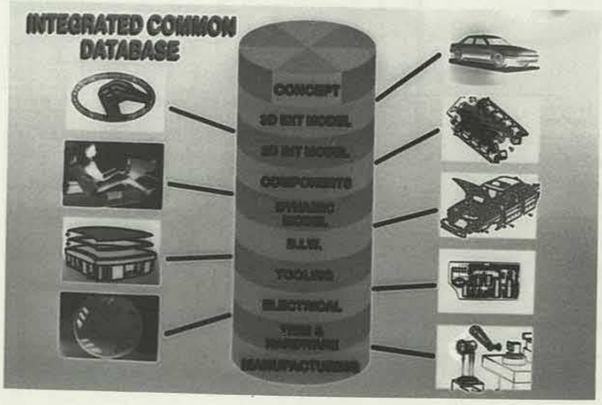

THE ELEMENTS OF THE TECHNICAL STRATEGY

The key elements of the technical strategy were:

- The use of a common three-dimensional database shared by the various operations from styling a new model through to computer-aided manufacture and inspection of the finished cars (see Figure 1).
- The attempt to standardise on common systems as far as possible for such things as computer-aided design systems, manufacturing robots, and planning and control systems.

The latter element has been a hard discipline to apply. In the 1970s, there were often operations where only one product from a particular supplier could provide the functions and performance required. Standardisation would have been impossible then. However, it was possible to establish a standard that could be enforced — in computer-aided design or office automation, for example. But there are still some areas where a mix of systems is inevitable — in manufacturing robotics, for example. Nevertheless, in this case Austin Rover now has systems from only five different suppliers, not the 200 or so that are in the marketplace.

Dr Butler advocated the discipline of standardising on the systems that do most for the company, and replacing the rest, despite the existing investment in the 'nonstandard' systems. Austin Rover now has more CAD systems in use (300) than any other European car manufacturer and it is the only manufacturer to use just one type of system.

OBJECTIVES OF THE TECHNICAL STRATEGY

To remain in business, Austin Rover needed to introduce an integrated but flexible means of designing, engineering, and manufacturing cars. New technologies had to be applied to:

- Computer-aided engineering, manufacturing, and inspection.
- Flexible manufacturing methods and systems.
- New support processes such as casting.
- Automatic manufacturing management and control.

All involve the application of computer power. By linking them together, using the common design database, the company streamlines the whole process from the stylist generating the shape for a new model through to the manufacture of the finished car. For example, previously it may have taken three months from the stylist's ideas being sketched on paper, through making clay models, to drafting hundreds of drawings to be used for production engineering and manufacture of the body. Now, the production engineer can use the same screen as the stylist to generate automatically within days the numerically controlled machine-tool cutter profiles used to make design models. Because of the time lapse before, there would often be several versions of a design in existence because the stylist would modify his or her ideas before the models were finished. Now, changes can be incorporated as they occur and there is only one design in existence at any one time. This change has greatly reduced the problems of design modifications and enabled new models to be put into production more quickly.

The technical database is used by the various functions shown in Figure 2. Not all the related systems in other functions are fully plugged in yet. For example, Austin Rover's 800 suppliers form a very important part of the production process. The smaller ones cannot afford their own computer-aided design systems compatible with Austin Rover's. Larger suppliers may wish to have systems compatible with the other car manufacturers' systems as well as with Austin Rover's.

Figure 2

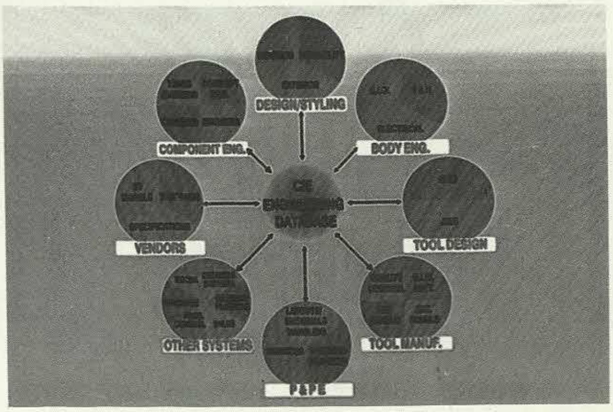

Figure 3

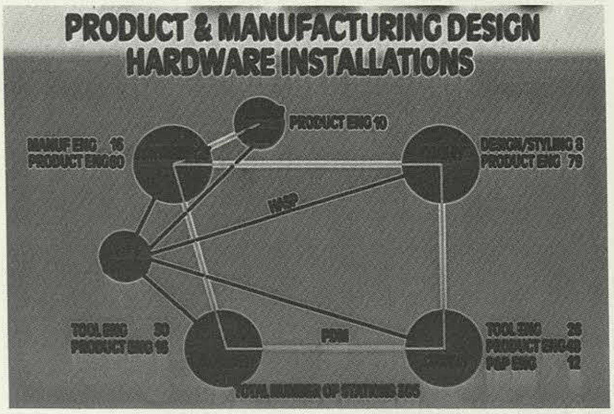

Today, about one in five of Austin Rover's employees now works with a computer terminal of one kind or another. For example, drawing boards are almost obsolete, being used only for the older car models still in production.

An outline diagram of the technical systems, showing how they are linked together, is depicted in Figure 3. Dr Butler showed many illustrations of the high technology now in place. They included laser-driven vision systems that are used automatically to position and check components in manufacturing and inspection. He agreed with Dr Stonham's opinion that image processing can be used in quality assurance to help select the cars that need detailed inspection by people.

IMPLEMENTATION OF THE STRATEGY

Dr Butler reinforced the advice given by Bob Alloway regarding the management of implementing such a technical strategy. Top-down support from the board was essential. The various facets of the project, such as training and the other people factors had to be consciously planned — not just left to happen. For instance, Austin Rover has organised two-week training courses in technology for all the board and senior management, and also provides facilities for staff to teach themselves about the new technologies being introduced.

He also agreed that the existing rule books for project management do not apply to sophisticated high-technology projects such as the one he was describing.

Figure 4

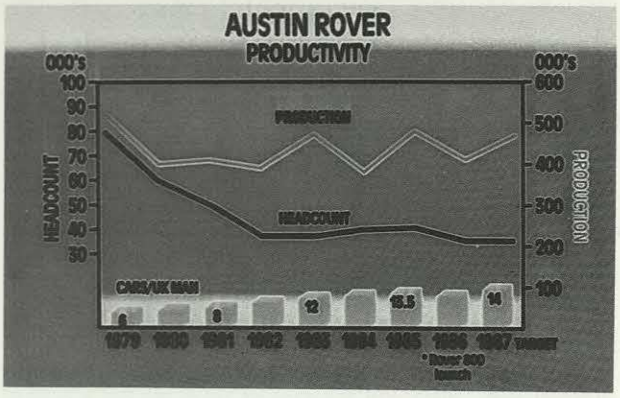

BENEFITS FROM HIGH TECHNOLOGY

Austin Rover is not trying to be the first to apply the latest technology. The company is happy to be second provided the technology works. Dr Butler indicated how Austin Rover's performance (in terms of productivity) had improved over the seven years during which the technology strategy has been implemented (see Figure 4).

Case History: Satellite Racing

Barry Stapley, Satellite Information Services

Barry Stapley is Technical Director of Satellite Information Services Ltd. He has 20 years' research and consulting experience in the fields of teleconferencing, networks, switching systems, and satellite-based systems. The company was started in early 1986 by a consortium of bookmakers, the Racecourse Association, and independent investors. The main objective of Satellite Information Services' business is to provide live televised racing and betting information to over 10,000 betting shops in the United Kingdom via satellite links, making the company the provider of one of the largest networks in Europe.

Dr Stapley began his presentation by outlining the driving forces behind the need for satellite broadcasting services. Significant changes had occurred within the betting industry in early 1986, and had led to increased dissatisfaction with the existing information services provided to betting shops. The most important change was that the legislation covering the facilities that could be provided in betting shops was relaxed, leading to live televised coverage of racing being permitted in the shops, as well as allowing the shop owners to provide a more 'up-market' and attractive shop environment. At this point, the possibility of a totally new communications package was recognised. Several communications media were examined, and satellite communications emerged as the most feasible and cost-effective method of transmitting a fully integrated service of sound, text, and pictures for the industry. Figure 1 shows the overall design of the Satellite Information Services' network.

Figure 1

Network design

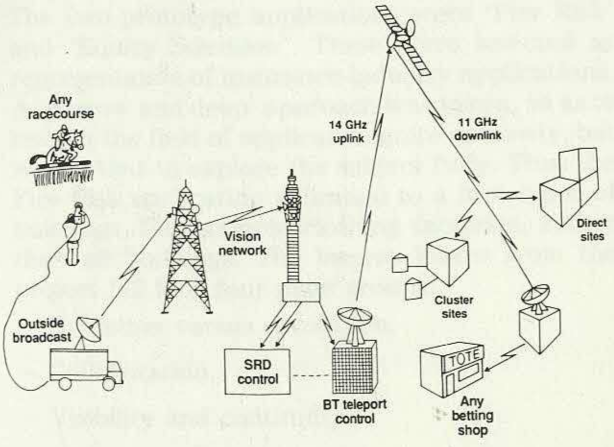

Satellite Information Services was set up to develop this project for the entire betting industry. The service that it offers will initially cover horse and dog meetings, but there are plans to cover other sporting events in the future. The service has five major components: information gathering, on-course television picture gathering, central collection, processing and editing, distribution of betting information, and in-shop receiving.

Dr Stapley concentrated his comments on the planning aspects of the project. He discussed the difficulties encountered in planning for the provision of each of the five key components of the service. In particular, unforeseen human difficulties were encountered in the installation of microwave receiving dishes on shop premises. He singled out planning-permission delays and uncooperative landlords as key inhibiting factors. Such difficulties led to a situation where traditional, formal project planning was virtually impossible, and certainly inaccurate. As Dr Stapley said, "We had a small project team and we forced things through." The result was that the project was delayed by several months. This was important as Satellite Information Services has high fixed costs, and delays meant that a build-up of revenue to offset these costs was also delayed. Figure 2 overleaf presents the overall project plan as it was originally drawn up. However, project costs were kept within budget.

Dr Stapley stated that several important lessons were learnt from the project that should be borne in mind by everyone when planning such an innovative, high-technology project. First, what appeared to be the most challenging aspect of the project, the use of satellite communications, was, in fact, one of the easier aspects. Far more difficult

Figure 2

Summary of key deadlines

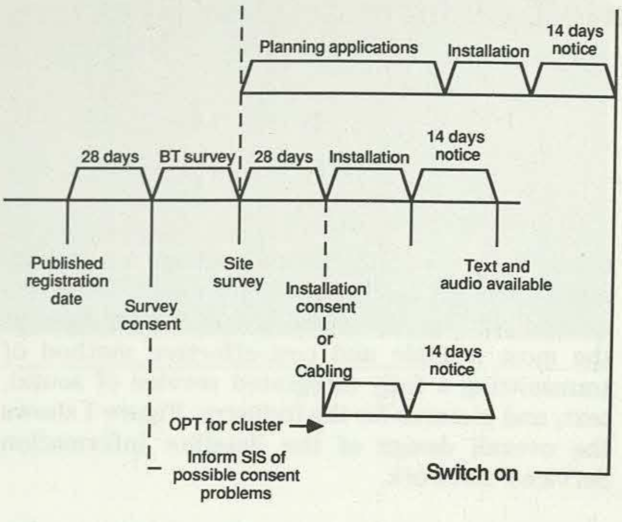

is the arranging for the installation of up to 8,000 microwave dishes, many of them on slate-roofed Victorian high-street shops with uncooperative landlords. There are also innumerable difficulties relating to legal and regulatory issues ranging from local planning permission to the very strict regulations in the betting industry itself. One of Dr Stapley's most important comments was that the introduction of such innovative business services was likely to generate more employment for lawyers and administrators than for engineers. However, despite the difficulties encountered, the eventual success of the project, and the pent-up demand for the services, will ensure commercial success, and a near monopoly position for Satellite Information Services.

Case History: Collaboration in Expert Systems for Insurance

Anthony R Butler, Colonial Mutual Group

Anthony Butler is Assistant Manager, Computer Services within the Colonial Mutual Group in London. He has particular responsibility for technology issues in the group.

In early 1985 he was invited to form, and to become chairman of, a community club within the Alvey Programme to investigate the use of expert systems in the insurance industry. The Alvey Programme is a government-funded programme of research directed at specific developments in information technology. One such development is the application of expert systems. Under the banner of the Alvey Programme, collaborative groups were established for a number of business sectors. The insurance sector group, called ARIES, has 31 participants, 19 of whom are representatives of insurance companies, the remainder comprising consultants and academics with specialised experience in insurance. The group worked in association with Logica who developed software to support the collaboration. The funding for the project was £350,000 of which £150,000 was a government subsidy.

Mr Butler described community clubs as a means of enabling people from the same business sector to learn together and to develop their experience of a new technology by collaboration. The specific objective was to ''deliver learning''. The vehicle used to achieve this within ARIES was to develop two prototype applications of the use of expert systems in insurance.

The two prototype applications were 'Fire Risk' and 'Equity Selection'. These were selected as representative of insurance-industry applications. A 'narrow and deep' approach was taken, so as to restrict the field of application quite narrowly, but within that to explore the subject fully. Thus the Fire Risk application is limited to a few types of buildings, for example clothing factories, rather than all buildings. The lessons learnt from the project fall into four main groups:

— Evolution versus revolution.

— Collaboration.

— Visibility and continuity.

— Pragmatism.

EVOLUTION VERSUS REVOLUTION

Mr Butler said that 90 per cent of high-technology projects are evolutionary rather than revolutionary. Evolutionary projects are low risk and have modest aims. You can learn from others, and previous experience is useful. It is also possible to use analogies as a guide. For example, although expert systems is a new technology, and there is little previous experience to draw on, in many ways the management of expert systems projects is not too dissimilar from complex database projects.

By contrast, revolutionary projects have dramatic aims. Previous experience and the experience of others can be ignored. In managing revolutionary projects it is important to avoid preconceptions and not to use the new technology in place of the old. For example, early cars had trafficators that simulated hand signals. It took 50 years for the motor industry to replace them with flashing lights, though the technology for flashing signals was available before the motor car.

COLLABORATION

Advantages of collaborating with others are the exchange of ideas and the development of centres of excellence. Collaboration also maintains impetus and continuity on projects. However, there are many problems with group working. Generally it takes more time and effort, in particular in technical coordination. There is a danger of the technicians making the key decisions. A particular advantage of the ARIES collaboration was that, because there was collaboration, the project had to be designed and planned in advance, an activity often omitted, particularly in high-technology projects.

VISIBILITY AND CONTINUITY

Projects must be linked to the business and be seen as business projects, not just technical ones. It is important to consider carefully the impact on the business. Mr Butler said it was important in this context not to give the project to a technician to manage, or to isolate the team from the business or to call the project 'research'.

In this type of project it is particularly important to involve users. However the users are not always easy to identify. Because of the nature of the project it was assumed (wrongly) that the user was the expert.

In practice the idea of the Fire Risk system was to enable branch-office staff to refer less often to the expert in head office by giving them access to the expertise via the system. However, the project team concentrated on the expert, neglecting the requirements of branch-office staff. Thus, one of the lessons learnt was to identify and involve the real users.

PRAGMATISM

It is important to start carefully and to set clear targets. Mr Butler agreed with an earlier speaker (Professor Feigenbaum) that it was worth spending a lot on the best tools at an early stage. Uncertainty is a feature of high-technology projects and you do not know at the start what level of complexity you are going to meet. Therefore, go for the best tools you can afford. If it turns out later that a simple PC-based tool would be adequate, the application can be transferred relatively easily. One of the ARIES projects had been ported to two different microcomputers at a cost of about 10 per cent of the original development cost. Other lessons learnt are to document and monitor the project carefully and to evaluate it thoroughly. This will enable you to carry out further projects using the same technology at lower risk.

So What? Implications for Users

Gordon M Edge, Scientific Generics

Gordon Edge is the founder of Scientific Generics, a research-based company specialising in optics, life sciences, and electronics, together with business disciplines such as economics and corporate strategy. In his view, the conference had set out to answer four questions:

- Should Foundation members be interested in high-technology projects?
- What are the implications of high-technology projects for MIS managers?
- Which business opportunities are worth exploiting?
- Which technological developments are worth watching?

These questions need to be answered against the background of a sea change in the business climate: service industries are becoming more capital intensive, and manufacturing industries are becoming more skills intensive. Managing high-technology projects is all about the management of the relevant skills to ensure that they produce the required output.

Professor Edge defined technology as "The application of science for commercial value". It is important to remember that business and commercial organisations are involved in managing applied research, not basic research. He identified the key issues as:

- Technology: its role and nature.
- Value added created by technology: its degree and location.
- Skills: the breadth and depth required.
- Management: innovation and interreactive.
- Financial: resources and rewards.
- Risks: forecasting and control.

The three key questions to answer about technology are:

- What technologies should be invested in today in order to be competitive tomorrow?
- Who will be the competition in the future? Does technology create new competitors?
- Where will the competitive markets be in the future? How will technology affect the marketplace?

In answering these questions, it is easy to be complacent and make dangerous assumptions. For example, the United States auto industry assumed, to its cost, that it operated in a closed domestic market. And the banks and insurance companies are now experiencing competition from unlikely sources such as the retail business. The key to answering these questions is to identify the value-added contribution that technology can make. Professor Edge's working definition of 'high technology' is that it provides a potentially high value-added contribution.

It is important to locate the contribution to value-added made by the successive operations performed by an organisation on bought-in goods and services. For some products (toothpaste, for example), most of the value-added contribution comes from marketing skills (manufacturing toothpaste requires little R&D and manufacturing skill). On the other hand, most of the value-added contribution for a toothbrush comes from manufacturing skills.

Sometimes, the value-added contribution changes with time. Initially, most of the value-added contributions for compact-disc players came from R&D skills. A couple of years later, it comes from manufacturing and marketing skills.

Thus, the strategic issues are:

- To identify the total value-added.
- To identify the location of the value-added components.
- To recognise that the value-added contributions from different components change with time.
- To be prepared to be flexible, so that the emphasis can be switched from marketing to manufacturing to R&D as required.

In this way, technology can be applied to provide competitive advantage in two different ways:

- Through product differentiation by concentrating on innovation at the materials, products, performance, or feature level.
- Through cost advantage by concentrating on innovation at the manufacturing-process level.

TECHNOLOGIES TO WATCH

Professor Edge identified the following technologies as the ones to watch, the ones that are likely to create 'discontinuities' that can be exploited for commercial advantage. However, he warned against the dangers of 'duality', where a new technology is used initially in the same way as the one it replaces. This means that the traditional suppliers and manufacturers miss out on the innovations made possible by the new technologies.

Surface technologies: Many new developments in surface technology are imminent, and these may well lead to new computing devices. (Recent advances in optical disc technology have come from research into the surface of a moth's eye).

Display technologies: New types of display technologies will be available "quite soon".

Optics: Optics (not optical computers), used in conjunction with silicon technologies, will lead to new devices. For example, special memories used for pattern-recognition applications, where the pattern is correlated optically. Optical cordless LANs will also be available, operating at 500M bit/s. Such systems will provide great flexibility in the location of workstations and the use of mobile and portable devices.

Mobile data communications: New types of communications techniques such as tropospheric scatter and meteorite scatter will revolutionise mobile communications.

Sensor technologies: New sensor technologies will improve the efficiency of sensor devices by two or three orders of magnitude.

Expert systems: Expert systems will be used in 'grey' decision support areas. For example, an auto manufacture could be provided with support to help decide whether to use a very expensive unburstable fuel tank for a new car instead of a cheaper alternative that (statistically) will burst a very small number of times.

Speech technologies: Speech technologies are not restricted to voice-recognition systems. Improved speech-coding techniques will lead to better quality telephony, and the effective use of bandwidth, and speech-synthesis applications will grow.

SKILLS ISSUES

A recent study of R&D organisations in the United States has identified the way in which the performance of technical staff varies according to various parameters. For example, they will be most effective if they spend only 50 per cent of their total time on their primary technical task. Performance reduces if the percentage is higher or lower. Overall performance also varies according to the main source of motivation (self/colleagues, colleagues, chief, or executive). The number of R&D functions performed also affects performance. The optimum number is about three or four. Age is also a factor. Performance seems to reduce at about age 40 to 45, but then increases again. It is almost as if organisations can promote the decline of their best resources.

All of these skills factors need to be taken into account when considering how to make innovative use of new technology.

Professor Edge summed up the skills issues as follows:

- Not all professional staff can be good professionals.
- Few professional staff can be good managers.
- Few managers can market professional services.

The keys to success are to build teams that provide:

- Technical excellence.
- Quality.
- Business awareness.
- A multidiscipline approach.

Technical innovation requires a wide range of skills. Organisations therefore need to create the right blend of marketing, sales, engineering, design, manufacturing, finance, fiscal, and quality skills. This implies a multidisciplinary and an interdisciplinary approach.

Finally, remember that a technical specialist who can both manage and market his or her speciality is a very rare commodity.

SUCCESS FACTORS

Professor Edge summarised the 'royal road' leading to success in technical innovation in the following terms:

- A high-quality science and technology base.
- High-quality research, engineering, production, and development.
- High-quality innovations in the product or process.
- High-quality marketing in research, launch, and implementation.
- Efficient scale-up of operations.
- High-quality management information.
- Good utilisation of external and internal resources.
- Commitment from the whole management structure.

Finally, it is essential to define what is to be done before starting out on a high-technology project.

List of delegates

ALLIED-LYONS	Darryl Whitcher
BARCLAYS BANK	Ernie Gledhill
BASS	Mike Charlesworth Kaz Diller
BOWATER	Dan Docherty
BRITISH AIRWAYS	Alan Ainsworth
BRITISH COAL	David Bardsley Keith Bodkin Paul Brown
BRITISH GAS EASTERN	Steve Williams
BRITISH GAS NORTH WEST	Gary Rawlinson
BRITISH NUCLEAR FUELS	Alan Harrison
BRITISH TELECOM	Nicola Dick-Cleland Andrew Stephens
BRITISH TELECOM (SOLENT DISTRICT)	Vernon Dover
BUPA	Richard Campbell-Carr
CARDIFF CITY COUNCIL	Mike Poulton
CORPORATION OF LLOYDS	Mina Gouran Des Lee Colin Talbot
CREDIT SUISSE FIRST BOSTON	Brian Woods
DALGETY	Colin Powell
DEPARTMENT OF THE ENVIRONMENT	Patrick Leonard
EAST MIDLANDS ELECTRICITY BOARD	Philip Champ
EAST SUSSEX COUNTY COUNCIL	Rob Briggs
ENGLISH CHINA CLAYS INTERNATIONAL	Bob Maddocks
EQUITY & LAW	Robert Hutchinson
GLAXO PHARMACEUTICALS	Iain Lee Bill Presley Mark Wadsworth
HONEYWELL BULL	Denis Chalk

List of Delegates

ICI	Dennis Connaughton
INLAND REVENUE	Don Brown
ISTEL	John Sansom
LEGAL & GENERAL	Fraser Mathews
LONDON RESIDUARY BODY	Ron Sims
MANUFACTURERS LIFE INSURANCE	Adrian Boyd John Gregory
MARKS & SPENCER	Jim Machacek
NATIONAL WESTMINSTER BANK	Geoff Birks John Hudson
PHILIPS ELECTRONICS	Andy Castle
POST OFFICE	Norlan McIntyre
PROVINCIAL GROUP	Mike Blaikie Stan Gardiner David Newton
RHONE POULENC	Manuel Bloch
ROYAL INSURANCE	Trevor Barraclough Graham Greene
J SAINSBURY	John Blake
SURREY COUNTY COUNCIL	Gurmel Bansal Colin Griffin Tony Lawson
SWISS BANK CORPORATION INTERNATIONAL	Richard Beal
SYSTEMS DYNAMICS	Brendan Doherty Bryan Mills
TOUCHE ROSS	Graham Norris Derek Pryce
UNION INTERNATIONAL	Mike Portlock Gwyn Stone
WIMPEY GROUP SERVICES	Andy Elliot
BUTLER COX STAFF	Tony Brewer David Butler George Cox Mandy Crouch Denise Dawes Fred Heys John Kinnear Rob Moreton Debbie O'Brien Jim O'Connor Irene Schenk David Seabrook

BUTLER COX
FOUNDATION

Butler Cox & Partners

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

Objectives of the Foundation

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

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

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

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

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

Membership of the Foundation

The majority of organisations participating in the Butler Cox Foundation are large organisations seeking to exploit to the full the most recent developments in information systems technology. An important minority of the membership is formed by suppliers of the technology. The membership is international with participants from Australia, Belgium, France, 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 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 FOUNDATION

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

Benelux
Butler Cox BV
Burg Hogguerstraat 791
1064 EB Amsterdam
☎ (020) 139955, Fax (020) 131157

France
Butler Cox SARL
Tour Akzo, 164 Rue Ambroise Croizat,
93204 St Denis-Cedex 1, France
☎ (161) 48.20.61.64, Fax (161) 48.20.72.58

Germany (FR)
Butler Cox Deutschland Ltd.
Richard-Wagner-Str. 13
8000 München 2
☎ (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) 486 1760 Fax (212) 319 6368

Australia
Mr J Cooper
Consultants (Computer and Financial) plc Australia
Level 5, 303 Pitt Street, Sydney 2000, Australia
☎ (02) 2870400, Fax (02) 2870450

Italy
SISDO
20123 Milano, Via Caradosso 7, Italy
☎ (02) 498 4651, Telex 350309, Fax (02) 481 8842

The Nordic Region
Statskonsult AB
Stortorget 9, S-21122 Malmo, Sweden
☎ (040) 1030 40, Telex 12754 SINTABS

Spain
Mr Sidney M Perera
Rosalía de Castro, 84-2°D, 20835 Madrid, Spain
☎ (91) 723 0995