Transcript

Management Conference

Torquay
January 26 – 29, 1981

The Butler Cox Foundation

The Butler Cox Foundation

EIGHTH MANAGEMENT CONFERENCE

Torquay, 26-29 January, 1981

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

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

CONTENTS

LIST OF DELEGATES

ALGEMENE BANK NEDERLAND NV	Dr J M Mudde
ALLIED BREWERIES LIMITED	M Connolly
BL SYSTEMS LIMITED	M Colin I Dewis B Johnson C Tristram
BP OIL LIMITED	D Brown M Hancock
BP TRADING LIMITED	H Belsey
BRITISH TELECOM	D Mildenhall P Polden
BRITISH NUCLEAR FUELS LIMITED	D Fletcher R Gunton
BURMAH OIL TRADING LIMITED	A Ford-Hutchinson D Hambling
CADBURY SCHWEPPES LIMITED	S Ross
CALOR GROUP LIMITED	C Gilbert I Hillan
CENTRAL COMPUTER & TELECOMMUNICATIONS AGENCY	E Conn V Dubery P Hall S Thomson G Watson
CIMENTS LAFARGE FRANCE	R Leger P Marchand
COURAGE LIMITED	J Ross C Montagnon G Wigginton M Vaughan
DEBENHAMS LIMITED	I Stoneman
ELECTRICITY COUNCIL	M Parish
ENGLISH CHINA CLAYS GROUP	H Patrick Dr T Thorneycroft
GRAND METROPOLITAN GROUP	D Crombie L Green A Knight J Mayger
GREATER LONDON COUNCIL	R Sims
IBM	J Brassart
IMPERIAL CHEMICAL INDUSTRIES LIMITED	K Edwards
INTERNATIONAL COMPUTERS LIMITED	G W Wilsher
ITT BUSINESS SYSTEMS	F Shubert
JEUMONT SCHNEIDER	J Petit
METAL BOX LIMITED	P Bennett J Davey
MIDLAND BANK LIMITED	D Whitworth

MOBIL OIL FRANCAISE	J Legrand
	M Mizrahi
MÖET HENNESSY	J Crepin
PHILIPS BUSINESS SYSTEMS	R Cotton
	D Tombs
	K Tweedale
PILKINGTON BROTHERS LIMITED	R McEwen King
PLESSEY TELECOMMUNICATIONS LIMITED	R Hart
	D Holloway
	J Hooley
PROVINCIAL INSURANCE CO. LIMITED	J Ross
RATP	H Sebag
RHM MANAGEMENT SERVICES LIMITED	A Davidson
	P Lorentz
	J Playfer
SCHWEIZERISCHE BANKGESELLSCHAFT	M Grouven
SCOTTISH & NEWCASTLE BREWERIES LIMITED	P Spink
SHELL UK LIMITED	W Sherratt
SKF	H Landahl
	B Rosenberg
SWEDISH TELECOMMUNICATIONS ADMINISTRATION	R Bernemyr
TSB COMPUTER SERVICES LIMITED	C Sommerville
UNILEVER LIMITED	P Hewson
	B Maudsley
UNITED BISCUITS LIMITED	R Talbot
WHITBREAD & CO. LIMITED	J Barclay
	N Redman
STAFF	
AKZO SYSTEMS	Dr H A M Verhelst
AKZO SYSTEMS FRANCE	A Loize
BUTLER COX & PARTNERS	L McKie
	J A Brewer
	J Broomfield
	D Butler
	A Caws
	G E Cox
	R Johnson
	R Moreton
	M Ray
	C J Woodward
STATSKONSULT	H Johannson
SPEAKERS	
BRITISH AEROSPACE	L Blonstein
BUTLER COX & PARTNERS	R Camrass
	D Seabrook
EASON CONSULTANTS LIMITED	N Eason
INTELMATIQUE	R Bright
INTERNATIONAL BUSINESS MACHINES CORPORATION	F Jelinek
MASSACHUSETTS INSTITUTE OF TECHNOLOGY	Prof N Negroponte
SYSTEM DEVELOPMENT CORPORATION	G Peel

CONFERENCE OPENING

David Butler,
Butler Cox & Partners Limited

Ladies and gentlemen, may I welcome you to this management conference of the Butler Cox Foundation, of which the theme is "New applications, new technologies". During the days which follow we hope that we will have the opportunity to catch a glimpse of some of the important changes which lie in the years ahead; some of the changes in computers, in office automation and in telecommunications — the three main themes of the work of the Foundation.

We have chosen, obviously on a very selective basis, the technologies which we will cover and the applications which we will draw out of them. With so much going on in the world of systems, we could have filled the agenda 10 or 20 times over; but we hope that the ones which we have in fact selected to include in the programme will give you an interesting and balanced picture of what is happening and, most of all, stimulate you to think about, and to discuss among yourselves and with us, the implications of some of these changes for the future.

First, we will set the scene against which the later speakers will develop their themes, describing their particular technologies and the lessons which they wish to draw from them. We will ask David Seabrook to paint that background picture for you in a session entitled "Future Trends in Systems Technology".

David is currently the Research Manager for the Butler Cox Foundation, which means that all the different pieces of research which we do flow across his desk in one way or another. He is responsible for controlling that research work and for the publication of the Foundation Report series. Before he joined us, David was in a senior management position with the computing subsidiary of the International Publishing Corporation.

So much for David's past background. Before I ask him to launch our conference on its way, let me also make a public announcement about his future activities. At the end of this conference, when Martin Ray completes his period of assignment as the Director of the Butler Cox Foundation, David will take over in that capacity from 1st February. So in listening to him, and in your response to his paper, I am sure that he is anxious in the job that he is about to undertake to be as sensitive as possible to all your desires and wishes about the future direction of the Foundation programme.

SESSION A

FUTURE TRENDS IN SYSTEMS TECHNOLOGY

David Seabrook,
Butler Cox & Partners Limited

David Seabrook is research manager for the Butler Cox Foundation, and as such is responsible for the research work and publication of the report series. His previous assignments with Butler Cox have included a study of the potential impact of videotex in Europe, and consultancy assignments for a United States corporation wishing to establish an online system in Europe, and for the European headquarters of a multi-national operation seeking advice on a pan-European computing and networking strategy. Before joining Butler Cox & Partners he was a director of the computing subsidiary of the International Publishing Corporation, where he worked for twelve years on the development and implementation of systems using computerised techniques for publishing airline timetables and large registers and directories. He holds a BSc degree in Applied Mathematics from the City University, London.

The theme of this conference is "New applications, new technologies". My task this morning is to set the scene for the conference by reviewing some of the trends in information systems technology. Some of the developments that I shall be speaking about will already be familiar to you; others may be new to you. My purpose in presenting any novel developments in technology will be not to titillate you, but rather to provide a perspective against which those developments can be evaluated.

Before I turn to the technology of information systems, I should like to begin by looking at the way that any technology develops. Any technology develops within a context; and, more often than not, the technology is developed to meet a need. That need may be commercial, economic, political, or perhaps military. I should like to illustrate this by an historic example.

About 150 years ago in this country, vast deposits of coal were being discovered. However, there was only one problem: that coal was not in the place where it was needed. It so happened that living on the edge of the coal fields there was a young engineer called Robert Stephenson, and the subsequent development of the railway engine owed as much to the commercial and economic need that existed at that time to shift coal about in bulk as it did to Stephenson's ingenuity.

What is the context and what are the needs against which information systems technology is developing today? I think that here there are two main influences. First, there is the shift to knowledge-based work; and secondly, there is the realisation that natural resources, and energy in particular, are finite resources.

Looking for a moment at the shift to knowledge-based work, the developed world in the last few decades has shifted from one in which most people were employed in manufacturing and agriculture to one in which most people are now employed in information-based occupations, and the growth in the number of office workers is an illustration of this shift. The response to this shift to knowledge-based work has been to turn to technology, to help us to handle information more efficiently and more effectively. That technology increasingly is based on microelectronics in whatever form it manifests itself: as mainframe computers, word processing systems, computer networks and so on.

As any technology develops, as it matures it passes through a number of stages. At the beginning of the development of a technology, the emphasis is mainly on how to make it work. But as the technology matures the emphasis shifts more to an emphasis on what we should be doing with the technology and why we should be doing it.

I have illustrated this shift on my first slide. Those of you who were present at the Foundation management lunch just before Christmas will recognise this slide as one that Gordon Scarrot used at that time. It illustrates the point very well indeed. The region on the left shows that initially most of the emphasis is on how to make the technology work, but as the technology matures the emphasis shifts more to what we should be doing with it and why we should be doing it. Along the bottom of the slide I have put the timescale of the development of steam engines as an example. The first practical steam engine was demonstrated round about the year 1700; but it was not until 150 years later that the formal theory of thermodynamics was developed. It was not until that theory was developed that steam engines could be used with any great certainty, and it led to a great proliferation in the use of steam engines, because once the theory of thermodynamics existed, engineers knew that they could build with confidence engines to meet the needs of the users.

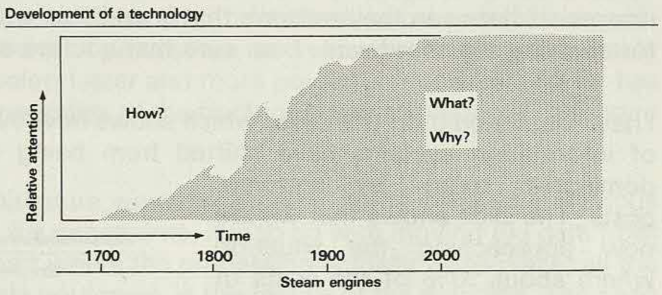

A lot of what we do in the Foundation concentrates on the what and why issues of information systems technology; and in this conference we shall also take a look at some of the how aspects of the technology as well.

The key to the development of information systems technology is microelectronics. I want now to review the major trends in the developments of this underlying technology of microelectronics.

Those who work in the field of microelectronics talk about the doubling rule. The doubling rule states that the complexity of chips doubles every year; that the speed of the circuit doubles every two years; and that the size of the chip doubles every four years. On the other hand, the design costs double every year. The net result of all of this doubling is that the cost per chip stays constant at about 30 cents per chip, which leads to a price per made up chip set of $300 to $400.

The effect of all of this doubling is that today we are able to see processors which are faster than ever before; memories which are able to store more information than ever before; but that both the processors and the memories become physically smaller, and that both processors and memories, in terms of their price/performance, become cheaper. In other words, the effect of microelectronics is to provide more 'bang per buck'.

The history of microelectronics has passed through a number of stages — through Medium Scale Integration and Large Scale Integration. Today, we are at the stage of Very Large Scale Integration; and the rapid advances in Very Large Scale Integration are challenging our ability to evaluate the technology, let alone apply it. But Very Large Scale Integration is not the end of the story; the next buzz word is Grand Scale Integration (GSI). Grand Scale Integration will bring with it the equivalent of a 370/158 on a chip by the second half of the '80s and one megabyte of random access memory on a chip in the same timescale. I am sure that tomorrow afternoon Norman Eason will tell us more about the development of microelectronics and the way in which he sees that technology being applied.

What I have said so far may seem to indicate that everything in the garden is rosy, that we will continue to get processors which are faster, memories which can store larger amounts of information, and better price/performance. But there is a problem. That problem is the gap which now exists between the hardware that is available and the software tools that we have available for applying that hardware. I am sure that you are all very familiar with this problem.

I have illustrated it on this slide, which shows how over the last 15 years or so the relative costs of information systems have shifted from being dominated by hardware costs to being dominated today by software costs. The slide shows that we are now approaching the situation where about 90% of the costs of information systems are related to software costs. What the slide shows is that the price of hardware is, in effect, tending to zero.

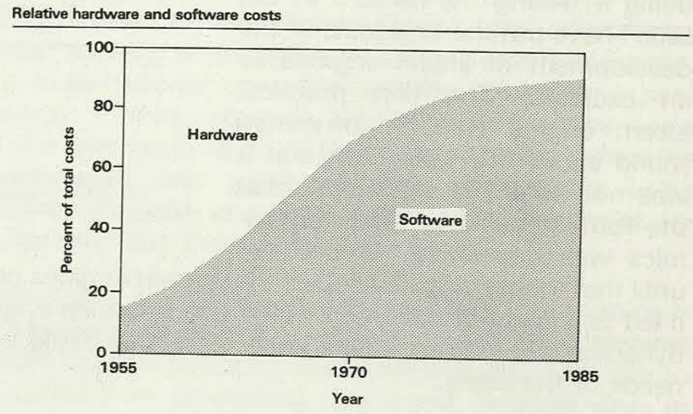

The problems that this dominance of software costs brings with it is an area which we have examined in the Foundation in the past, in Foundation Report No. 11 which dealt with systems productivity. It was also largely the theme of the conference in Venice last May. It is a theme to which we shall return later this year, in Foundation Report No. 25 which will look at alternative methods of developing systems.

I now want to move on to look at the way in which the basic underlying technology of microelectronics will be applied. It is necessary to do this at two levels: first, to look at the different types of equipment in which microelectronic technology will be used; and then to look at the way in which that equipment will be put together to form the hardware environment in which applications systems will operate.

The equipment that I want to look at is equipment for handling all types of information: voice, data, text, and image. I want to consider four different types of equipment: equipment for processing information; equipment for storing information; equipment for transforming information; and equipment for distributing information.

With equipment for processing information, I think that the effect of microelectronics will manifest itself in two extremes. At one end of the scale we will see a proliferation of small, cheap, portable computing devices — microprocessor devices. These devices will be used in homes, offices and schools. They will be afforded by people with anything but the smallest budget for office equipment, and increasingly will be bought by people out of their domestic budgets.

This proliferation of microprocessors will introduce many people to the joys and the frustrations of using personal computing for the first time. I think that this in turn will lead to a change in the attitudes of the end users that the corporate information processing function has to deal with.

If I may recall for a moment my first slide, which showed the change in emphasis from how to make the technology work, to an emphasis on what and why — as users become more accustomed to using computer-based equipment they will be in a position to concentrate on the what and why issues themselves. This means that, probably for the first time, the corporate information processing function can look forward to dealing with a knowledgeable and informed user base. That could well require a fundamental change in the attitude of the corporate information processing function.

At the other end of the scale, the impact of microelectronics will manifest itself in the development of very fast and very powerful processors. There are undoubtedly some specialised applications which require very fast and very powerful processors, and voice processing is probably one such application. But the need to develop faster and more powerful processors so far has been due more to a need to overcome the restraints of the traditional von Neumann architecture of computer hardware.

The constraints of the von Neumann architecture were discussed in some detail at a previous Foundation conference, but to summarise the constraints, it is to do with moving the data from the memory part of the processor to the part where the processing actually takes place; all the data has to pass individually through a single bottleneck in the middle of the machine. The need to develop faster von Neumann type machines is in turn due to a need to protect the software investment that large organisations now have in software designed to run on that type of machine.

There have been a number of alternative architectures proposed for overcoming the von Neumann bottleneck. But before I mention a couple of those, I should like to consider how much faster the existing, traditional architectures can be made to process. It turns out that the limiting factor is in fact the speed of light. I am sure that you are all well aware that nothing can travel faster than the speed of light, including electricity in electrical circuits. This means that if we want to build a processor with a one nanosecond cycle time no signal can travel more than six inches, because that is the distance that light will travel in one nanosecond. It is theoretically possible to build a very fast, very powerful computer in a six-inch cube of silicon. However, there is one great difficulty with doing that: it would generate so much heat that the silicon would melt. There is no known coolant that would work because any coolant known would boil. So what is required is a technology which provides very fast switching and also has very low power consumption.

Such a technology is under development by IBM, and by at least one major Japanese manufacturer — the technology of Josephson junctions. IBM have built fragments of computer circuitry, using this technology. Their predictions are that it is possible to build a processor which has a 3 nanosecond cycle time. That translates into a processor capable of operating at 100 million instructions per second, in other words about 10 times faster than IBM's current largest machine offering, the 3081. But there are quite severe problems with Josephson junction technology. It requires cryogenic temperatures, that is temperatures just above absolute zero. Operating at those sort of temperatures is likely to produce quite severe manufacturing constraints and operational constraints.

The technology of Josephson junctions is still very much under development; but one day perhaps it could lead to processors which are very small, very cold, very fast, but initially very expensive.

A few moments ago I mentioned that a number of alternative architectures have been proposed for overcoming the von Neumann bottleneck. The next slide illustrates one such architecture which uses parallel processing techniques. It is a diagrammatic representation of ICL's Distributed Array Processor, which consists of a 64 x 64 element matrix where each element in the matrix itself consists of a 1 bit processor and a 4k memory store. Each processor can communicate with its immediate neighbours in a north-south-east-west configuration and can also access the memory module of its immediate neighbours.

ICL originally designed the Distributed Array Processor with physical field problems in mind which require very large arrays of numbers to be processed. But they have since realised that there are many other types of problems which require several processors all to obey the same common instruction stream simultaneously. One example would be retrieving information by complex content criteria, where simultaneous searches could be initiated.

Distributed array processor

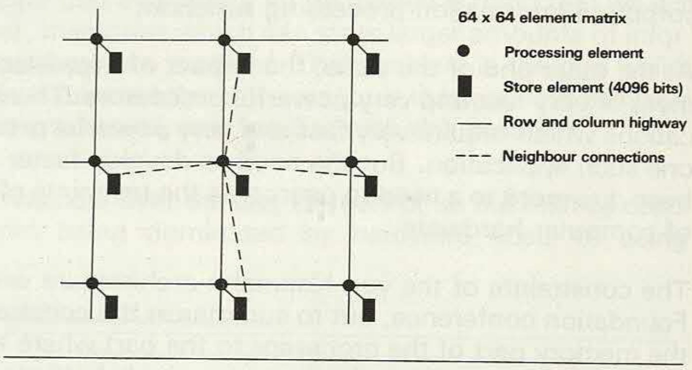

What I have just described sounds very much like another innovative device that ICL has — the contents addressable file store device. This illustrates a problem which will become increasingly apparent in the future: it will be ever more difficult to distinguish between equipment for processing information and equipment for storing information. The concept of active storage will become increasingly important in the future.

Two other alternative architectures are data driven architectures and pipelining techniques. With a data driven architecture the instructions are taken to the data to operate on it. With a pipelining architecture a task is broken up into a number of sub-tasks and dedicated, specialised processors carry out each individual sub-task.

Having reviewed the impact that microelectronic technology will have on equipment for processing information, I now want to move on to look at the impact that microelectronic technology will have on equipment for storing information. A lot of what I have said about microelectronics in general and about equipment for processing information applies also for equipment for storing information. Solid state storage devices will continue to store ever larger amounts of information whilst becoming physically smaller and, in terms of the price per bit of information stored, cheaper.

This slide shows the general trends in the reduction in the cost of memory over the last 15 years or so. Each individual curve on the slide shows the way in which the cost per bit stored has reduced over a period of time and has then been superseded by a subsequent technology which further reduces the cost. It is interesting to note that although the changes between core memory and solid state memories are quite fundamental in terms of the technology used, the older technology continues to exist for some time after the newer technology has been introduced. The straight line on the slide represents the general trend. It is perhaps unfortunate that the way in which this slide has been drawn might make it appear that the cost of memory will drop to zero in about 1983. I do not think that is quite true. The cost trends shown on the slide will continue into the future, without doubt, but the cost of memory will never quite reach zero.

Memory cost trends

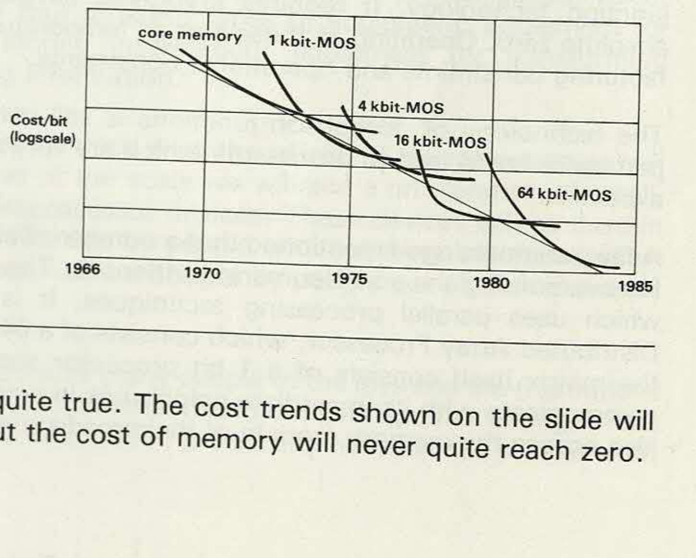

In a couple of years time it will be possible to add to that slide the cost curve for 256k random access memories, and by the second half of the '80s also the cost curve for 1 megabyte random access memories.

Solid state memory devices consist of random access memories, read only memories and, somewhat in between, programmable read only memories and erasable programmable read only memories. Solid state random access memories typically have an access time ranging between one-tenth of a microsecond and one microsecond. It is, however, possible to build solid state memories which cost less than random access memories, and these alternative types of solid state memories require the use of serial techniques. Two examples of solid state memories using serial techniques are charge coupled devices and bubble memories.

By serial techniques I mean that the information stored is circulated within the chip around a number of storage locations, and the data becomes available to be read only once in each cycle. So if the data is shifted through 64 locations, the access time of the charge coupled device will be about 64 times slower than an equivalent random access memory.

Charge coupled devices require about two to three times less area of silicon to store a bit of information. The control circuitry is also cheaper and less complex. This means that high density storage in solid state form will appear as charge coupled devices before it appears as a random access memory. So 256k memories and 1 megabyte memories are likely to appear as charge coupled devices before they appear as random access memories.

Bubble memories operate on exactly the same principle as charge coupled devices, except that instead of electrical charges being moved around it is magnetic bubbles being moved around, or, to put it more formally, microscopic domains of magnetic polarisation shifted about in a thin magnetic film by applying the appropriate magnetic fields. One advantage of bubble memories is that they are non-volatile — rather like a game of magnetic "pass the parcel": you switch the power off and the bubbles stay where they are until you put the power back on again.

Bubble memories will be used in situations where they can be used as a replacement for tape and disc memories, typically requiring between 1 million bits of information and 10 million bits of information. Mention of tapes and discs serves as a reminder that not all equipment for storing information are solid state devices. Tapes and discs of course are examples of moving surface memories, and drums are a third example. So systems designers today have available to them a hierarchy of memory devices.

Memory access time and price

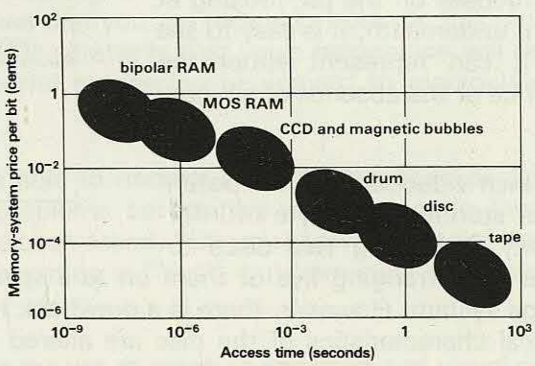

This slide shows the relationship between different types of memory devices. The vertical scale shows the price per bit of information stored on the different types of devices shown on the slide; the horizontal scale shows the access time of getting to the information on those devices. On the top, left-hand side we have the more expensive solid state memories with faster access times; on the bottom, right-hand side the moving surface memories, which are cheaper to store information on but slower to get at the information. Moving surface memories are non-volatile, like bubble memories.

Let us look in more detail at one of those moving surface memory technologies — discs. A few years ago, a number of people were prepared to say that, perhaps by now, discs would be obsolete. It is now well accepted that discs will be around for some time. There are many situations where storing information on discs will not be replaced by solid state storage devices for the foreseeable future. Disc technology has made very steady progress over the last 20 years. That progress will continue for the foreseeable future. Today we have discs capable of storing 1,260 million bytes of information, in other words more than a gigabyte of information. The disc manufacturers are confident that, within two to three years, the density will have been doubled again, to 2,500 million bytes of information per spindle of discs. By the end of the decade they believe that it can be doubled yet again, to 5 gigabytes.

To achieve these higher densities on discs will require the use of microelectronic fabrication techniques, both to create thin film platters on which to record the material, and to create smaller read-write heads so that more information can be stored on a smaller area of the disc. The disc manufacturers are also experimenting with techniques of recording information vertically into the magnetic film, rather than horizontally as at present.

Microelectronic technology will also be used in disc controllers to provide more intelligence. That intelligence will be used, for example, to enable the disc to anticipate the next access that it will be asked to make and also to provide greater error correction and detection than has hitherto been possible. What I am beginning to talk about again is the concept of active storage devices.

There is one type of moving surface memory which is not shown on the slide, and that is the process of using optical data storage techniques, that is videodiscs. Videodiscs have the potential to increase by an order of magnitude again the amount of information that can be stored online to a computer system.

On this slide I have shown the principles of recording information on a videodisc. The slide shows a cross section through the disc. The disc itself consists of a thin sheet of clear acrylic and information is recorded on the disc by burning small pits into the top of the disc. Once the pits have been burnt into the top of the acrylic layer, they are covered with a thin coating of a reflective aluminium material and the whole lot is then sealed in a protective coating. To read the information back from the disc a low-powered laser is focused from underneath on to the pits — or rather what is left after the burning process took place originally. If the laser focuses on the pit, looking at it from underneath, it is easy to see that it can represent either the presence or the absence of a binary digit.

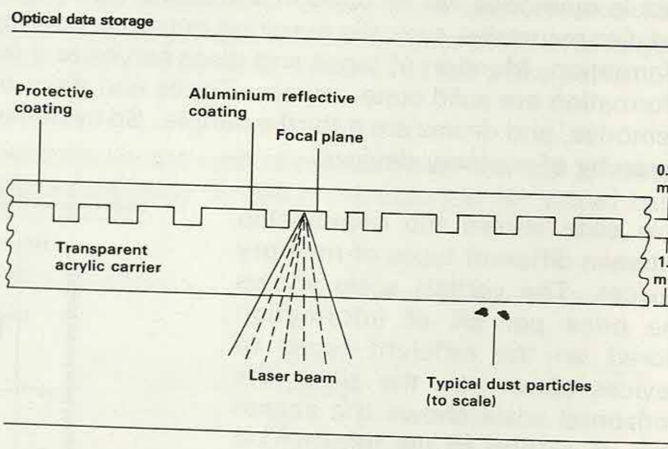

A 12-inch videodisc has the potential for storing a gigabyte of information, so putting two discs together and arranging five of them on to a spindle, you have the potential for a 10 gigabyte storage system. However, there is a drawback to optical data storage techniques: because the physical characteristics of the disc are altered when the data is recorded, it is not possible actually to overwrite it. But with such large amounts of storage available I am sure that it is possible to arrange techniques where data can be updated by writing it elsewhere on to the disc — in other words by introducing a certain amount of data redundancy into an optical data storage system.

So optical data storage is likely to be ideal for archival storage, perhaps for use in electronic filing systems.

I have talked about equipment for processing information and equipment for storing information. I now want to move on to look at equipment for transforming information; that is equipment for getting information into and out of computer systems. Input devices consist of keyboards, key-to-disc systems, visual display units, optical character recognition systems, mark sensing systems, laser scanners, and voice input; the telephone also can be used as an input device to computer systems. Output devices consist of printers in their many forms, line printers, character printers, dot-matrix printers, ink-jet printers, laser printers, daisywheel printers. Visual display units are also used as output devices as well as input devices. I also think that increasingly in the future typesetting equipment will be used as output devices. Voice output devices: again the telephone can be considered as an output device from computer systems.

I do not intend to give a detailed breakdown of the developments in the technology of each of those types of devices, but I should like to make one general point, which is that input and output devices, increasingly in the future, will have associated with them their own processing power and their own storage. So once again the concept of active storage emerges. For example, the days when a 100 megabyte terminal will be available are not far away.

The equipment for inputting and outputting information will normally form the interface between people and equipment, and the ergonomics of such equipment becomes increasingly important. It is an area that we looked at in some detail in Foundation Report No. 20.

To provide a truly usable interface between people and equipment it is necessary to use concepts which we all use in our everyday life to make the equipment as familiar and friendly as possible. This afternoon Professor Negroponte will be talking to us about developments at MIT, using the concepts of space and spatiality for storing and retrieving information. Professor Negroponte argues that we must use the same concepts of space and spatiality that we use every day, for arranging papers on our desk for example, for arranging information in information based systems. Until we do use similar types of concepts we will not provide a truly usable user interface.

I want to mention briefly two technologies used in equipment for transforming information. One is an input technology and the other is an output technology. The input technology is that of voice recognition. In the next session this morning, Fred Jelinek, from IBM, will be updating us on the state of the art of voice recognition technology. As I understand it, voice recognition technology today is capable of recognising isolated words, but the aim is to recognise continuous speech. In fact IBM's stated aim is to transcribe spoken English into typed text. However, there seem to me to be formidable problems to be overcome. For example, one second of spoken speech, if it is stored in an uncompressed digital form, requires about 50,000 bits of information which, on a quick calculation, means that you are receiving from me this morning about 150 million bits of information. What that illustrates is that voice recognition will require very large amounts of storage and very powerful processing equipment to manipulate the information which has been stored.

To be truly useful, voice recognition must be able to understand what the speaker actually means. That could present some very difficult problems. Let me give you an example. I wonder how many times porters at railway stations have been asked, "Do you know the time of the next train to London?" The logical answer to that question is "Yes", but I am sure that is not the answer that the questioner wants to hear.

The output technology that I want to mention is the use of graphics and colour on visual display units in a normal business environment. Graphics and colour have been used for a long time for special applications, but once again the developments in microelectronics are reducing the price of those facilities to a level where they can be used more widely in commercial environments.

Prestel and other videotex systems have demonstrated very well the potential for using graphics and colour. The next slide is a photograph taken from a normal television screen in Canada displaying information retrieved from a videotex data base using the Canadian Telidon technology. This demonstrates the potential for combining text and graphical information using colour as well.

If videotex terminals are to become widespread - and despite the slow take up of Prestel in this country there is still a good chance that will happen — it means that there will be an increasing demand for graphics and colour facilities in a normal data processing commercial environment. The use of graphics and colour in that environment is limited only by our imagination.

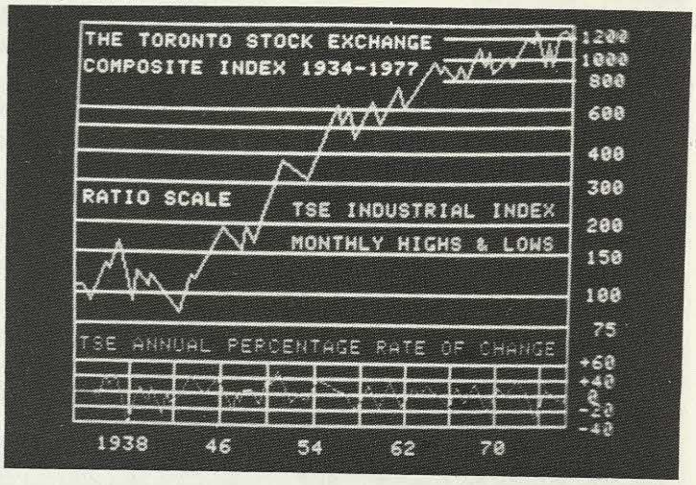

It is interesting to think of the way in which visual display units have been used so far. In many ways they have been used just to transform existing computer printouts directly on to a screen in their existing form. I think it was Confucious who said that a picture is worth a thousand words. There are many situations where financial information can be displayed as a histogram, a graph, or a pie chart; and statistical information also. Graphical displays can be used also for showing deviations from expected norms.

The fourth type of equipment in which microelectronics will have an impact is equipment for transmitting information, that is equipment used in telecommunications. Information has been communicated to date largely using analogue techniques, but developments in microelectronics will mean that in the future digital techniques increasingly will be used for telecommunications. This is an area that we looked at in some detail in Foundation Report No. 21, published just before Christmas.

The use of digital transmission techniques will mean that it will become cheaper and easier to move information about. There will be many situations where it will be cheaper and easier to move information about rather than moving people about. I think that this explains the current interest in systems such as electronic mail systems and teleconferencing systems. I am not suggesting that all business travel, or indeed most business travel, can be replaced by electronic forms of communication, but there are some areas where it is certainly possible.

Developments in equipment for transmitting information depend very much on the plans that the PTTs have for digitising the telecommunications infrastructure. Most PTTs have plans which eventually will provide a channel in parallel to the voice channel for transmitting information. That data channel is likely to be 8k bits per second, in parallel to the voice conversation, so that while you are speaking on the phone you can transmit data simultaneously over the same circuits. But that sort of facility will not be available at the local level, from local exchanges out to subscribers' equipment, until well into the 1990s. However, PTTs are making good progress with digitising the high-level telephone network. For example, by 1983 or 1984, in this country most of the major business centres should be interconnected with digital transmission facilities. That means that by that timescale it should be possible to lease from British Telecom digital lines at the highest level between the major centres. It would then be necessary only specially to engineer connections from the local exchanges out to the subscribers' equipment to provide end-to-end digital working.

The constraint so far on providing high bandwidth transmission facilities has been on the physical bulk of the coaxial cable necessary to accommodate the higher bandwidth; but there are some technologies under development which will help to overcome that constraint. One of those technologies is the use of optical fibres. A very thin pair of optical fibres should be capable of carrying several thousand simultaneous telephone calls. This means that in the same physical cabling conduit it would be possible to provide very much higher transmission capability using optical fibres than it has been with coaxial cables.

To use optical fibres for transmission of information also requires the development of optical data receivers. To date the development of optical data receivers has lagged behind the development of optical fibres. But IBM announced last November that they have developed a relatively inexpensive optical data receiver, which can be built on a single chip and which is capable of receiving data at a rate of 200 million bits per second. That type of development brings nearer the feasibility of the integrated office where voice communication, stored program controlled switchboards, store and forward message and data switches will all use compatible communication circuits and processors. I am sure that tomorrow morning my colleague, Roger Camrass, will be telling us more about this sort of development when he speaks to us about developments in private automatic branch exchanges, which are evolving from being basically voice switches to devices with non-voice functions as well.

A second technology which is helping to overcome bandwidth restrictions and constraints is the technology of satellites. After lunch today, Larry Blonstein will be telling us more about developments in satellite technology. But I think that the use which large organisations in Europe will be able to make of satellites depends very much on the European regulatory environment. For example, it is most unlikely that the PTTs in Europe will permit the equivalent of Satellite Business Systems to be established in Europe. SBS is the American corporation jointly owned by IBM, by Communication Satellite Corporation, and by Aetna Life and Casualty Insurance Company.

SBS launched its first satellite on 15 November last year, and that satellite is due to come into commercial operation soon. SBS sees the market for its services as being all forms of business communication, ranging from facsimile to electronic mail to teleconferencing.

The third technology under development for transmitting information is the use of the infra-red spectrum for transmitting information on a local basis. Using infra-red techniques provides the possibility of making local-area networks much more flexible, because the equipment will not have to be physically interconnected. We know that local-area networks are currently a hot topic and that a number of Foundation members are experimenting with equipment such as Ethernet, the Cambridge Ring, and Xi-Net.

But networking is not just about local-area networks, it is also about wide-area networks. Wide-area networks have their own set of constraints and problems. Many of those constraints are tied up with the PTTs' plans to digitise the basic telecommunications infrastructure, and also with the PTTs' plans for developing public data networks. It may seem that the changes in the telecommunications infrastructure are painfully slow. System X in this country is talked about reverently as being available in the late '80s, or the middle '90s, or perhaps into the 21st century. It is interesting to think that in this respect some of the underdeveloped countries may leapfrog over some of the more developed countries. Some of the Arab states may well go straight to using digital techniques for telecommunications before developed countries, because they have nothing to replace at the moment.

Nevertheless, by the end of this decade digital telephony will be the norm or, if it is not, it will be soon after then. The availability of cheap, ubiquitous digital transmission facilities will have a revolutionary impact on the demand for transmitting information in all its forms of voice, data, text, and image.

I have now summarised the impact of microelectronics on the four types of equipment which will be used in information processing systems: equipment for processing information; equipment for storing information; equipment for transforming information; and equipment for distributing information. I want now to look briefly at the way in which that equipment will be put together to form the hardware environment in which application systems will operate.

The basic point that I want to make is that information systems technology will become all-pervasive. It will invade all parts of our business life and our private lives. Let me illustrate this with a few examples. In factories, information systems technology will be used increasingly for production control, and also for controlling robotic systems. The Mini Metro production line at Longbridge is an example of the way in which things are moving. In offices also, information systems technology will be used increasingly. Office automation is currently the focus of much attention, but so far nobody seems to have made much progress with moving forward in a big way into office systems. I think that I may have an answer as to why this is so.

Let me refer back to my first slide which showed how the emphasis in technology shifted from an emphasis on how to make it work to an emphasis on what we should be doing and why we should be doing it. I suspect that with the application of information systems technology in the office we are still concentrating very much on the "how to make it work" area, and that we will not make dramatic progress until we concentrate much more on the what and why issues of information systems technology in offices.

But developments are taking place, novel developments of applying information systems technology in the office. Tomorrow morning, Gregory Peel will be speaking to us about the development of electronic filing systems and the impact that that type of system will have on storing information, transforming information and distributing information.

Information systems technology will also be used in a retail environment. I am thinking here particularly of point-of-sale systems. This is an interesting area because it brings us right up against the interface of information systems technology and the general public. One of the supposed benefits of point-of-sale systems is that it should no longer be necessary for the retailer individually to price each packet on the shelves; by reading a bar code on the packet the price can be retrieved from a central data base at the checkout desk. But consumer pressure in the United States is forcing retailers to continue individually to price each packet on the shelves.

Finally, information systems technology will also be used in the home. The emergence of personal computers and videotex systems is an example of this. France has very ambitious plans for introducing society in general to information systems technology. The French plans for the information-based society are extremely ambitious. The French, of course, have invented a word for it: Télématique. It is instructive that I have heard a number of Foundation members use the anglicised form of that word recently — telematics. Tomorrow afternoon, Roy Bright will be reviewing for us the French Télématique developments. He will concentrate on two aspects: on their Télétel experiment, which is very similar to the Prestel system in this country; but perhaps of more immediate interest, the French plans for automating the directory enquiry service. Those plans call for every home in France to be provided, free of charge, with a videotex type terminal. That terminal will be used to access the directory enquiry data base. The use of the terminal will not be free because the subscriber will have to pay at least the cost of a local telephone call each time he uses the terminal. Nevertheless, once those terminals are installed in homes it means that they can be used to access any other videotex data base.

I hope that I have demonstrated that information systems technology will become all-pervasive. During this conference we will hear details about a number of the developments that I have mentioned; some of the sessions will be about developments in the technology itself; others will be about developments in applications. Some of what we will hear about is already available; others are still at the research stage.

To conclude, let me refer again to my first slide. A few months ago, I believed that information systems technology had passed the point where there was now a greater emphasis on the what and why issues — what we should be doing with the technology and why we should be doing it. But I am not quite so sure now. I suspect that information systems technology is at about the same stage of development as the steam engine was just before the formal theory of thermodynamics was developed. I suspect that what we require today is a formal theory of the dynamics of information. I am not suggesting that I can even begin to say what such a theory should consist of, but until it does exist I do not believe that we can move forward in any great way into the what and why areas of information systems technology.

This means that the applications gap to which I alluded earlier will be with us for some time. Nevertheless, I hope that some of the things that we will hear about during this conference will show us the way in which we can begin to fill that applications gap.

SESSION B

DEVELOPMENTS IN THE RECOGNITION OF SPEECH

Frederick Jelinek,
International Business Machines Corporation

Frederick Jelinek was born in Prague, Czechoslovakia. He received bachelor, masters and doctorate degrees in electrical engineering from the Massachusetts Institute of Technology in 1956, 1958 and 1962 respectively.

Since June 1972 he has been with the Computer Sciences Department of IBM Thomas J. Watson Research Center, Yorktown Heights, New York, where he manages research on automatic recognition (transcription) of speech. He had been an instructor at MIT (1959-1962), a visiting lecturer at Harvard University (1962), a Professor of Electrical Engineering at MIT Lincoln Laboratory (1964-1965) and a visiting scientist at IBM Thomas J. Watson Research Center (1968-1969). His principal interests are in speech recognition and information theory. He is the author of Probabilistic Information Theory.

Dr Jelinek was the president of the IEEE Group on Information Theory in 1977 and was the recipient of the 1969-1970 Information Theory Group prize paper award.

I define speech recognition as the automatic transcription of speech. Metaphorically, you can think of it as an attempt to construct a typewriter which you talk at, and it types out what you have said. That, of course, is quite far away in the future, but there has been research along these lines, and there are some partial products on the market at this moment which I will mention.

Before that, however, I should like to say something about some related areas and products which are not speech recognition, but which have recently become the subject of interest. First, there is the area of speaker recognition/verification which is about recognising the identity of a speaker. The aim is to use a speaker's voice as a security key that would unlock doors, open accounts, or allow a bank to give you money.

Next, there is the area of speech synthesis where you try to synthesise utterances out of text. This area has different degrees of synthesis. Sometimes, what is referred to as speech synthesis is simply putting together sequences of words which were pre-recorded by a human being. Sometimes it is more advanced, and the idea is to synthesise every sound without a human being involved at all. Finally, there is the area of speech coding/compaction. Some of you may know that to record speech and transmit it digitally by PCM or delta modulation takes a lot of bandwidth and a lot of bits. Attempts are being made to compress the bandwidth and the number of bits. You may have heard about vocoders, which are machines in this area. I will not talk about any of these related areas today. I just wanted to distinguish between them and my real topic, which is three degrees of speech recognition.

The first degree of speech recognition is isolated word recognition, which is saying words in isolation, with pauses between them, and trying to recognise those words. There are actual commercial products on the market for isolated word recognition. These products are able to distinguish between about 50 to 100 words, and they work in various noise environments — maybe even over the telephone.

Next, there is the area in which I am involved, that of continuous speech recognition, which is an attempt to transcribe words spoken naturally in the cadence in which a normal person would speak.

Finally, there is the area of speech understanding, that is not just recognition but also interpretation of the meaning of what was said. That is an area which in some sense is even more difficult and awaits some further breakthrough. Some people may say that we will never have continuous speech recognition until we have speech understanding, because the correct transcription of words and sentences depends on whether one understands what was meant. This concerns homograms, such as H-E-A-R, H-E-R-E, and so on.

Why do we need speech recognition at all? My first slide shows some of the advantages of inputting speech into machines. I have marked those that I consider most important with arrows. First, the applications of speech recognition right now are almost exclusively industrial applications where it is very important that eyes and hands are left free. Speech input allows you to move around and not be tied to a keyboard, because you can use a microphone and a transmitter with no wires, and you can give your commands as you move around the workplace. For instance, General Motors tried to see whether they could carry out their quality inspections using speech recognition, by walking around the car and saying, "This door is dented", and "This lock doesn't work".

Advantages of speech input

- → Eyes and hands free
- → Mobility — not tied to a keyboard
- Speed
- Accuracy
- → Natural modality
- Cost effective
- → Rotary telephone becomes a terminal
- 24-hour, 7-day input capability
- → Little or no user training
- Permits simultaneous communication with machine and other humans
- No panel space or complex apparatus
- Compatible with radio

Speech is a natural modality, you do not have to train anyone for it. Another very big advantage for speech is that you can bring databases to the people, because almost everybody has a telephone. Therefore, if you develop speech recognition and speech understanding sufficiently, you can input information direct into a database. You can order merchandise and so on. Also, speech requires little or no user training. Of course that is an exaggeration and it is an important thing to remember when we talk of speech recognition and ask why we have not made as much progress as we might. After all, a child takes maybe 10 years to train before he can speak. So, when I say there is no training, that applies for an adult. But, of course, there is always evolutionary training. So the amount of training is very large, and, like us human beings, you can expect that a machine, in order to understand speech, will require a lot of speech to be trained.

Why should speech recognition be successful in the first place? My first slide on the next page shows an example of the speech waveform of an utterance used as part of a speech recognition programme at Carnegie-Mellon University which tried to recognise chess moves. This sentence says "Bishop moves to King Knight 5". You can see the regularity of patterns which recur at regular intervals. They are not identical from cycle to cycle, but they are very similar, and it is this similarity that gives rise to the hope that we will be able to recognise speech.

We have made progress in isolated word recognition rather than in continuous speech recognition for three major reasons. First, because of the small vocabulary of isolated word recognition, we can recognise the words by asking the speaker to say every word in the vocabulary once or twice. We then use the samples of his speech as prototypes or patterns to compare with what will be spoken next. For a 30 or 50-word vocabulary, it is not too tedious to have the speaker say all 50 words. On the other hand, natural speech recognition would require all of the millions of words in the Oxford English Dictionary, so obviously we cannot approach natural speech recognition in the same way. Even a moderate dictation system would require about a 10,000-word vocabulary, and again we cannot expect a user to go through such a training period.

The next advantage of isolated word recognition is that the pauses between words separate the words and makes their pronunciation rather standard. So you eliminate all the edge effects where words adjoin other words. For instance, you do not say "handlabelling" — you would put a pause in between and say "hand . . . labelling". You do not swallow your R's, and you do not say "onthetable", you say "ON-THE-TABLE".

Waveform of the utterance showing the actual word and phoneme boundaries.

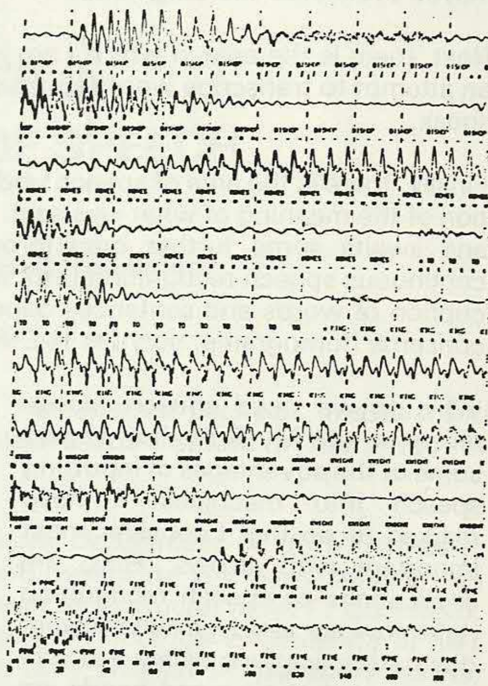

Finally, because of the pauses, the word pronunciation is distinct, stable and uninfluenced by context. If you have pauses between words, the recognisers will synchronise with what you have said. In other words, you know that you are about to try to recognise the 17th word because it comes after the 16th pause. On the other hand, if you are speaking continuously, since there are long words and short words, you never know exactly where you are.

The next slide shows some of the isolated word recognition products on the market and some of the firms researching into this area. In the availability column, an "N" means that there is no product on the market, and a "Y" means that you can buy it. The price ranges vary widely. You can buy very inexpensive recognisers for hobbyists and so on, such as the one from Heuristics for $300 (the hobbyists' market seems to be the biggest market these days). On the other hand, Nippon Electric has a recogniser for $70,000. The slide also shows whether the equipment works over the telephone, or in a normal environment, or in a relatively quiet office, and so on.

State of the art in discrete (and connected*) utterance recognition

		Available	Unit price	Applic. system?	Vocab. size	Resp. time	Claimed accur.%	Talker indep.?	Entry environ.	Standalone?
BTL		N		N	50	Slow	95	Y	Lab	N
Centigram		Y	3.5K	N	32	Fast	95	N	Norm	Y
Verbex (Dialog)		Y	70K	N	30	Fast	98	Y	Tel	Y
Heuristics	1–	Y	299	N	30	Slow	?	N	Norm	N
	2–	Y	2-3K	N	64	Fast	?	N	Norm	Y
IBM – STS		N		Y	50	Fast	95	N	Lab-Tel	Y
Interstate	1–	Y	18K	Y	50	Fast	95	N	Norm	Y
Electronics	2–	Y	2K	N	40-100	Fast	98	N	Norm	Y
	3–	Y	5-900	N	16	Fast	98	N	Norm	Y
Threshold		Y	10-80K	N	50+	Fast	99	N	Norm	Y
Fujitsu		N	¥5M	N	?	Fast	?	N	Office	Y
Sanyo		N	?	N	32	Fast	90	N	Home	Y
Nippon	1–	Y	¥18M	N	120* 1000	Fast	99	N	Norm	Y
	2–	Y	18.2K	N	4	Fast	99	Y	Tel-Norm	Y
	3–	Y	22.7K	N	16	Fast	99	Y	Tel-Norm	Y
	4–	Y	68.2K	N	128	Fast	99	Y	Tel-Norm	Y

The capabilities of the products shown are very large, and the vocabulary sizes on which they can operate differ by quite a lot. Actually, I have never seen the 1,000-word vocabulary system from Nippon in use. In general, 100 words is about the limit. That is not because word recognition could not go beyond 100 words — I think that it could go to about 200 or so — but in actuality there are very few isolated word applications that need more than a 30-word dictionary. The next application is an office dictation system, and that requires thousands of words. So there seems to be a natural dividing line.

This next slide shows a list of the applications for isolated word recognition. For instance, baggage at Chicago Airport is handled by speech recognition. That means the handler says "Kennedy" and presumably the baggage goes to Kennedy Airport. Anyway, it provides a good excuse for mishandling.

Present applications

- Parcel sorting systems
- Baggage handling
- Numerical control systems
- Warehousing and distribution
- Inventory control
- Quality control
- Distributed data processing data capture
- Command control interaction
- Machine control
- Aid to the handicapped
- Banking and credit card transactions
- Cartography
- Training air traffic controllers
- Aircraft cockpit communications
- Keyword spotting

Two major applications areas are inventory control and quality control. Today, you can see entire warehouses being run with isolated word recognisers. Speech recognition is also used for parcel sorting. Normally, somebody has to straighten out the package for another person, who then keys in the zip code, and that is how it is sorted. Now, with speech recognition, it is only one person who does the sorting; he reads in the zip code and the machine then does the sorting.

The next slide gives some other future applications for voice recognition. Some of these are applications for continuous speech recognition. I have divided them into three types of usage. First, public use, where the public would use the system. In this case, of course, isolated word recognition is probably unacceptable, because, if you have customers who are not trained, you will not want to insist that they put pauses between words, and you cannot take a sample of their voices. The requirement for public use is essentially for use over the telephone — i.e. for voice recognition over the telephone, in an undisciplined speaking environment, and for a speaker-independent recognition system. These are very heavy requirements, and we have a long way to go before we can meet them.

Some applications

I. Public use

- Catalogue ordering
- Pay-by-phone
- Stock market and other financial transactions, queries and instructions
- Access to data base
- Aid to the handicapped

II. Business use – occasional

- Speech filing commands
- Credit card verification
- Flight control log

III. Business use – intensive

- Direct data entry from questionnaires
- Small business billings
- Language instruction
- Help for the handicapped

The second type of use is occasional business use, where, possibly, you have some trained operators, but the use of the system is not intensive. Speech filing is where you send messages over the telephone to your friends or colleagues, which are then stored. Your friends can ring up to retrieve their messages, and the messages can be routed, and so on. In the United States, speech filing requires a touchtone telephone to input all the commands, and you have to remember people's dialling numbers and so on.

The third type of use is intensive business use for direct data entry from questionnaires, small business billings, and so on. In this case, you could expect to have a person who would do nothing else but use a voice recognition system, and you could afford to train him highly, and you could afford to train the machine very much to his voice. When continuous speech recognition becomes a business reality, I think that maybe data entry will be in this area.

The most likely office applications for voice recognition are for dictation machines. The idea is essentially to replace the typist and construct your office correspondence by direct dictation, with a display screen in front of you. The system would also provide an editing system and a formatting system and a hard copy facility. Thus, when you have dictated your letter, it is typed out and you are ready to send it. Perhaps most of you have secretaries and do not appreciate the problems of those of us who do not have personal secretaries. For instance, I share a secretary with about 30 people and when I dictate a letter I do not know when it will be sent out, and I worry about it. It is typed and it comes back to me for proof reading three or four days later. If the secretary has made any errors, I have already forgotten what I meant to say so I cannot correct it very well. So one of the major advantages of voice recognition systems is the instantaneous creation of documents for a low-level manager.

On my next slide I have shown two choices. It is not clear that the first choice is acceptable, although, personally, I think I would be willing to pay the price of dictating the words in isolation because the availability of the hard copy would be very useful to me. But there is the question of whether managers in general would be prepared to do the editing and correcting. It is a well-known myth — I was going to say that it is a well-known fact, but I do not believe that fact even though all the research seems to substantiate it — that managers hate the keyboard and will never touch it, because it is socially demeaning or whatever. Personally, I cannot believe that if something was made really useful, available and easy, managers would not use it. But all research says that they would not, so perhaps keyboards are only for low-level managers!

Office applications

Isolated phrase dictation machine
20,000 word vocabulary
Advanced man-machine interface

Memo and letter dictation machine
1000 word standard vocabulary
1000 word tailored vocabulary
10,000 word potential vocabulary available through isolated word recognition and spelling

So generally what I am talking about is a situation of the type shown on my next slide. Of course, you would not want to hang anything like that type of equipment on a manager. Perhaps he would want an array of microphones so that he would not need to sit at the terminal but would be able to pace about. Perhaps he would like to eliminate the keyboard altogether, so that the editing commands could also be made via speech. But that would be the general idea. The lady said "women" and the system has made a mistake, so there is still a need to correct it.

Why do we not have continuous speech recognition systems now? It is because there are great difficulties in doing continuous speech recognition. First, the acoustic signal is continuous, and it has no boundaries between words. It does not have boundaries between sounds either. Clearly, if you are to have a continuous speech recognition system, your building blocks cannot be words — they must be sounds, and the sounds run together. Furthermore, sounds do not sound the same way in every context. A "T" at the beginning of a word, for instance in the word TIP, is usually aspirated, but a "T" at the end of a word, like PIT, is not aspirated. So even though we are trained to hear sounds as being the same, acoustically they are not the same, which provides us with difficulties.

Human beings can easily distinguish speech sounds from non-speech sounds, and can have a lot of conversation going on around them but focus on one speaker only. It is very difficult for a machine to filter out speakers other than the one in which it is interested. Also there are a large variety of dialects, mannerisms and voice qualities, which again humans do not seem to have difficulty with. I think that you are perfectly able to understand what I am saying, and yet I have a Czech accent, and few of you have heard a Czech accent before.

Human understanding of speech depends on context and what the speaker has said previously, so that is another area where one has to have some kind of a tracking of what was said and what the meaning is. Finally, even if you write with perfect grammar, you cannot rely on spoken grammar, because most speech is irregular and ungrammatical. Also, people rephrase what they have started to say, they sometimes "um" and change in the middle of a sentence and try to scratch the whole conversation. So this is the problem we are up against.

In order not to be overwhelmed by all these difficulties at once, in research one minimises the problems and imposes limitations. First, you try to train the machine for a single speaker. By this I do not mean that it will forever work only with me, but you require from the person who is to be understood a large sample of his speech before you let him try the machine. You need to do the recording with very good microphones, in a relatively noise-free environment, and that means decidedly not over the telephone.

Initially, you may make the speech more regular by doing reading recognition rather than spontaneous speech recognition. In other words, you give somebody some text to read, and you attempt to transcribe what he reads. The idea is that he will read it at a regular pace, and that the text that he is reading is grammatical. To start with, you do not insist that the recognition is carried out in real time — that is you allow the processing time to be much larger than the speech time. During research, one second of speech might take 50 seconds to recognise. Finally, you do not build any machines, but you use general purpose computers to do your recognition.

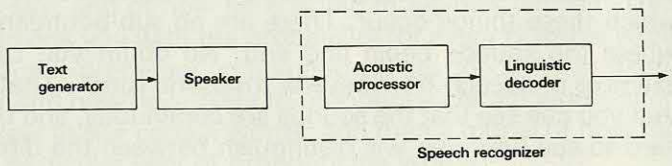

Speaker reads text presented by generator

Acoustic processor converts speech into a sequence of symbols suitable for the linguistic decoder

Linguistic decoder attempts to deduce the text read by the speaker, it chooses that word sequence which accounts best for the acoustic processor output symbols

I should like to continue by telling you a little more about how we carry out speech recognition research and then show you some examples of how far we have got and to show you what machines can do these days. This slide shows the point of view that we take in the design of a speech recogniser. All the experiments that have been done so far in continuous speech recognition have been with a speaker reading a text, although you can consider this diagram metaphorically and think of the speaker as generating text in his brain and then reading it off his brain. So this is not a bad diagram, even for spontaneous speech, but it is more easily interpretable when the text is being read by the speaker.

The speech recogniser usually consists of two parts — an acoustic processor and a linguistic decoder. The acoustic processor converts the speech into a series of distinct symbols which are suitable for processing by a linguistic decoder. Originally what people had in mind was to have an acoustic processor which would emulate one who goes among savages and transcribes their speech without knowing their grammar and without knowing the meaning of their language. As you know, missionaries are highly trained to do that, and this is how grammars for some languages have been developed. So it is possible to do this acoustic processing and segmentation into sounds without knowing anything about the linguistic content of the speech.

Anyway, it turned out that this approach was a dream. It was asking for too much and it is now known that an acoustic processor should not try to recognise phonemes and their segments. An acoustic processor should simply be a data compressor that takes the wave form spoken by the speaker and compresses it into a sequence of distinct symbols.

Finally, you have the linguistic decoder which tries to deduce the text from the string of symbols produced by the acoustic processor. The linguistic decoder decides what was spoken by using models of the speaker pronunciation, of the performance of the acoustic processor, of the text itself, and of the propensity of the speaker to formulate various sentences.

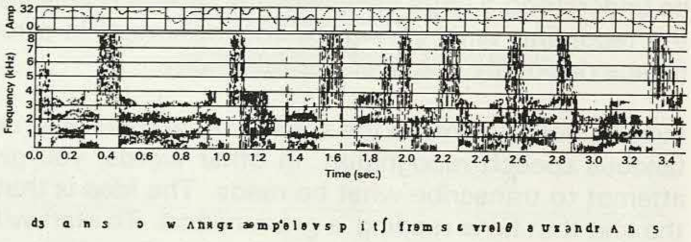

My next slide shows a typical output from an acoustic processor. This is called a spectrogram. It gives you a short-term picture of the energy content of different frequency bands as the speaker speaks. Here you have a spectrogram which is discrete in time — that is where the content in energy bands is different for each 10 milliseconds of speech. Wherever there is a dark area there is a lot of energy, and wherever there is a light area there is very little energy. Spectrograms used to be referred to as visible speech. It is from this spectrogram that the acoustic processor tries to put out the sequence of symbols. Previously, when the acoustic processor tried to put out a series of phonemes, the ideal thing would be to transcribe the spectrogram in the phonetic sequence in which these things occur. There are no sub-boundaries between phones but I have indicated where the sounds begin and end. No doubt you cannot read it but it says "John saw an example of speech from several thousand runs". That is irrelevant, but the whole point here is that you can see that the sounds are continuous, and that boundaries do not exist — so it is very hard to see how you will distinguish between the different sounds.

Let us assume that we have a perfect acoustic processor — we still need a linguistic decoder — for example:

Wimeɪkɔlavɚčɪldrɪn

We make all our children.

We may call our children.

Let us assume that we have a perfect acoustic processor, which means that a machine would put out the sequence of sounds perfectly as they are. I then claim that, even if you had that, you would still need a linguistic decoder, and my next slide illustrates why this is so. This is a phrase written in the way in which Websters' Dictionary gives the pronunciation for every word. If you try to read it where the words run together you will have a very hard time. But even after you have decoded it, there are many examples where the meaning is ambiguous. The slide shows two alternative meanings. You can see that both sentences have exactly the same sound pattern. Obviously I need something to distinguish which alternative was said. But in most cases, sounds do not have a perfect homonymy, which means they do not sound the same even though they have the same sound pattern. But as soon as you have an error in the transcription you can easily be sent from one interpretation to another. You need the linguistic decoder because you need to have some kind of measure of whether it is more likely that a person would have said one thing or another in one context or another context.

The task of the linguistic decoder is described in technical terms on this next slide. It says that the decoder sees the acoustic processor output (string A) as a string of symbols, and tries to find the word string W which maximises the probability of word string W being correct given the evidence of string A. Because of a well-known formula in mathematics (Bayes rule) this probability can be written as the product (or ratio) of three factors. You are trying to find the word string W which makes this term a maximum. W is not even involved in this term so you can forget about it as you vary the W. The term involves two probabilities. One is the probability that that word string will be uttered in the first place *a priori*, without any evidence of anything that has gone before. In other words, how likely is anybody to say that sentence. The other is the probability that the output of the acoustic processor will be A, given that the speaker read the word sequence W.

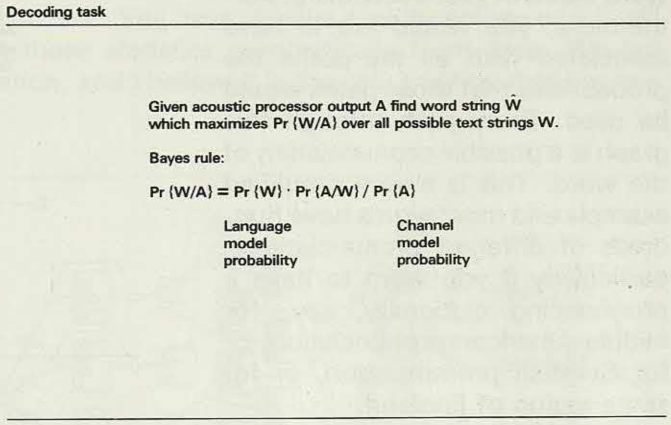

So, in order to do speech recognition, the linguistic decoder has to use a model of the language itself, of the text, or of the area of discourse that you are trying to attack. For instance, if you are trying to recognise business correspondence, you have to have a model of what a person who dictates a letter or memo is likely to say, and of how he is going to express himself. That model, in a statistical way, would then provide you with the probability for every possible word string. The probability of the speaker/acoustic processor combination is best decomposed into two parts — the phonetics of the speaker and how he is likely to pronounce certain words such as tomato, or either, and whether he is likely to slur certain words; and then a model of how the acoustic processor is likely to react to a particular way of pronouncing words.

My next slide shows the basic design of the linguistic decoder. There is a word sequence hypothesis search and it works with three models: a model of how frequently a word sequence is pronounced; a model of how the acoustic processor performs; and a phonetic model, which is a pronouncing lexicon fitted to the particular speaker. If you changed speakers, then you would have to provide a different phonetic model for the new speaker; and if you changed topics or applications, you would have to change both the language model and the acoustic processor model.

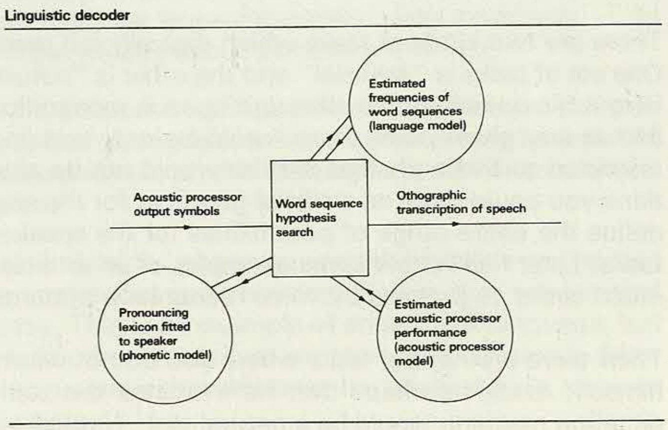

Here is an example of what a speaker model might do. The slide represents in a graphical form the ways of pronouncing the word "two". This simply says that "two" is a T which can be aspirated or not, and then the "oo" can either be short or long. What you need in your machine is a pronunciation graph for every possible word that is in your vocabulary. Furthermore, you would like to have associated with all the paths the probabilities that those paths would be used. Every path through this graph is a possible pronunciation of the word. This is a very simplified example and most words have hundreds of different pronunciations, particularly if you want to have a pronouncing dictionary, say, for middle-American pronunciation, or for Scottish pronunciation, or for some region of England.

Speaker model

Assigns probabilities to various possible word pronunciations specified by phone strings.
To each word there corresponds a pronunciation graph

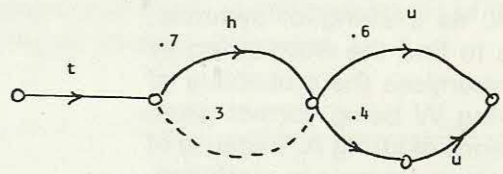

Finally, you would say that the speaker pronounced words by pronouncing a sequence of phones — phones are the basic units of pronunciation — and then you would like a model of how the machine is likely to react to these strings of phones.

So you need a model of how a machine reacts to a particular phone, and then to all phones. There is a problem of segmentation because the sounds run together. Then there is the problem of whether, when the phone is pronounced, you will recognise it as one phone or two phones, or perhaps no phones at all. So you have three possibilities — correct pronunciation, or deletion of the phone, or incorrect pronunciation where one phone is split into two. Then you need a list of probabilities. Suppose you say a "T". What is the probability that the machine puts out a "T" or an "L" or a "K" or what-have-you? Ideally, you need a complete statistical description for every possible reaction of the acoustic processor to the pronunciation of every given sound.

You also need a language model. In principle the language model is very simple. You would like a description that gives the probability of the next word being any particular word, given that a certain string of words has already been pronounced. How you arrive at this language model is a deep question. If you had a grammar of a very ambitious type, then you would say that the string of words already uttered would put you in a certain frame of reference, and through the grammatical analysis of what was already said, you as a human being would have some kind of expectation of the word that would be said next. We would like to give the same kind of feeling to a machine.

There are two kinds of tasks which typically are carried out by continuous speech recognition. One set of tasks is "artificial" and the other is "natural". Consider an insurance broker trying to fill out his questionnaires through speech recognition. He could have a menu in front of him, and at any given point there would be only certain phrases that he could say. He would be restricted to those phrases and he would not be able to say other phrases. For some applications you could write an artificial grammar for the speaker — artificial in the sense that it would define the entire range of possibilities for the speaker. These tasks would be called "artificial" tasks. Later I will show some examples of an artificial task and I will argue that those tasks are much easier to process by voice recognition systems.

Then there are natural tasks where you do not want to restrict the way the speaker expresses himself, except perhaps that he indicates the topic or the application. Certainly an office dictation machine would be a natural task. Natural tasks are much more difficult to implement.

It came as a surprise to many researchers that when we developed these language models of the pronunciation and of the speaker, they all required probabilities. The question is: how do you get these probabilities?

People assumed that you can derive the probabilities by counts, and by observation. But the question is: how are you going to carry out these counts? Surely you are not going to hire a million phoneticians to listen to the speech of radio announcers and count how people pronounce various words? You just cannot do that. You have to have a methodology which will be self-organising and which will extract these statistics automatically from data. We are striving to develop this sort of self-organisation, and I believe it is the only method that has any chance of success in the long run.

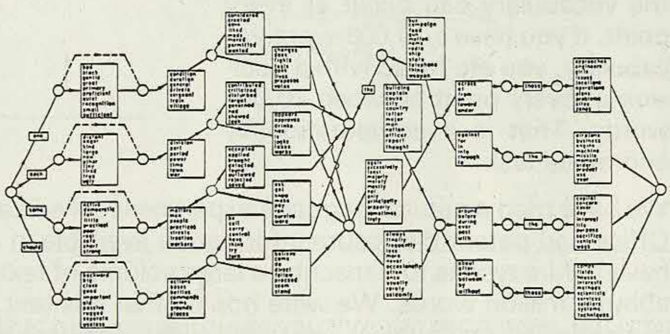

The next slide is an example of a very simple experiment of artificial grammar. This is known as finite state grammar, and it is called the New Raleigh Language because it comes from a project in Raleigh, North Carolina. All of the sentences that can be produced are described fully on the slide. The sentences are generated by computer. You place the computer in the initial state and it can select what state it goes to next. As it goes to the next state there are certain states that can follow, and this selection proceeds until the computer reaches the last state. Whenever the computer makes the transition from one list of words to another, it is free to choose any word from that list.

Most of the possible sentences that can be constructed are totally nonsensical. One possible sentence might be "Each large division criticised the captain over the vehicle". Whilst it is theoretically possible to say such a sentence, it is not a very probable sentence in any language in real use. But this is just an experimental example. You can see that I could put into this example quite a lot of restrictive situations which would mirror menu-type applications, such as answering questionnaires, or enquiring into a database about aeroplane schedules, or requesting stock market quotations, and so on.

The important feature that you see here is that the vocabulary is fairly small. This is a 250-word vocabulary, which has attractions for us. But the thing that makes this particularly easy compared to natural tasks is that the 250 words cannot be uttered in any spot, and that at any possible point there is only a very small number of words that can be uttered. That is what makes artificial tasks — even though they might be large-vocabulary tasks nominally — so much easier. The biggest choice shown on the slide is about 24 different words.

Lest you think that this is a very easy case, the number of different sentences that you can pronounce in this example is 14 million. So the speech recogniser is faced with a choice between 14 million possibilities. That is quite a respectable choice and, if it is going to be 100% accurate, it has to make the decision correctly every time.

We have achieved 100% success with the New Raleigh Language, and the first slide on the next page tells you how arduous the progress in continuous speech recognition is. You can see that it took us four years to achieve 100% success. That is an example of an artificial language, but now we are interested in natural languages. In other words, we are striving to develop a dictation system. The main point about natural languages is that whatever the vocabulary is, you can say a word from that vocabulary at every point.

That might sound strange to you, but in English and in any other language that I know of, most of the words are nouns or verbs. The number of prepositions, conjunctions or articles is very small (the total number may be something like 300). Yet the language has millions and millions of words. Also, when you speak a sentence and you are about to say a noun, you can say any noun from your entire vocabulary of English, and you can then continue the sentence in such a way so that the sentence makes sense. So in any natural task most of the words that are in the vocabulary can occur at every point. If you have a 10,000-word vocabulary, you are faced with a decision at every point between 10,000 words. That is a computationally enormous task.

New Raleigh language

Date	% Correct Whole sentences	Single words
11/74	73.0	96.4
8/77	95.0	99.4
8/78	100.0	100.0

Test results on 100 sentences
single speaker
sound room recording

We have tried a natural language experiment called Laser Patent Text. The United States Patent Office had patent disclosure information available in a machine readable form, and we did not have to hire typists to transcribe a large volume of text. The Patent Office has a very large text of about 2 million words. We were hoping that this text would be confined in terms of syntax and semantics to one subject only. We were also banking on the lack of originality of lawyers in expressing themselves. We were richly repaid in this expectation! In fact, too richly, because many patents for which people no doubt paid a lot of money to the lawyers differ only in one or two words from other patents. So we had a big problem with patent duplication, and somehow we had to remove the patents which were essentially duplications of each other. In any case, there is a stylistic uniformity across this text. This is different, for example, from a text of music reviews in newspapers. A music reviewer does not have a solid subject to write about, so he writes in a very flourishing language, and he is proud of the way that he expresses himself compared to his colleagues. Unfortunately, we were disappointed with the size of the vocabulary which, at 12,000 words, was much too large a step from a 250-word vocabulary. So what we did instead was to examine a subset of the sentences in the patent text, and that subset consisted only of words from a 1,000-word vocabulary. That means that we found the thousand most frequently occurring words in the whole text, and then found those sentences which are made up entirely of those 1,000 words. This is still a very large text, with about 72,000 words.

This slide shows some examples of that experimental text. No doubt the sentences seem very complex and long, but grammatically they are not as complex as they could be. For instance, there are no questions, and there are no exclamations. So whole parts of English sentence structures were not included in the experiment. The text is somewhat simplified to restrict it to a 1,000-word vocabulary. For instance, there is the prevalence of the word "eighteen". That is because the possible numerics that can occur in the text are infinite, and to construct a 1,000-word vocabulary, we had to keep the numerics that could occur finite, so we decided every number was

Sample laser sentences

After a period of time, the dye in the switching cell will decay and not be as effective in its switching operation.

Rod eighteen may be any of the known crystalline lasing material such as ruby, or alternatively, the active material could be any of the lasing gases, or any of the injection-type lasing media.

In operation, the rod eighteen is pumped from the light source eighteen to effect the usual inverted population state of the laser ions between given energy levels.

"eighteen". Also, every formula, no matter how complex, was translated to "sigma equals rho". We took other liberties of that type, but the sample still represented quite a respectable task.

Once we had this task, we were faced with deciding on the kind of language model to put together. In the artificial task the language model was given *a priori* because there was a limited set of sentences that could be pronounced, but how do you take a natural text such as this and construct a language model for it? Given a string of words, how do you estimate the probability of the next word?

We took a very simple-minded approach and said that we would predict the next word on the basis of "trigrams". In other words, we will try to find out what the next word is given the two previous words. This is stated more formally by the equation on this slide, which says that the probability of the occurrence of a string of words can be expressed as a product of the Kth word, given the previous (K-1) words. We have simplified the problem by saying that we will not consider the entire past, but we will consider only the last two words.

Laser language model

$$\Pr(W_1 \ldots W_n) = \prod_k \Pr(W_k | W_1 \ldots W_{k-1})$$
$$\approx \prod_k \Pr(W_k | W_{k-2} W_{k-1})$$

It seems extremely simple-minded to think that trigrams will guide you through sentences, but I can show you evidence which I think would convince you that a very good grammarian could not do much better.

We can now achieve about 92% success on recognition of words, and I will show you examples of how we do it. The price that we pay is that we use 70 CPU minutes on a 370/168 to decode a minute of speech. That is a gigantic number, and it seems like a very expensive proposition. However, digital computers are not really suitable for this type of task because a task such as this requires a lot of parallel processing. You can do a lot of work in parallel, you do not have to do the work in sequence. I am sure that this task can be speeded up enormously. We have carried out sizings which show that, when we finalise our algorithms, we will be able to produce not-too-expensive hardware working in real time.

Here is a slide to give you a feeling of what it means to achieve 92% recognition success. In a moment I will play a recording for you of how the person pronounced the sentence shown on the slide. I am sure that you will agree that he did not make any concessions to the machine as he spoke. In fact, I think you will find that he speaks much too fast for your taste. The reason that he speaks so fast is that he had to record about two hours of speech, and he was enormously bored and wanted to get it over with. He will read the sentence as it is printed out on the lefthand side of the slide, and on the righthand side is what the machine came up with. You can see that the machine made one error — it changed "the" into "a". As you listen to the recording you will hear that there is very slight evidence for the "the". I am not saying that you will not recognise that he did say "the", but you will see that recognition is quite difficult. (Taped example).

Sentence key 0000054

True	Decoded	True	Decoded
advantageously	advantageously	the	the
the	the	optical	optical
acoustic	acoustic	wavelength	wavelength
wavelength	wavelength	generated	generated
is	is	by	by
of	of	•the	a
the	the	coherent	coherent
order	order	light	light
of	of	source	source
magnitude	magnitude		
of	of		

19/20 words correctly decoded

In the second example there are several errors. "Silver" is decoded as "silvered", and the next error is an example of an error that an isolated word recogniser could not make, but it is an error which is only too prevalent in continuous speech recognition. One word (dielectric) was said and it was decoded as two words (by electric). (Taped example).

Sentence key 0000064

True	Decoded	True	Decoded
the	the	other	other
particular	particular	suitable	suitable
reflective	reflective	material	material
coatings	coatings	for	for
on	on	example	example
the	the	* dielectric	by
mirrors	mirrors		electric
may	may	* coatings	coating
be	be		
* silver	silvered		
or	or		

15/18 words correctly decoded.

The second slide on this page shows a homonymic example where "hear" is decoded for "here" and "as" is decoded as "is". The third error on this slide is due to the language model, because unfortunately it changed "if" into "be", and then the phrase "may be desired to be" is much more probable than "may be desired be". So the decoder is forced into that error by the language model once the first error is made. (Taped example).

Sentence key 0000077

True	Decoded	True	Decoded
although	although	may	may
flashtubes	flashtubes	* if	be
eighteen	eighteen	desired	desired
are	are	I	to
* here	hear	be	be
illustrated	illustrated	connected	connected
* as	is	in	in
connected	connected	parallel	parallel
in	in		
series	series		
they	they		

15/18 words correctly decoded.

The third slide is the final example, and it illustrates the kind of catastrophe we come up against. But again, *ex post facto* in research, everything can be explained. The "l" in "available" comes from "will", and "the" and "v" are similar sounds. This is the type of problem you run into with continuous speech recognition. (Taped example).

Sentence key 0000093

True	Decoded	True	Decoded
because	because	the	the
D of		gaseous	gaseous
the	the	discharge	discharge
losses	losses	is	is
associated	associated	* available	of
with	with		the
such	such		tube
an	an		will
arrangement	arrangement	for	for
however	however	pumping	pumping
not	not	D in	
all	all	the	the
of	of	ruby	ruby
the	the	rod	rod
light	light		
emitted	emitted		
by	by		

25/28 words correctly decoded.

I think that you would agree that if I could achieve this kind of recognition in real time and I could do it for office dictation, 92% reliability would be entirely satisfactory. Presumably you would see the text appearing on a screen in front of you as you were dictating, and therefore you could either say an incorrectly decoded phrase again, or let it pass and correct it later. So, in one sense, we are right at the threshold of being able to create an office dictation system as far as reliability is concerned, but we have two problems that we have not tackled yet. One is the problem of decoding in real time, and the second is the problem of spontaneous speech. At present, continuous speech recognition works when the person reads the text, but in real life it will have to deal with the vagaries of a person dictating — with the fact that he does not speak regularly, and so on.

On the other hand, spontaneous speech gives us an opportunity for improvement because a recognition system where the person is reading does not allow, through feedback, any teaching of the speaker. The speaker reads the text, and five months later the computer produces the output. So he has no way to adjust the way he speaks to the way the machine recognises his voice. We are hoping very much that speakers are malleable, and that when they realise that the machine is not very good at recognising swallowed articles, they will say the articles distinctly, because the penalty for not saying them distinctly will be rather high.

My last slide shows the major remaining problems that must be overcome before we can think of voice recognition products. We would like to be able to adapt to multiple speakers. This may not be necessary commercially but certainly it would be desirable if we could recognise many speakers at once rather than designing the machine for one speaker only. We would like to be able to reduce the bandwidth requirements so that lower quality microphones can be used.

Major remaining problems

Adaptation to multiple speakers

Reduced bandwidth and increased noise input

Treatment of spontaneous speech

Special purpose hardware design to attain real time response
by exploiting parallelism of recognition algorithms

Large vocabularies and incremental addition of new words

Human factors of dictation:
- Acceptability of isolated word input
- Error rate tolerance
- Vocabulary requirements
- Auxiliary aids: key entry
  software formatting

Design of interface facilitating interactive error correction

In a real application, it is very hard to restrict a person to a certain vocabulary, because clearly it is not practical to have the person learn a vocabulary and know which words are and are not in the vocabulary. We do not know the extent to which a person can memorise the allowed vocabulary, but it will be very interesting to see whether his corrrespondence will be impoverished because he will subconsciously not use the words which he would like to use, but will instead use the words that are available to him. Undoubtedly, we will need to make provision to put in words that are not in the original vocabulary, and we have not solved how to do that yet. On the other hand, everybody has his own vocabulary, and one way of proceeding might be to provide a basic 5,000-word vocabulary for everybody and let the other 5,000 words be built up through usage. In other words, provide some kind of machinery that collects statistics on how likely a person is to use particular words, and always keep in the vocabulary those words that he is currently most likely to use. Thus the vocabulary would be dynamic. One month the person might be corresponding about a widget problem and another month he might be more interested in his tennis game, so that vocabulary might become more prevalent. Perhaps there is scope for the machine to be provided with an adaptive learning capability.

Finally, the human factors of dictation have not been studied — in particular, how acceptable isolated word input might be. It is not just a question of isolated word recognition versus natural language. There are all sorts of gradations that you can have between different modes of input. Perhaps you could allow continuously spoken phrases with articles or adjectives or nouns, but with pauses between noun phrases. I do not know what you could allow, but certainly there is a large area of research there. How tolerant are people to errors? Is there a threshold? It is interesting that in isolated word recognition there is a very sharp threshold at 97% accuracy. I do not know why, but everybody reports that if the machine is more than 97% accurate, the users are satisfied. But as soon as the accuracy drops below 97% they walk away from it and do not want to have anything to do with the machine. Researchers do not understand why this is so and how it comes about, but in the voice recognition industry 97% accuracy is a cut-off point. There is no justification for it, but everybody recognises it and tries to get over that hump.

Other problems concern the vocabulary requirements and the kind of auxiliary aids that are needed. Is key entry acceptable? What kind of software formatting must you provide?

If you want the machine to behave like a secretary, it must have quite a lot of intelligence pre-programmed into it, such as an understanding of paragraphing and layout and so on. Finally, what kind of interface facilities should you have for interactive error correction?

Let me end by saying that I think that dictation machines as commercial products are still some way off in the future. I should not be surprised if we did not see any earlier than 1990.

SESSION C

SATELLITES FOR BROADCASTING AND BUSINESS COMMUNICATION

Larry Blonstein,
British Aerospace

Larry Blonstein was in charge of the design and construction of Europe's first satellite, UK3, which was launched from the USA in 1966. Subsequently, he worked on satellite earth stations in Europe, and on associated radar in the USA. He spent five years in Paris, working with an international organisation on space systems, with particular emphasis on Third World needs. He re-joined British Aerospace in 1979 as sales manager of the space and communications division.

This afternoon I shall try to bring you up to date on the European scene and the imminence of satellite communications, and to overcome some misconceptions, of which there are many when it comes to satellite communications. I will take you through what is happening now, what will happen very shortly, and what will happen inside this decade.

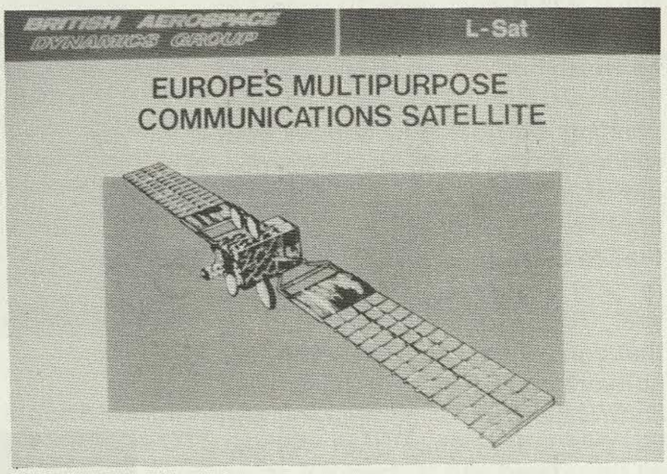

The first slide shows a picture of L-SAT. I am sorry about the name which is a temporary abbreviation for 'large satellite'. We are not allowed to give a name to a satellite until it is about to go into orbit. I want to call it OLYMPUS which means Operational Large Heavy Multi Purpose Satellite, but Rolls-Royce are not very happy about that and we are having a gentle fight with them over their rotten Olympus engine. So regrettably, for the moment we have to call it L-SAT. It is large. When those wings carrying solar cells are erected they will be three times the width of this conference room, and from tip to tip it will be nearly 170 feet. Its height from the base of the body to the top of the tower carrying the antennae will be 22 feet. When it is launched, hopefully at the end of 1984, it will be the biggest commercial communications satellite in the world. And it is being built here, in Britain! How about that?

However, it costs $60 million. To overcome some more misconceptions, my next slide has on it some very simple sums to show you why people want to spend $60 million on a satellite. That is a talking price, at 1980 prices, so it is subject to escalation.

You do not buy one satellite, because it might fail; you have to buy two in orbit. That means two in orbit and they both have to be launched. In addition, you have to build a third and keep it on the ground. For example, our OTS satellite was dropped in the sea by an American rocket failure. It was insured for $23 million; it was a tiny satellite. But you have to keep a third one on the ground in the event that one gets dropped in the sea. So as a manufacturer you have to build three satellites to put a two-satellite fully redundant system into orbit. So there is $180 million gone.

A launcher for something the size of L-SAT has to be a complete Ariane rocket, which is being built in Europe, or a large lump of shuttle which I will mention briefly later. The typical cost at today's prices for that size and weight of satellite is $40 million for each launch. Now the insurers, mad as they are, at the moment are offering a 10% premium on launches. They are getting a little worried because Ariane has worked once and failed once. They should be operating on a 50% premium but they are not, thank goodness. In the United States they are offering 10% insurance on the total thing launched into the sky. So for each $100 million launch you have to find 10% in advance of launch, which is another $10 million. So your system capital cost before you borrow any money means that you have to find $280 million. You, as a Post Office or as an independent operator, have got to find $280 million.

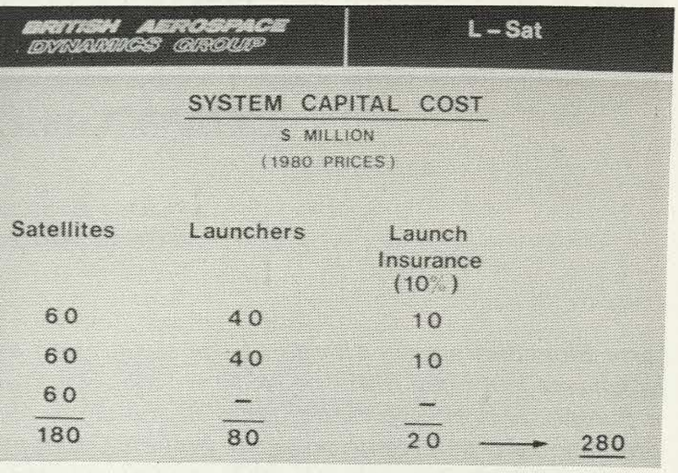

Satellites	Launchers	Launch Insurance (10%)	
60	40	10	
60	40	10	
60	–	–	
180	80	20	→ 280

What will it cost over 10 years? In my next slide, I have taken the capital which you have to pay back, and looking at interest typically, depending on how you borrow your money or how you pay your dividends, I have nearly doubled it, to another $250 million. We know that to operate and keep a satellite in orbit costs roughly $7 million a year. This is for the ground stations that control the satellite in orbit, the network control centre, the satellite control centre, the computers, and the things that keep it on the station.

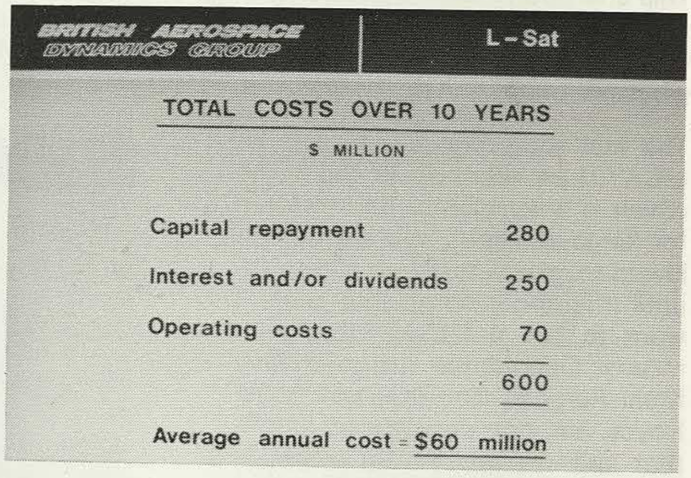

Capital repayment	280
Interest and/or dividends	250
Operating costs	70
	600
Average annual cost = $60 million	

A typical price is $7 million a year. The life of a satellite is designed for 10 years, so you will have to find another $70 million, which should come out of revenue but I have taken it as initial capital. So over 10 years you have to pay a total of $600 million, an average of $60 million a year.

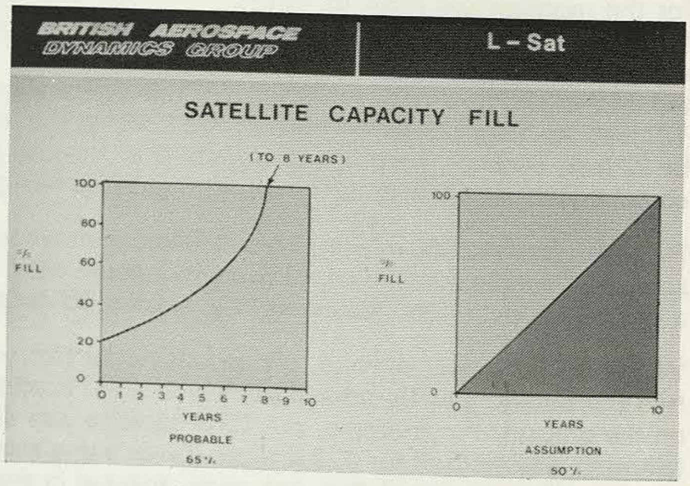

What do you get for your $60 million? First, you have to estimate how your satellite will be used and my next slide illustrates this. We have great fights with the Post Office on this because they are beginning to use satellites. They assume the righthand curve, which is the most pessimistic possible assumption. In other words, when you start the system nobody knows the system exists, and you start with a satellite that is empty. You have done no selling in advance. You start with an empty satellite and, hopefully, you fill it up after 10 years. What we say is that, with adequate selling beforehand, at the time you launch you have customers waiting for it, as SBS have in the United States. They launched in November and they have about 11 customers. Then you follow the normal growth curve of communications and you fill your satellite before the end of the life. So you could assume that 65% of your satellite will be used in its lifetime. But we will take the most pessimistic assumption — the Post Office assumption — and take only 50% usage, so you are using only one half of the capacity of your satellite over the total period.

So you have to find $60 million a year. The L-SAT that I am talking about in the particular configuration which I will show you later, for business use over Europe, can offer 80,000 half circuits, that is 40,000 simultaneous telephone conversations. Take the pessimistic 50% filling assumption and you will be able to sell typically only 40,000 half circuits. The cost per half circuit then is $60 million a year divided by 40,000, which is $1,500 per year.

For a two-way circuit, assuming that you lease a two-way circuit between where you live and where another office is located, the cost is $3,000 a year. At current rates of exchange that is £1,250 per annum. If you are the Post Office who have bought that satellite, you could sell it at £2,500 per annum. But the current Post Office price for a leased circuit between London and Oslo is £22,000 per annum. That is why Post Offices buy satellites. That is why your local telephone bills will go up by a factor of 10 in the next five or six years, because the Post Office accepts that their long-haul prices are totally unrealistic.

In their press conference last week about our new ECS satellites, the Post Office has accepted that satellite link prices must not be tied to terrestrial link prices but must be economic. Up to now it is the long-haul circuits that have been subsidising your local telephone calls to a very great extent. The Post Office, pushed by us, are accepting that long-haul telephones — which are not all that long-haul, they could be London/Edinburgh, or London/Oslo, or London/Paris — should be priced in accordance with a reasonable return on capital of the satellite system. But that means that a subsidy will no longer apply to your local telephones. I regret that part of my selling activity will result in your telephone bills going up — but that's life.

So $60 million for a large satellite and $30 million for a small one. That is a lot of money — why does it cost $60 million?

What I should like to do is to take you for a quick walk through our factory to show you why it costs $60 million for a large satellite. We launched the OTS satellite in 1978, and it has been up in orbit for nearly three years. It is standing over the Equator, at 10° east. It has been used by the Post Offices and television authorities of Europe since 1978, for trials in communication and television broadcasting. We were lucky when we launched that satellite because the rocket put us within a few degrees of its final position. So we did not have to use much fuel to get it into its position in orbit, and we therefore have two to three years' life left. OTS is becoming available for trials to people who want to use it in April this year.

My next slide is a picture of one of the views inside the satellite body. A satellite is not just an empty dustbin with some solar cells on it. It is in fact a telephone exchange. That is one picture of the inside of a satellite showing some of the structure and wiring.

The next picture shows the central tube of the satellite which contains the final apogee motor. That is the motor that finally fires you in apogee into high orbit, and the live motor lives in that tube. Round it are those balls made of titanium which contain the gas — hydrogen — that keeps you in orbit. Satellites are subjected to perturbations due to solar pressure, dust and gravitational effects, and float around. We have to try to keep them within plus or minus 0.1 of a degree to the observer on the earth. So about once a week you let it float about 60 or 70 kilometres, then you fire it back again and stop it. This uses gas which is stored in those containers. How much gas you put in determines the life of the satellite. We designed OTS for about five years and there is plenty left. But for L-SAT and large satellites of the future, we (and the Americans too) are designing for 10 years, which means more gas, more weight, more mass into orbit, and more effort to get it up there.

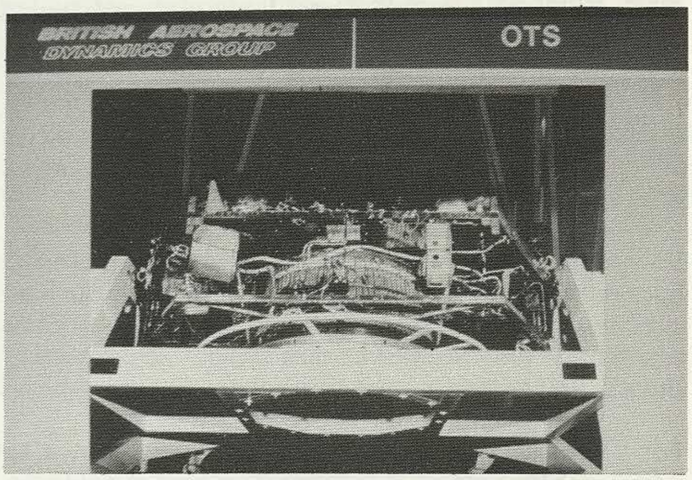

Here is a typical picture of a worried workman in our workshop. You get very worried looking guys in this particular shop; you see them walking round with a bundle of wires, with seven ends coming out of one end and six out of the other. They walk up to the satellite with this bundle and it does not fit. In that particular loom alone there are about 2,500 ends; that is one side of one base of the satellite. There are four major looms of that nature, so it is about 8,000 to 10,000 connections. All the equipment has to be space-proven and suitable for use in high vacuum.

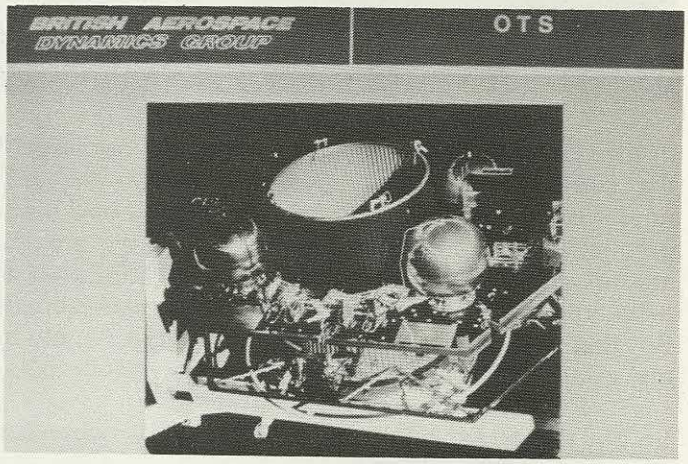

We have to build satellites in clean conditions. What you see on the slide on the next page are satellites in construction at Stevenage now. We have two going through on what is effectively a production line. Two satellites is a production line, out of a total of eight. These are MARECS, the maritime ECS satellites. On the left is a picture of a clean room showing the MARECS satellite being put together. It is a development of OTS and being used by the Post Offices. On the right is another view of the complexity inside the body.

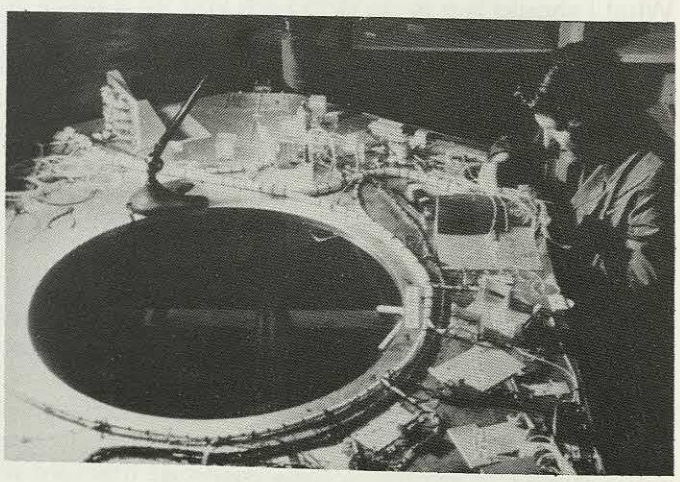

We have to look after the temperature of the satellite. Once it is in space it is subjected to direct sunlight on one side, and almost absolute zero on the other, looking into black space. So you have to control the thermal environment internally by the use of thermal blankets, heaters and special surfaces. The surface which faces the sun has a particular absorptivity / reflectivity ratio to sunlight, so that the amount of sun falling on the surface and entering the satellite is totally controlled, and the same applies on the other side which is looking into space. So the surface of the satellite and all the equipment in it has to be designed for a controlled temperature, which takes quite a lot of software and control.

The next slide shows the solar cells. This is the structure that hangs on the side of an ECS or MARECS. They span a total of about 50 feet from tip to tip. On the one shown on the slide there are 12,000 solar cells, and the price for each cell was about \$20, so on that picture there is about \$¼ million of solar cells, which is expensive, but of course the power is free. On L-SAT, with 80,000 solar cells giving about 7 kilowatts, the cost to us of the solar cells is roughly \$2 million, which is another part of the cost.

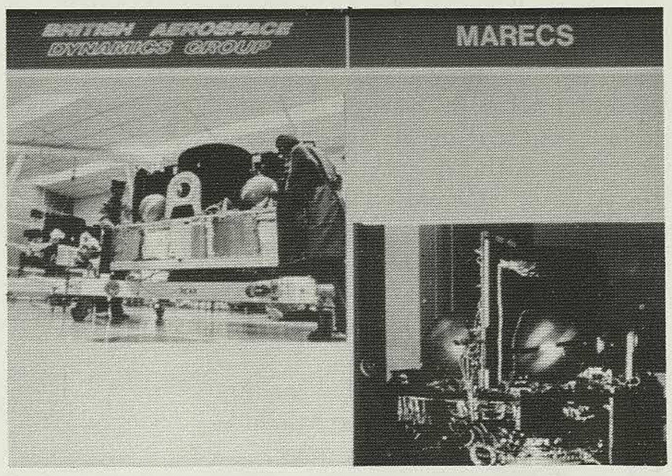

There are some beautiful mechanisms inside satellites, and one delightful device keeps the satellite pointing at the earth accurately. This is done by infra-red sensors that look at the earth's edge. They can either see the warmth of the earth or the coldness of black space. There are four of them, looking at four points on the circle. If the satellite tilts, one of the sensors will see black space. It signals a central microprocessor and says, "I can't see the earth any more". The microprocessor waits for a given time, because it might be a temporary perturbation, but if the sensor keeps on saying, "I can't see the earth," the microprocessor finally says, "OK, I know it's you." It then, with pre-set programs, sends signals to four wheels. They weigh about 2 lbs each, and are mountable magnetic wheels, totally frictionless in a vacuum. There are three of them, in each axis of the satellite, and one spare at an angle. The microprocessor instructs the wheels to spin. It knows which ones to call up to spin. It might call one to spin for 4½ seconds and another for a few seconds. What is happening here is that the satellite feels inside its belly some electrical power being converted into mechanical energy, and it has to conserve the angular momentum. So when it feels these little wheels spinning it has to correct because of the conservation of angular momentum. It will shift one-hundredth of a degree, until the sensor sees the earth coming back, when it signals the microprocessor and everything stops. With the accuracy that we will need in the future we cannot depend on thermal views of the earth. They are good to 0.014 of a degree, which is the accuracy with which they can sense the earth. We will have to be more accurate. We will have to point satellites to radio beacons on the ground. That costs money, too, as you can imagine.

We also test satellites in our Anechoic chamber. It is black and the size of this room. It is surrounded with radar-absorbent material, made by Plessey, which works from the VHF range up to the gigahertz range. We can mount a satellite in the chamber and make sure that its aerials are working properly and that there is no electromagnetic interference in the satellite. The chamber costs £1 million, which is another price that has to be paid off. That is just a quick view of what goes on at Stevenage at the moment.

This slide shows a chart which gives you a better view of what is happening right now, at the beginning of 1981. The OTS satellite was launched in 1978. In April 1981, it comes to the end of its formal trials and will be available for users who would like to try satellite communications. We already have several companies, including oil companies and large multinational companies, who would like to try videoconferencing and save travel costs. We expect it to last for another couple of years, but we are not absolutely sure of that.

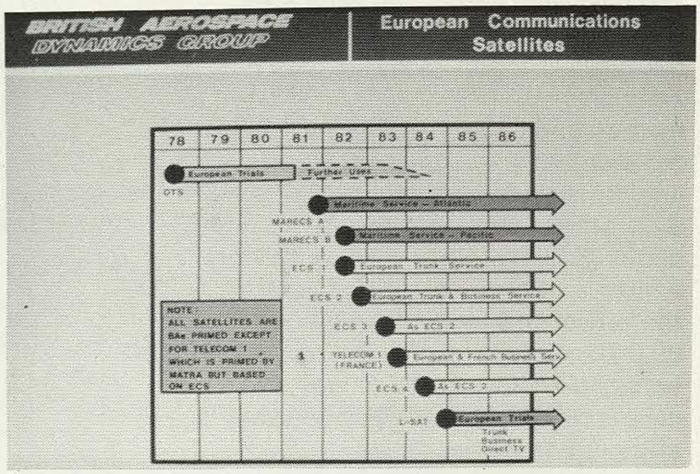

OTS was a test satellite which gave us the opportunity to test in orbit the various equipment that we are now putting into the seven satellites that we are building right now.

Let us look at the various satellite coverages. OTS's coverage diagram is of the shape shown on the slide. It has a spot beam over Europe and some Eurobeams that cover much larger areas. The spot beam is the one that offers the facilities that I have mentioned, for trials by people who want it via the Post Office.

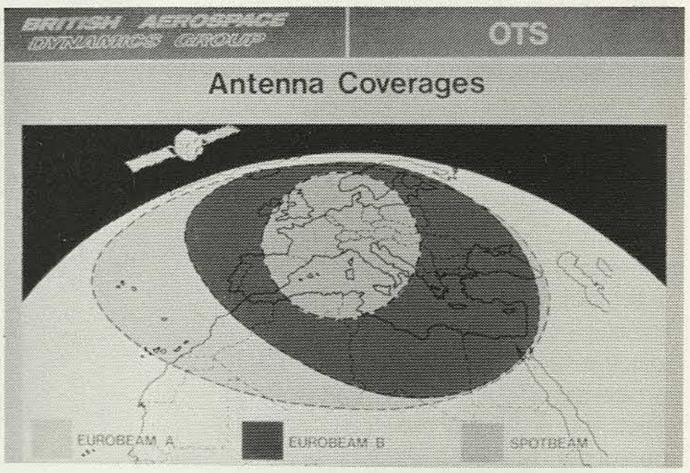

The next slide is a picture of our own station which we are operating via OTS at the moment. This is a 3½ metre dish. Notice, regrettably, that it is a Nippon antenna; in fact we are going to change it for an Andrews because we are going to transmit from British Aerospace. It is to our regret that there is no British manufacturer who can help us there. We have talked to Plessey and Marconi. Plessey and people like them do not seem to be aware or ready for this enormous market. It does shame us to have to use a Japanese antenna at British Aerospace.

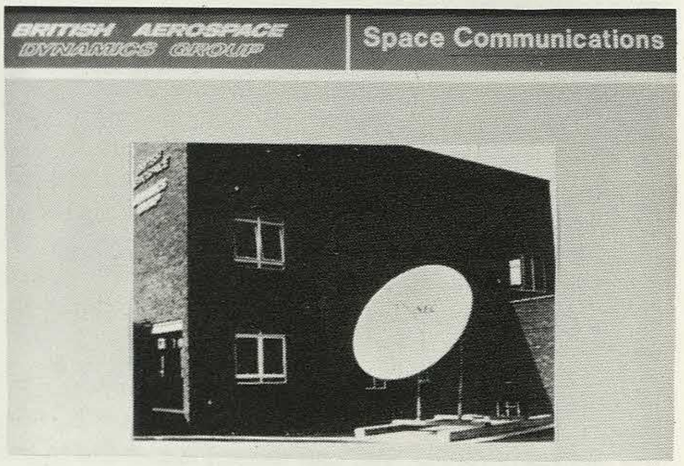

We use the dish at the moment to take trials continuously. The French put out news every night to Algeria, so we pick up French news at 6 o'clock. We have picked up some of the Italian porn which has been retransmitted via Turin. We get the amusing sight at night, around 7 o'clock, of a crowd of dedicated engineers standing in front of all this equipment. The dish is the size of station that people like you will be using, about 3½ metres. It arrived just before lunch a few months ago, and at 6 o'clock in the evening it was on site and we were locked on to the satellite. That is an indication of what you can do with satellite communications: you can be online in half a day. Assembly of the station did not start until about half-past one and we were locked on to OTS at about 6 o'clock, and we have been ever since.

Let us look at ECS. This is the follow on development of OTS. OTS was the satellite which told us how to build things and how not to. We have had some failures in OTS and we know what can go wrong, but it is still working. We have now developed ECS, and we have sold five of them to the European Post Offices, under the auspices of EUTELSAT which is centred in Paris. EUTELSAT is the international consortium of 17 European PTTs, which is all the European nations plus Greece, Turkey and Cyprus. They have recognised the potential profitability of buying satellites of this nature, at roughly $30 million each because these are small satellites. The intention is for launch in 1982, and there will be four over the Equator, covering the entire European continent from Turkey up to Norway, Greenland, Portugal, and North Africa. Already, the EUTELSAT people are forecasting that, depending on how much traffic they put through these satellites, they should be saturated by about 1986 to 1987.

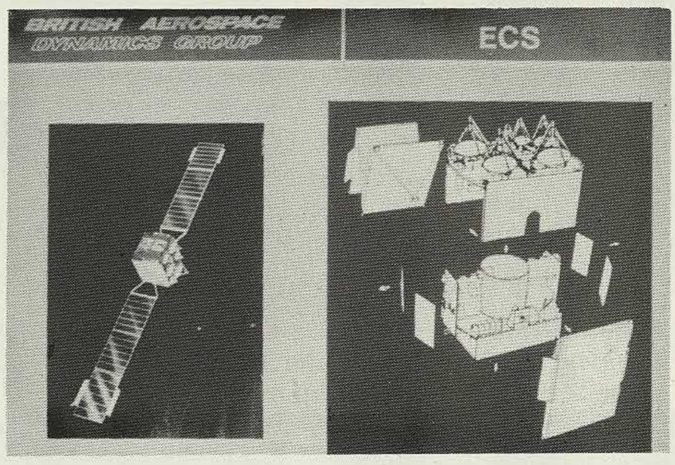

It is worth mentioning in passing that it is our normal practice now to build what we call a "modular" satellite. The righthand side of this slide shows how the satellite splits down into effectively four main sections. The bit in the middle is called the bus — the one with the tube going down the middle. That is standard for whatever satellite application you are using. For ECS it is like that; for MARECS it is like that. If we sell a satellite to the Arabs, which we may, it will look like that. It contains all the controls to keep the satellite stable and in orbit, and to get it into orbit. Above it is the meccano looking bit, which drops on to there and carries the payload. On ECS there is that great mass of antennae which I will show on the coverage diagram, and below it all the communications equipment associated with that particular mission. That is the bit that varies from mission to mission, and on each side are the solar cell arrays.

ECS is relatively small, but you are already talking about 12,500 simultaneous telephone circuits plus two TV channels. 120 megabits per second is the operational rate of the transponders on the satellite working with fairly large earth stations owned by the Post Offices in Europe. ECS can carry 12 transponders; nine of them work in sunlight, three on standby, and five in eclipse. I shall not go into the geometry of the eclipse, but twice a year the satellite passes into the earth's shadow and it lasts a maximum of 72 minutes. It peaks up over a two-week period. During that time you have to run on batteries for those 72 minutes, and you try to position your satellite so that the eclipse occurs in the early hours of the morning. On the second stage of ECS we are now making it fully active in eclipse, which means carrying more batteries, which we can do. So ECS, which is a relatively small satellite, can already handle 12,500 two-way circuits for Europe.

The slide at the top of the next page shows the coverage of ECS. Those antennae that you saw in the drawing on the right are designed to give spot beams of fairly high power density at the ground, the major one being called Spot West (the one over central Europe). It also covers the North Sea, which is of interest to the oil companies operating out of the North Sea and the North Atlantic. The Spot Atlantic one does cover the Canaries, but it is very much a wasted beam. Spot East is most of Eastern Europe and the Mediterranean, Cyprus, Greece and Turkey. There is one Euro beam which is a much larger beam, and because it is larger it is much weaker; but it is used mainly for Eurovision transmission of television from any country in Europe to any other country, via the Post Offices.

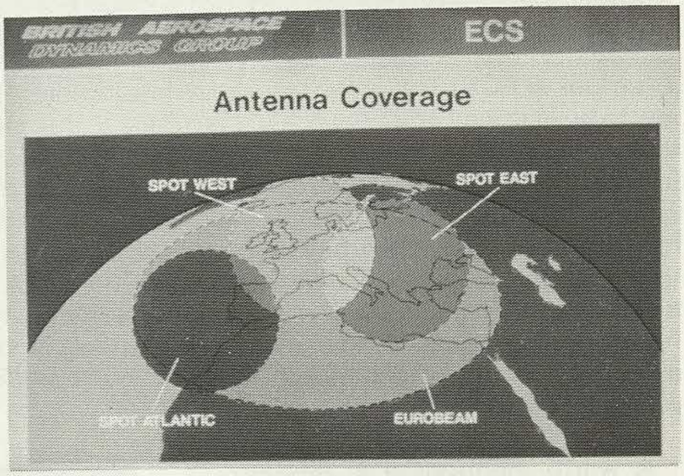

What we are aiming to do now, and what the Post Offices have asked us to do, is to add to our second ECS, which is in construction, a business package. The Post Office announced this last week because they recognise, finally, partly due to our poking their eyes out, that businesses do want communications via space. So we shall be carrying business packages working in a different frequency for use with rooftop or car-park antennae, from 1982.

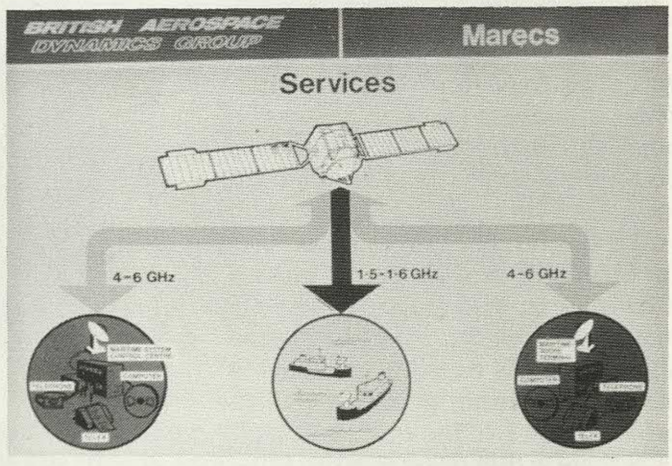

Before ECS itself is launched, a development of ECS — MARECS — is due to be launched at the end of 1981. Two MARECS are scheduled and they are Maritime ECSs. We have sold three of these to the International Royal Marine Satellite Organisation which is centred in London. Its job in life is to provide communications to and from ships. About 70 nations are members of it, including the Russians and the Japanese. It has ordered a development of ECS with a particular antenna on it. You can see from the top diagram on the slide that the satellite is the same shape as ECS except for the module that carries the antenna and communications equipment.

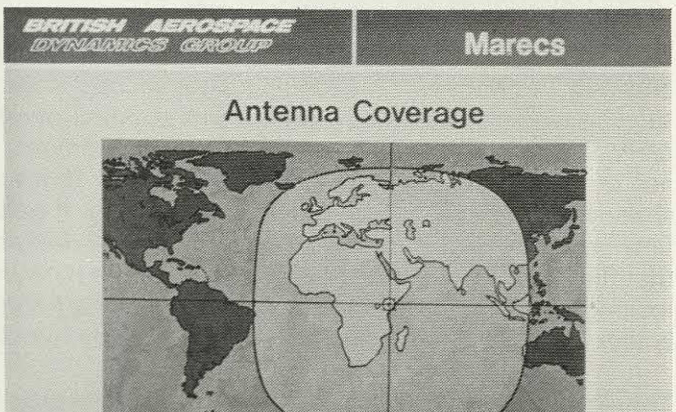

MARECS has one large antenna to cover oceans. If it were centred over the east coast of Africa, as shown in the slide, that would be the coverage of a MARECS satellite. In fact the first MARECS satellite, MARECS A, which is due out of our works in a couple of weeks' time for final tests in France, will be placed at 26° west, which puts it over the Atlantic. So the Atlantic traffic will be covered by MARECS A. MARECS B, our second one, will be right round the other side of the world, at 179° east, covering the Pacific. For the moment the Indian Ocean will be covered by the Americans. Then we have a third one available as a spare. This gives communications to ships carrying antennae only 4 feet in diameter, fully stabilised because the ship rolls and pitches and the antenna has to go on looking at the satellite. It is a very small antenna, with very poor performance; and because of this the number of channels that can be handled is very limited. You remember that ECS could handle 12,500 telephone channels, working into big, 15-metre Post Office stations. This satellite, which is identical, is working into tiny antennae on ships. As a result, it can handle only 35 channels. But these 35 channels mean that one ocean can be shared between perhaps a thousand ships. Several hundred ships can have on demand access to a channel, so that a ship can order a channel on demand. We are now offering that service to the oil companies in the North Sea. We expect the typical waiting time for a channel to be not much more than two or three minutes, and in the early days there will be no waiting at all. But eventually the big satellites will take over.

Another development of ECS is the French Telecom 1 satellite, providing European and French business services. With the exception of Telecom 1, British Aerospace is the prime contractor for all the satellites I am talking about today. Matra is the prime contractor for Telecom 1, but it is based on ECS.

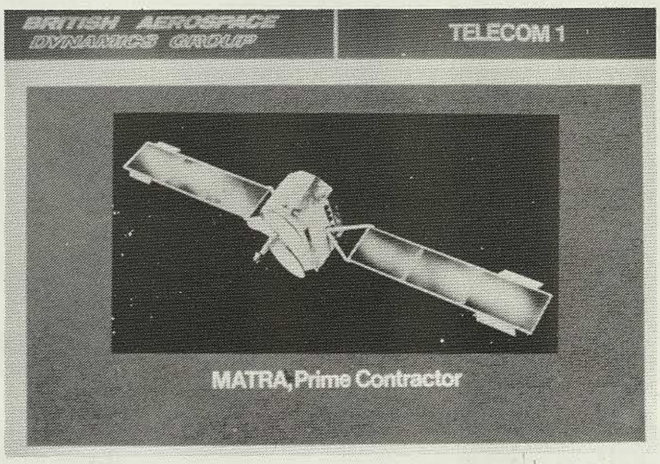

Whatever the French say, Telecom 1 is a British Aerospace satellite. This slide shows their drawing of it, and they have drawn it back to front to make it less like ECS, but it is in fact an ECS structure with their own module on it. That is being built by Matra to our design. The French have charged ahead here.

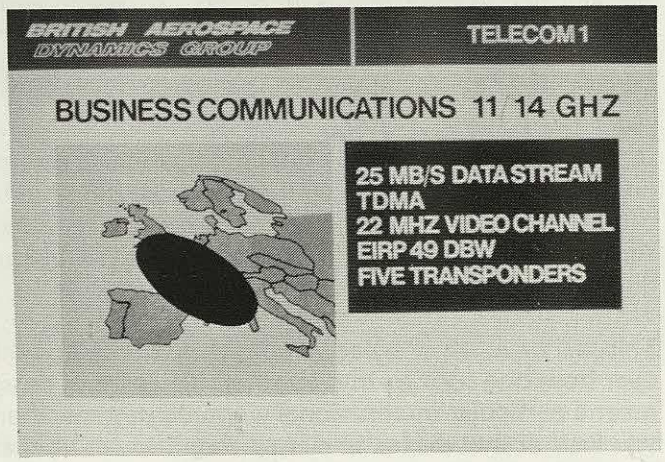

The French have decided to offer a business system, for use into 3 metre antennae, mounted on rooftops and car parks, for the whole of French industry. This slide shows the coverage pattern, which covers southern England too. Telecom 1 has a fairly high data rate of 25 megabits per second, works in time division multiplexes (which is a fairly complex way of doing it), provides a video channel, and has an EIRP (an equivalent isotropic radiated power) at the ground of 49 DBW (which is high) and it carries five active transponders, just like ECS.

In addition, the French have decided to put in a global beam. You can see on the first slide on the next page the French intentions on Africa. They have covered the entire African continent with a fairly low level beam, but nevertheless, there are three transponders covering Africa. Telecom 1 will be the first active satellite available to the Africans for intra-African telecommunications traffic. They have also thrown out a beam to north of South America, to cover what was French Guiana. Quite right too. It means that the French have communications to their own launch site. But please remember, despite what the French say, that Telecom 1 is a British satellite.

So, we are now building the MARECS and the ECSs. MARECS is due for launch at the end of this year, and ECS4 will be launched in 1984. After that comes the "large one", which is the next test satellite for Europe. OTS was the test satellite which led to the existing series. L-SAT, the "large one", is coming into being now.

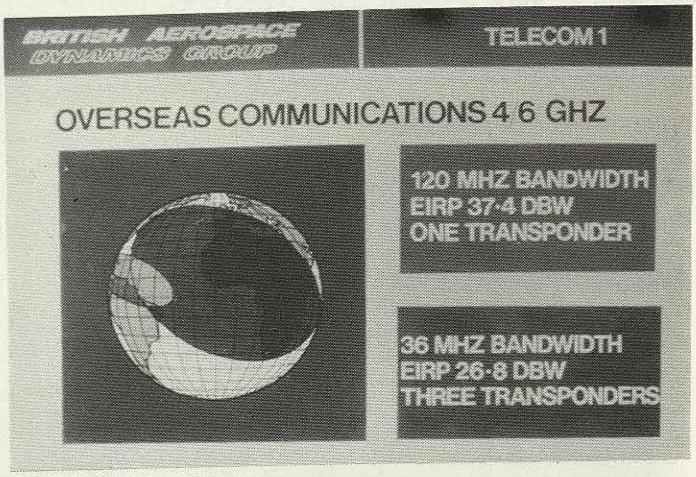

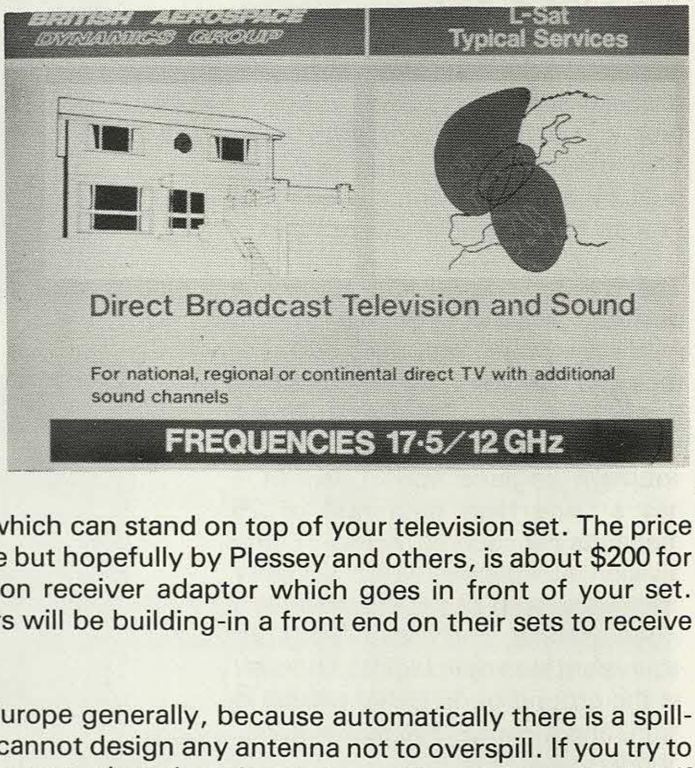

The second slide shows us what L-SAT can provide. The most impressive service that will affect the public is direct broadcast television. When our first L-SAT goes into operation, it will have a direct television service to Italy, because the Italians are very large contributors to the L-SAT programme. They have ordered one of our antennae to point at Italy, to provide direct television over the Italian Peninsula. But we have put in, in addition, a steerable antenna for direct television. I have plotted out on the slide the shapes that those television patterns will take. The British one is that funny shape because it is not a right shaped antenna, but it will cover the UK, Holland, Belgium, Luxembourg, Sweden, Denmark, Spain, France, Germany — in fact every European country. The steering ability of the antenna is enough to cover the whole of Europe.

Because of the power that we are transmitting in that frequency of 12 gigahertz, you can receive the picture in an 80 cm. antenna, shown on the front of that house, with a very cheap receiver. The receiver that we have at Stevenage is a little box which can stand on top of your television set. The price being talked about, again by the Japanese but hopefully by Plessey and others, is about $200 for the antenna, its mount, and the television receiver adaptor which goes in front of your set. Almost certainly, television manufacturers will be building-in a front end on their sets to receive direct television.

This will have some massive effects on Europe generally, because automatically there is a spill-over from one country into another. You cannot design any antenna not to overspill. If you try to meet a particular multinationally agreed pattern, then there has to be a spin-off. It means that if you look at that ellipse that is spilling into south east England there, which I think is the Dutch beam, that is effectively the size of the beam for Holland in which you could receive with an 80 cm. dish. If you live in Edinburgh, the signal is certainly weaker, but if you put up a 1¼ metre dish, which is not much bigger, you could pick up Dutch television. Equally, you see the size of the Italian beam: that could be picked up with a 1 metre antenna. Equally, British television programmes can be picked up in Spain. You can imagine the social concern that is going on at the moment about the Spanish picking up Swedish television, and about the British picking up the nightly material from Rome. There are 600 stations operating in Italy at the moment, pouring out the amateur material — and it will all be available in Britain. We had the BBC up only a few weeks ago and they were filming our own television set. We had the French news on and then some other French programme. They were quite happy, and they stopped filming. Then the Italians came on next, transmitting the most gorgeous bit of porn, but the BBC people had run out of film. They wrung their hands in despair, and we were delighted because Mary Whitehouse did not see what came over that evening. But people like that will get very concerned.

Then there is the problem of advertising. For instance, Belgium has no advertising, but it will get it. The Luxembourgians are suddenly given the possiblity of increasing their coverage from their little town of 100,000 people to 10 million viewers. So there is real money involved, and the advertisers are having orgasms over what is going to come in Europe, and the multinational advertising agencies are already beginning to design advertising programmes which are multinational and have minimum language content. So that will have a tremendous effect, but it is only one aspect of what a large satellite can do.

Here is what the large satellite can do for the Post Offices of Europe. It is not dissimilar to what we can do with ECS, except that here we have done typically four beams, working in the frequencies in which we are already working. This is for international exchange between large antennae in the capital cities and main provincial cities of each of the member nations. You can see the coverage, again with a big Euro beam for television for subsequent surface broadcast.

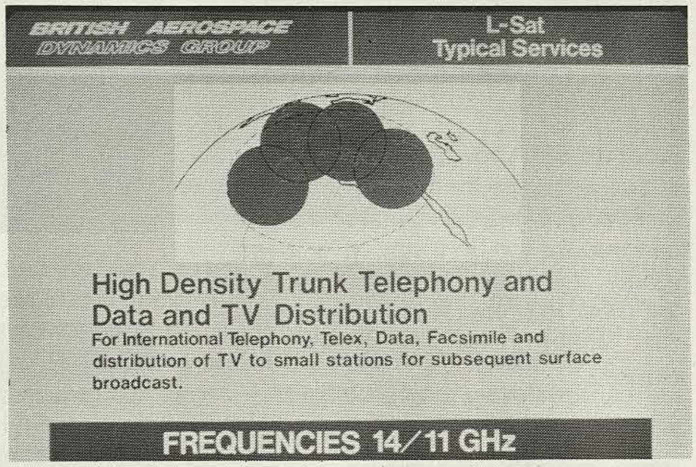

Also, within L-SAT we can add an overlay of the very high new frequencies which we are testing on L-SAT 1, in the 20 and 30 gigahertz range. We are limited in bandwidth down in the 11 and 12 gigahertz range because of interference with ground systems. In the 20 to 30 gigahertz range, we are much less limited and have much wider bandwidth. Using part of the 20 to 30 gigahertz range, we can handle 4 gigabits per second, working into large stations, which is a very large capacity. So that is what L-SAT can do and will be planning to do, to replace the ECS satellites in 1986/1987.

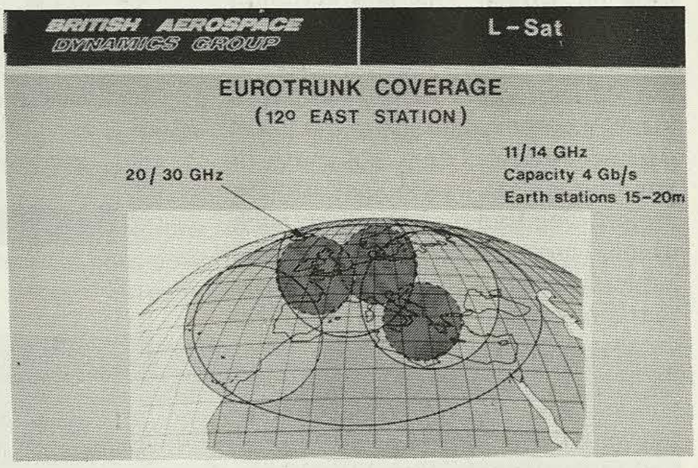

The first slide on the next page shows an L-SAT beam configuration that the Italians are interested in. You probably know what the Italian telephone system is like. They have decided that it would probably be worthwhile to overlay the entire system with what we call "national intercity". The slide attempts to show that for national, high density trunk telephony, you can concentrate beams on to your major cities, and because of the size of satellite and because you are working in very high frequencies you can produce beams as small as 150 kilometres in diameter.

The other thing that the large satellite can do is provide communications to/from mobiles. We have already talked about ships, and that will continue. The second slide on the next page is a more accurate picture, and up on the bridge of the ship to the left of the arrow you can see the size of antenna that will be mounted on the ship. The aviation industry is not really interested, despite their rotten communications once they are out of reach of VHF; they do not want to spend money. But it will come; we will force them to take it. But the people who really want the service are the international truck lines. They are fascinated by it, and they have shown us how much they can save if they can communicate with truck drivers going from, say, Sweden to Saudi Arabia. They need to know where the truck is and when it is due to arrive. If they can tell the customer the exact date of arrival of that truck, he can organise all his lifting gear and not have it waiting for a week. They can divert trucks, as ships already do, to pick up loads. They can be in constant communication with the trucks.

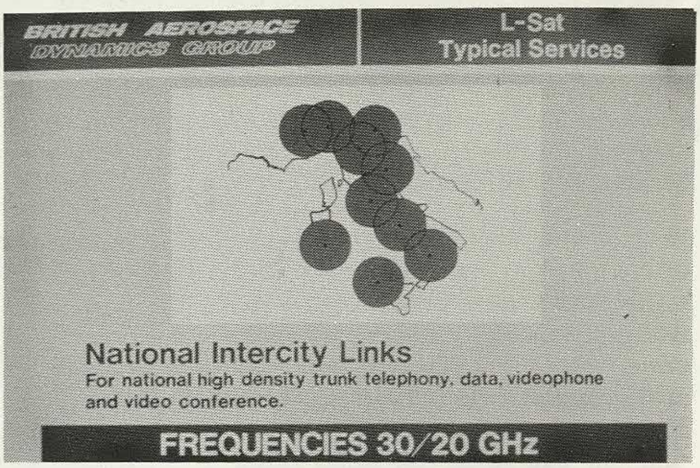

Now we can provide that service. It is bad enough on a ship, but you can imagine the sort of antenna you can mount on a truck. It is even worse. It is a very poor antenna. But we think that we can handle beams that would cover several thousand trucks.

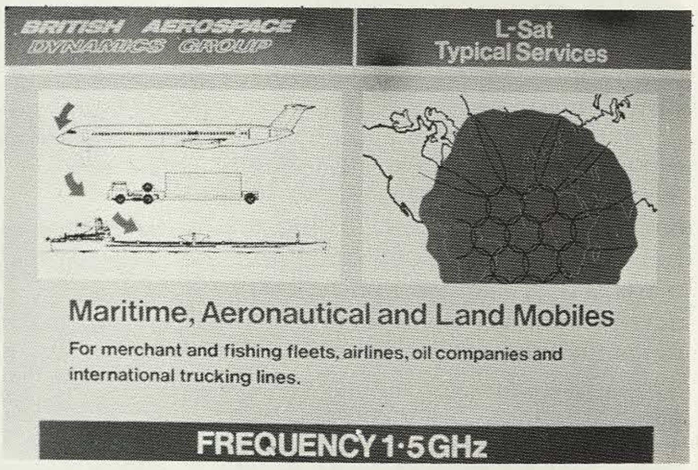

Once you are in the large satellite range you can cover many services, such as those shown on the next slide for Africa. This is a situation where we show typically the sort of beams that might be required for the African continent. There is a mass of communication in the West African region. There is a mass of communication in the East African region. The two do not talk to each other. The only link between them is a microwave link through Zaire, and that broke down in 1978. There is no east-west communication. You cannot get to the telecommunications link in Zaire because they run on diesels, and you have to get fuel to the diesels. The 20,000 kilometres of road in 1972 is now down to 7,000 kilometres, because they are overgrown. They just cannot keep a microwave link going. In this case we would suggest an overall beam for east-west communications in Africa, with high density beams working into small stations all over East Africa and all over West Africa.

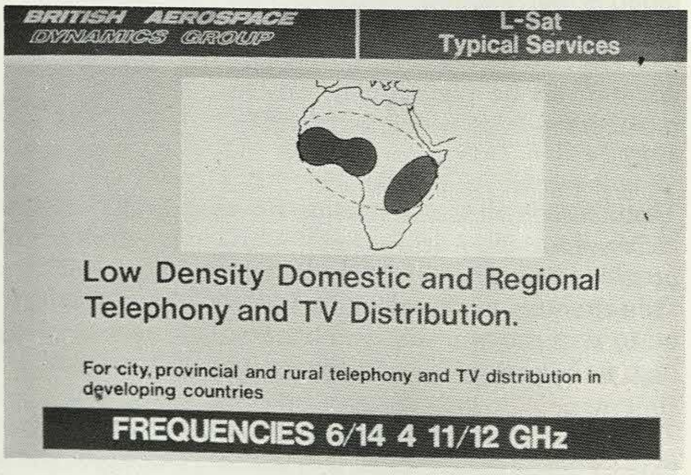

Let me come to what really is most fascinating and more interesting to you — Specialised Comercial Services. Here you are talking about business to business, car park to car park, or rooftop to rooftop. Already, because of people to whom we have spoken and because of the imminence of the service, the Daily Telegraph, who are putting up a new building in London, are designing their roof to carry their own 3 metre antenna. Several other companies are designing the same thing for London. People are preparing spaces for those sort of antennae in car parks.

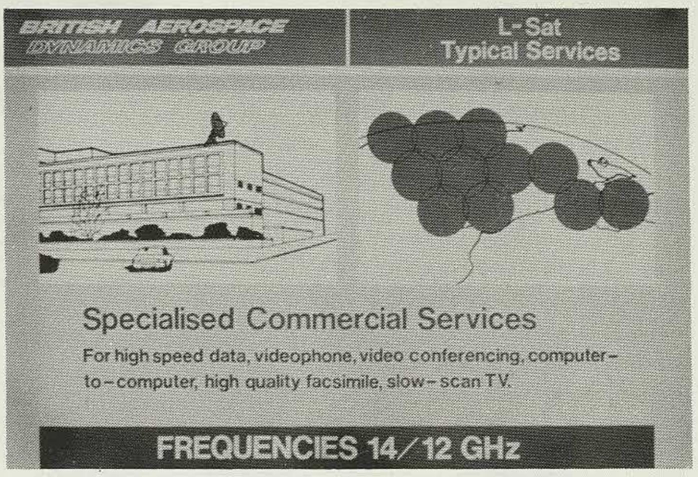

We are now talking about an antenna of the size that I showed you outside our building in Stevenage — 3½ metres. The sort of coverage that we can provide is shown on the right-hand side of this slide. We are now talking about communications that are independent of distance. In any of those beams there is interconnectivity in the satellite. It acts as a telephone exchange. So if you are in Scotland and you want to talk to Athens, you can: up to the satellite, across the switch, and down to the beam that is going down to Athens. If you are in Edinburgh and you want to talk to Bristol, you can: up one beam and, in that case, down another. We are talking about independence of distance here. We are talking about the essence of space communications, which is the message that we are trying to get across — the flexibility of a system like this. Once you get that station in, on site — and it can be done quickly, it can be put in on the back of a truck in a day by the Post Office — you then have the flexibility of all of the bandwidth of that satellite that you wish to use. In other words, if you wish, you can have a video conference in the morning, using say 8 megabits which you pay for, for a couple of hours. You then do not require that bandwidth any more and you revert to, say, 24 telephone channels. In the afternoon you may want some data. The flexibility that comes once you are online is what matters and the speed at which you can do it.

Equally, if you have a multidrop requirement (for example where you have many buildings that belong to your organisation), and you want to transmit a lot of data to each site — if each site has a ground station, you multidrop by broadcasting. You can set up a mesh network, you can set up a star network, you can set up point-to-point connections, quickly, and you can change them quickly. If you are in a nasty situation and have to shut down some factories, you do not waste the money that was laid out in bringing your wideband terrestrial link to your premises — you just drive the antenna away.

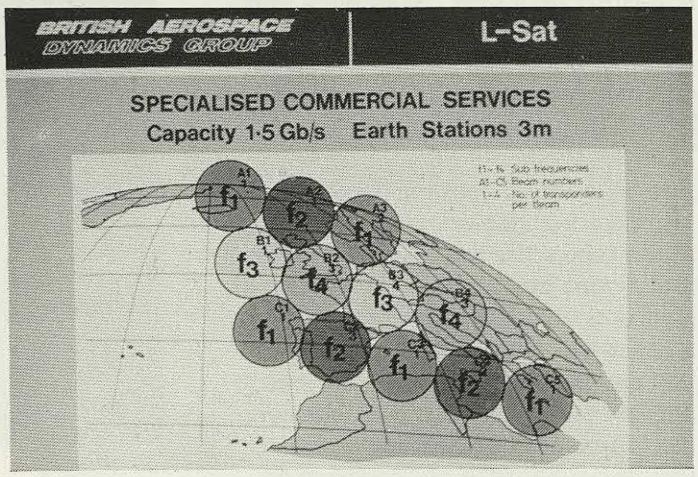

This is where we think it will have a significant impact — Specialised Commercial Services. Typically the L-SAT can offer a pattern like the one shown on this slide. We show here how we re-use frequencies, because we are short of frequencies. Those circles with "F1" in them represent a sub-frequency which we can re-use five times, because anybody within one of those beams cannot see anybody else in a similar beam because he is too far away. So we can re-use frequencies time and time again. With the limited bandwidth we can handle 1½ gigabits in that situation, working into small stations, 3 metres diameter. That is the sort of service that will become available by the mid-'80s by L-SAT, round about 1987-88.

The next slide gives you an idea of what the satellite looks like at the top of the rocket. We spend a lot of time looking at the installation of a satellite in, say, the Ariane rocket, which is the prime rocket in which we intend to launch the satellites. This gives you an idea of the antenna designer's problems. It is not much of a problem. You can either have antennae mounted rigidly on the front of the satellite, inside it, or you can play with them and erect them once you are in space.

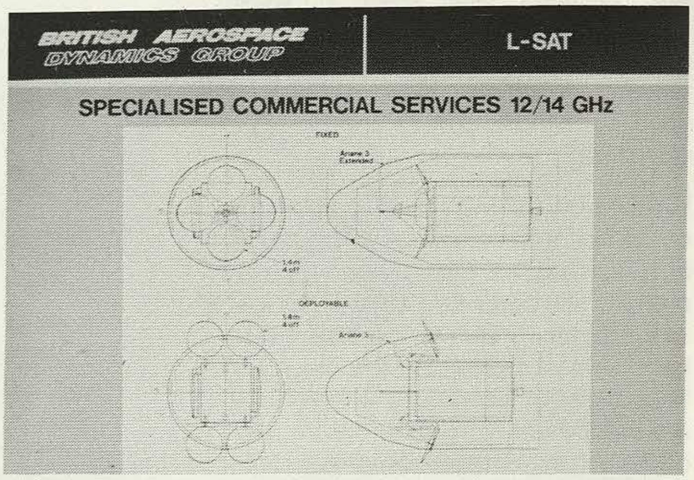

This is a more accurate picture of what I was talking about earlier — Italy. Because of the size of the satellite and the frequency at which you are working, we can put antennae on that give you beams that small. This beam pattern is not trying to cover the country, but it is coupling Turin to Bologna, to Ancona, to Rome, to Naples. Here we can handle 6 gigabits, which is several hundred thousand telephone circuits across Europe. This has been required by Telespatzia which is an Italian organisation. They want about 150,000 circuits. We can offer through the two antennae nearly 200,000 circuits, plus of course the use of slow scan TV, data, telex, and high speed TV, by using an appropriate number of channels. That can be squeezed into Ariane.

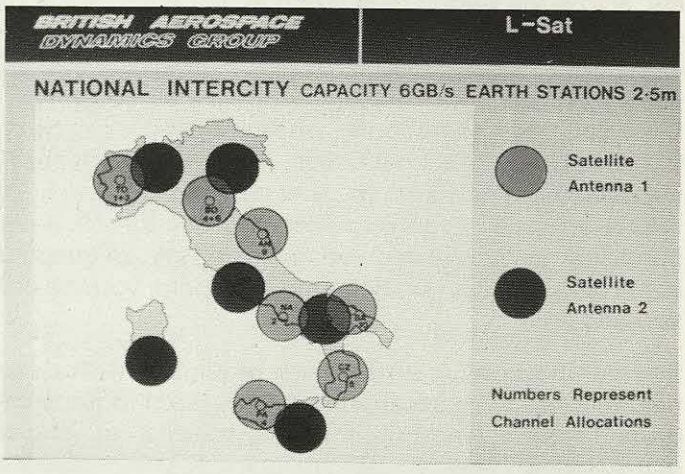

On this slide you can see much larger antennae — this one is about 4 metres in diameter. But we shall be making antennae that can be squashed up inside the satellite and inside the rocket. After the rocket loses its nose, the antennae can be erected in space. The antenna design is a fascinating problem.

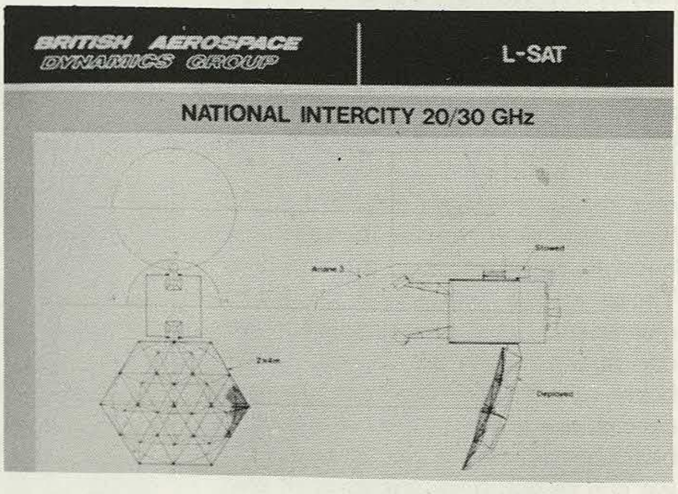

The other thing that we can offer is multinational television. On the slide at the top of the next page we have drawn two beams, one for Switzerland and one for Luxembourg. The smaller the nation the more power you can pour into that small area, and in something of L-SAT size we can carry enough transponders to give two services to two separate nations, that is five channels per nation. That gives you an indication of the interference that will occur between two nations, between Switzerland and Luxembourg. You can imagine overlaying those on each other; and you can see what can be picked up in the United Kingdom with larger antennae. A little mark, called a 6dB/k contour on the map, shows you the limit of where you can accept the signal to an 80 cm. dish. As you go out across the pattern a larger and larger dish will give you the same picture, but you will get more and more interference.

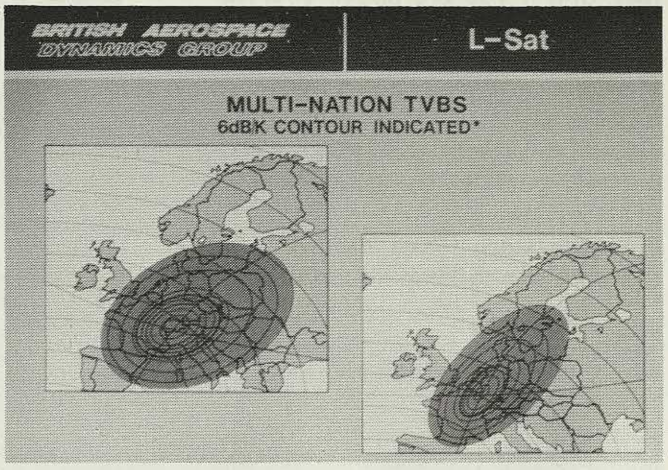

The next slide shows a way out thing that we can do, but this is for the late '80s and early '90s because we have not done the antenna design for this yet. But what you can offer here with an L-SAT is a capacity of 15 gigabits. Now we are talking about half a million telephone circuits across the whole of Europe; by using large antennae on the satellite, by re-using frequencies time and time again, and working into 6 metre earth stations, that is the sort of coverage you could put over the whole of Europe, and in fact we have particularly covered the North Sea.

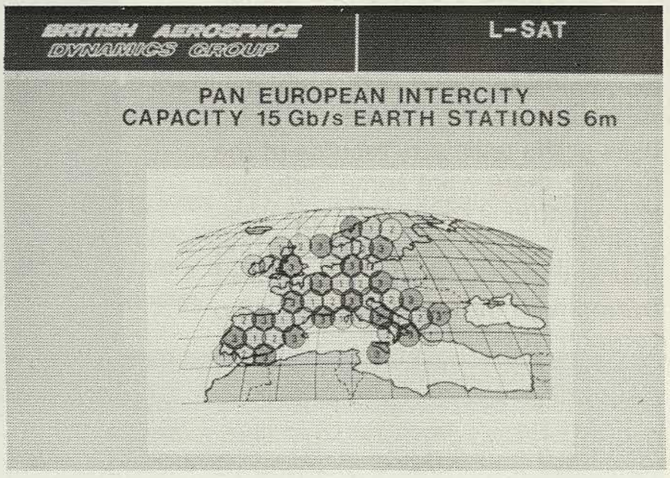

The third slide on this page shows a slightly more way out antenna. It is the largest antenna that we have thought about, which is 15 metres in diameter and working at a lower frequency. Because of the size of the antenna and the frequencies allocated, we can work into land mobiles.

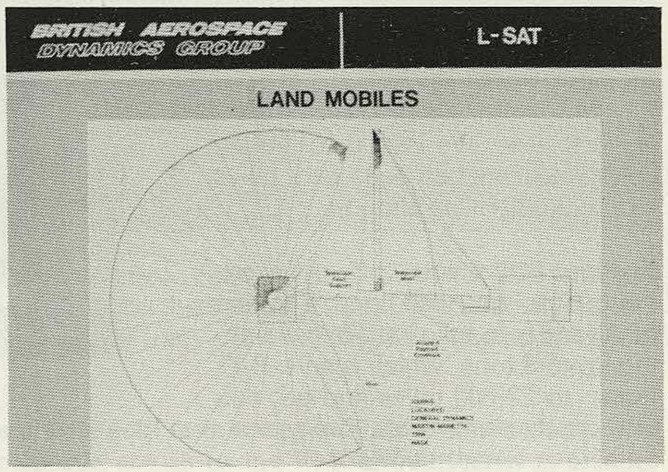

The slide at the top of the next page shows a typical layout to cover the whole of France, with sound broadcast and with land mobiles. With that size of antenna and the size of L-SAT, about a thousand trucks could be covered in that area. That is not good enough at the moment, but it is the way that we are working.

It is probably worth mentioning what you have to do to get the satellites up. We are talking about the geosynchronous orbit. The second slide on the next page shows the typical process of injection into geosynchronous orbit, launching into what is a circular orbit, above the Equator, at 22,000 miles altitude, injecting the satellite with the central apogee motor which lives inside a central tube, de-spinning it, because it is spun to keep it stable; when it is de-spun erecting the arrays of solar cells; getting the attitude right by using gas; getting the station right; and then sliding round the orbit by going slightly above or slightly below. If you are at exactly 22,000, you will circle the earth once every 24 hours and remain stationary to an observer. If you are below 22,000 you run too fast, and if you are above you will run late. So that wherever you finish up with your rocket in orbit, you then go up or down just a few dozen kilometres. You slide round the orbit, running late or early, until you get your station. At the moment you get your station, you fire into orbit and then stop. From there on you use your gas to keep you in orbit.

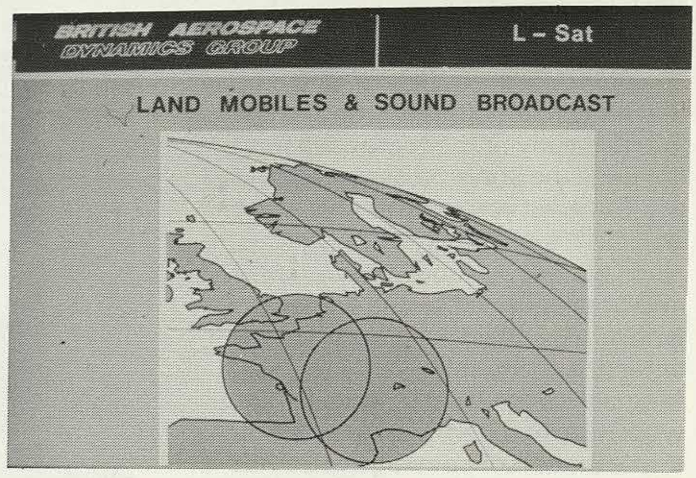

That is 22,000 miles, and that is where misconceptions abound.

Ariane is the European rocket mainly driven by the French. We make large numbers of parts of it, but it is essentially a French rocket. This is a picture of it on its launch pad at Kourou, in French Guiana. They have had some lousy luck. Number one went off perfectly. Actually it was not perfect. Looking back at the telemetry records of the motors, there were some fairly serious vibrations seen in the gas pressure in one of the motors during that launch; but in the euphoria of the moment the reaction was, "Who cares? It worked like a dream". However, Ariane number two blew up because of the same problem. They had a combustion instability. They now have a nightmare of trying to stop a combustion instability in their engines, and they delayed the next launch. This is worrying the insurers somewhat because the success rate of Ariane is only 50% at the moment.

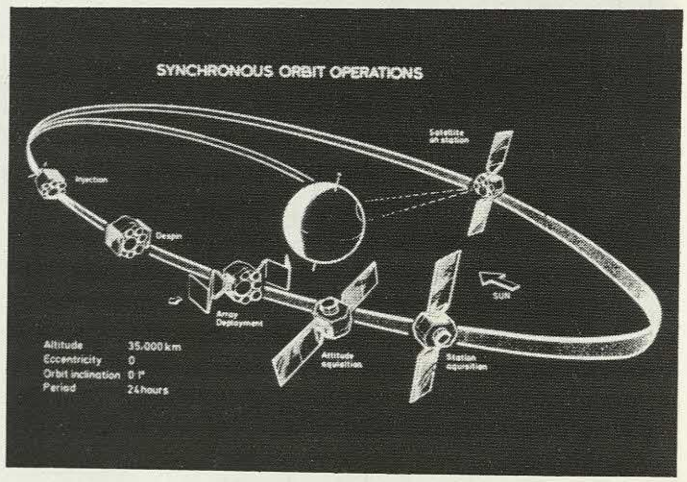

Nevertheless, we intend to launch an Ariane if we can because it takes us to 22,000 miles altitude.

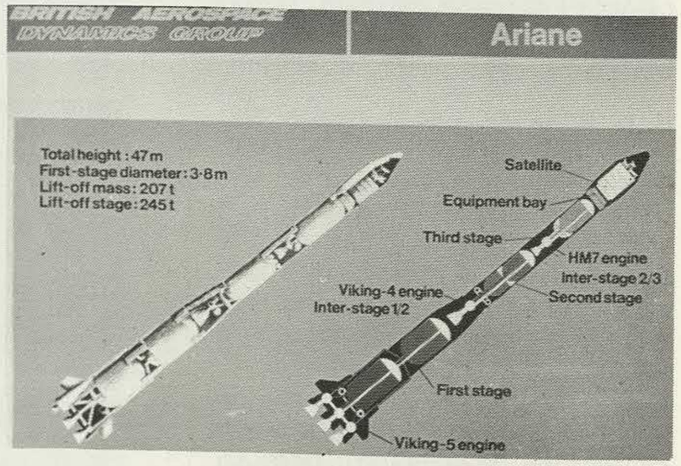

The idea is that you stuff your satellite up into the nose cone. The aim of Ariane is to take satellites weighing up to 3 tons, which is an L-SAT mass, to 22,000 miles and get you into synchronous orbit. The typical cost for a total launch is $40 million, which is not a large sum, looking at what it led to.

It is said that shuttle will be more economic. It is built by Rockwell in California and it is a gigantic, complicated beast which will work eventually. It is having terrible problems at the moment with its thermal transfer. It will be carrying all sorts of payloads.

This slide shows an artist's impression of the shuttle, carrying Space Lab (which is built in Europe), and behind it all the experiments that are subjected to the outer vacuum of space. We build the equipment that lives at the back. The problem with shuttle is that it goes to only 300 miles altitude. If you take up a satellite of L-SAT size, you have to carry with you another big rocket motor. You have to shunt the thing out of shuttle, and then light up and get out of the way quickly, to get this thing from 300 miles to 22,000 miles. In addition, you have to carry up with you all the test gear that you would normally have on the ground. We have always talked about ground support equipment, we are now having to talk about space support equipment. Obviously when you are preparing a satellite for launch it has to be right. This means a mass of test gear. So that when you do fire via shuttle you must make allowance for the cost of all the extra equipment that you must carry, and the rocket motor which is lost. When you do the sums you find that Ariane and shuttle are roughly the same price. Eventually shuttle may be cheaper, but we aim to operate out of both, preferably with Ariane.

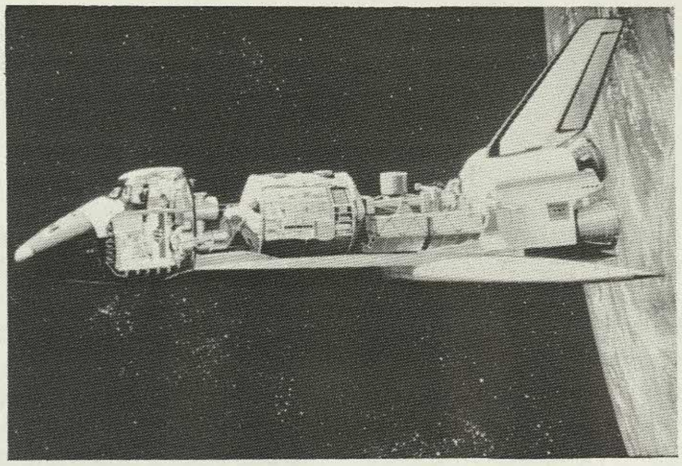

L-SAT is the end of a development line so far, starting with OTS which is available soon, ECS, MARECS, Telecom 1, and into the next series of test satellites. This will give us test transmissions to Europe for television, to Italy and television; very high frequency communications for Italy; and propagation experiments in the high frequencies.

And this leads to the essence of the whole requirement: business communications via space.

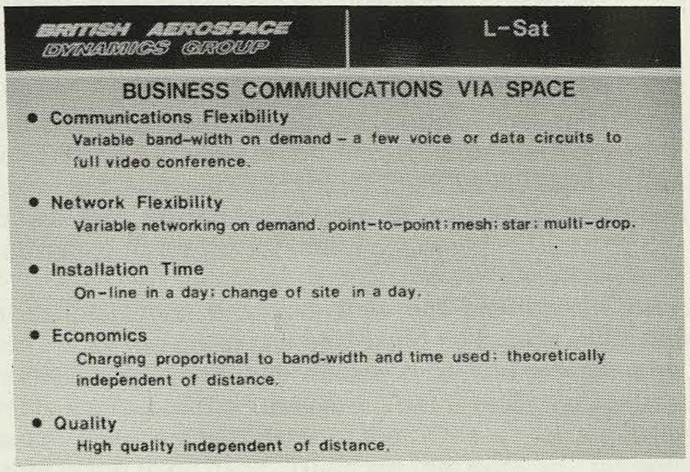

To summarise, it gives you the flexibility; the variable bandwidth on demand. Once you are online you can have a bandwidth of a few kilohertz up to megabits. You can have a flexible network on demand where your stations are. You can go point-to-point; you can build up a mesh network, a star network, or you can broadcast multidrop.

Installation time is literally a day. The Post Office right now are ordering some stations from Ferranti for use with our ECS satellites when they are launched in 1982. They will be mobile, on the back of a truck. Theoretically, potentially, you can order them from the Post Office now for trials with the ECS. They should be on-site and online with the satellite in a day.

The economics are independent of distance. It should save money, but that depends on who owns the system. But we have got the message across to the Post Office that it should be charged economically. The charging is proportional to your usage, not to the distance. The quality via satellite is not dependent on wet, noisy, mechanical switches, or even electronic solid state switches. That is the message of satellite communications as it exists in Europe today.

SESSION D

AN INTRODUCTION TO SPATIAL DATA MANAGEMENT

Professor Nicholas Negroponte,
Massachusetts Institute of Technology

Professor Nicholas Negroponte is Professor in the Department of Architecture at Massachusetts Institute of Technology. He holds both bachelor and master degrees in Architecture from MIT. Since receiving his degrees he has been responsible for undertaking and also supervising a wide range of research activities concerned with the use of technology in graphics of all kinds. Among his major presentations are ones on the future of television, computers in architecture, spatial data management and the human interface. He is also the author of a number of books and papers concerned not only with spatial data management, but also with the use of computer techniques in architecture and urban design.

There is some misrepresentation in the title of my session in that Spatial Data Management is the name of one project that we do in our laboratory and all of our projects are a big, fat excuse to work on the human interface. Spatial Data Management, which we call SDMS, is certainly our most comprehensive and largest system and I will show it to you in gory detail. But before I do that, I should like to show you other things that are going on. I want to give you a slightly broader panorama than just that one application of an area of study that, as recently as three or four years ago, people thought was cissy stuff. They thought it was not the real world of computer science, but was soft and not particularly germane. I am talking about topics such as computer graphics, voice synthesis, all of the stuff at the interface between people and machines, which was tolerated but was not thought to be important. Today, four or five years later, people are realising that in some sense it is the whole ball game and that it is an area that has been ignored for too long.

I have a story which, by analogy, is quite telling. If you do not believe me, you can try this yourself. Go to a department store and watch people select hand calculators. Hand calculators are rather interesting, because for between $6 and $12 you can buy a calculator that adds, subtracts, multiplies, divides, and does the per cent sign work. Some of them have memory, but most people do not even know how to use that, so the price from $6 to $12 is your basic calculator. There are usually up to two dozen of those for sale on a counter.

If you watch people select these, it is an interesting process. They will pick up the hand calculator, hit a few keys, and set it down. They will go to the next hand calculator, hit a few keys, and set it down. Finally, they will make a choice and they will buy one. You wonder why, and there are usually two criteria: the first is the feel of the keyboard. That is really very important. The second has to do with appearance. There are two parts to that. One is the readability and the nature of the display, and the other is the thing itself. There are thin ones and fat ones and so on. Now they even offer fringe benefits: they tell you the time and do other things. But it has nothing to do with the function of the calculator.

What has happened in the calculator market is that the function of calculation has been reduced to such a low cost that what you are selling people is the interface; and when you sell people the interface you are, all of a sudden, selling them something very subjective. One person says, "This keyboard feels great," and another says, "No, this one is better." They are both right.

I think that there is a message here for computer manufacturers, that in maybe five or ten years from now — it does not matter which, the difference is insignificant — that is what will happen with computers. People will be buying only the interface with the computing system, and the number of mips or bytes, or the number of gigabytes of secondary storage will be something between irrelevant and meaningless in the process of selecting a computing system. People will sit down in front of terminals, where there are terminals to be seen, and say, "Yes, that feels good". Managers will buy them because they will say, "Yes, that looks good. I will be very happy to have that on my desk". If you talk to manufacturers and tell them things like that, they think you are absolutely nuts to be worrying about the appearance of a terminal. But it is things like that which are becoming increasingly important.

I shall do four things this afternoon. First, I should like to position a new area that we are trying to start at MIT, which is my own pet project and what I do most of the time now. Secondly, I should like to show you a handful of experiments that have to do with different styles of interaction and modes of presentation; things that I hope you have not seen before. Thirdly, I will show and describe the evolution of the Spatial Data Management System. Finally, the *piece de resistance* of this presentation, I will show you some of our latest work which combines speech recognition and gesture recognition.

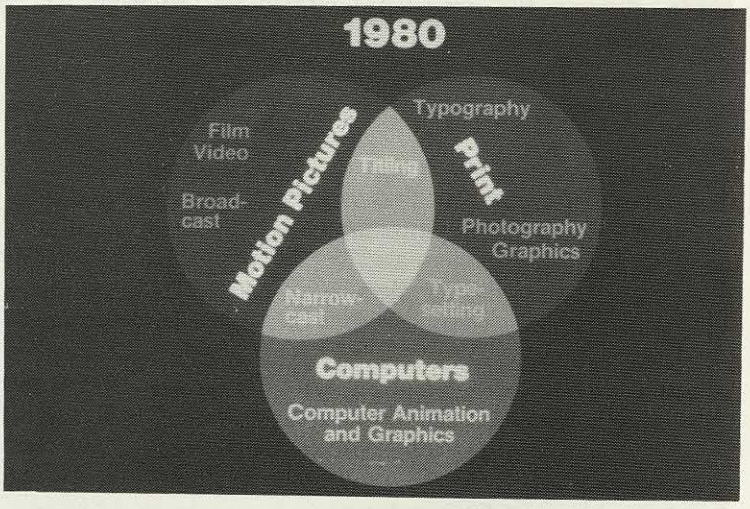

This is probably, intellectually, the most important slide I have with me. What I am trying to do here is to describe three industries: the publishing industry, the broadcast industry, and the computer industry, and show areas of intersection. There are obvious couplings between print and computers that have been going on for years, and other couplings; but there is very little at the centre, at that white triangle. There is very little intellectual activity in that area. Worse, from my point of view, there are very few people who are trained in that area. It is a previously very well-formed compartment, with very well-formed intellectual pursuits. People have their own journals, their own conferences, their own technical communities which really do not overlap very much.

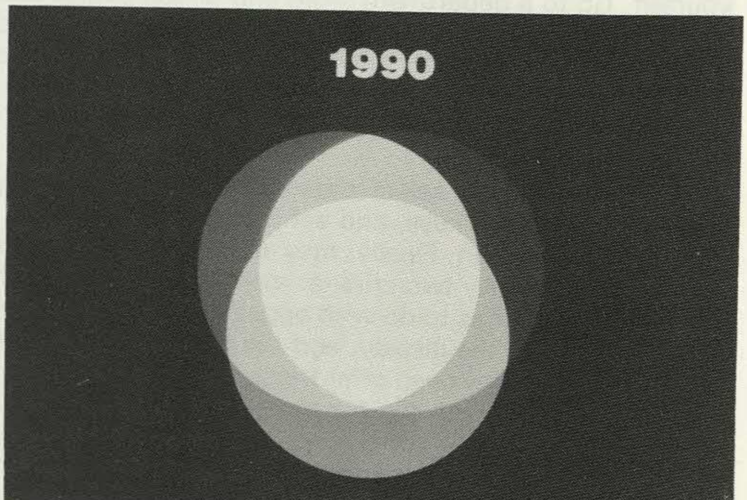

I am very interested in the white triangle because I think that what will happen over the next few years is that it will look something like that. It will get bigger and bigger; and the boundaries that separate those intellectual domains will start to disappear.

There are some good examples of what is happening, say, in the area of Prestel. It is a publishing medium. It uses the video domain, and it is certainly a computer resource. So Prestel is a specific example and lives right in that white triangle. I will argue later that optical videodiscs are another excellent example that are at the intersection of publishing, broadcast and computers. But what really worries me is that the style of thinking of these people is very different and they come to problems like Prestel and optical videodiscs with very different backgrounds and, in some cases, prejudices which are almost dormant and subconscious. It is an area that we would like to fill and we affectionately call it "Media Technology".

Let me show you half a dozen small projects that have been going on over the past three or four years, which have to do with the stuff that separates people from machines. I should like to start with touch sensitive displays, partly because we have been working with them for at least six years now, looking at different ways of doing it and looking at different technologies for achieving touch sensitivity on a display.

When we started doing that we found a real misconception. We found people who had done some work in this area, who assumed that the finger as a graphical stylus of some sort was a very low resolution device — rather like writing a postcard with a cigar. That is not true. It turns out that it is a very high-resolution input device. Admittedly, at a couple of stages you might touch the screen, but if you use your finger very gently you can position a cursor with greater accuracy than any CRT can display it. In fact you have a very high-resolution input medium.

We thought that we would explore that, and people posed questions about why it was so good. Fingers turn out to be very good. They have a couple of properties that are interesting. One is that you do not have to pick them up. You type away, and you point, and you type: it is an input device that is with you all the time. Another thing is that you have 10 of them.

With 10 of them there are certain things that you can do that you cannot do with other graphical input media. The area of multiple fingers is one on which we are still working. There are a number of ways of doing it and it is rather interesting.

At this point I should like to share with you as much a style of thinking, or at least what drives our laboratory, as showing you a project. When we were working on touch sensitive displays we found that the coefficient of friction between your finger and glass, if glass was the substrate that you were touching, was such that there was just a little too much friction for you to draw with your finger. You get an effect that precludes freely using your finger to draw on a CRT.

We turned that "bug" into a feature. We made a pressure sensitive display. We thought that it would be interesting if you could not only touch the display but introduce forces on the face panel, and interact with a model; and when you touched the screen you could pick up forces in the plane of the screen and in the Z axis, basically getting five degrees of freedom.

The arrow on the slide represents the direction in which he is pushing; its length is how hard he is pushing; and the sides of the square represent how hard he is pushing in on the screen; the XY location is at the tail of that arrow. I offer this as an example of the kinds of projects that we do.

It turned out that this was useful for virtual face panels. There are a lot of industries which have a need for virtual face panels of one sort or another, because they have run out of real estate. There is no more room for the dials and knobs. With your average oscilloscope you either have to poke at the knobs with tweezers or have very tiny fingers these days, because there are too many knobs and buttons. So people are interested in "soft" face panels, ones where the right knobs and dials come up for the occasion. But they tend to be rather unpleasant to use; the feel of those things is just not there. People like grabbing that knob, turning it and feeling it go "click, click", and watching the numbers. Sight, sound and touch are all wrapped into how the person interacts with that control panel in that example of an oscilloscope.

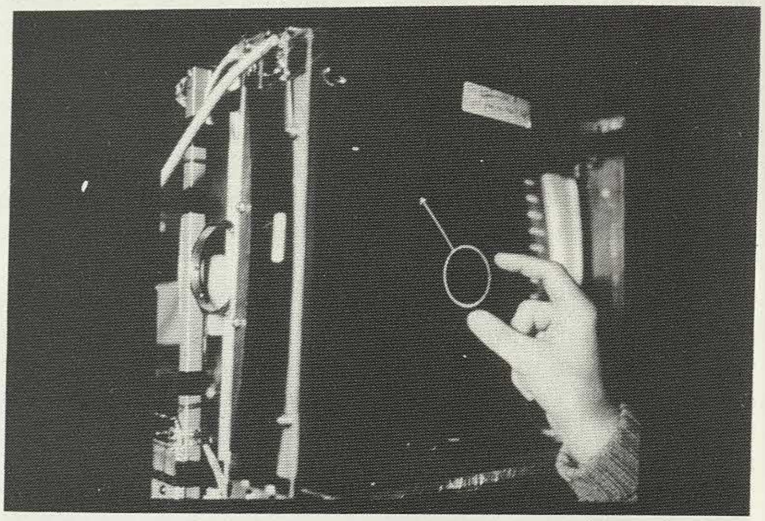

Let me show you a movie. All this movie shows is what I just showed you in slides, in operation. This is to show you X. X is the horizontal, Y is the vertical, and Z is the force into the screen. Z cannot go negative. We have figured out how since, but at the time Z could not go negative and that is what is being shown here. The nature of our laboratory is such that we build things and all they have to do is to work once. This was certainly one of those cases. In fact this particular machine was stepped on in a move recently, and no longer exists.

This is a torque around a Z-axis which you can see down below. What happens is that you can feed back to the user physical properties that would otherwise be unavailable to both you and him — you as a programmer or designer of a system. The arrow is the direction in which he is pushing and with the size of the square here, you can track it over time; he pushes a point and the area he pushes just follows. This is great for the games market!

This next one was the most important and goes by very fast. When he touches an object he pushes it, but the physical properties of that object now can be fed back to him. If it is a heavy object moving on a surface with a high coefficient of friction, then it would be hard for him to push. If it was a ping pong ball on a sheet of ice, he would just need to touch it. You can feed back physical properties to the user because it is harder for him to push. In this case our sponsor was military, and this is our token militarism. The number of bullets is how hard you push and so on. I was very disappointed in the programmer who did this because the ducks did not even fall over.

The next one is equally important. It is the feel of the knob. There was a sound track to this and the machine literally went, "click, click, click" as you turned it. You start to get a feel for that knob or that square moving across, in a very literal sense. (End of movie.)

It is amusing and it is nice to talk about it, but there is a real message there that has to do with actually giving people a sensory involvement with the display that is significantly richer than what we are accustomed to. As problems that we tackle with computer systems become more and more sophisticated, we will need interfaces that are equally more sophisticated in the sensory domain. I am not trying to argue that we have been incompetent in our modelling and our simulation, but I think that the incompetence and tragedy has been that we have dealt with computer systems that are parsimonious at best in terms of how they interact with people.

I ought to digress for a moment to give you a brief station identification. We were an old computer aided design group in the School of Architecture and Planning. We started in 1966. Around 1970, computer aided design was a bad joke. It was not computer aided design at all.

Designers would not touch the stuff. We felt that the problem was that there was an intrinsic incompatibility with the early design stages and what people had as computer systems available to them. We were much more interested in the human interface and trying to make that a subtle, smooth and creative place to be. In 1972, we made a decision which everybody thought was foolish: we moved all of our vector graphics into what is now commonly called "frame buffers" and we raster scanned displays. One of the reasons that we did that was to stay in the television domain. There are a lot of reasons to stay in the television domain which I will enumerate later, but all of our colleagues — some of them still today — said, "That's silly, because television doesn't have enough resolution. You can't use television properly because of its lack of resolution".

That is a lot of nonsense, and it comes from the fact that people have not looked at it in the right way. I think that everybody will agree that a common denominator of Prestel, Telidon, Ceefax and all the others is absolutely ghastly text. It is so ghastly that it is absolutely inhuman to sit down and read a Prestel display. In fact in the United States it should be an OSHA (Occupational Safety and Hazard Administration) violation. It is an offence to ask people to sit down and look at those fonts.

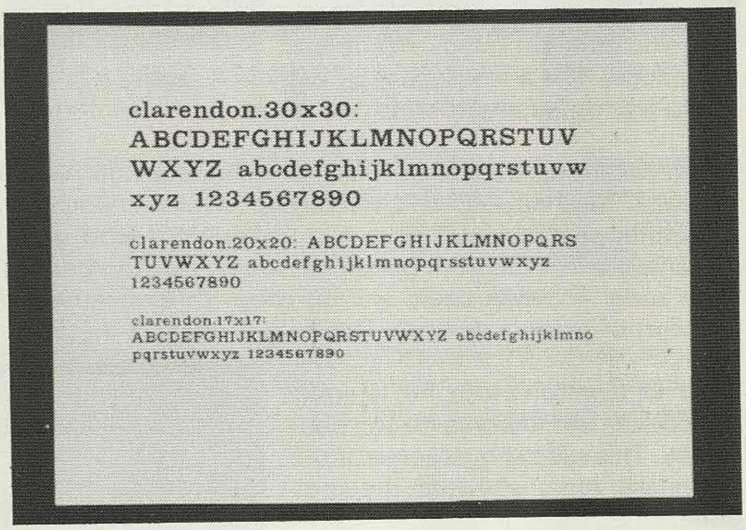

So we decided to look at fonts differently. This slide shows a standard 525 line television display with 3 size fonts displayed in front of you. These fonts are not the greatest but they are pretty good. They are better than you have ever seen before, and they are worse on this slide than they are in reality. One of the things that we cannot do is the gamma curve and our slide reproduction equipment does not quite reflect what is on the monitor. This is a colour monitor displaying black and white text. We can get up to 125 characters per line and it is still perfectly readable. There are some tricks and some caveats. I do not want to go into all that, but let me show you how it is done. What I should like to impress upon you is that no geniuses have been working on this. Three people worked on this project. One of them was a professional graphic designer, in fact a typographer; one was a broadcast engineer; and the other was a computer scientist. Jointly, those three people came up with a simple and stunning idea. That was that you can do something on television that you cannot do on paper: you can trade off acuity for resolution. You can introduce purposeful blurriness in the definition of a character and the net result would be a higher perceived resolution.

That is parodoxical, but it works. I will prove it by example. If you do not believe that, you will have to come to MIT. To me, the interesting part of that is that the way of thinking about the problem was very different. Let me talk about the computer graphics community. I can be rude about that because I am part of it. People came to the problem thinking squarewaves. They are squarewave thinkers. They like crisp edges. They like ones and zeros to have sharp rises and falls. Television hates that. It is exactly what television does not want.

If you introduce in the definition of a character purposeful grey tones — take this "Q" shown on the next slide — this is what our "Q" looks like on the computer. If I had made a slide and shown you that blown up full size, that funny little array of greys, you probably would not even have recognised it as a "Q". But in its actual size down here it is a beautiful "Q". It is gorgeous. It does not scintillate; it gets rid of all the artefacts of television. The way it was done was to look at the "Q" in terms of a sampling theory instead of character matrix and introduce grey tones, not one-to-one, but let us assume for the sake of brevity that the grey tone is reflected by the percentage of coverage of that area.

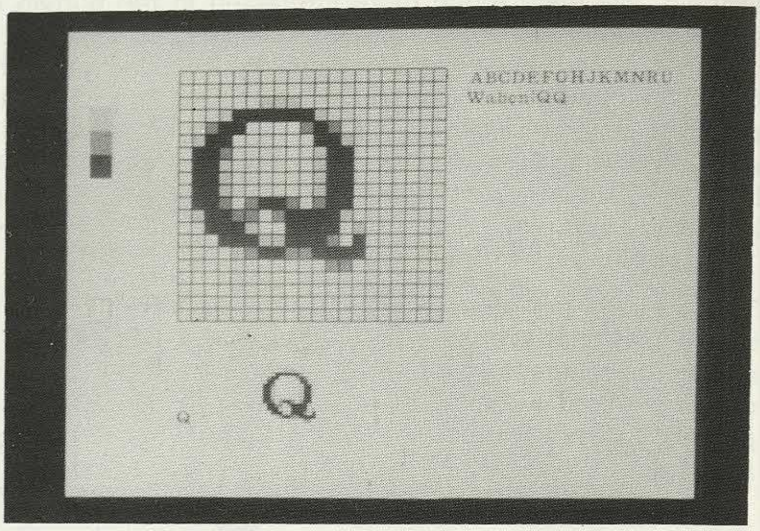

This is not absolutely new. People have been trying to do this for vectors, and there is a long history of it. It turns out to be quite fantastic for characters. In the horizontal direction it lowers the bandwidth of the signal because you do not go from black to white; in the vertical direction it gets rid of scintillation because every part of each character is in both interlaces. So you get a pretty nice character. These fuzzy greys do different things in different directions. It is a small piece of work but, again, it is an example of the kinds of things that we do to try to pay attention to the quality of that interface.

I believe that, as people start using management information systems, you will find a new class of users. This will also be true in homes, with children, who will not tolerate the junk that our computer programmers have been tolerating, who are much more amenable as they see the advantages in a different perspective. We are now getting to a point — maybe not this year but very soon — when users will be distinctly concerned about the quality of displays and the comfort of using a system. This sort of thing will surface as a very serious problem.

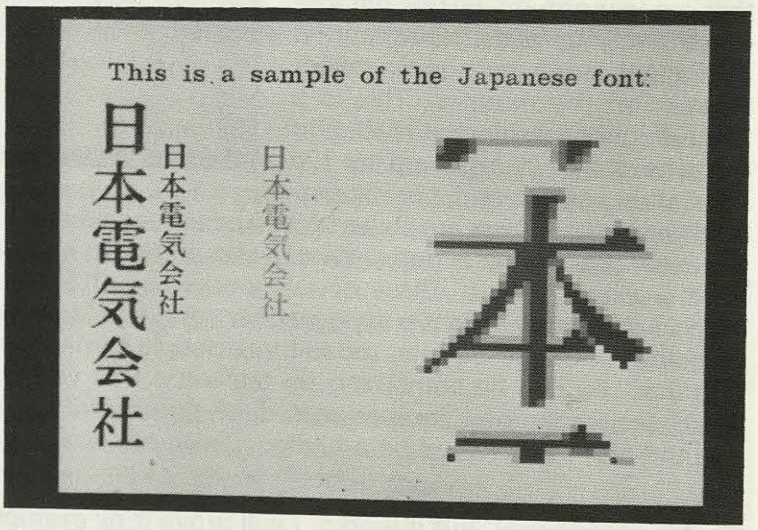

I will not dwell on these, but it works for colour, and it works for the Japanese. This slide shows some of the Japanese fonts showing the grey tone.

What is shown on the first slide on the next page is much older. There was a memo recently that we distributed amongst our various friends in the ARPA community called "Don't throw away the message with the media". It was a plea for paper and to go back to look carefully at some of the qualities that paper had before we discarded it completely. Not that we should go back to doing everything on paper, but let us look at some of the messages and signals that are on paper. When you get a letter and there are two creases in it, you know that it has probably come in an envelope, through the mail. So you have a couple of bits of information there, namely it came through the mail.

You can look at other things. Annotations on paper are a wonderful example. When you make annotations and you file a paper document and retrieve it, it turns out that what you really look at more frequently is the annotations than the document itself. You can tell something about the reading habits or the occasion, from a whisky stain versus a coffee stain on the piece of paper. If you have ever tried to uncrumple a letter to put it in a file, it is pretty hard to do; and that tells you something again.

We looked at this problem in relation to terminals which scroll text. In reading text that scrolls, whether it scrolls and then you read it, or whether you try to read it as it is scrolling, you really tend to have no sense of place, and you do not know where you are vis-a-vis the whole. Yet when you buy a cheap paperback book and you read it, you have a random access medium which has built into it a wonderful sense of place. When you are reading the book you can tell whether you are at the beginning or the end because the pages that you have read are in your left hand and the ones to go are in your right hand, so you know where you are. A page is in fact a syntactical chunk that should not be discarded as carelessly as we seem to have done it. Frequently in language you will say, "I received a three-page memorandum from so-and-so", and that means something to the person you are talking to.

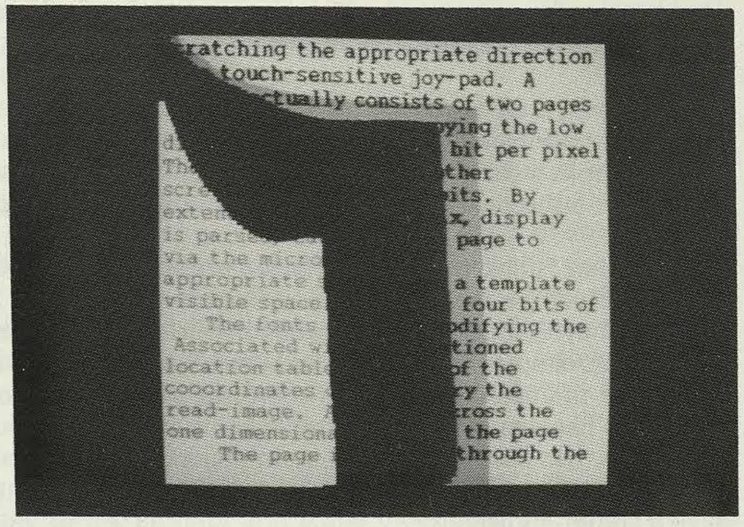

So we decided that maybe what we could do in text and character displays is to reintroduce the page as a primary chunk, or at least as a modular page. We introduced a little bit of animation, which you will see in the next film, where the next page actually follows it. It flips by, it does not scroll by, but you get a page 1 that is revealed, and then you get your second page. This happens smoothly and floats by, and you see page after page.

We are not very good at doing any psychological testing that would validate this as the right way to display text, but again it is an attitude to the problem, at least to try to be concerned about things like this. Whether or not you agree with the solution is in some sense less important than agreeing that the area needs some work.

Another project we have is to look at television as a display medium for books without pages. As you start getting a lot of storage capacity in your home — I am thinking of videodiscs — and as you start getting computing capacity, how can you best use a standard 525-line television economically and still come up with reasonable images?

The map shown on this slide is a standard television image and the memory bit-map is being used to support this. We looked at this picture in a very different way from the way that anybody had ever looked at image compression before, by basically doing the opposite of television. For those of you familiar with television technology, you will know that most of the bandwidth is an illuminant signal and then a small amount of bandwidth is overlaid and applied to the two colour components.

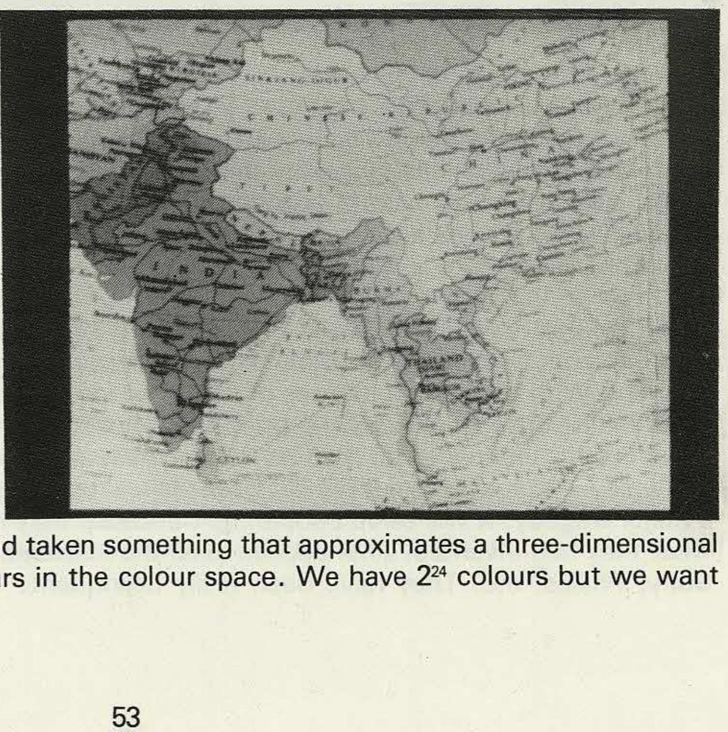

Here we have sampled this picture and taken something that approximates a three-dimensional histogram of the population of colours in the colour space. We have 2^{24} colours but we want to pick 2^7 or 2^8 that are the most indigenous to this picture, and use those, hence reducing our memory from 2^{24} to 2^7 or 2^8. This works quite effectively. In fact we have not encountered a picture that needs more than 2^8 colours, and it has worked quite successfully. It works for maps and a lot of other things.

Another area that we work on is printing colour hard copy from computer systems. This is motivated from a number of different directions. First, we would like to work more in colour soft copy. There is a story that I used to tell a lot. It is about the Xerox Corporation, and then I heard them tell it publicly and I thought maybe I had better stop telling it.

We got what I think was the first or the second Xerox 6500 colour copier. At the time that Xerox put this colour copier on the market they could not lease them, and they could not sell them of course. They went back and asked themselves why they could not sell a colour copier. The answer was simple: in the world of business there is no colour copy to copy. Then they puzzled and asked themselves, "Why is there no colour copy to copy?" Obviously one of the major reasons is because of Xerox. That is being hoist with your own petard, in the sense that their technology has contributed greatly to making the world of business practice very much a black and white world. There is indeed very little colour copy to copy and very little justification to have a colour copier.

Then they invented a slide adaptor because there were a lot of colour slides to print, and so on. Having said that, it is clear that people are working more and more online. I agree with the comment that people are getting much more comfortable at typing and working in front of displays. What is also clear is that these displays in a few years will all be colour; it will not be in anybody's interest to install a black and white display. A few years beyond that it will probably be more expensive to put in a black and white display, because you will have to buy it from an antique dealer. Basically people will have high-resolution shadow-mask displays in all of their terminals, and whether it is five years or ten years from now does not matter.

But when you work in colour soft copy, if you want to go out on a futuristic limb, you can argue that you will not need any hard copy whatsoever; and that we will come up with flat panel, thin, flexible, waterproof displays that can be used in the shower. Maybe that is the case, but in-between, and for various reasons, to take it seriously and really work in colour you need colour hard copy. Our project started way back when we got the copier. We can now produce direct Xerox originals that come directly out of the system, not with a laser but with a CRT. This was work done under contract to Xerox. They promptly went off and did exactly the opposite, and came out with the dreadful laser option which produces, in my opinion, very poor images. Again, this project shows you the kind of experiment that we call Media Technology.

Let me go on to a favourite example that has to do with the use of optical videodiscs. They are very much in that triangle I showed at the beginning, unbeknown to some of the manufacturers, including Philips. People like Philips turn out to be their own worst enemies, because there has been a serious failing in looking at this as a new information medium that can store digital data as well as picture in the same domain. In the long run it will all be digital, but in between why not just hide some digital data in the analogue signal and put this out as a new publishing medium? But that is not being done.

If you look at optical videodiscs — and this will be particularly true in future office and home environments — disc playing time is not necessarily a good measure, but let me use it for a moment. With disc playing time as one axis and viewing time as another, a 45° line is a movie as we know it. Because of focusing on that 45° line, there will be people like RCA who will soon introduce the world's worst product. When that particular videodisc hits the market — if it does sell which it should not — it will set back the real opportunities by about five years. That will be a major tragedy. It will serve people right because they have been very poor at understanding what they had invented. That applies to all the manufacturers, including IBM, MCA, Pioneer, Sony, Philips, etc. They have been really delinquent. There are new movie types and new book types that have not been explored. One of the reasons, with all due respect to these giant corporations, is that there is nobody who knows how to think about these problems. There is no culture of people who are at once in the business of authorship, but on the other hand have made films, and also know something about computer programming. We are missing that kind of person, so this is quite explainable.

SDMS is deeply dependent on videodiscs. Let me refer to them as a new book type. Before doing that I should like to show you some new movie types. I could not bring them because they are all on video tape. One of the international problems right now is that our video standards are not compatible so that I could not bring you current video tapes of some work in that area. But I can show you statically and you can try to image them. They all have to do with videodiscs connected to computers. You have an absolutely incredible storage medium. If you hide digital data on a videodisc, crudely you get 1 billion characters per side. That is 10^{10} bit and you get an error rate of 10^{-11}. If you use it as a video storage medium it is 54,000 frames, which is a lot of data. Ironically, you have this beautiful storage medium that is controlled by a handheld panel, which is a disgrace.

The first thing we do is to connect it to a computer because the computer in it is limp, to say the least. We add a touch-sensitive display to the monitor and some other controls. Then we find that, lo and behold, you have to mix computer signals, so we add a second player. Then you have a sound system. This is basically the configuration of what we are working with. Let me now quickly illustrate two projects.

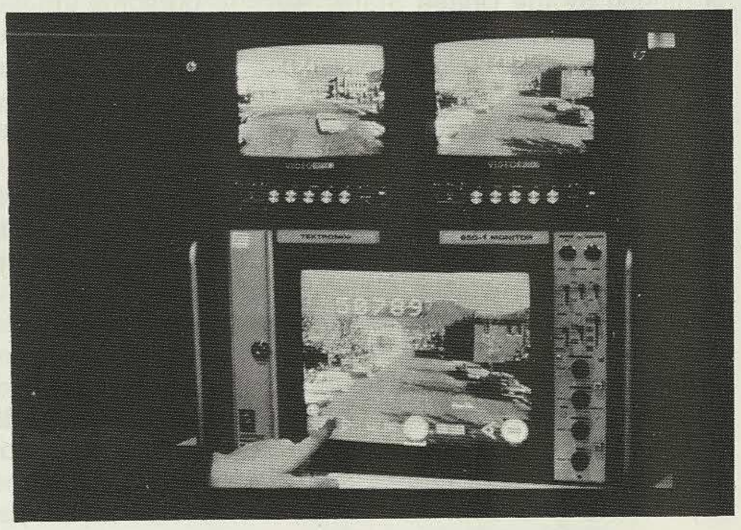

One is a movie map. It has received a lot of publicity so you may have seen it reported or discussed. It was a very simple idea. The task that we were funded to do was to get involved with the experiential aspect of mapping. One of the things about maps and cartography in general is that it is focused on structural relationships. You look at a map and you know what is to the north or south, but you have no sense of what it would be like to be there. Could we make a system that approximated to being there?

We took a small town and we photographed every 10 feet, travelling down every street, in every direction, taking every turn in every direction, during every season, night and day. The images reside on an optical videodisc. You have two players. Player Number 1 is playing for you as if you are travelling down a street, and the numbers are purposely superimposed. The user of the system has indicated that he wants to go left, so he has touched the screen — the touch sensitive monitor in this case — to go left. Player Number 2 which is not being used right now has gone ahead and is waiting for him at that intersection, so when he gets to the intersection it just cuts and smoothly shows the turn going through to the left. While Player Number 2 is doing that, Player Number 1 is free. It goes and gets the street, and waits for him; and when he comes out of the turn, cuts to that, and the user freely drives around the town.

In fact you freely drive around the town and the two players take over. There is a cute little detail, which is that when you get to an intersection you have three options: straight, right, and left. But they are just spliced head to tail, so if the person decides to go right you play Player Number 2 forward, and if he decides to go left you just play it backwards. You can play all these things backwards as well as forwards. You have all the options sitting there simultaneously, and the person travels around.

There are also side views. You can look out of the side window. On our TV set, instead of having a channel select knob we have a season knob. You can dial-in in winter, summer, fall or spring. We have a special season called animation, which is typical computer animation. I wish that I could show you a video tape of this, but the season knob is quite spectacular because one of the things about any kind of experiential mapping system is that seasons make a dramatic difference, in some places.

I remember as a child going to a school that had a campus in the winter that was up in a ski resort. I went back there, a few years later during the summer, and I could not recognise the place. The skating rink was a tennis court; trees were there that I never knew about; and it was a whole different place. I have two slides that show the same house in winter and summer. When I show these slides, there are people who swear that it is not the same house. It is a whole different place in winter. We even have it now so that you can have 50% winter and 50% summer, which is an amusing trip.

In some places, where we could, we put in a time knob where you could wind back, and see what the scene looked like 70 years ago. The computer database behind that is enormous. It constitutes a new kind of movie in the sense that there is more footage out there than you would ever look at. There is no way that a single user would explore that whole space.

The next project I want to mention is the use of the medium to make personalised movies. This is funded by the United States Navy to do some innovation in maintenance and repair. The United States Navy has a problem in that their equipment is getting more and more sophisticated and their personnel are getting dumber. They are very frank about it. They have stopped preventative maintenance in many situations because people did more damage than good. They are really concerned about it, witness our helicopters most recently.

What we tried to do was to make a personalised movie with two axes. One has to do with the user's previous experience, and the other has to do with the user's cognitive style. Previous experience can probably be illustrated best with an example from cooking. I do not know how many of you have seen the *Larousse Gastronomique* which is printed in English. It is an incredible reference cook book. It has the most complex recipes in it, and then the last line is "Cook until done". That is what we would call an unexpandable macro, in the sense of what do you do with it. To an expert cook that is perfectly complete information. To me, and I suspect most people here, you would say, "Cook for an hour at 350°". For a beginner, you would say, "Preheat the oven, put the shelf at this location, wait for the light to go off, then put it in". What you see there is a degree of elaboration that is a function of the person's previous experience, and hence it has to do with the model of the user.

The other axis is one of presentational means. If you show some people a plan or a section, they are just hopeless; they cannot read them and they do not understand them. Yet if you showed them a perspective drawing with shaded surfaces, they would find that much more understandable. Other people are the reverse, and so on. So there is an axis of presentational means.

What we try to do is to make very interactive movies; and make the film media, which is historically, and almost by definition in people's minds, a sequential one with a beginning and an end, into a very random access, interactive medium where you expand macros. You can have alternate sequences, which is changes of mode. You can combine those. We did it for the maintenance and repair of a bicycle.

This is the beginning of the advertised topic: Spatial Data Management. It is just one of our programmes, and was a nice excuse to work on the human interface. Back in 1975, we were asked to build a management information system that senior management could use, and would want to use; but most importantly, that the senior manager could learn how to use in less than 10 seconds.

We have a little jingle in the laboratory that says that we build computers for generals, presidents, companies, and six-year-old children. One common denominator between a six-year-old child and a senior executive is that they want things immediately. If there is something on the screen, you touch it and out it comes. There is a directness and simplicity. If it is too complicated they will get frustrated; if it is too simple they will get bored. There is a very delicate boundary, and to capture the person's interest and enthusiasm is, to say the least, challenging.

Somewhere along the line, I am not sure how, we discovered that there was something about spatiality that had not been explored as an organising element. There are two incidents that illustrate this. I get to the airport and I have forgotten a telephone number. So I telephone back to my secretary and say, "Beth, I've forgotten a phone number, but if you go to my desk, just to the left of the telephone, underneath the calendar, there is a pile of papers about three inches deep, and two-thirds of the way down there is an orange sheet of paper, and on the back of it, in the lower lefthand corner, in blue Pentel pen is written a phone number. Can you read it to me?"

That may be a bit extreme, but even people who are not what you would call visual thinkers or people who are very attuned to spatial memory, still are very good. You see examples of this on bookshelves, how people tend to organise their books. If you put the book there, there is even more reinforcement and you remember where it is — it is to the top of such and such, or it is to the right of such and such. Most of us remember the colours on the spines of books or how thick it was. There are lots of cues that we will use to retrieve it.

A real incident that we uncovered was an admiral who was selected for a test site of a very complex command and control system, and he refused to use it. Instead, he had a bulletin board with his ships on little thumb tacks; and the enemy's ships were also on little thumb tacks. You can imagine that theirs were red and his were blue. He sat there, putting the ships on the thumb tacks on the bulletin board. He kept moving them around, and then he would stand back and then move others over. People tried to discuss with him why he was not using his very advanced, full colour graphics display command and control system. Everybody came up with different reasons that had to do with the standard reasons such as he was intimidated, or it did not work well enough. But what nobody really considered was that maybe there was a modal memory reinforcement. Maybe his own body movements were important to him, and when he picked up that ship and put it down there, he was interacting with the display in a way that should be taken seriously. That is what brought us into the idea of making, as part and parcel of our work in the Spatial Data Management System, a media room. In other words, a room that was a terminal. So it is not really a terminal in front of which you sit, but a terminal into which you go. Call it an office of the future if you want, but when you are in it, this terminal has the authority to engage every sense, and you have every opportunity to interact with it.

The first media room no longer exists, but the newer one is a derivative of it, quite directly. The floor-to-ceiling, wall-to-wall display in the current one is in some measure a simulation of what we are expecting to happen in the near future. There is a conspicuous absence of a keyboard in this system, probably as much for metaphorical reasons as anything else. I am certainly the last person to hate keyboards. I cannot write any more, I even typewrite postcards. The absence of a keyboard and any set theoretic type of retrieval is more metaphorical than any position we are taking on the topic.

The new media room shown on the slide on the next page has three monitors. One allows the user to keep track of his information space which we call Dataland. The other, on the right, is a key map which allows a person to explore particular data types, and then on the large screen is what the person is attending to.

I will now show you a movie of this in operation. It will be difficult because we will be looking at just one screen and moving between these. So let me talk you through it.

This is Dataland, which I shall describe in more detail in a moment. These are key maps which come up. In this case it is a telephone that happens to be routed in Dataland, because he needed a telephone at that moment.

Here is a little window. Dataland is the fabrication of the user. In very much the same way that you put papers on your desk or put books on a shelf, you put data in Dataland. You can create neighbourhoods. This neighbourhood is non-technical reading; there are four books stored here. There are five books stored up there, which are technical documents. There is some electronic mail in that neighbourhood. There is a personnel file down here; there is a calculator and a calendar. There is a telephone missing from this picture. There are some TV sets which are live broadcast or recordings. There is a zoo, because the guy whose Dataland it is has a child, so there is a very large zoo stored in this section. Then there are some maps; some satellite imagery; some maps of Suez and the Far East.

It looks sparse, but it is not. What is happening is that there is a Z-axis to this world. As you helicopter over it, you can go down and gain resolution. SDMS is very much like an Advent calendar. Once you get to a book or a document, you can open it up and look in. When the monitor on the right gets engaged, you are into that key map.

Dataland is created by the user so he remembers where things are, because he put them there. Everything is up front, nothing is hidden around the edges of this bezel, that is the full extent of Dataland. However, there is a Z-axis. Think of the Z-axis in the following example. You are the personnel officer in a company. You are going to build up Dataland; you are going to have a neighbourhood, which is all the companies that you buy and sell electronic components from. You will have another neighbourhood which is the furniture and office supply people. You will spread out on your surface, as data, the various companies that you deal with as a purchasing agent, and your Dataland is scattered that way. Then for any single company, maybe the recent three years of purchase orders with them is the Z-axis of that data.

Somebody else may be the president of this corporation, and that person, as president, will probably have in his Dataland just one little box called "purchasing office". All of your data as purchasing agent which was spread out over your X-axis and Y-axis, in his world is in the Z-axis. So Dataland is not only enormous in extent, but things go into the Z-axis as a function of your immediate needs and interests.

If this picture represents 3,000 miles of physical extent the current system can get down to a picture element of 10 feet. So you have an enormous extent. You can keep zooming in to this thing. If that were the United States you can get down to images that are smaller than streets and not exhaust more than one optical videodisc. You have built into this an assumption that facsimile-style data types will become increasingly easy to deal with, so not everything is stored as structured data.

There is a sub-theme, which you will see in the film, which has to do with the richness of data types in general. You will get movies stored in there, operable machines, dynamic processes of one sort or another, colour photographs, text and so on. So let us look at the film.

Our user will travel over Dataland in three different ways. At first, we will see him move the joysticks which are pressure sensitive and allow you to fly over the surface, and also to go down in the Z-axis. As he is moving, there is a window moving over Dataland. As he goes over an area, it appears on the large screen at increased size and resolution. First, he will go down to this personnel file and what we will see in a moment is what he sees on the large screen, which is this array of people drifting by underneath his window. As he zooms in, at a certain point, the chunkiness is replaced by a higher resolution image. So now he is even closer down into the personnel file. If he now wanted to look up his records, he would address the righthand monitor to which I referred. It now becomes the Advent calendar, and you go through the face and find out what you want on the other side.

Here, with a double exposure, it may be a little clearer. He is driving due north in Dataland and going to look at some electronic mail, which will appear here. This is what he sees on the large screen, totally unreadable. But the logo tells me that I had better read it, as it is my sponsor; and I read my sponsor's correspondence. It is a bad example because they really do fit on the screen.

All you have done, to this point, is drift over the surface and go up and down. We have flown over this information space, we have flown north/south primarily in Dataland, but now we are going to poke through. We are actually going to go in and look at a book, in this case. As soon as we do this, the righthand monitor will appear with a table of contents. This looks like a book. It has chapter headings. The whole thing is touch sensitive. You touch a chapter heading and it tells you the sub-heading. You touch the sub-headings and, if there are any, it will show you the sub-sub-headings. This red line is a marker of where you are in the book. When that display comes up, the thickness of the book reflects the number of pages stored in the database. So we have gone to every extent possible to build in cues to help you randomly access, in this case, a book of 150 pages.

The artefact of page flipping may be a bit extreme, but it certainly works. Here he goes back to the title page. He flips through pages by stroking the side of the chair. Again, this is sort of on the lunatic fringe, but it reflects interest. If nothing else, appreciate it. Here he is just reading pages. One student wrote a program so that as a page is flipped the room actually made a fluffing sound, from upper right to lower left. You heard the page float through the air, because there is an octophonic sound system in the same room.

> "SHOW ME THE FOXES" (from the film soundtrack).

He is talking his way around Dataland.

> "TAKE ME TO THE MUSEUM" (from the film soundtrack).

There is a whole museum that is stored there.

> "GO TO THE MAP ABOVE AND TO THE LEFT OF THE CALCULATOR" (from the film soundtrack).

This is a calculator which we will see later, and that is the map above it.

> "I'D LIKE TO MAKE A PHONE CALL, PLEASE" (from the film soundtrack).

The "please" is clearly unnecessary, but up will come this machine which has a key map that looks like a telephone. He is going to dial the phone and place the phone call.

(At this point in the film, the user makes a telephone call by "touching" the push buttons of the telephone on a touch-sensitive screen. This is followed by a telephone ringing tone.)

This telephone has a relatively intelligent roller deck, one of those little rotating address books. You can find names, pull up these little cards, automatically touch telephone numbers, and it places the call for you. Again, this is not particularly interesting. What is interesting is that it is buried in the same medium.

(There then follows a dialogue on the film where the user is using the SDMS as an intermediary to make a telephone call.)

What he has been doing up till now is talking through the system, but now he is going to talk to the system.

(There then follows a dialogue on the film where the user asks the SDMS several questions, and as a result of the spoken replies, the user asks the system to place a call to someone's office. In this case, the call is not dialled — the system already knows the number to call. The scene concludes with a telephone "busy" signal being heard.)

We did not know how to get out of that scene, so we had it ring "busy".

The third way of going around Dataland is just touching it, which is a little boring; you just touch the surface and go round. So I want to take the opportunity of the slow part of the film to describe this. What started to happen there is absolutely critical, because a new notion emerged. We are very accustomed to when this window in graphics moves around and we think of a graphical window. What we developed here was, to my mind, very significant: we developed the idea of a vocabulary window. As the person drives over Dataland and gets near the telephone, we load words that have to do with dialling telephones into our speech recognition system. As he goes near the calendar, we load words that have to do with January, February, Monday, Tuesday and so on. When he goes to the calculator, we load words that have to do with plus, minus, divide and so on.

So on the one hand you can say that here is the world's most expensive hand calculator, but on the other hand it is one that can be driven by both voice and touch. I agree with the question to Frederick Jelinek this morning, namely that mapping spoken words into typed correspondence is maybe not the correct motivation; but to use that channel of communication when you need to use it might be a more noble beginning because, if nothing else, it is easier. What you are seeing here works. The reason that it works is because there is so much redundancy at the interface that you can disambiguate. I have another movie which will show you that in some detail.

Let me hold the rest of that paragraph in abeyance and go back to SDMS where all of a sudden the data itself is spatial. The only reason to show this, except for the fact that that is MIT and that is Logan Airport and one gets homesick, is that these picture elements are 50 metres on a side. That is about 400 miles and represents a tiny bit of Dataland.

This is the last thing on the film where the data type itself is a movie. If you are at all into videodiscs you will have seen this movie many times and are probably nauseated by it.

(At this point of the film, a sequence demonstrated how the user can react with the system to show the same visual sequence twice, using a different soundtrack on each occasion. In other words, multi-lingual films.)

This key map is important. One revolution of that clock is the duration of that film. What happened there with that little wedge is that at the instant he touched it he put a marker into the database, so that by touching that marker he can get back to that exact frame. We refer to that frequently as a film "dog-ear" in the sense that you can fold these pages and get back to them accurately. That is the end of the film.

There were a lot of things that happened in that film, even though the database illustrated is sparse. We certainly did not look at what could be in all of the books, all the letters and all the correspondence. But probably the single most important thing is the cohabitation of these different worlds. So often people say, "Oh, that's television technology" or, "That's a movie" or, "That's text". Here they are all in one place and all in one medium. When the user in the film switched from looking at the satellite photographs to data or to a film or to a letter (whether as a facsimile image or as ASCII data), it was all transparent to the user. You get these different data types that are mixed, very conveniently, in the video domain. I think that is a very important part of the system. It has nothing to do with spatiality *per se* but has a lot to do with the future of using these kinds of computer systems.

When we finished SDMS, which was about 18 months ago, there were many frustrations. One of them was how could we interact with the large format display. How could we engage the user more with that display? You had the two on the side that you could touch, but what could you do with the other? That was where we started working with speech and gesture. We wanted to point at it and talk to it.

Our experience in speech recognition is only about two years old. We used the Nippon Electric connected-speech recognition system. We certainly have made no contribution whatsoever to the state of the art of recognising words in the way you heard this morning. What we have made a contribution to is a very simple addition, which is not to try to work with just speech and process it, or just graphics and process it; but to try to work with them together so that one mode of communication disambiguates the other mode of communication. It is so obvious. We do it all the time as human beings when we talk to each other.

We say things, and sometimes you do not even hear what the other person has said, but you know what they mean not just because you know them, but because you know in which direction they are looking. Without even going into the world of artificial intelligence, which unjustifiably seems to scare people, but just staying with the surface problems, the interface, you can disambiguate a great deal.

The Nippon Electric speech recognition system is the top of the line, and you pay an arm and a leg for it; but one of its strong points is that when it fails, it fails gracefully. It does not crash. It will say, "Can't quite understand that," and it returns what it cannot understand. What I want to show you in this next film is a program that has been completed and now works quite successfully, which combines a magnetic technology which we implemented and advanced from where we found it, which allows you to track six degrees of freedom, X, Y, Z, azimuth, elevation and roll, with magnetics instead of optics or acoustics. It is nice because it goes through people; and unless you march into the room with an anvil you will not disturb the system. It is extremely accurate. We can track the arm movement which, since it is six degrees of input, is literally a vector; so we extend a vector from the person's arm to where he is pointing to on the screen and the screen knows where he is pointing.

The speech recognition system can deal with up to five utterances concatenated. An utterance is defined as anything less than two seconds of speech. When you concatenate these it disambiguates them in a way that is rather interesting. On the film you will see a very limited interaction that has to do with moving objects around. But when you look at the film, notice the free use of words like "this" and "that". If I say, "Please pass me that," and you did not hear the word "that" but you saw me pointing to the water jug, you might pass me the water jug. So my pointing is literally quite redundant with the word "that". If I say, "Please pass me that," and I am pointing to the table, you might be a little taken aback by the fact that it is a rather large thing to be passing around. At that point you might ask why. There is nothing wrong with asking. It is a very interactive system. It does not go off and process and come back a few minutes later. If you cannot disambiguate, you want to ask the right and timely question. If you ask the right and timely question, the whole scenario is extremely comfortable. We have a system where these people come and sit down, move ships around and so on. The speech recogniser in connected speech mode does not work more than 85% of the time. That is a good performance. But with the gestures combined you get a perceived response of 95% to 96%, because it does the right thing; not because it has understood what you said, but it has understood enough of what you have said combined with what you are pointing at and a little bit of knowledge about ships, land and so on, to enable it to respond, and its performance looks pretty good.

Then when it fails, it fails gracefully. It says, "Where?" or "Which object?" It will ask you questions, which is much less offensive than its just going off and doing the wrong thing; that is what really gets offensive.

The user wears a little magnetic sensor. The transmitter is on the arm of the chair. It came out of helicopter head-mounted siting systems where it was mounted on the top of the cockpit and on the helmet of the pilot. Obviously, azimuth, elevation and roll are very important when you are siting; X and Y are almost meaningless; and Z is meaningless. So that is what we are using.

Let us look just at the beginning of the next film. Remember the Spatial Data Management System and think of the person shuffling things on his desk.

(The film showed a user seated in front of a wall-sized screen, pointing at the screen and speaking instructions, such as:

"CREATE A YELLOW CIRCLE THERE."

"CREATE A RED TRIANGLE THERE."

"LAY A BLUE DIAMOND THERE."

"MOVE THAT THERE."

"MOVE THAT BELOW THAT."

"MOVE THAT WEST OF THE DIAMOND."

"CREATE A LARGE GREEN CIRCLE THERE."

The images on the screen appeared and moved in response to the spoken commands and the pointing gestures.)

SESSION E

TRENDS IN VOICE SWITCHING DURING THE 1980s

Roger Camrass,
Butler Cox & Partners Limited

Roger Camrass is a consultant with Butler Cox & Partners specialising in voice and data communications. He holds an MA degree from Cambridge University and an MSc degree from the Massachusetts Institute of Technology.

Prior to joining Butler Cox, he worked for Plessey and was concerned with the introduction of the PDX telephone exchange. Since joining Butler Cox & Partners he has carried out a wide range of assignments in the telecommunications field. These include: a study for the British Post Office on the effect technological developments will have on their marketing plans, several equipment selection and planning studies involving integrated voice and data communications and a study of the technical problems of designing a European network for a major US company.

He has written reports for the Butler Cox Foundation on public data services and on the selection of computerised PABXs, and has carried out market studies on facsimile and electronic mail for an international market research organisation.

Before I came to Torquay, I was hoping to entitle my talk today "Trends in Voice Switching During the 1980s", but after Professor Negroponte's performance I decided that any such title would be extremely mundane in the context of his futuristic views. So instead I will take you aboard my space/time travel machine for a guided tour of Voiceland.

More seriously, the telephone is a very important part of our commercial lives. It is a device which has hardly changed its characteristics in the last 50 years, but despite this it is worth serious attention today. So I should like to review the background to what I believe will be the major events and developments in voice communications in the 1980s. Then, having looked at the background to these developments, I want to give you a personal view of where I believe voice switching and voice communications may go in the 1980s.

I am concentrating my talk around the PABX, the private branch exchange which I believe is the key element to voice communications in a private, corporate sense. The PABX is the device at the centre that provides the switching between all the telephone instruments. In Europe at least, it is the one box which is subject to the greatest degree of change, barring the telephones and the wires themselves.

But it is at the PABX that I believe the major developments and changes will initially take place. That is why I should like to concentrate my talk about the PABX.

I should also like to extend the geography of my talk not only to Europe but to the USA, because I believe that any talk on this subject has to include the developments that are now taking place in the USA. By and large what happens today in the USA may either happen in Europe or have a direct influence on what happens in Europe in five to ten years' time.

When I have given you the widest possible geographic view over the next ten years, I will then conclude my talk with some positive forecasts about what I believe will happen in Europe. I hope that these forecasts will be of some use to network planners and to people in this room who are concerned with large, corporate networks, and will give them some guidelines or a feeling of confidence when they begin to lay down plans for the next ten years.

The progress of voice switching

Pre-1940	Semi-automatic only
1940s	Strowger systems offering automatic features
1950s	Crossbar technology introducing common control
1970s	Stored program control and solid state switching
1980s	?

The logical place to start my talk is to review the last 50 years of PABX technology. If we look at the first slide, we will notice that in consecutive decades we have witnessed fairly major changes in the technology and the fundamental design of the PABX. In the earlier part of the century all systems were entirely manual. The manual systems were originally replaced by Strowger switching equipment that provided automatic switching not only within the public telephone network but also within the private telephone network. Although it was automatic, Strowger switching provided only the very basic features and facilities to the telephone user. Progress in technology has tried to improve the speed and function of the telephone exchange, to give a higher level of service than was available with the traditional Strowger technology. So in the '50s we witnessed a change to crossbar switching, which, although it still used electromechanical technology, provided faster dialling and push-button dialling, but it did not provide a great deal more in terms of features and facilities.

But by the '70s we saw the introduction of computers to the telecommunications field and, for the first time, we began to have some degree of power and sophistication in the switching of traffic in the voice network. As a consequence, the facilities available to the user and to the management of the network were becoming more sophisticated.

Today, I will concentrate my talk on what will happen in the 1980s. If we look back over the development of the PABX, many of you who have used the various generations of equipment may say that nothing very much has happened in the last four or five decades in terms of what the user sees at his desk. It is still the same old telephone. I suspect that very few of you have become accustomed to using the new facilities that are provided by a modern telephone exchange because they are only a small add-on to the very basic and fundamental function of the telephone. The real question today is: will technology in the 1980s take us in a fundamental sense beyond the traditional telephone function that we are all accustomed to and to which all these technologies were directed? I believe that the answer is "Yes" and I shall provide some pointers today to just where that technology will take us.

To provide a perspective on the developments of the 1980s, I should like to look at the various factors which are now building up to promote change in the telecommunications industry. I think that there are two pressures providing the momentum for change. One is a straightforward commercial pressure — a marketing pressure; and the second is technology.

Let us start by looking at the commercial motivation for change. If we look at the US market, we will notice that SPC (store program control) exchanges have been around rather a long time. In fact, the majority of US companies now rely entirely on SPC technology. That has meant a wonderful boom in the interconnect industry for the suppliers of SPC technology during the '70s. Unfortunately, it rather dries up the market, and in the '80s the market is approaching saturation. Those suppliers who have built up enormous revenues on first-generation SPC exchanges are now scratching their heads and looking for a genuine new product with which to approach the '80s, and which they can sell into the installed base of SPC exchanges.

Europe is a little different. We started five to ten years after the Americans in the introduction of SPC technology, and therefore Europe is really an embryonic market. Perhaps 20% or 30% of organisations throughout Europe now have SPC technology exchanges in one form or another, but there is still plenty of scope for selling more first-generation SPC exchanges.

Given the different stage of development of the American and European markets, there are two options available to suppliers. In the USA, the most logical option is to develop a new generation of electronic telephone exchange, which offers fundamentally new features and facilities. The hope is that all first-generation SPC PABX users will throw them out into the garage, and immediately race out and buy the new technology. However, in Europe that option really does not provide us with a logical strategy. Although the market is young in Europe, the competition is very fierce. There are a lot of SPC PABX suppliers. Naturally they are all looking for a technological lead and for some way of getting ahead in the competitive race. They may be able to do that by bringing out second-generation SPC technology. However, given that the market is still young and that relatively few organisations have first-generation SPC PABXs, it is perhaps more logical that the suppliers would look for some degree of added value to their first-generation products that will put them a little bit ahead of the competition in a market which is still growing and which is still intensely profitable. So, even in Europe there is a need for suppliers to look ahead and to try to provide new products.

The commonality, therefore, between the USA and Europe is a commercial driving force to find new products or new functions within existing products to put the suppliers out ahead in the marketplace.

Telecommunications used to be a little island in a very large sea of information technology. It was a very traditional industry and rather segregated from all the other industries. But I am sure that the major characteristic of the '80s will be an awareness within the telecommunications industry and within the surrounding industries that telecommunications should fit into this jigsaw puzzle much more closely than ever before. If we ask ourselves, "What is this total environment in which telecommunications is placed?" the answer must be the office systems market. Anyone looking to add functions, or to come out with a new generation of technology in the PABX market must surely be looking to the office systems market as the overall target for their product plans.

If you ask anyone to define office automation or the office systems market, they will have a great deal of difficulty. The easiest place to start trying to define that market is to look at what people are doing now and try to imagine what they may be doing in the future. Recently, I spent some time with a large, multinational company at their headquarters' office. They provided me with a down-to-earth and very common picture of just what office systems means to a large number of traditional European companies.

In the slide on the next page we have a fragmented picture. In the far left, upper corner we have the telephone exchange and the PABX. This may still be a manual system, but it is a separate system providing telephone facilities. Down below, in the lefthand corner, we have a few local minicomputers or terminals linked to computer bureaux, or even mainframe computers, but nevertheless they are used for specific data processing applications, in isolation to all the other information needs of the organisation. In the far, righthand corner we have the communication of text and images. This organisation had traditional telex machines located in the basement, which were fed by a manual courier service taking pieces of paper from offices to the basement. They also had facsimile machines, again fed by people taking pieces of paper down to the basement. That was their electronic mail system, so to speak, which again operated in total isolation to all the other systems. Finally, in the bottom, righthand corner we have new developments such as word processing, and communicating word processing, which again are growing up in isolation to every other system in the office. So the total picture is very much a fragmented one of a variety of different systems handling separate needs of text communication, image communication, data communication and voice communication in very different ways, and without any thought of co-ordination.

The office systems market

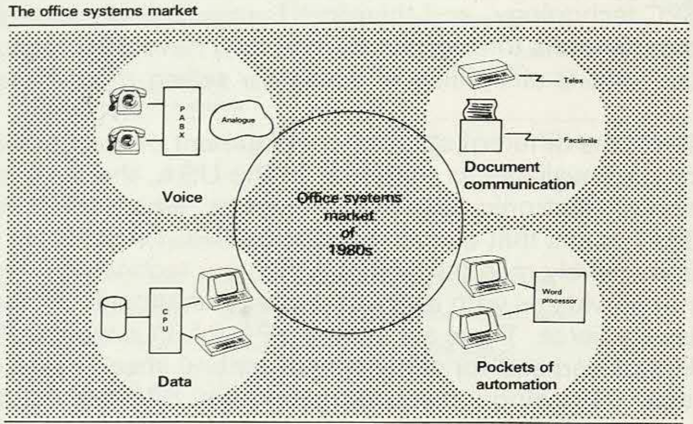

Naturally enough, people are trying to look for a centrepiece to bring this fragmented picture together in the '80s or the '90s. From a telecommunications perspective that is where I believe the office systems market lies; it is a coming together of these functionally and physically separate systems, to try to build a more coherent office systems approach around these systems.

That sounds wonderful in theory, but if I designed a box and took it round to many of your offices, saying, "I've got a means of pulling together all your fragmented systems", you would probably show me the door. It does not really do anything for you until you can talk about payback. In this respect, IBM are one of the leaders; if IBM can find some way of talking about payback, it can justify the sale of virtually any type of hardware for any application.

We do not need to look too far to see just where the payback may come. First, if we were to bring these different systems together, perhaps using common wiring and common switching equipment, we are saving physical resources. There is a cost saving in terms of internal wiring, space, air conditioning and all the rest. That is one payback. That is a fairly unsophisticated area that everyone is aware of. The more intangible and exciting area of payback is in staff productivity. So how can we achieve payback in that area?

Improved communications: the pay-off

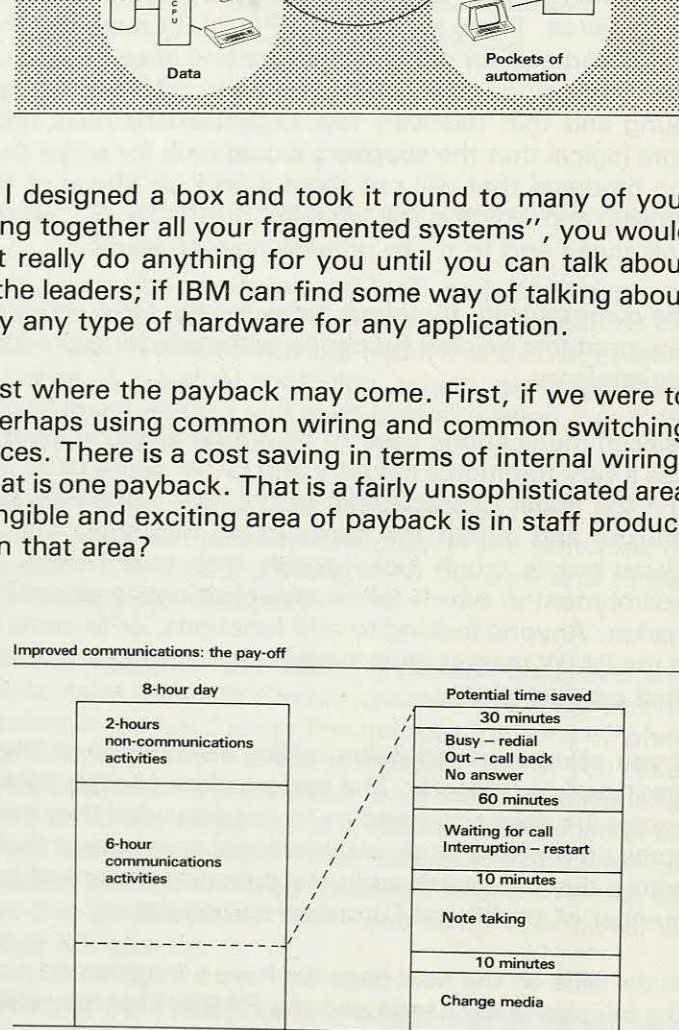

Source: SRI

Looking at this rather overcomplex slide I should like to direct your attention to the lefthand side. According to SRI, the office worker spends two hours of his day in non-communications activities, shut away from the world. For the other six hours of the day he is communicating with people, either in his own office or outside to the world. It is this six hours a day in which we, as communications specialists, are interested.

Let us see how efficiently the office worker uses that six hours of his day devoted to communications activities. On the righthand side of the slide I have broken down the time wasted at present due to inefficiency in the communications systems, in particular, the telephone communications system. I am sure that the examples that I am going to quote are commonplace to you, even though you may be surprised at the actual time spent in these various areas.

The first 30 minutes of the six hours is spent making abortive calls. For example, you ring someone and his number is busy; or you may ring someone and he is out of the office altogether; or someone else will say that he is going to call back, but he never does. So you have 30 minutes of your six hours already wasted because the telephone exchange and the telephone switching system itself is not doing its job properly.

The next 60 minutes is rather more difficult to eliminate, because this is much more of a human problem. You have called someone and his secretary has answered the 'phone for him and said, "He's in a meeting but I'll go and call him", and you are sitting on the line while he hums and haws about coming out to speak to you. So 60 minutes can be wasted just trying to get hold of the right person and to induce him to come and talk to you on the 'phone.

Again, if you are in the middle of a very intense piece of work and the telephone rings, that creates a disruptive problem for you. The interruption creates an inefficiency in your own timescale, because you have got to respond to someone outside, and possibly search around in your office for some piece of information which you have hidden in Dataland. You have got to come back with a sensible answer, satisfy him, and then go back to your work. It can take you five or ten minutes just to remember where you were and to regain your line of thought. So there is a major wastage factor in that element of communication.

Then there are some fairly trivial aspects such as taking messages for other people, or going from a telephone to a telex machine to confirm a telephone call. Those aspects are to do with shifting from one communication media to another, but they constitute a waste of time. So even if you do not believe my figures here, I am sure that you will all acknowledge that there is a major case for improving the efficiency and productivity of communication. I know that many of the first-generation SPC exchanges were sold on the basis that they saved six minutes of executive time per week because of the speed of abbreviated dialling. We are not talking about that sort of saving but about something much more fundamental. We are talking about the basic nature of the telephone system — its interactive nature and the problems that it poses in terms of disruption on time and inconvenience.

So that is one very interesting area of payback. If we can provide products, and functions within products, that can tackle those areas, then we are talking about a major impact on office productivity.

To be able to tackle such complex and sophisticated aspects of office behaviour we will need a lot more power and technology than the traditional PABX provides us with at the moment. Yesterday we were told about many of the developments in component technology and voice processing technology that could give us economic processing capability in the next 10 years.

First, I should like to look at the PABX as it is today and to see just where the impact of component technology and the reduction in the cost of components will affect today's PABX. Then I should like to describe where new functions could evolve from. If you pull off the front covers of your PABX and look inside, you will probably see a mass of spaghetti; but if you are a discriminating observer, you will notice that it breaks down into four different components, which I have illustrated on the next slide on the following page.

There is the component that I have called an interface. It is the interface that takes the signals from the outside world, the analogue world which connects to your telephone or to the public telephone exchange, and brings those signals into the inner world of the PABX. In fact, today, increasingly that inner world is a digital world, quite different from the outer world of analogue signalling. Once inside the PABX you have the big box in the centre called the switching circuits. That may not consist of a whole lot of physical relays or selector switches as it did with Strowger technology. Today, it may be something much more akin to the bus or highway of a minicomputer. Nevertheless, that box provides the basic switching function that connects one subscriber to another within the exchange, so it is a fundamental part of the exchange.

Supervising those two components is the common control circuits. These are the circuits that allocate switching facilities to incoming subscribers, monitor the level of traffic in the exchange, and generally supervise the running of the exchange. It is a very important part of the PABX. Last but not least is the power and the ironmongery that keeps this PABX alive and well.

PABX hardware components

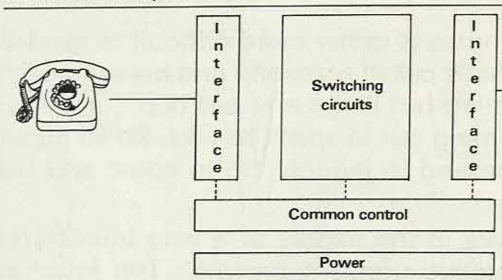

Relative costs of:

	1970s	1980s
Interface	30%	70%
Switching	30%	3%
Control	30%	17%
Power	10%	10%

If we look at the relative costs of these functional components in exchanges designed during the 1970s, we will see that the cost components break down quite evenly between the cost of the interface circuits, the switching circuits, and the control circuits — approximately 30% for each of those elements, with the remaining 10% accounted for by the power supply element. The shift from analogue circuitry in traditional space division switching to time division technology, which exploits the digital world, means that we can expect a rapid decrease in the price of the switching circuits and the control circuits because of the impact of digital circuitry. The forecast by one major PABX manufacturer for machines designed in the 1980s leads us to believe that the cost of the switching circuits themselves — which may become just a digital bus — will virtually fall to zero. Also, the cost of the common control element, which will probably be a number of microprocessors, will begin to fall dramatically from a hardware point of view (the cost of developing the software is another matter). Compared to the switching and common control elements, the interface circuits will stay at a very high price because they need to interface to a very messy, analogue world, which has high currents and requires a considerable amount of circuitry.

What are the implications of these cost changes? One obvious implication is that it does not really matter if you over-design your switching circuits. Why not build the switch to handle as much capacity as could possibly be forced through the exchange? Traditionally, PABX designers have designed them to handle an activity level of about 10%. If your telephone is off the hook for about 10% of the day, the PABX will cope with that traffic. But what happens if we have our telephone off the hook for 100% of the day? This could happen if we are using a data terminal through the telephone, or if we have a four or five-hour conference call? Today, such usage of the telephone by everyone in the organisation would totally disrupt the traffic patterns of a PABX. But, if you increase the switching capability of the exchange to the point where there is unlimited switching capacity (what we call a non-blocking switch), then we can cope with as much traffic as you care to push through any particular terminal.

So, it makes a lot of sense to provide a non-blocking switch, and the incremental cost of building such a system will be smaller and smaller, to the point where it really will not make much difference. On the control side, we have built our exchanges to handle simple voice communication patterns. Why not add some extra processing power there, just in case we need to make use of it for, say, the processing of non-voice traffic? Again, the obvious consequence here is that it will not cost us much more to double the size of our processor or the main memory in the processor. We will add a little bit of cost to the total PABX system, but on the other hand we will have a lot more flexibility and, as a consequence, we can begin to think about handling not just the traditional voice traffic but other traffic as well.

I have now looked at the traditional PABX. If we look at the more far-reaching aspects of voice processing that have not yet revealed themselves in the market today, we notice that voice, when it is turned into digital signalling, is fairly messy and requires a large number of bits for representation. The number of bits required is rather overwhelming if you are used to the data processing world. Therefore, you need a lot of storage and a lot of processing power to be able to handle voice in a digital form — to process it, to file it, and to transform it.

The first thing that we are in sight of doing today is to take digital voice at 64,000 bits per second and compress it down to a more manageable bit rate. As you can see in the slide, we can compress 64 kilobits per second (which is the the top level of toll quality voice communication) down to 16 kilobits per second without any loss in the quality of the speech. 16 kilobits per second is still higher than our traditional data processing transmission rate, but it is more manageable. This compression can be performed by a single microprocessor in a telephone so that the cost of doing it will become almost trivial.

Digital compression of speech

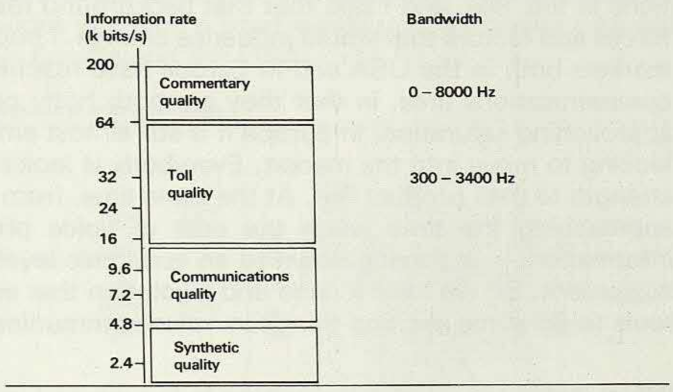

If you combine the compression of digital voice into a slightly more manageable bit rate with the rapidly reducing cost of random access storage, then you have an increasingly attractive formula for store-and-forward voice systems. During the 1980s, the economics of voice storage and, ultimately, of voice recognition and of voice synthesis, will gradually move nearer to what we can afford.

To give you an example of where voice store-and-forward is rearing its head as a service, we need to look at the United States. Bell Telephones, who provide the majority of the public telephone facilities in the States, have just introduced a new service called advance calling. This is based on the digital storage and compression of voice. I should just like to give you a feel for the facilities that they are offering, and the price.

The facilities are divided into two parts. The first part is either a business or a domestic facility where, as you leave your house or your premises in the morning, you ring up the central public exchange and leave a voice message in the exchange. When you ring the central exchange you are talking directly to a computer. You put the voice message into the exchange, and you then tap out on your push-button telephone the destination number to which that message has to be delivered. Finally, you tap out the time at which the message has to be delivered, and you then put your telephone down and go off to work. The computer takes care of the rest. It makes the call for you at the appropriate time, and it delivers the message. If the line is busy or there is no response, it keeps on trying until it has delivered the message.

You can also ring the public exchange from any different telephone that you want (for example, when you have reached the office), and check on the status of the call delivery. The exchange will respond and tell you whether the call has been delivered.

The second example is much nearer to home. It is really a telephone answering machine service, whereby somebody calling your home or office when you are out can leave a simple recorded message at the central exchange. You can then collect that message either from your home or office telephone or from any other telephone.

You would think that those two facilities, given the cost of storage and digital voice compression in the 1970s, would cost a fortune. But I should like to point out that those two facilities are provided virtually on an unlimited usage basis for $3 per month, which is a fractional increase in the monthly telephone charge of about $20 a month. So you can see that the economies of scale and the voice processing technologies really have brought new functions into the realm of the domestic and business user.

I said that I would provide you with a background for the developments of voice communications in the '80s, and I said that that background really consisted of commercial and technical forces and factors that would influence change. I hope that I have convinced you, first, that the markets both in the USA and in Europe have reached a fairly critical stage in the PABX voice communications area, in that they are both hotly competed for. In the States, the market is approaching saturation; in Europe it is still almost embryonic, but there are a lot of competitors looking to move into the market. Everybody is looking for a new product alternative, or a new strength to their product line. At the same time, from a technology point of view, we are rapidly approaching the time when the cost of voice processing — actually manipulating voice information — is coming down to an economic level that we can exploit in the design of new equipment. So we have a need and a solution that are coming together to provide us with the tools to do some exciting things in voice communications during the next 10 years.

When looking at what is happening in Europe and in the States at present in the development either of second-generation PABXs or of new product ideas associated with voice switching, I have tried to sift out two fundamental trends. I should like to use examples today to illustrate those trends so that I leave you with a picture of two particular directions in which suppliers are beginning to move. I think that these two trends form a useful basis for discussion of just what will happen to us in Europe in the next 10 years.

The first noticeable trend, which is coming out of the USA, is that some of the suppliers have grabbed the bull by the horns and are involved now in designing second-generation SPC telephone exchanges from scratch. You may ask, "What on earth does that look like?" I would be loth to recommend any particular supplier as being the ideal example of a second-generation SPC PABX, but I do want to use one particular company and one particular product as an example of what I believe a second-generation PABX may contain.

The product may be familiar to some of you. It is called the Integrated Business Exchange. Notice that is an IBX not a PABX, and we can see the subtle change from voice to information, or integrated business. But there is a subtle change of emphasis there, even in the name. From the point of view of design and technology, naturally enough we see all the latest components of technology incorporated in the system — for example, ample use of fibre optics. Voice is not switched in the traditional circuit-switching sense, it is packet switched. Again, that looks very much like getting on to the voice-plus-data bandwagon, and making the switch compatible not only with voice communication but with data communication.

The cost of switching and of processing is coming down, so naturally enough the IBX is designed as a fully non-blocking switch. It can handle unlimited traffic from all terminals, regardless of duration or activity. It has a massive processor compared to anything that has been built into a PABX before. It has two 32-bit microprocessors, 4 megabytes of main memory and a lot of distributed memory around the machine.

Functionally, the first thing to notice is that it covers a very large range of lines, 100 lines to 4,000 lines. I prefer to call these ports because, again, we are not talking about a traditional telephone exchange. For reasons that I will demonstrate soon, a port can handle both voice and data and is no longer specific to an extension line or to an exchange line.

Every port on the IBX can handle both voice and data concurrently. In fact, each port has a transmission capacity of 128 kilobits per second, which accommodates 64 kilobits per second for voice, so that one single voice conversation can be made down any of these ports, and at the same time a data conversation of up to 56 kilobits per second can be carried out concurrently from the same port.

The IBX is functionally equipped to handle all sorts of data communication and can process data on the fly to provide data facilities that are much more at home in data networks than they ever were in PABX networks. Finally, the IBX is not just a box that gives you a front window to the public telephone network; it also provides you with a front window to the public data networks that are coming on stream, both in Europe and the USA.

Integrated business exchange

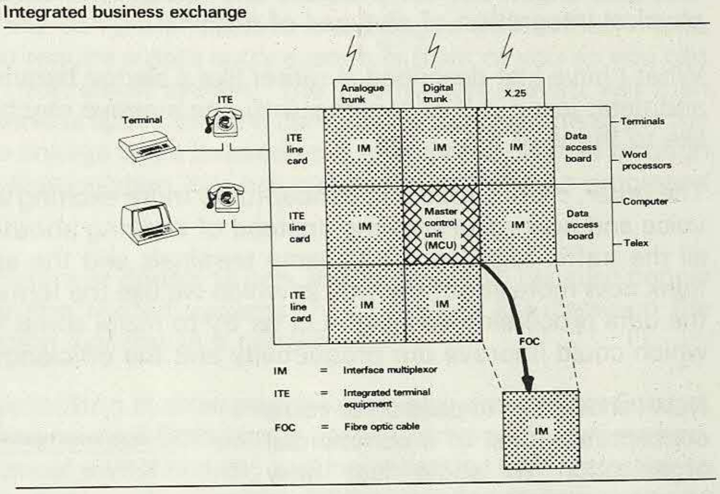

Looking at the IBX in a little more detail, we find that it is no longer one big box that fits downstairs in the basement. In fact, it is a whole kit of parts consisting of small modular units that can be assembled floor by floor, across the building, and that provide local switching, rather than wiring all the way down to the basement.

However, what it does have is one master control unit that represents the nerve centre of the exchange. The master control unit is a packet-switching unit that will handle voice or data transparently. Modular switching units, which each handle 256 ports, are bolted on to the master switching unit. The modular switching units can be assembled around the building and are linked to the master control unit by using fibre optics.

The IBX is also modular in the sense that it will accommodate both voice and data. The telephone is no longer called a telephone but is called integrated terminal equipment. This equipment provides you with a handset and a variety of data features and data ports for accommodating different types of terminals. So with the IBX you really are moving out of traditional voice-only systems into voice plus data in a big way. On the slide you can see data access ports to provide communication into a variety of public data networks, as well as to the public voice network. The IBX also provides a variety of different terminal interfaces for specific applications such as telex, word processing, computing and so on. So the IBX is a modular unit, and is quite different from the traditional single-box PABX. It is very much intended to handle voice and data together.

The office configurator

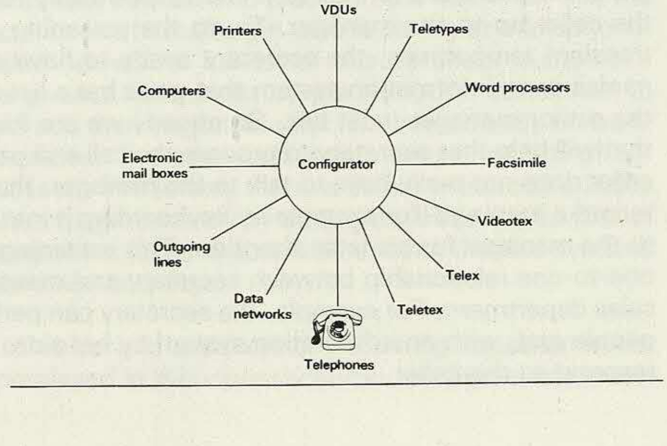

What the suppliers of the IBX are trying to achieve is a coming together of the separate communications systems that I identified earlier into one unified network. It is physical integration of all the terminals shown on this slide, to provide connectivity between different types of terminal as well as between terminals of the same type. Instead of calling it an integrated business exchange, it would be more appropriate to call it an office configurator that provides you with connectivity, on demand, between any of the different types of terminal (telephone or non-voice) that you are likely to find in the near future throughout the office. Obviously, an office configurator requires ample traffic handling capacity to connect all these devices together as and when required. It also requires the processing capability to undertake the delicate transformation between one terminal type and another. For example, when a word processor logs into the configurator, the configurator recognises that terminal as a word processor — in fact it is recorded in its class of service — and then provides the appropriate transformations, say, to the telex network, so that you can use a word processor to talk directly through the private network to the telex network as though the two were entirely compatible. The office configurator, therefore, provides the maximum amount of connectivity, and it provides a physical integration of all types of terminal.

What I have just described is rather like a sledge hammer approach to the integration of voice and data: just put it all together into one massive machine that will handle everything that you like to throw at it.

The other, more subtle and conceptually more exciting approach is the functional integration of voice and non-voice systems. Instead of thinking about the physical wiring and about bringing all the traffic together in the same terminals and the same pieces of equipment, let us try to think now more about the way in which we use the terminals, the telephone, the telex machine, the data processing terminal. Let us try to make some functional links between these systems which could improve our productivity and the efficiency with which we use these systems.

Now I should like to pass on to you a conceptual model of a commercial organisation to show just how important the telephone is, and how well related the telephone is to a number of other information systems, which may at present be manual systems, or may already be embodied in data processing equipment.

Voice: the key to business communications

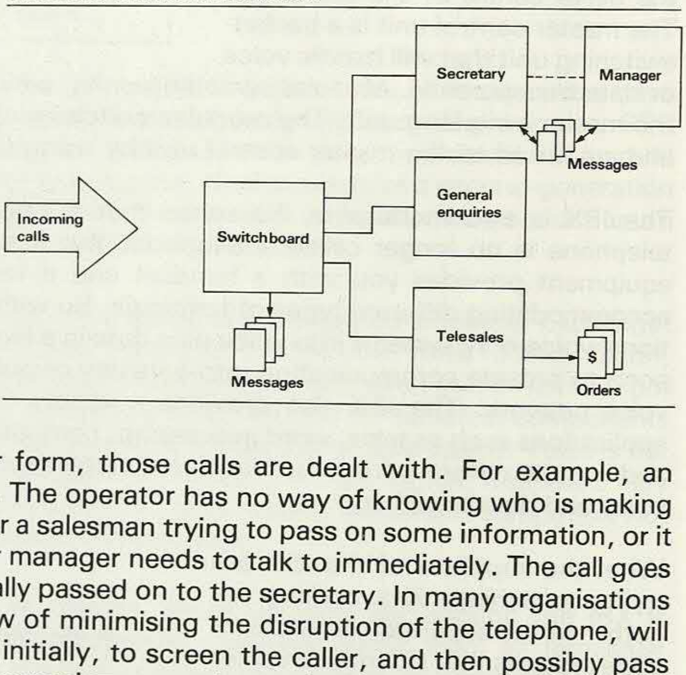

If you look at this very simplified company, incoming calls may be of a variety of types. They come to the switchboard operator. The operator then decides who to pass these calls on to, and then, in some way, shape or form, those calls are dealt with. For example, an incoming call may be to a senior manager. The operator has no way of knowing who is making that call; it may be an advertising agency or a salesman trying to pass on some information, or it may be a high-level contact that the senior manager needs to talk to immediately. The call goes through the switchboard and is automatically passed on to the secretary. In many organisations the senior manager, from the point of view of minimising the disruption of the telephone, will prefer his secretary to interview the caller initially, to screen the caller, and then possibly pass the caller on to the manager. To do the screening properly, particularly if the secretaries are transient temporaries, the secretary needs to have some form of database, be it just a list of names or an information system that gives her a list of important clients and contacts to whom the senior manager must talk. So already we are involved in some related information system that will help that secretary to process the call and pass it through to the manager. Again, if the caller does not really have to talk to the manager, that same information system can be used to record a message (for example by keyboarding it into the system), which can then be passed on to the manager for his later attention. This screening of calls does not need to be handled by a one-to-one relationship between secretary and manager. It could be one secretary for a whole sales department. For example, the secretary can perform the job of operator for many different people and, with an information system by her side, she is much better equipped to be able to respond to the caller.

Another call arriving at the switchboard may be a general enquiry for a product of the company, and it will be routed to a general enquiry office. This office may be staffed by some specialised staff, or it may be staffed by some uninformed people who need a very elaborate database to provide them with accurate knowledge of the company's products and services so that they can respond effectively to the telephone caller. So again, if you want to achieve a high level of service, you have a link between an information system and the caller.

Last but not least, a large number of calls of a general commercial nature arrive at the switchboard. These calls might be potential sales coming into the company, or you might have telesales girls making repetitive calls out of the company. You are in a dialogue with someone who wants to place an order, so, ideally, you require a data entry system in front of you so you can enter that order immediately into an information system. The information system will then pump the order through at maximum possible speed into the commercial organisation to get the order out on time. So there is a definite linkage there between the enquiry that comes through the telephone, and some form of data entry system that will enable the order to be processed and acted on.

If all that sounds like a scenario of the future or double Dutch, it is not quite so because people are genuinely looking at and implementing related systems that provide a link between the telephone and the commercial information systems in a company.

I should like to exemplify these trends by looking at three positive examples drawn from Europe. The first one is the Pharma Bauer Pharmaceutical Company in West Germany. Five years ago, they had 40 telesales girls sitting at manual switchboards, making repetitive calls each day to obtain orders from thousands of retail outlets — pharmacists. All they needed to do was to ring up the chemist, ask him what he needed for the following day, record his response on a computer terminal in front of them, and pass on that information. If you think about that process, two-thirds of the girl's time was spent getting through to the pharmacist and copying down his order. One-third of her time was spent in saying, once she had the order, "How would you like to take advantage of our new bulk offer on Disprins today?" or, "Have you thought of buying some Senecot?" Then she would be able to expand the order and really do her sales talk rather than just act as a sales clerk.

Surprisingly enough, it was possible to eradicate two-thirds of her work through a very simple form of automation. The company had just bought a new SPC PABX. The company also had a mainframe that was used for entering the orders and had customer files. It also had the telephone numbers of the pharmacists to help the girl when she was making the calls. Someone said, "Why not ask the PABX to dial those calls automatically?" So a link was made from the mainframe through to the PABX, to initiate automatic dialling to each customer on the customer list. At the same time, each customer was given a terminal in his own sales outlet where he could enter his daily requirements on to the terminal, early in the morning or late at night. When the automatic call was made to his shop, the computer was immediately connected to his terminal and read off all the orders contained in the terminal before involving the telesales girl. As soon as the orders were read, verified and placed in the customer file, the computer instructed the PABX to link the call to the telesales girl who is sitting at a telephone. In front of the girl was a VDU that provided her with a full list of this particular customer's requirements for the day. The customer is given a little shout by the computer to tell him to come to his telephone, and the telesales girl can then negotiate with him immediately on bulk-discount packages or new product opportunities. The net result is that only one-third of her time, the productive selling time, is now utilised. In this way Pharma Bauer produced a 200% improvement in the productivity of their telesales girls, merely by adding a wire to the control of the PABX. There was no clever physical integration, just a simple link between the two systems and a little bit of systems thinking between voice and data.

A second simple but effective example of voice and data systems coming together is the telephone answering service, which is underplayed in this country but in the States is a major way of life. Traditionally, you have a girl sitting at a manual switchboard with a lot of little pigeonholes in front of her, one for each subscriber to the service. As a call comes in, she gives the single number belonging probably to the answering service itself.

She then responds to the enquiry, writes out a message for the subscriber, and puts it in his pigeonhole until he rings in to collect it. These services were very expensive, because they were labour intensive with girls sitting at manual switchboards. Also, the telephone system itself was very inefficient, and the level of service was poor because the operator never knew which subscriber was being referenced as the call came in.

What has been achieved recently is the linking of a much more modern telephone system — an automatic call distribution board — with an information system. A call that is forwarded from a subscriber's telephone to the telephone answering service bureau initiates a request from the information system for information about the particular subscriber being called. That information is displayed in front of the girl, on a VDU, and, as she takes the call and talks to the person, she can respond immediately with the salutation provided by the subscriber. She knows all about the subscriber, his hours of work, perhaps even where he is today because he rang in earlier to leave a message. She can make a very intelligent reply to the incoming caller. She can record the message quite simply, and put it back into the information system. Later in the day, when the subscriber rings in, she puts in his personal code, and up comes the information. In this way, you have not only radically improved the efficiency of the girl because she is able to handle more calls per day. You have also improved the quality of service that you are providing to your subscriber, by personalising it, by exploiting the information system to improve the amount of information available to the incoming caller.

Functional integration: electronic mail

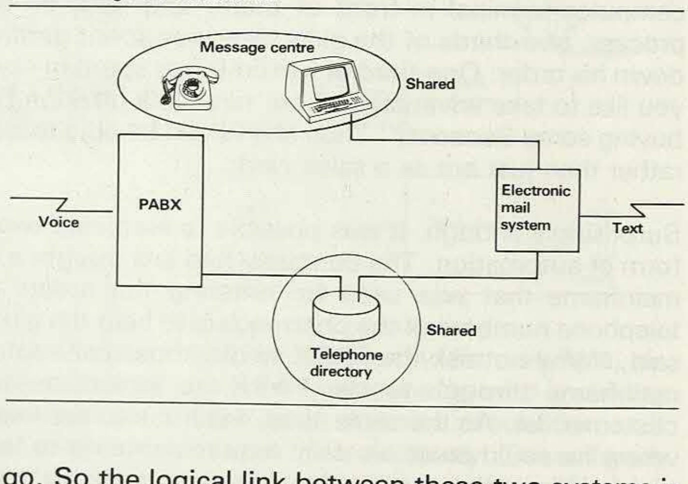

The third example of functional integration is the elusive area of electronic mail, which everyone is planning to introduce but has not quite done so. How can electronic mail possibly link up with the telephone system? This slide does not show a real live example, but it is an example which is being worked upon. There is one common element that an electronic mail system should or maybe could have with the telephone system, and that element is the physical address of each subscriber within the organisation. That address is where his mail goes and that is where his telephone calls also go. So the logical link between these two systems is the telephone directory. If the telephone directory is shared between the telephone exchange and the electronic mail system, it does not matter if you are using either a voice electronic mail system with voice store-and-forward which can be accessed by the telephone itself, or a terminal which is sitting next to the telephone but addressed by the extension of the telephone. In either case, the telephone directory defines the subscriber.

The other common point of reference shown on the slide could be — and this is just a hypothesis — some form of message centre within the building. An external call that is made into the building which is unable to reach a particular person can be re-routed to the message centre. This centre may be one girl at a keyboard. She will take the call, record it in the electronic mail system, either verbally or by typing. The message then sits in the electronic mail system, and when the person returns to his desk, he either picks up the phone or looks in his electronic mailbox and the message is waiting for him. So the whole problem of abortive call making, both external and internal, could be overcome by a simple conceptual link, rather than by a major physical re-work of these two systems. This link can be a shared database, which may be shared physically or conceptually, together with a telephone located next to a data terminal and keyboard.

I hope that those three examples show you that by providing some sort of control link — not a major physical link but a simple control link — and some system unification between the telephone network and the information systems that are already available in most commercial organisations, you can, in general, improve the productivity of individual functional groups or staff by a substantial amount — up to 100% or 200% in the examples that we have seen. Productivity improvements of that order are something that you can go out and sell without having to produce some form of elaborate new PABX design.

To review the blue-sky product opportunities in the 1980s I have talked firstly about the second-generation SPC PABXs. Second-generation PABXs throw technology at the switching problem and, in doing so, create a vast amount of flexibility, with great reserves of switching and processing capacity that you can channel into the handling of non-voice communication as well as voice communication. But I have also tried to show you that you do not need to take this heavy-handed approach. By merely linking systems in a control sense you can achieve not a physical integration, but a functional integration between systems which will give you equal, if not greater gains in productivity, without a dramatic change to the hardware and software that you employ in your organisation.

I emphasised the blue-sky approach. What happens in the marketplace in reality is somewhat different. I will now take you quickly through the US and the European marketplaces to form a perspective against which we can gauge possible trends in the 1980s. The USA is a market dominated by one single carrier — AT&T or Bell — which, rather like the PTTs of Europe, has been the traditional supplier of PABX and voice-switching equipment. In the 1960s, the interconnect industry emerged with licenses granted to private companies to provide voice-switching equipment for private use. You had a sudden surge of commercial activity and some real "get-rich-quick" boys got on the bandwagon. One interesting aspect of the interconnect industry is that it has demonstrated just what the US market is looking for. Bell spent many years pondering that question but never really came up with the answer, judging by the spectacular success of some of its competitors.

To give you an example of the interconnect industry and what qualities these people have built into their products to achieve success, I should like to take the Number 1 company in the US interconnect industry, Rolm, and show you its staggering success. From 1976, this company which had been ruggedising minicomputers suddenly turned its hand to telecommunications, and clothed its minicomputer with a few analogue interfaces, and called it a digital PABX. As you can see on the slide, it did the company no end of good because it has increased its revenue by tenfold in the last four years, and there are not many companies around that can do that. It has since gone public, and everyone is laughing. Rolm did not base its miraculous rise to power on fantastically modern pioneering technology like the IBX that I have demonstrated to you.

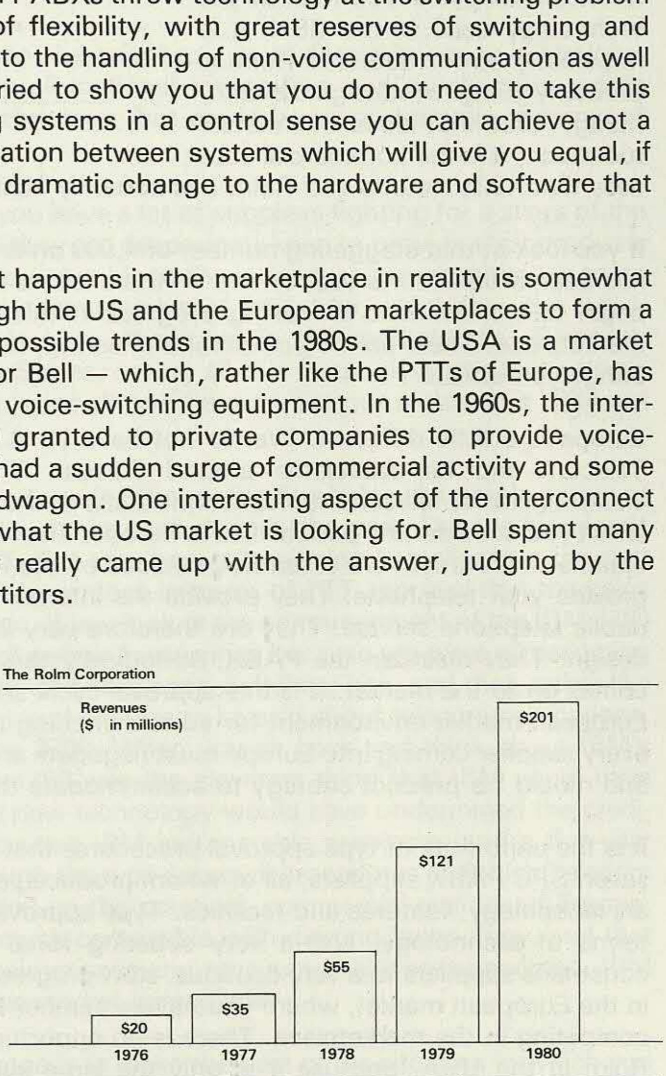

On the slide on the following page you can see that deliveries, in terms of systems, are now up to 3,800 systems in the USA. All Rolm has done is to carefully package up a very simple box which they call a CBX (Computerised Branch Exchange) and that box did not go wildly beyond the expectations of the users. Rolm's strategy, which was a direct response to Bell's traditional offerings, said "We can do all that Bell can do. We can do it at the right price. It is 100% reliable. It will not throw you into a great spin of technology and confusion. It is simple; we install it for you; and you can forget about it". In fact, their success was really commercial packaging, not technology. If you look at the technology used in the CBX, it is hopelessly out of date and Rolm is probably going to redesign the technology. But that does not matter too much; it is the commercial success that really counts and Rolm has certainly achieved that.

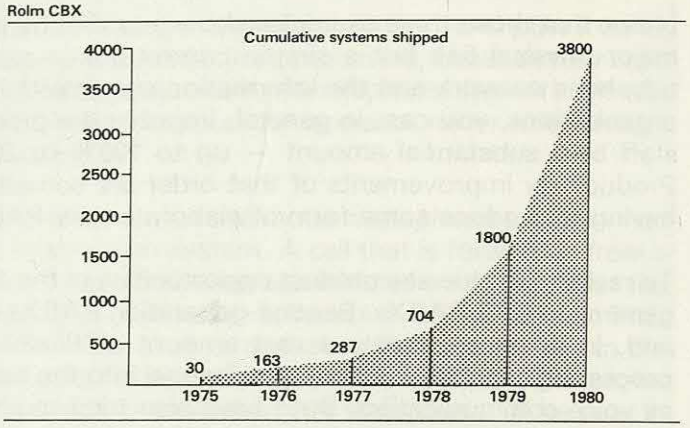

If you look at that staggering number of 4,000 units delivered, and compare it to the IBX, which has been around for a year or two — there is one IBX system installed and one on order. You begin to wonder who is making progress in this world. That is the reality of the American marketplace. There are no points for pioneering; it is just good commercial commonsense that brings you ahead.

Europe is a little different. We do not have an AT&T, but we do have the PTTs — British Telecom and its associates around Europe. They have produced an entirely different environment for telecommunication products to the USA, and it is important to understand that when you consider the possibilities in Europe. The rules that govern the European telecommunications environment are essentially decided by the PTTs, because it is the PTTs who install and provide your telephone. They provide the internal wiring to a building, and they provide the public telephone service. They are therefore very interested in the PABX, its function and its design. They maintain the PABX, so logically they have to approve every new product that comes on to the market. It is this approval cycle and the cost of approval that determines the European market environment for voice switching technology. This is the major obstacle that every supplier coming into Europe must negotiate and comprehend. Each supplier must change and mould his product strategy to accommodate this important constraint.

It is the uniformity of type approval procedures that has produced today in the UK about six or seven SPC PABX suppliers, all of whom produce very similar products in terms of state of the art technology, features and facilities. Type approval procedure is a powerful levelling force in terms of technology, and a very sobering force in terms of technological innovation. It constrains suppliers to a very cautious, slow progression in technology. The result is to be seen in the European market, where you have a number of well-matched, uniform, large companies competing in the marketplace. There is no opportunity for the entrepreneur, as there was for Rolm in the USA, because it is only the large suppliers who can play the cautious, slow, deliberate game that the type approval process dictates.

In contrast to the USA, the European marketplace is not fast-moving; it is not highly commercial, it is well controlled, and it is very uniform. If we review the different ingredients of success in the USA and in Europe, the US presents us with a very fast-moving picture where virtually anyone in a garage can design a second-generation PABX and be on the market within a year with no problems. In Europe, however, you need to have a lot of money to be able to negotiate the PTT approval cycle just to be able to establish your product — not just in one country, but perhaps in seven different countries. This costs many millions of pounds, and is obviously a disincentive to escalating the process of technology.

What does all of this say to us about the future of PABX developments in Europe? I have given you the background of, and an introduction to, the two fundamentally different areas of development that I believe are beginning to take place in the market. The first development is an emphasis on building a new generation of switches that address the question of physically integrating all your information systems together into one uniform network that will handle all sorts of traffic, and will provide the maximum connectivity.

The second aspect is functional integration, which does not really begin to address the physical wiring problem, but looks at the way in which one system is used in relationship to another. For example, how is voice used in relationship to data processing, or in the generation of text or images.

Also, how do we use telex in relationship to the telephone or to facsimile, or to a straightforward data enquiry or data entry system. Functional integration is a much more subtle approach. It is less dramatic in terms of technology, but could be more dramatic in terms of productivity. That is the background.

From a commercial point of view in Europe you have a lot of suppliers fighting for a share of the market, but they are all evenly matched and they are desperately keen to provide their products with a leading edge. So they are beginning to unjumble all these product opportunities, to try to grab for themselves some part of the office systems market of the future, which they can individually exploit and build, hopefully to their own advantage in the marketplace. European suppliers are constantly looking over to the States, seeing all this rampaging new technology, and saying, "Is there anything in this for us, or do we have to set our sights in a different direction?"

I have tried to summarise for you the basic options in Europe, the first being the simple option of throwing technology at the traditional PABX, as we have done in the past. Decade by decade we have progressed the design of the PABX. But now we are beginning to realise that the cost of throwing technology at the problem is so enormous in terms of PTT approval that the gains are not necessarily worth making at that price. If you look at the announcement of the IBM 1750 telephone exchange, it caught quite a lot of people by surprise because we were all assuming that it would be some fantastically new advanced fibre optic switching box, and that, rather like the IBX, it was going to do and say everything there was to be said about telecommunications. In fact it was a scaled down version of the 3750. When you sit back and think about IBM's telecommunications product strategy, the 1750 was the cleverest thing that IBM could have done, because to have brought in a totally new technology would have undermined the credibility of the 3750, and also would have caused IBM innumerable problems (and a five-year procedure) to get the new technology through the type approval procedures of British Telecom and its equivalents overseas. So IBM settled on the practical, commercial realities of Europe, and brought out a product that was very much compatible with existing technology, and that provided no cause for concern as far as maintenance and type approval were concerned. IBM were therefore able to get the 1750 into the market immediately.

As a contrast, the alternative option in Europe is to make the best of the situation in which you find yourself. Most of the major suppliers in Europe have now launched their first-generation SPC exchanges. None of those exchanges represents scintillating, exciting, new technology — it is all pretty standard, down-to-earth technology. But now that their first-generation SPC PABXs are in place and their customer bases are gradually being built up, the obvious next step is to incrementally improve the capability of their existing PABXs. This incremental improvement can be achieved either by establishing the sort of links that I showed earlier with other forms of information system, or by improving the effectiveness of the PABX itself (for example, by providing a greater emphasis on networking). If a supplier has brought out a medium-sized switch he can then bring out a couple more switches — a small switch and a very large switch, with the appropriate software that enables him to sell either to existing or to new customers. The supplier is then in a position to improve his revenues whilst he is still selling to his existing customer base.

If you are a supplier, you can take advantage of non-voice communication, possibly in the functional sense, and you are at liberty to add on hardware modules to your existing product range to accommodate non-voice functions outside the PABX, unlike the integrated approach of the second-generation exchange, which causes you too many problems and frustrations with the PTTs.

Eventually, a supplier has to take the difficult plunge into second-generation technology as events speed up in the States. I know that Rolm, Northern Telecom, Itel and all the major suppliers will bring out second-generation PABXs within a year, and there will be some pressure on European suppliers to do the same. But I cannot see this being a viable commercial proposition in Europe for another five years, mostly because of the cost of type approval. Also, in Europe we make a careful separation between voice and data, since the voice network is heavily controlled by the PTTs and the data network is, to a certain extent, free of PTT control, at least for the moment. Those two networks have been developed almost with total indifference to each other, because of the fundamentally different effects on the marketing of products in these two areas. So there will be an inducement to the bitter end both to keep voice and data separate in a physical sense, and to prevent the data communication markets coming under the scrutiny and careful control of the PTTs.

What a supplier has got to do to steer his way through the type approval process is to bring out new systems which incorporate, to a large extent, existing technology and existing design criteria, but which also provide improved functions by adding new boxes, new software and new capabilities, without causing a major disruption in the type approval procedure.

So what does all that mean to the network planner who leaves this conference today, and has another look at his private voice communication network? I have emphasised that in Europe it is my firm belief that the European suppliers have got to maintain a slow, cautious approach towards enhancing technology. So you are unlikely to find your IBM salesman greeting you at the door, saying, "Now about that 1750 that you've just installed — I can sell you a fantastic new 4750 that completely obsoletes all the technology that you've bought from us in the last five years". If you do not believe that IBM would do that on the voice side, you have plenty of evidence that it is already doing that on the data side. Why cannot IBM do that for voice as well? There is one simple answer: IBM cannot afford to do it. They just cannot afford to take on the PTTs again so recently after they have gone through the painful and costly first-generation type approval process. I think that also applies to all the other established telecommunications suppliers in Europe. So the suppliers have got to provide an incremental approach to upgrading their present systems. They have got to provide you with new software, new add-on bits and pieces to improve the voice function, to provide low-activity data functions that will not disrupt the traffic handling capacity of your existing installation. Thus suppliers are likely to add on data support boxes, and they will also become more involved in the total environment in which the telephone finds itself. In other words, telecommunications suppliers will involve themselves with the office systems environment, so that they can plan and understand the relationships between the telephone and non-voice systems. As communications specialists, their aim will be to really address the wicked problem of inefficiency that the telephone has always represented in the office, and that no one to date has ever been able to tackle and overcome.

So, it is an incremental strategy that will allow you, the network planners, gradually to improve your voice communications, based on your existing investments, or on the investments that you are about to make in current day technology. It is my belief that you are not about to be unseated by some frightening news of new technologies that will make your existing investments obsolete, despite all of the pointers that we see in the USA.

So I end my presentation on a note of confidence and reassurance to you that the future will see a cautious, gradual change that will be to your benefit, rather than a quantum leap into the future without any pre-knowledge or warning of just where that leap will take us.

SESSION F

DEVELOPMENTS IN ELECTRONIC FILING SYSTEMS

Gregory Peel,
System Development Corporation

Gregory Peel joined System Development Corporation in 1972, after receiving his masters degree in Business Administration from UCLA. After three years as a systems analyst in both military and commercial areas, he served for two years as the corporate assistant treasurer. Since 1977, he has held positions in the electronic publishing systems market area, first as divisional controller, then as manager of operations, and currently as marketing manager. In this position he has responsibility for marketing, sales, and planning worldwide for product lines for the publishing industry.

One of the problems in talking to people about electronic filing is to define the problem that we are trying to solve. Technology has got to the point where we are able to offer a variety of different devices for storing and retrieving information, but how will those devices be used both in the office of the future, which is one of the things that people talk about constantly these days, and in the office that exists today? I will start by defining the problem as I see it, and I will then examine the use of electronic filing in the automated office.

I will follow that by a look at the origins of electronic filing systems and at the particular technologies that have evolved into the current filing systems that are available electronically. I will take a deeper look at some of the current technologies which I believe will be useful to you and I will try to evaluate the systems that are available today.

I will also take a look at future developments. It turns out that most of those developments have been covered in previous talks in terms of voice processing, increased storage capacity, and image processing, but I will delve into them a little as they particularly apply to electronic filing. Then I will summarise what I have been trying to put across today.

I should mention at the start, and then not mention it again, that the reason I was invited to speak at this conference was that my company and myself have been involved, in some depth, with electronic filing systems for the past few years. We have a new product that we shall introduce shortly, and we believe it is the state of the art product. I will try to avoid the crass commercialism today and deal with the subject on a more general level, but I invite any of you who are interested in electronic filing systems to join the session tomorrow morning when I shall be happy to give you a preview of what the system will be like.

The problem that we are trying to address is today's information explosion. One of the key buzz words now in the United States is the problem of management productivity. Booz Allen & Hamilton have just come out with a major, multi-thousand page report on that topic. They have said that there are indeed increases in productivity to be had in the management world, although a lot of them are very difficult for people to put their finger on. But they believe that in the next ten years various technologies will come along that will allow at least a 20% improvement in management productivity in that time frame. I should like you to keep that in mind as we go through the rest of my presentation.

There are two particular areas that I believe will make both the cost and the results from electronic filing systems particularly justified in the next few years. One area is the ease with which managers run these systems. Ease of use is essential to overcome any resistance that managers currently have to using the technologies available.

The second area is the reduced cost of filing or data storage. This was talked about during a number of the presentations yesterday. The cost of data storage has come down significantly. However, it has still not reached a point where it makes it cost effective for businesses to put in filing systems that can store massive amounts of data.

Let me now relate the use of electronic filing systems to the way in which current office procedures work. Word processors are becoming quite widespread in the United States and are starting to be used in Europe also, but even so there is normally a draft created by a manager or professional, which is then input by a secretary or an administrative assistant on a word processor. However, once the document has been input to the word processor, we go back to the usual manual functions of running off eight to ten copies of individual documents. Those copies are distributed to various people throughout an organisation, and then filed in long rows of filing cabinets. The problem comes six months later when you try to find the information.

If you take a look at the cost trends over the last 20 years shown on the slide, this does not make a lot of sense. The average annual increases of paper, of supplies and of people have been between 7% and 9% per year, while the cost of communications, data processing and memory have been dropping at an extremely rapid rate.

Profiles of cost trends

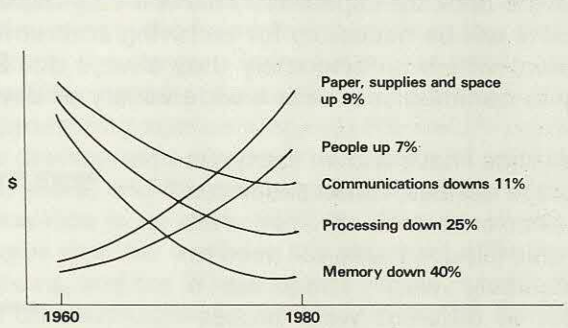

The real use of an electronic filing system is to replace the manual functions with digital storage of information in the form in which it comes out of either stand-alone or shared-logic word processors. For the next few years, managers and professionals will probably, in most cases, continue to prepare a manual draft. But from the word processor onwards, the document will be retained in an electronic form in a filing system. This will allow automatic distribution of documents within an organisation, with the storage and retrieval functions hopefully eliminating the lost and misplaced documents. Also an electronic filing system can incorporate an electronic mail system. As Roger Camrass mentioned, electronic mail could be handled by the phone system — by an automatic PABX. However, another potential way of handling electronic mail is through an interface to the electronic filing system. From there, information can be accessed either by word processors, or by executive terminals which would be fairly dumb terminals that managers and professionals could use for electronic mail and for the storage and retrieval of information.

Electronic filing systems use in automated office

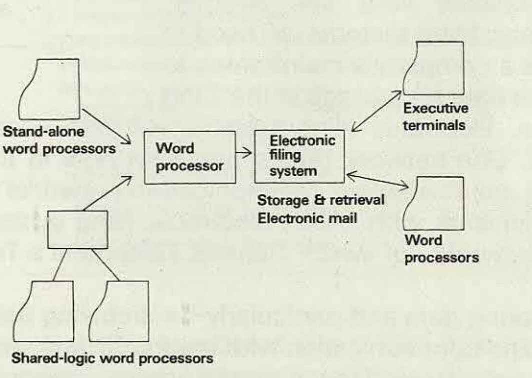

I have examined the requirements for an electronic filing system. We believe that the storage capacity would have to be 100 million characters or more. This is equivalent to 20 file drawers with 3,000 sheets of information per drawer and an average of about 1,500 characters a sheet. Storage capacity of 100 million characters would be a minimum size that would be really useful within an organisation. The system should be capable of retrieving documents on a full text search in less than 20 seconds. If the retrieval times are more than that, managers and professionals will not use the system to the extent that it should be used. Automatic document indexing would be an absolute necessity within the system, and the storage of full text would be required. The system should accommodate spelling variations, and the capability to specify synonyms for search terms by particular individuals should be available, so that a constant researching of the file is not necessary for particular words and particular variations in words. The system should also permit additional terms to be added manually to documents. For example, in the United States, if you had documents that referenced Ronald Reagan you might want to go back in the last couple of months and add "President" to some of those documents so that you could still access the documents that you are looking for. The system should also provide for the easy retrieval and modification of stored documents so that documents can be retrieved, have changes made to them, and be restored — hopefully in an online situation.

I also see the availability of electronic mail, or at least the interface to electronic mail systems, as being a requirement for an electronic filing system. Also, I believe that a filing system will need to provide back-up capability. This will vary depending on the application, but some back-up capability will be necessary for archiving and retrieving systems, for example when computers go down, which unfortunately they always do. Electronic filing systems will also require the ability to communicate with a wide variety of devices.

On this slide I have shown the types of storage devices, retrieval devices and electronic mail devices that electronic filing systems will need to communicate with. There are a number of different word processors on the market today, and we believe that electronic filing systems will need to communicate with them. Also, as various computer devices (minicomputers and large mainframes) acquire the ability to communicate with one another, electronic filing systems will need to access a company's mainframes to retrieve data for storage in the filing system. Electronic filing systems will also need to communicate with communications networks. One network that is prevalent now in the States is Ethernet, and various other companies are developing communication networks which it will be necessary for filing systems to communicate with. Also, electronic filing systems will need to communicate with electronic mail networks, of which General Telephone's Telemail is an example.

Electronic filing systems communication requirements

An electronic filing system should be able to communicate with:

	Storage	Retrieval	Electronic Mail
Various word processors	x	x	x
Minicomputers	x	x	x
Large mainframes	x	x	x
Communication networks (e.g. Ethernet)		x	
Electronic mail networks (e.g. Telemail)			x
Magnetic tape	x		
Other electronic filing systems		x	x

For storing data and particularly for archiving data, by far the cheapest method is still magnetic tape, and communication with magnetic tape would be an absolute necessity for filing back-up. Finally, electronic filing systems need to communicate with other electronic filing systems, so that organisations that are located in various countries around the world will be able to communicate with one another through electronic filing systems, and also use those electronic filing systems for electronic mail.

Talking to a number of people interested in the office of the future and in electronic filing, we have found that a much wider variety of applications is being discussed than we had ever believed possible when we started looking into this area. Work group filing, project management, and activity management are all areas within an organisation that will be applicable to the future uses of electronic filing, in addition to personal filing and electronic mail.

This slide shows that practically every area within a major organisation can be used in the electronic filing arena. The 100 million bytes of information that I quoted earlier as being an absolute minimum requirement within a large organisation could easily be used up by any one particular area within a company. Particularly in the legal area, one individual would easily be able to fill up that amount of information. This is another reason for the absolute necessity to have filing systems communicate with one another.

Applications

Engineering Specifications Change orders	Personnel Skills inventory Requisitions Applications
Marketing Proposals Account profiles Competitive evaluations	Financial Taxation analysis Budgets Credit reports
Administration Policies and procedures Library card catalogue	Purchasing Contracts Purchase orders Vendor selection
Legal Litigation support Patent application	

I want now to look at the origins of electronic filing, so that you will have an understanding of where filing systems are coming from, why they exist in the forms that they do today, and hopefully that will indicate the path they will take in the future. There have been three separate developments that have resulted in electronic filing systems: database management systems — those systems coming from traditional data processing applications over the last 20 years; from word processing which has been a fairly new development, particularly in terms of proliferation, having taken off only in the last two or three years; and from micrographics which I will touch on briefly. Although in many ways micrographics is an antiquated technology, it still offers many cost advantages over other technologies that are available. These three different technologies have resulted in electronic filing systems, and the filing systems based on each of these technologies have different advantages and disadvantages, on which I will touch briefly.

Let us look first at database management systems and the development of electronic filing within those systems. One of the first, and biggest developments that happened in my company, System Development Corporation, was the invention of the inverted file system. At the time my company was a public, non-profit company in the United States, so the inverted file technology became public property, and it has been used in most of the filing systems that have been developed to date. It has been very important in the development of electronic filing. Bibliographic retrieval systems started to be offered about 1969 by a variety of organisations, primarily headquartered in the United States. I have listed a few examples on this slide. Most of those systems still exist today, essentially in their original form, and they are available worldwide. The system that is offered by my company is Search Services, which is offered all over Europe as well as North America and Japan. Most of the other systems that I have listed are also available worldwide.

Data base management systems history

Year	Development
1967	First inverted file system developed (System Development Corporation)
1969	Bibliographic retrieval systems offered ORBIT – System Development Corporation INFOBANK – New York Times DIALOG – Lockheed Corporation
1971	STAIRS Developed by IBM for internal use (International Business Machines)
1973	STAIRS Offered as customer software package (IBM)
1981	Bibliographic retrieval services/systems Search services/ORBIT – System Development Corporation INFOBANK – New York Times DIALOG – Lockheed Corporation LEXIS/NEXIS – Mead Data Central INFOGLOBE – Toronto Globe & Mail Software packages STAIRS – International Business Machines QL – QL Systems Docu/Master – Turnkey Systems, Inc. AFPS-1-Datafusion Corporation

The first real full-content retrieval system was the STAIRS system, which was developed in

1971 by IBM. It was developed by IBM for internal use in their law suit with the US Government, the anti-trust suit that is still continuing and for which a court date has yet to be set. Back in 1971 IBM had already accumulated such massive amounts of data that the lawyers could not possibly cope with all of it.So IBM had their software development people put together the STAIRS system which is based on the inverted file system that I mentioned earlier. STAIRS turned out to be a very good system for a number of IBM's customers who had similar problems. IBM started offering STAIRS as a customer package in 1973, and they are still offering it and supporting it today.

Today, in 1981, a variety of bibliographic retrieval services are offered. Most of the ones that are mentioned on the slide are offered around the world. There are a number of different software packages that are offered in the database management arena. These packages tend to run on large mainframes, and they practically all use the inverted file system that I mentioned.

Data base management systems

Strengths	Weaknesses
Multiple databases on-line	Data processing orientation: Requires specialized language
Most allow search on content	Expensive: Requires medium to large mainframe computer
Offers full text retrieval on-line	Limited access to information by appropriate individuals
Powerful query language	Requires structured files
Allows very large databases	Off-line storage and updating from central files
Extensive available communication	No electronic mail capability

If we examine the strengths of database management systems, there are many advantages that primarily come from the fact that you are usually running these systems on large mainframe computers. They provide the ability for multiple databases, due to very large storage capacities that are capable of being accessed by large mainframes. Most of them allow full search on content through terminals from the various databases that are available. STAIRS and the other software packages that I listed on the previous slide offer full online text retrieval. They tend to have powerful query languages which have been developed by programmers to allow access to data, no matter how obscure it is. They allow very large databases due to the range of disc drives that are available today from most manufacturers. Also, because they are on a mainframe, there is extensive communication available with all kinds of devices.

However, there are also a number of weaknesses. One is that because these packages were, in general, developed by programmers as part of a database management system, they often require a specialised language to be learned by the user. This has largely restricted the use of these packages to either libraries within organisations, or to very specialised users who have the need to learn that language to retrieve documents. As I mentioned, this is one of the two key areas that will be dealt with in the next few years — simplifying the language and encouraging managers and professionals to use systems. Because a medium to large mainframe computer is required, these systems tend to be fairly expensive, although the cost of large mainframes has come down dramatically in the last five years, and will continue to do so. If an organisation already has a large mainframe available within its environment, it tends to make the installation of a software package cheaper, but there is still an inherent cost in using that mainframe computer.

Another weakness is the limited access to information by the appropriate individuals. This is primarily due to the specialised language that is necessary. Because of the technology utilised by these systems, a very structured file is required. This may not be such a bad thing for people to start using. They will become as accustomed to using structured files within an electronic filing system as they are to using structured files in their manual filing system. However, in the future as technologies develop, I believe that unstructured files will become the everyday norm that people will be dealing with.

Off-line storage is the way in which these systems are used and so updating from central files is particularly difficult. Files or particular documents have to be deleted and then re-entered with updated information. The final weakness is that, although some electronic mail capability is available on large mainframes, it is generally of a limited nature.

I would like now to look at the word processing industry. A number of developments have taken place in the last couple of years which have enabled word processing manufacturers to surpass the electronic filing systems that are offered by the large mainframe manufacturers. The first real storage device offered on a word processor was magnetic card storage, offered by IBM in the mid-1960s. There are still thousands and thousands of these devices out in the field, and they still have quite a bit of use left in them. The next development, which also came from IBM as well as from Wang, was the use of floppy discs for storage from typewriters. This occurred in the early 1970s, followed by hard disc storage being offered as an option on word processors. This was first developed by Wang Labs in the mid-1970s.

In the last couple of years there have been some very interesting developments. One example is AIM (Associative Index Method), offered by Datapoint. AIM is an extension of the hard disc storage device and it is able to retrieve data in various formats from word processing input. Practically all of the major word processing manufacturers now offer a variety of hard disc and floppy disc storage systems that interface directly to their word processors.

Word processing storage

Strengths	Weaknesses
Office orientation: Most require less specialized language than DBMS	Currently require specific manufacturer's equipment
Offers both word processing and data processing capabilities	Limited database size due to available word processing equipment
Automatic filing as part of word processing capability	Limited query capability
Some offer integrated communication	
Electronic mail available	

On this slide I have shown the strengths and weaknesses of filing systems based on word processing devices. One of their major strengths is their orientation to the office. Most word processors require a much less specialised language than the database management systems. In many cases, we believe that the language is simple enough for managers and professionals to use easily. Most of the word processing manufacturers offer data processing capabilities on their word processors, as well as word processing, so that both data and text would be available through the electronic filing system. Automatic filing is accomplished on the systems through the word processing interface, and that enables filing to be carried out as an automatic part of the word processing environment. Some of the manufacturers are currently offering integrated communication. This means that in addition to the disc storage and the word processor, a number of other devices can be communicated with, either directly through the word processor or through the electronic file. Also, electronic mail facilities are now available from most of the word processing manufacturers in some version or another, so the electronic filing system has access to electronic mail facilities as well.

The slide also shows the weakness of filing information in word processing systems. One of the evaluation criteria that I mentioned for communications devices was that it will be absolutely necessary to communicate with a variety of different manufacturers' equipment in the future. All of the current offerings from word processing manufacturers require the use of their specific equipment. I believe that they have intentionally designed their systems so that it makes it very difficult for users to try to access somebody else's equipment. The limited size of database is another problem. This is due to the available word processing equipment and the fact that the retrieval mechanisms are fairly slow. Increasing the database size would slow down those mechanisms to an intolerable rate. Word processing storage systems also have limited query capability, although some of the latest offerings are starting to expand significantly in this area.

Let me now tell you briefly about the developments in the micrographics area. I believe that, in general, microfilm is a fairly antiquated technology. However, it does have a lot of uses, and for the next few years it will continue to be particularly useful in systems where it is necessary to store an extremely large amount of data. Microfilm was invented by George McCarthy in 1928, and after that there were very few developments in the basic microfilm technology until the late 1960s. In the late '60s people started linking computers to various microfilm devices, and Stromberg-Carlson developed the first computer-output-microfilm system so that computer-stored data could be output on microfilm. The microfilm was then stored in its own mechanism, and the data was retrieved using that storage system.

In 1977, 3M introduced their 715 COM device that greatly reduced the price of computer-output-microfilm and made it applicable to today's environment. Now, in 1980, about 65% of the market is in 16mm film, and 35% in 35mm film. There are computer-assisted retrieval units available from a variety of manufacturers. All of these units have the capability of accessing a great deal of stored documents.

Indeed, that is the greatest advantage of micrographics: it offers the least expensive storage per document, and probably will for a few more years. I believe that optical videodisc technology will probably do away with a good deal of that advantage of micrographics, but currently micrographics does offer the least expensive storage per document. Another strength of micrographics is that an extremely large number of documents can be stored while still retaining relatively good performance.

Micrographics storage

Strengths	Weaknesses
Least expensive storage per document	Initial document input is manual and expensive
Very large document storage capabilities with good performance	Limited communication available
Answers problem of handling non-digitized information	Query limited to manually-input information

In taking a look at micrographics, one of the interesting stories that I heard from one of the manufacturers was that the Three Mile Island disaster in the Pennsylvania area, a couple of years ago, was really the result of people not being able to get at data fast enough. It turns out that after the nuclear plant was built there, all of the documentation that went along with it was delivered in several trucks to the plant on the day that it was turned over to the people who were going to run it. Of course, they had no way of processing all of those documents and storing them in a system that was adequate for retrieving them on a timely basis. On the morning that the accident at Three Mile Island happened, they were just not able to get to the proper document in time to find out exactly what was going on in the system, and within an hour they had made the decision to close the plant down. If they had spent a little extra money, even on a micrographics system, to have the builder of the nuclear plant store those documents properly, they probably would have been able to contain the incident, and we never would have heard of it.

The other strength of micrographics is that it answers the problem of handling non-digitised information. That is particularly difficult right now with images. On the other hand, it also has an inherent weakness in that when you want to work with digitised information you have to find a way of re-inputting the documents to transform them to a digitised form.

Looking at some of the other weaknesses of micrographics — initial document input for use in a computer-assisted system has to be a manual exercise, which is very expensive. Also, micrographic systems have limited communication facilities, although some of the manufacturers are now making available the ability to transform the microfilm into systems that can be easily communicated through a facsimile methodology. Finally, queries to a micrographics storage system are normally limited to manually-input information, which means that full content addressed queries are almost impossible.

Let me now spend a few minutes reviewing some of the current technologies that are used in electronic filing systems. I want to deal particularly with inverted file systems because it is used in practically all of the full text retrieval systems available today. I will also give a brief overview of how the bibliographic retrieval systems work and also take a brief look at micrographic retrieval.

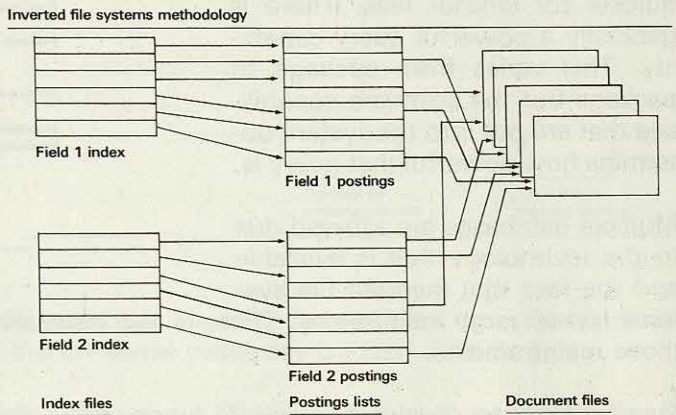

In the inverted file system methodology, which is illustrated on this slide, there are three different files set up within the computer system: an index file; a file which contains a list of postings; and the files containing the documents themselves. The index files, broken down into several fields within each document, contain a list of every word that is likely to be accessed by the users. Each of those words then has a pointer to the postings list. The postings list contains the indices of exactly where each of the documents contains each of those words. The documents are filed in a digital form within the system. The inverted file methodology requires documents to be duplicated a number of times, and since the words generally exist once within the documents themselves, most of the words will exist again within the index file. In addition, there will be the list of the postings, again containing pointers to all of the documents. One of the weaknesses is that quite a bit of storage is required for inverted file systems.

Inverted file systems types

Full inversion
postings lists contain every term (word or number) in documents

Partial inversion
entire documents are stored, but postings lists contain only partial information (e.g. no location within document)

Incomplete inversion
vocabulary is limited by excluding some terms ("stoplist") from indices and postings lists

There are several types of inverted file systems. If you were talking to particular manufacturers, or to software package vendors about systems available, they might use some of the terms shown on this slide. Full inversion means that the postings list will contain every word that is in the documents. The advantage of full inversion is that you can retrieve documents using full text keys. However, the disadvantage is that a number of words — 600 to 800 usually in the English language — appear many times and have very little meaning within a document.

Partial inversion is when the entire document is stored but the postings list contains only partial information. For example, you will not be able to retrieve documents with key words in a particular order; you only know that those words appear somewhere within the document. Incomplete inversion is when the vocabulary is limited by excluding some of the 600 to 800 words that I mentioned, which are usually called a "stoplist" or "stopterms". Those stoplists would then be excluded from the indices and the postings lists, and this greatly reduces the amount of storage that is required for the inverted file technique.

This next slide shows the strengths and weaknesses of inverted file systems. One strength is that they offer full-text search on text. The technology allows fairly fast query times since you go directly to the list of fields, find the locations of words in various documents, and retrieve those documents. Generally, even with very large files, the systems can offer retrieval times of about 15 to 25 seconds and they are much quicker for smaller files. There is generally a powerful query capability. This varies from package to package but the software capabilities that are put into the system determine how powerful that query is.

Inverted file systems

Strengths	Weaknesses
Full-text search on content	1.5 to 4 times document size required for storage
Fast query time	Does not allow real-time updating
Powerful query capability	Expensive: Requires large mainframe computer
Allows multiple data bases	No electronic mail
Extensive communication available	

Multiple databases are allowed due to the technology that is available and the fact that inverted file systems run on large mainframes. There is also extensive communication facilities available with those mainframes.

Turning now to the weaknesses of inverted file systems — between 1.5 and 4 times the document size is required for storage. 1½ times would be required when partial inversion is used. Full inversion would require up to four times the size of the document, depending on the number of vocabulary words in the documents themselves. Inverted file systems do not allow real time updating. Updates have to be done through a process of deleting documents and then re-entering them with the updated information on them.

Inverted file systems are expensive. They require a large mainframe computer. In general, there is no electronic mail system available as part of the system itself, although some of the mainframe manufacturers now offer interfaces to electronic mail systems.

I want now to look briefly at bibliographic retrieval systems. Both full text and abstract searches are utilised by particular packages. Most of the early bibliographic retrieval systems that came out in the late 1960s offered searches merely on abstracts, looking at key words to retrieve documents. Some of those early vendors, and some of the newer ones, have converted now to full text so that full text search is available, generally done on a key word environment. The big problem with bibliographic retrieval is that it is really only a partial solution, and an organisation's own files are not easily put into the system. Some large companies now ask the vendors that offer bibliographic retrieval systems to input their own files, but this is a process that takes a certain amount of time, and generally the files will not be available on a real-time basis. The bibliographic retrieval systems usually are charged by the various vendors on a connect-time basis so that you are charged for the actual time that you are connected to their mainframe and accessing information. This makes the initial cost of getting into bibliographic retrieval very small and makes it attractive to a variety of types of users.

The third methodology is to use a computer-assisted micrographic retrieval system. Generally, a central processor is used and in most cases this is a minicomputer with software prepared for it such that an inverted text type system is used. Index files and postings lists, as used in inverted file systems, are input into the central processor. However, it is necessary to do this manually in the case of micrographics because the terms are not readily available in digital form. Usually the user will have two terminals in front of him. One terminal will be for entering the requests into the minicomputer, which will search the index files, find the postings list, and in this case the postings list will refer not to an electronically stored document, but to a location of a microfiche or microfilm unit. A microfiche cartridge or a microfilm tape can be inserted into the unit, and an automatic key will be sent across from the minicomputer processor to the microfiche storage unit which retrieves the microfiche itself on a separate terminal. This does have the advantage of offering the user a very large number of documents — several millions of documents are available in some of the systems today. Some of these are very elaborate systems. However, they tend to be very mechanical in retrieving either fiche or film documents from a stored location.

Micrographic retrieval systems methodology

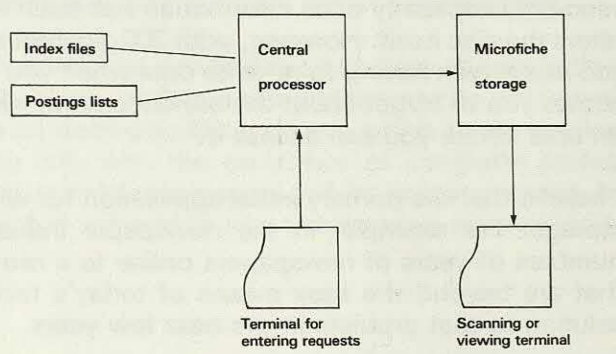

Let me now take a look at some of the future developments that are going to help greatly the implementation of electronic filing systems within organisations. Two of these were covered to a great extent by previous speakers so I will not spend much time on them. Increased data storage capacity is one, and the other is voice storage which was covered in a state of the art sense yesterday by Fred Jelinek.

Image storage is another area that I will talk a little about. This was touched upon by Professor Negroponte, and it will become more and more important for electronic filing systems.

Data storage capacity

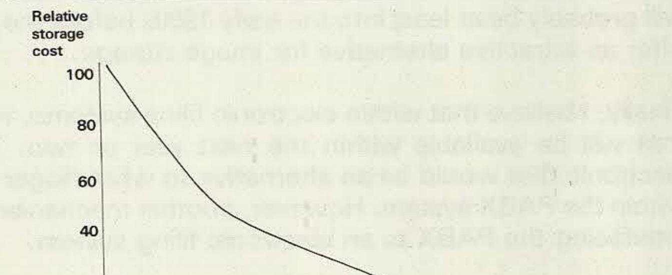

Storage costs have come down significantly over the last 20 years to the point where it looks as though they will almost go down to zero. Of course, that will not happen, but for the foreseeable future, at least for the next ten years, there are developments in the wind that will continue to reduce the cost of storage at almost the same rate as the last 20 years. I imagine that by the end of this decade other developments will come along to continue to reduce it. So within the next few years we will still continue to see a dramatic reduction in the cost of storage. This will make the storage of documents for subsequent retrieval in the office an extremely attractive proposition.

Even with current technology, there are a lot of new developments in disc drives. One of these is the Winchester disc drive which is a sealed disc technology. In addition to reducing the cost of storage, Winchester drives also provide high reliability, since the disc drives are sealed and are not subject to a lot of the things that happen in the computer room environment when disc drives are removed, returned to shelves, and then put back into disc drives. Another developing technology, which is an extension of the current technology, is thin film technology. This technology will continue to enlarge the amount of data that is available on rotating disc drives. Currently, products that provide up to about a gigabyte of data are available and, through the use of thin film technology, this storage capacity will undoubtedly expand by many factors in the next few years.

Optical disc is a very exciting technology that is currently being developed by a number of vendors. Commercial products are becoming available, although generally, in thinking about optical discs, we are considering products made for the home, for online movies and so on. But if you apply the same technology to data storage systems, the amount of data that can be stored on an optical disc will increase dramatically. By 1983, I expect that computer manufacturers will be offering drives that can store up to 200 gigabytes. Currently, optical discs offer a read-only capability once information has been stored, since the storage process permanently alters the disc itself. However, with 200 gigabytes available you do not have to concern yourself too much with having to re-write data when you want to update it. That size of storage system allows you to forget about the old document, so you can re-write the changed document into an area where you can access it.

I believe that the primary initial application for which optical discs will be used is mass archival storage. For example, in the newspaper industry, newspapers are anxious to get several numbers of years of newspapers online to a retrieval system. This requires storage capacities that are beyond the easy means of today's technology, and optical discs will offer a quick solution to that problem in the next few years.

A related area to optical discs is storage of images. The reason why optical discs will become very important in image storage is the current requirements for storing an image. A single page of 2,000 characters of data would take 2,000 bytes of storage, but a black and white facsimile image requires 20,000 bytes (even with the data compression techniques used in facsimile transmissions). However, a single-page size grey-scale photograph requires 200,000 bytes of storage capacity, and a colour photograph will require between 800,000 bytes and one megabyte of data. Using today's technology disc drives it is very expensive to store that amount of data, but optical discs will offer a solution to that problem in the next few years.

The current retrieval technology that is available for images is similar to micrographics, and it will probably be at least into the early 1990s before the optical disc manufacturers will be able to offer an attractive alternative for image storage.

Finally, I believe that within electronic filing systems, verbal message storage will be an offering that will be available within the next year or two. The storage of verbal messages within electronic files would be an alternative to what Roger Camrass mentioned as a storage device within the PABX system. However, another mechanism for handling voice storage would be by interfacing the PABX to an electronic filing system.

By voice communication I really mean the elementary form of voice communication that Fred Jelinek talked about as "isolated word recognition". But I believe that within five years manufacturers will be offering this so that, for example, a manager could call on the telephone or communicate with the office where his terminal to his electronic filing system was. Instead of having to sit down at the terminal, he would use the natural modality of speech and say, "Please give me my messages", and the system would give him any verbal messages that were stored there.

Another area that we are looking at is verbal document storage where documents can be input and stored as verbal messages. Initially, this will be particularly useful for the visually handicapped. However, there are several applications that will expand significantly over the next few years. Within electronic filing systems, the storage of verbal documents will be possible once optical discs are in common usage.

Dr. Jelinek spoke of voice dictation systems as being available in about ten years. That was a little more optimistic than I realised, and I shall be very excited if that happens in the next ten years. But voice dictation is definitely another development that will be very useful within the environment of electronic filing systems.

Let me now briefly summarise. There are two things which I believe are holding up the use of electronic filing. First, whilst the technology exists today, we still have the problem that storage costs are high relative to word processing. Storage devices for the 100 megabytes that I talked about cost in the region of $50,000 or so, while word processors can be bought for under $10,000 now, and prices are dropping every day. However, I believe that storage costs will drop to a relatively low level in comparison to word processors within the next few years, and that will help to enhance the implementation of electronic filing systems within organisations.

The second item that I mentioned several times was the availability of an easy-to-use, office-oriented interface. This is necessary so that managers, professionals and secretaries will require very little training, so they will be able to use an electronic filing system which today requires specialised training, and is useful as a device only with the assistance of computer professionals. By the end of this decade I believe that most businesses will be communicating by digitised information, and the real mechanism that will enable that to happen is the availability of electronic filing systems.

SESSION G

FUTURE APPLICATIONS FOR MICROPROCESSORS

Norman Eason,
Eason Consultants Limited

Norman Eason has worked in computer systems since 1964. He initially worked on high-integrity dual-processor message switching systems and computer-controlled warehousing systems. Subsequently he spent five years as European systems engineering manager for a US minicomputer manufacturer before joining ERA Technology Limited where he was manager of the computers and automation division, responsible for CAD and microprocessor systems development. Mr Eason is now an independent consultant. He has a BSc (Eng) degree, is a Chartered Engineer and a Fellow of the Institute of Electrical Engineers.

The first problem regarding microprocessor applications is where to start and where to finish. As the title of the talk is "Future Applications for Microprocessors" I could say that the applications are anything and anywhere, and then leave the stage. In fact the problems are much wider than that and I could spend all afternoon talking about applications in depth. However, what I will try to do is to look at the application trend and give some indication of the limitations to that trend.

In order to try to look at where we should go from here or where the technology will go, first let us look at where we have been. Microprocessor technology started way back in the early 1960s when people decided that they should replace a semiconductor device by a package system. They replaced the one device right at the bottom of the first slide — a transistor — with a gate made out of resistor/transistor logic. The RTL gate consisted of something like four components on a chip. That was the start of integrated circuit technology.

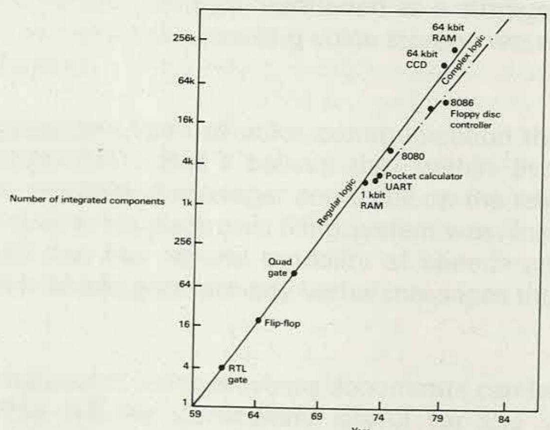

We then advanced fairly slowly till about 1964 and put four devices on a chip, which was a quad gate. So you had the same thing that you had before, an RTL gate, but four times over. We then expanded this a little further and produced the ability to store one bit of information. This one-bit storage facility was called a flip flop and required about 16 devices on one gate. (Sorry, the quad gate came later).

So we had the ability to store memory on an integrated circuit. The complexity of chips expanded, until you get to round about 1970, when you had 1,000 integrated components on a chip. There is a fair time gap between the quad gate and the 1k bit RAM. I should like to fill in some of the information on that gap and the reasons why the technology started to take off.

You are all aware that, during the 1960s, America had a race on its hands: Kennedy's race to get to the moon. Computing technology then had a target to aim at. That target was to build a device which was capable of processing computing information on a space craft, on the moon rocket, which had high reliability, small physical size and did not consume much power. We also had to try to make sure that the device would encompass all the surrounding activity. In other words, if you can reduce the interconnections round your central processing device you can start to bring up the reliability, because reliability is very much dependent on the interconnections. Prior to integrated components, interconnections were made by human operators soldering joints together. So round about 1970, primarily because of the space race, we had 1,000 devices on a chip.

There were two main reasons why we got to this stage. First, there was the space race; so we had the motivation. The other thing was that at that stage America was not aiming to try to build a microprocessor; it was trying to get all the devices of a computer into a small package and to build up a configuration of packages together — to produce in fact a minicomputer, not a microprocessor. The motivation of the space race is an important point to remember throughout this talk, because I will make the point later that microprocessor technology as it goes through the 1980s will be very much dependent on people's motives for increasing chip density and so on, rather than the ability to do it.

If you look at the graph on the slide on the previous page you will see that once we get to around 1974 two separate graphs start to emerge. The upper graph is regular logic and the lower graph is more complex logic. The upper graph means that for a particular point in time you have a higher packing density of devices. The reason for this is due to the number of interconnections. Regular logic is typified by memory devices which have a fairly regular pattern of cells. Complex logic, typified by the pocket calculator chip or the 8080 microprocessor, has a complex interconnection structure. You pay for complexity, and the way you pay is in silicon area.

The UART (Universal Asynchronous Receiver/Transmitter) is a communication chip which transfers power from the microprocessor or computer out to a serial data bus. So you have two trends. You have a memory trend or regular logic trend which will be much more device intensive than the complex logic trend. These areas are important in considering where we go from here with micros.

One device that is not on that chart is the 4040 from Intel, which was really the beginning of microprocessors. Around 1970 you had a 4-bit microprocessor which really provided the first programmable function on a chip. If you extend the complex logic line you get into a lot of different areas such as microprocessors, TV games, floppy disc controllers and so on. The complexity there is interconnection.

If we look at this trend for complexity, we can see that the trend is for more complex devices on chips and, furthermore, the ability to bring the functions together in many ways. The Intel 8080 has 76 different instructions. If you combine all these instructions you can program-in 2^{608} different ways. So the complexity is important, but the complexity combined with the combination of devices and instruction types is even more important. What we are getting is the ability to combine these instructions into building blocks and combine the different devices into building blocks for system development.

Let us try to predict the trend. If you look at the regular logic line, you get more and more devices per chip, and bigger and bigger memories. Everyone has heard a lot of talk about 64 kilobyte RAMs and even higher packing densities than that. So it is possible to predict where this line will go, and it is all to do with number of bits per chip. Round about the early 1970s you could start to predict where programmable logic would go. You could say that since we had a 4-bit device around 1970, we really ought to make that an 8-bit device eventually, and then 16-bit devices — we ought to be able to increase the number of bits on a microprocessor from 4 to 8 to 16 to 32, and we should be able to increase the number of instructions. Furthermore, we should be able to increase the addressing capability of the chip, and hopefully the speed. In the early 1970s you could have predicted where microprocessors would eventually go, although you could not predict when and in what way.

These four areas, that is increase in speed, increase in the number of instructions, increase in the memory capability and increase in the number of bits, have all been achieved in the last decade — but increases in speed have been a little bit less than the others. So you now have 16-bit devices, typified by, first, the Texas Instruments' 9900, which came out in 1974 or thereabouts, and then Intel's 8086, Zilog's Z8000, and now Motorola's 68000. Memory addressing for these devices is in megabytes rather than kilobytes. Now we have single chip micros with a lot of these other functions already on the chip; you have ROM and RAM on the chip. If you take a device like the Intel 8022, you also have analogue to digital conversion on the chip. You are starting to get more and more functions inside the silicon.

An interesting side point about these devices is that if you look at the Zilog Z8000, although it has less packing density than the 64k RAM, it has about 17,500 instructions. It is an extremely complex device and the interconnection of these instructions is complex. The 17,000 active components on a Zilog Z8000 happens, by sheer chance, to be the same number of valves that the old ENIAC computer started with.

The evolution of integrated circuits

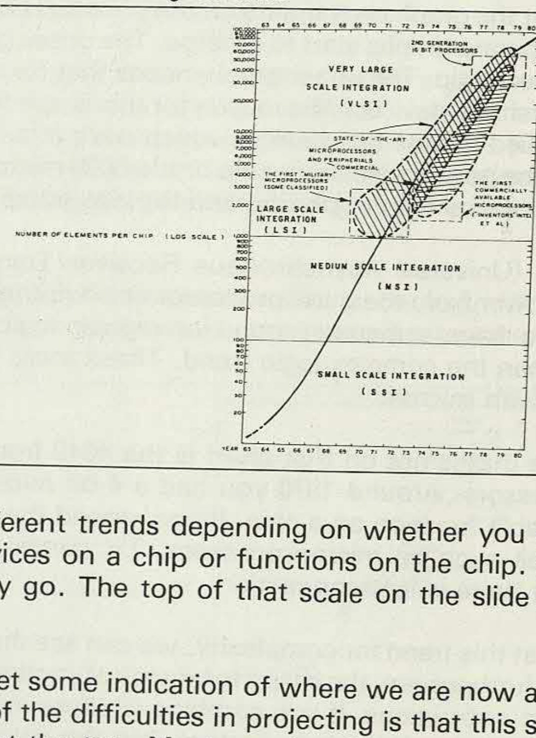

I have put some of the information from the previous slide in a slightly different form. We have gone through the different categorisation of packing densities, from small scale integration with one or two devices up to 100 devices on a package, through medium scale integration, right up to large scale integration and very large scale integration, and into grand scale integration after that. I have marked the areas where these devices were used. There were military microprocessors round about 1971/72; commercial microprocessors coming out about the same time. Then you have different trends depending on whether you are talking about state of the art with respect to devices on a chip or functions on the chip. So we are beginning to see where we will eventually go. The top of that scale on the slide is 100,000 components.

What I should like to do now is to try to get some indication of where we are now and predict where we will go in the near future. One of the difficulties in projecting is that this scale is not linear, it is logarithmic; and it tapers off at the top. Most people who want to use previous information and predict into the future like to put a ruler through the graph and project upwards. We cannot do that in this case. So let us look at a different method of projecting into the future.

On the next slide on the following page we have the cost of processing functions on a chip. You have seen this slide in many forms in other presentations, but in each case device cost is coming down. The first line shows a combination of chips and the cost of that combination; and the second line shows how the price of a single chip configuration is coming down. You will see that the cost of the devices is getting to a stage where we can predict that it will cost nothing in the near future. But that is not really a good indication of the usage of the devices, because the testing and the development labour will be more critical and more fundamental to the cost of the systems than the actual cost of the devices themselves.

So let us try something else and look at device sales. There are a fair number of indicators here. I will convert it for you so that you do not need to double up. We come back to this doubling up situation that we heard about earlier in David Seabrook's presentation. You can see that the trend in each case — for 4-bit, 8-bit and 16-bit devices — is for the number of sales to grow to about twice the number of the previous year. That makes it look as though we have a nice indication of the use of these devices in the future, by projecting all these things linearly, by doubling.

Cost of processing function no longer significant

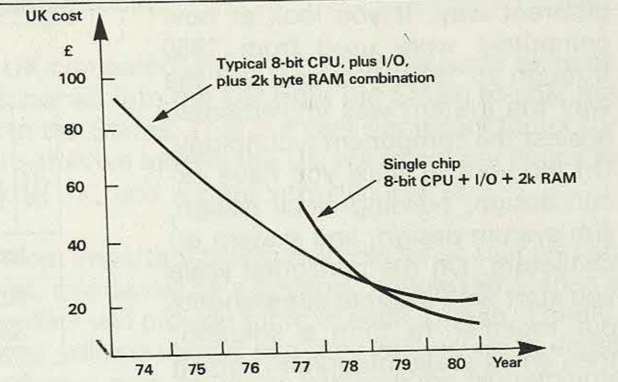

But let us look a little more cautiously at each of these devices starting with 16-bit devices. We have only three years' information to go on, so anyone who predicts future trends on the basis of such sparse information knows that they might just as well not bother, because there is not enough information.

Worldwide microprocessor and single chip microcomputer sales (unit in millions)

	1973	1974	1975	1976	1977	1978	1979	1980	1981	1982
4 Bit										
Annual	0.1	0.3	1.0	2.0	4.1	8.0	15.0	34.4	95.4	207.0
Cumulative	0.1	0.4	1.4	3.4	7.5	16.5	31.5	65.9	161.3	368.3
8 Bit										
Annual	0	0.01	0.6	1.5	3.0	6.0	8.0	13.0	20.0	35.0
Cumulative	0	0.01	0.6	2.1	5.1	11.1	19.1	32.1	52.1	87.1
16 Bit										
Annual	0	0	0	0	0.2	0.3	0.8	1.6	3.6	6.0
Cumulative	0	0	0	0	0.2	0.5	1.3	2.9	6.5	12.5
Total										
Annual	0.1	0.31	1.6	3.5	7.3	15.3	23.8	49.0	119.0	248.0
Cumulative	0.1	0.41	2.0	5.5	12.8	28.1	51.9	100.9	219.9	467.9

If we look at 8-bit devices, we are starting to see a growth which less than doubles every year. It is tapering off. It is still growing a lot, but it is not growing at the same rate. You would think that 8-bit devices would be an indicator for other devices — by taking where we have gone from 1975 to 1980 and seeing whether that is relevant also for 16-bit devices. If you look at the 8-bit devices in relation to 16-bit devices, they are starting to taper off because 16-bit devices have been introduced with more processing capability. So the 16-bit devices are starting to take away some of the market from 8-bit devices. The question is whether or not these 8-bit devices will taper off.

Let us go back a stage further and look at 4-bit devices. The trend goes the opposite way: they are increasing at more than twice the rate every year. So you cannot predict the ratio of 16-bit devices to 8-bit devices by looking at the ratio of 8 to 4. Why are 4-bit devices increasing when you would think that they are old hat and people ought not to be using them? Four-bit devices were used initially for computation when you did not have anything else around at that price and size. They were also used as a basis for controllers, for games and little devices that get information, for example, thermal information on a car or the speed of a wheel or something like that. They are devices which are quite sufficient for doing that sort of function, and they are very cheap.

I know that some of the audience is from the car industry and you know that the cost of every fuse is important. So if you can perform the function with a 4-bit device, you will continue to use that and there is no need to go to higher architecture.

What about 8-bit devices? It is fair to say that these will not grow at the rate of the 16-bit devices, but they will not die off. You will get quasi 8-bit devices where you have 16-bit processing capability inside the chip, but an 8-bit data bus which makes them compatible with previous systems. So we cannot use that table as a prediction of trends.

Let us try to examine the trends in a different way. If you look at how computers were used from 1950 through to 1978, you can plot the way the system was implemented against the component technology. On the vertical scale you have circuit design, building block design, sub-system design, and system architecture. On the horizontal scale you start with discrete components put together to form a solution; then small scale integration, going through medium scale to large scale. The trend has been for the traditional system suppliers' share of the total design and manufacturing activity to become less. The system suppliers' share is being eaten into by the microprocessor manufacturer. As the microprocessor manufacturer has put more and more devices on a chip, he has performed functions that previously had been performed by someone who glued things together with transistors and resistors. So you get to the stage where the device manufacturer is starting to take over a lot of the systems engineering functions that systems houses are currently doing; and the systems houses are concerned with systems design rather than circuit development.

The semiconductor device manufacturers are becoming the small-system manufacturers

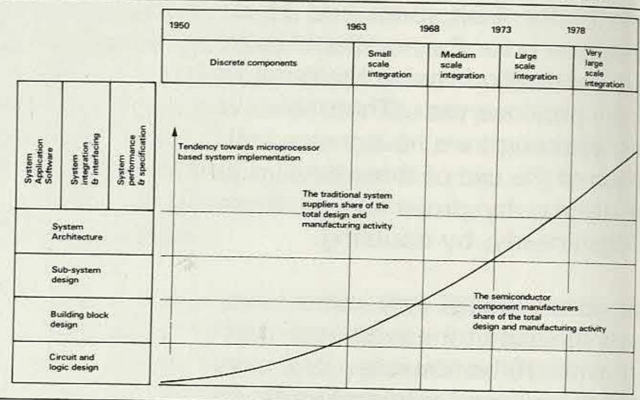

We can extrapolate this curve into the future, and if we go off the screen you will see that the integrated circuit manufacturers will take over systems by 1985, according to the graph. But I do not believe that will happen for two reasons. One reason is because of the law of diminishing returns — it is just not worth their effort. Secondly — and this is very important — as you get nearer the problem-solving area, the activity becomes very applications dependent, and systems suppliers do not have that expertise. So unless they can identify a market where they can put all the functions on a chip, or a significant amount of functions on the chip, and unless this market is broad enough, they will not get into that area. The main problem will be with the degree of applications knowledge of the systems houses.

So where does the UK stand in comparison with America? This chart is for 1977 but it is a good indication of where we are now because today the differences are worse. On the lefthand side we show the annual production of electronic equipment production. In the middle is the annual production of integrated circuits, and on the right the number of installed computers costing over £15,000. The lefthand columns show that the annual production of electronic equipment in America was 10 times that of the UK in 1977. For the righthand columns, the ratio is 13:1. But if you look at the middle columns — integrated circuit production — it is round about 20:1.

Value of electronic production – USA and Europe

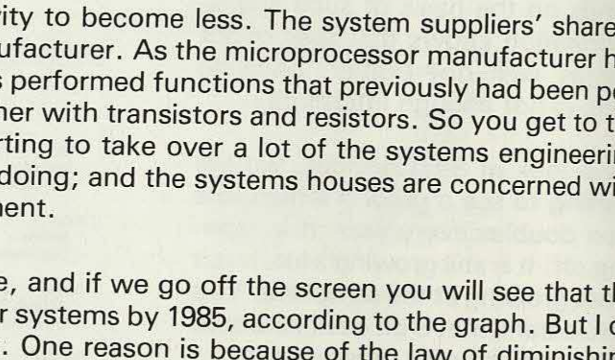

This brings me back to the point that I made previously: a large percentage of the systems sales value will be from the integrated circuit manufacturers. So, in the UK, we are not doing too well so far as that is concerned. One thing I should point out is that whereas the righthand columns show the number of installed computers costing over £15,000 in 1977, we now ought to be talking about £5,000 and above for the same power today.

Another problem is the cost of systems in the UK compared with the cost of the same systems in the States. It is very common for systems imported into the UK from the States to cost the same amount in pounds as they cost in dollars in the States. That is a very significant factor for the future, because the number of installed units that we have in the UK will affect our ability to expand the knowledge base of computing in the UK, and we are very bad at this.

At this stage let me just underline some of the most important points made so far so that we can look at the next stage with respect to them. First, micros will be a systems module rather than a component. Secondly, the semiconductor suppliers will provide more of each system. Thirdly, present systems suppliers will be chasing a lower volume market and they will require in-depth applications knowledge, because the semiconductor manufacturers will take the high-volume part of it. Lastly, systems will be much more complex and more will be expected of them as the building blocks become more complex.

The market size

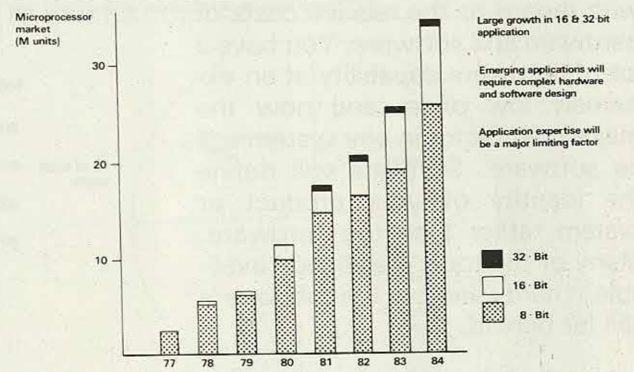

I should now like to look at the market for different microprocessor types. First, let us look at the traditional computer market, normally programmed by users or systems houses. On the slide you can see the growth for 8-bit, 16-bit and 32-bit applications. Obviously, there is no significant usage of 32-bit devices yet, but those devices will start to have an impact from this year onwards. There have been 16-bit devices since about 1977, but the dominant area has been 8-bit devices. But the growth in the future will be in applications for 16-bit and 32-bit devices. The limit to the expansion of these areas is the ability to understand specific applications. The limit will not be the technology — it will be the applications.

An indication of just how critical this applications area is to the industry was shown last May when I attended a meeting arranged by Intel for major consultants in the UK. Intel was very concerned at that meeting about the ability to expand into many of the applications areas, because they needed a growing facility between the applications and the devices. That facility is very much dependent on individual knowledge of applications areas. It is an area which is not really worth a manufacturer's effort to get into, but he needs to do this to increase his sales. So they are very concerned about applications.

Let us look a little bit more at trends in microcomputer applications. Round about 1975, very early days so far as micros are concerned (although, even today, they are still extremely immature devices) applications were fairly basic, and included CRT terminals, cash registers, scales and electronic instruments. The system requirements of the devices were fairly limited.

Today, in the early 1980s, the applications require more complex hardware and software configurations for word processors, retail business systems and communications processors. The systems requirements of today's applications are extensive and include network communication, high level languages, peripheral controllers and different types of memory.

In the future, VLSI will find its way into many complex applications. In business data processing, it will go into the small business computers area, database management systems and graphics display terminals. Office equipment will include word processing systems, microfilm systems, facsimile equipment and teleconferencing products. It will also be used in financial systems such as automatic tellers, electronic funds networks and point-of-sale systems.

VLSI will also be used in data communication and telephone networks. We have covered some of these areas already in the conference. It will also be used in process control applications, and I will come back to this later in the talk. Engineering and scientific applications will also use VLSI technology. The spread is phenomenal, but it is applications dependent.

By 1985 microprocessor applications will be even more widespread. We have already talked about electronic offices. Perhaps, by 1985, there will be 10 processors per employee on the factory floor — not under the control of the employee, but used in different processes. What these processors will need are things like fault tolerant computing, transparent multiprocessing and fairly sophisticated network architectures. So we will be asking these microprocessors to do a lot more than we are asking at the moment. The systems are capable, but we are asking more from them. So, tomorrow's applications will reach a crisis level in system complexity. If we cannot provide these facilities in the UK, then Germany, Japan and America will produce them.

One of the problems is the trend with regard to the relative costs of hardware and software. You have a lot of hardware capability at an extremely low price, and now the major cost factor in any system will be software. Software will define the identity of your product or system rather than the hardware. Many of the tools are already available in hardware, but the software is still far behind.

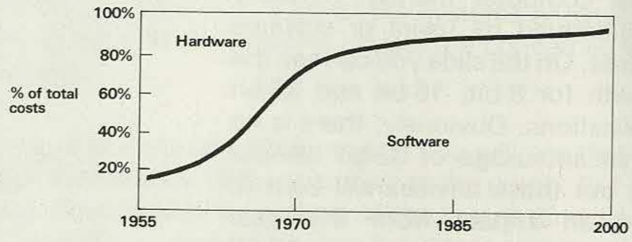

Let us look in a bit more detail at three different applications areas—byte handling, computation, and low-cost control. Low-cost control in the process control area is handled by 4-bit and 8-bit devices and it is a fairly significant part of the market. Byte handling, word processors and so on, account for the same percentage of the market. Computation, which mainly uses 16-bit processors, is a small percentage of the market, but that percentage will grow.

The three major market segments for microprocessor components

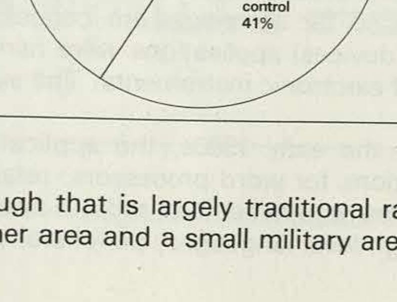

Let us look (slide on following page) at the 4-bit and 8-bit low-cost control market. As you would expect, it is mainly industrial applications (40%). Telecommunications applications (25%) are important, and you also have some computing usage (20%), although that is largely traditional rather than being the best way to do it. There is a small consumer area and a small military area.

The second slide on this page shows the 8-bit byte handling market. Once again, this market is dominated by industrial applications, telecommunications applications and computer applications.

The third slide on this page shows the breakdown of the 16-bit computation market (there is a mistake in this slide: consumer and computer should be transposed). 16-bit devices are mainly used in the computing area. That pie diagram represents only 18% of the total usage of micros. I said before that there is only two years' history of 16-bit devices, so you take the figures shown with a pinch of salt, because it will change dramatically.

The 4-Bit and 8-Bit low-cost control market (41% of world total)

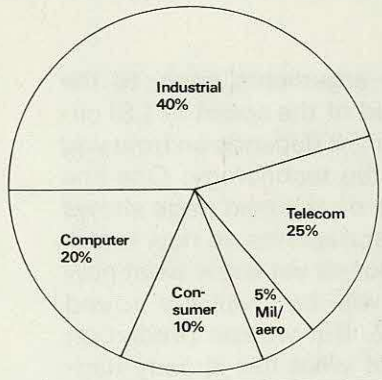

I want now to look into future trends. The slide on the following page shows the trend in memory bit-density, and it shows magnetic bubbles starting to be the dominant area. Moving head disc drives are important and will be significant in micros, but we will see different trends within different technologies.

The 8-bit byte-handling market (41% of world total)

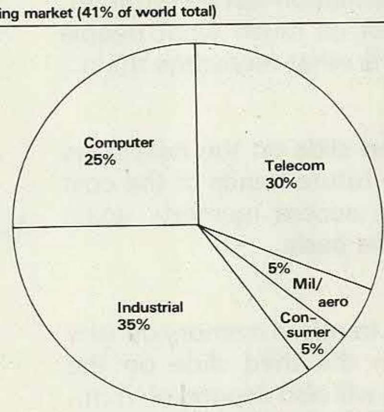

If we look (second slide on following page) at dynamic RAM trends, throughout the 1980s we will get more and more devices on a chip. This is what people predict, and this in fact will happen. The motivation there is to provide cheap memory.

But if we look (third slide on following page) even further into the future, we have to change our whole attitude to prediction, because what actually happens will depend not only on what people can do with technology and particular devices. It will also depend on what motivates them to do it. So there are three lines on this prediction and it goes right on to the year 2000.

The 16-bit computation market (18% of world market)

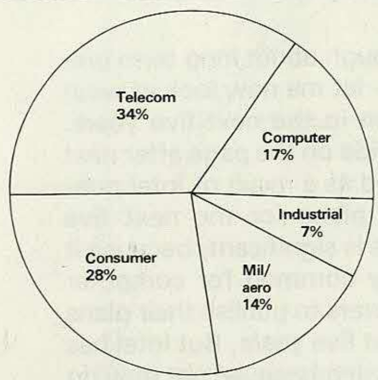

The middle line is the best prediction of what will happen from present knowledge. But there are two bands on either side of the best prediction. One band depends on the industry having the motivation to reduce the cost of the packaging. The other band, the top one, assumes that there is no such motivation. If the suppliers muddle through, they will go along the top trend. If there is something like another space race or, heaven forbid, a war, then they will go along the bottom trend. These are the factors that will affect how dense memories will be, how much it will cost, and the capabilities of particular devices.

Trends in memory-bit density show magnetic bubbles as today's densest devices

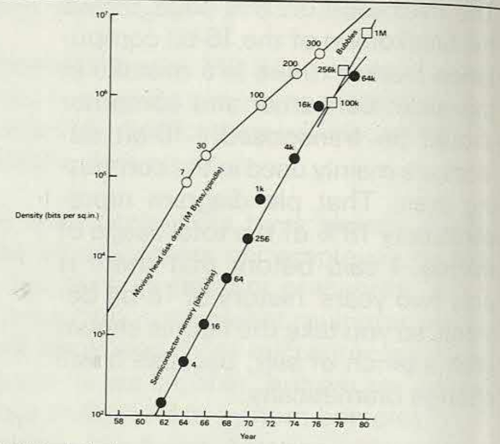

(source: National Semiconductor)

The same arguments apply to the future trend of the speed of LSI circuits. Again, it depends on how you approach the technology. One line of the slide on the next page shows a trend that depends on new materials. We do not yet know what new materials will be available round about 1992. But we can predict on the basis of what has already happened in the last decade that there is likely to be a breakthrough where the materials will cause an increase in speed. But if we do not do anything the situation will be different. So it is not so much what people can do, it is what motivates them.

Dynamic RAM Bit Density for State-of-the-Art RAM's

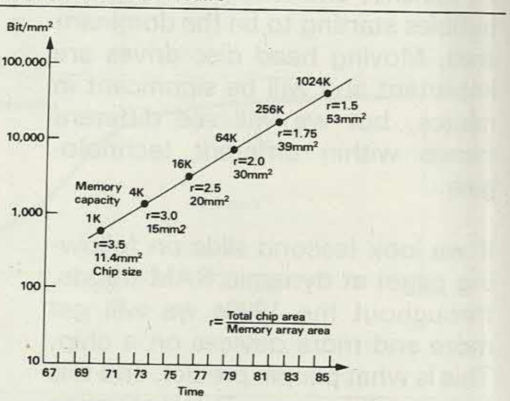

The second slide on the next page shows the future trends of the cost of random access memory, again on the same basis.

Long term trends in memory density (shown on the third slide on the next page) will also depend on motivation. Referring back to the space race in the 1960s, the technology would not have advanced at the pace it did if there had not been the space race.

Expected cost per packaged gate in LSI circuits

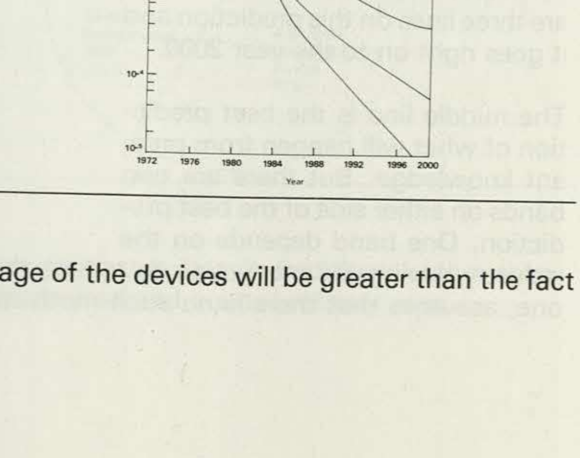

That's enough about long term predictions — let me now look at what will happen in the next five years. The first slide on the page after next is produced as a result of Intel publishing its plans for the next five years. This is significant, because it is not very common for computer manufacturers to publish their plans for the next five years. But Intel has taken this step because not only do they feel that it is important to identify where they are going, but also because the problems associated with the usage of the devices will be greater than the fact that they can actually build them.

The chart starts in 1974 with the 8080, and that device is the basis for all the other families. At each arrowhead we have the processor of the family, and behind that are the supporting devices. So we have the 8080 with an 82XX and 825X, which are devices that have not been announced yet. So the 8080 will not die off, it will expand in the near future, but not for that long.

The 8085 has similar functions to the 8080 but it is a faster device. The 8088 is a 16-bit microprocessor with an 8-bit data highway. Then there is the 8086 which is a 16-bit device. With the 8086 you have the 8087 and 8089, which are an arithmetic processor and I/O processor respectively. Towards the bottom of the slide you have single chip microprocessors with less capability. I should like you to consider these three devices on the righthand side: the micro-mini, the micro-maxi, and the micro-mainframe. There is some information on the second slide on the next page about them.

What I have done on that slide is to relate the three devices to a standard microcomputer and a microcontroller. The microcomputer is the 8088 and the microcontroller is the 8022. As the functionality increases the number of bits per device increases, and obviously the price increases also. The performance relative to the microcontroller shows you that the micro-mainframe has a CPU performance between 20 and 70 times greater, and the I/O ability is 6 to 45 times greater, depending on how you use it.

If we look at the memory addressing capability, whereas there is 1k addressing on a microcontroller there is between 256k and 8 megabytes of addressing on the micro-mainframe. Memory management types and how they are implemented are shown on the righthand side. The important thing to remember is that the first member of that family, the micro-mini, will be available this year. So you are now beginning to get fairly complex and significant capability on silicon.

Speed per gate at constant power in LSI circuits

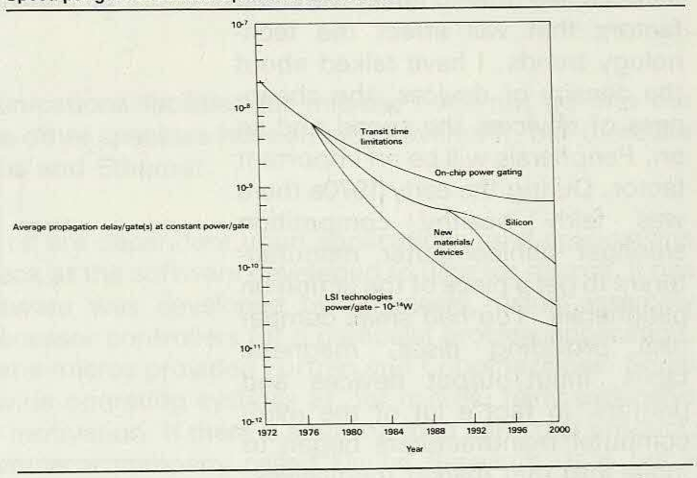

Forecast of cost per bit for random access memory (RAM)

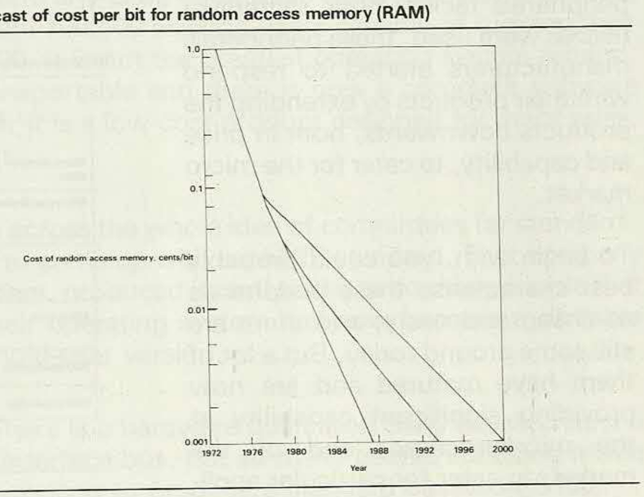

Forecast of developments in microcomputer memory density

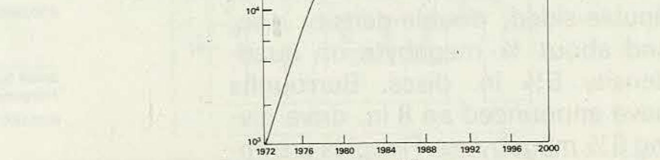

On the third slide on this page I have a comparison of performance for the same price between 1978 and 1985.

I should like now to cover the other factors that will affect the technology trends. I have talked about the density of devices, the cheapness of devices, the speed and so on. Peripherals will be an important factor. During the early 1970s there was fairly healthy competition amongst minicomputer manufacturers to get a piece of the action on peripherals. You had small companies providing discs, magnetic tapes, input/output devices and printers. In fact a lot of the minicomputer manufacturers began to move into that market themselves. DEC, for example, has a fairly large peripherals facility now. Wherever micros were used, these peripherals manufacturers started to respond with their products by extending the products downwards, both in price and capability, to cater for the micro market.

To begin with, you could probably best characterise these peripherals as cheap and nasty, and there are still some around today. But a lot of them have matured and are now providing significant capability at the microprocessor end of the market, to cater for particular applications. There are people like Qume and Xerox producing printers, primarily for the word processing market; but now DEC, Data Products, and Racal are starting to impinge on that market.

Traditionally the storage on a micro has been via floppy discs. The current capacity of a floppy disc is about 1 megabyte on an 8 in. double-sided, double-density disc, and about ¾ megabyte on quad-density 5¼ in. discs. Burroughs have announced an 8 in. drive giving 6½ megabytes. Floppies are important, but, for micros, Winchester discs will be really important. They will be very significant because of their reliability, their cost, and the packing density. You are now talking about between 8 and 60 megabytes on a Winchester, which costs between

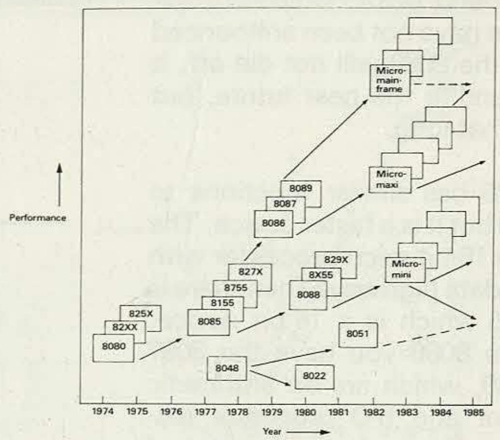

Microprocessor functional classes

Micro system class	Level of functionality	Typical price	Performance relative to microcontroller		Memory	
			CPU	I/O	Typical size (bytes)	Management
Micromainframe	32 bit	$400-$3,600	20-70	6-45	256K-8M	dynamic addressing, segmented or paged; adaptive virtual support
Micromaxi	16 or 32 bit	S100-S500	25	12	128K-1M	structured addressing, segmented or paged; virtual support
Micromini	16 bit	S20-S150	8-10	6	32K-256K	static addressing, segmented or paged
Microcomputer	8 or 16 bit	S10-S50	1-5	3-5	4K-64K	segmented or direct
Microcontroller	8 bit	S2-S20	1	1	1K-2K	direct or absolute

System evolution, 1978-1985
At a constant price, performance greatly increases

	1978	1985
Home computer $1000	16 KB RAM Cassette CRT	128 KB RAM 256 KB Floppy disk CRT Thermal printer
Personal computer $3000	32 KB RAM 80 KB Minifloppy CRT Thermal printer	512 KB RAM 1 MB Floppy disk CRT Impact printer
Work station $10,000	64 KB RAM 0.5 MB Floppy disk CRT Thermal printer	1 MB RAM 10 MB Fixed storage 2 MB Floppy disk CRT Impact printer
Small business computer $30,000	128 KB RAM 10 MB Disk CRT Impact printer	4 MB RAM 80 MB Fixed storage 8 MB Floppy disk CRT Impact printer

$1,000 and $5,000. That is obviously significant for the future of storage on small machines. The problem with the Winchester is that it is fixed and you cannot use it for archiving. Another device that has recently been announced is a cassette that can store 35 megabytes in seven minutes. Also, you are now getting removable cartridge Winchesters and 5¼ in. Winchesters, so the technology is not static.

There are also improvements in communications facilities for micros. I will not go into the communications area very much because other speakers have already covered it, but there are obviously things like ring networks, Z-bus and Ethernet.

Limitations to growth of the use of micros are dependent upon applications, and applications are dependent on software. If we look back at the software developed to date for micros, it has been pretty pathetic. Initially, the software was developed by engineers, using assembly language when they developed microprocessor controllers for a particular process application. Later on, Basic was introduced, then some micros provided Fortran and Cobol facilities. Some microprocessor manufacturers now provide operating systems at, for micros, fairly expensive prices. But there again we come back to motivation. If there is a reason for a particular product, someone will provide it. A case in point is a company called Digital Research. It is not a computer manufacturer; it is not linked to Intel, Zilog, or Motorola; but it produces an operating system which will sit on any Intel computer and any Zilog computer, either 16-bit or 8-bit. The Digital Research people estimate that there are 100,000 existing users capable of running their operating system, and it sells for £200 or £300. It is not the greatest operating system around, and anyone can dig holes in it. But it is transportable and there is now a complete software package industry sitting on top of their CPM. It is a low-cost product designed for mass sales, and that is what the micro business is about.

The approach taken by Digital Research cuts across the whole idea of committees for standardising software and hardware, that take years to come up with an ANSI standard. That company found the need for a particular operating system, produced it, and sold it to so many people that people now have to pay attention to it. Their operating system has now become a *de facto* standard, and they have now produced a multi-user version.

The same thing happened with hardware. There is a hardware bus called S100 which is now a *de facto* standard. Again, it is not the best interface bus, but so many people are using it that the IEEE in America now has a committee working to produce a standard based on S100.

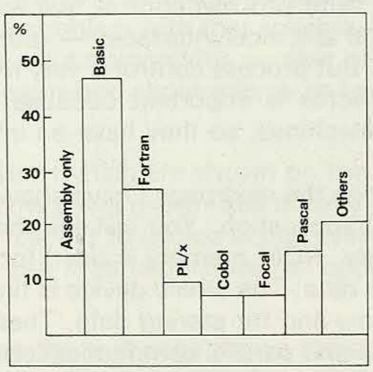

This slide shows the percentage usage of languages on micros. At the moment Basic is the dominant language. Basic will continue, but the language to watch is Pascal.

I have now looked at facilities, at software, and at trends within the hardware structure. I would like now to look at some applications, beginning with data processing. Microprocessor devices are developing at a tremendous rate and there is an indication within data processing circles that perhaps you ought to be considering them. However, there is a problem because most people want to be sure that the systems they purchase are mature. They do not want to burn their fingers on untried systems — they want someone else to use it first. But the incentive for using micros in data processing or, more likely, for connecting data processing installations together, is price.

The trend in the technology is for devices to become obsolete in 2 to 3 years, and that is your first major problem. Wheeled transport existed for thousands of years. Then came the industrialised society, and Watt's steam engine lasted for 102 years. Capital equipment in the 1930s was written off in 30 years. Mainframes today are written off in five to eight years, depending on the installation. Micros will become obsolete in two to three years. An analogy with this is someone who bought a pocket calculator in 1975 with an LED display and four functions, for £20 or thereabouts. Nowadays they probably do not use it because they have to replace the batteries every two or three months. It is not obsolete, it is still perfectly capable, but there is something better on the market, with more functions, with LED displays, and you replace the batteries every 18 months.

The same thing is happening with micros. They are not becoming incapable of being used, but something more cost effective has come onto the market. That is a difficulty for any company which is basing a product on microprocessor development. One of the problems is that you have to get the product onto the market as quickly as possible to make sure that you come within this two to three year period. That problem impacts on the whole development process of the product. If you do not have much time to develop the product, you do not spend too much time on designing for maintainability, or reliability, or testability. It does not happen in every case, but there is a tendency to get the product on the market at all costs.

If you look at data processing, smaller machines will affect the decision making process in data processing. You are all familiar from your trade journals of trends both in smaller machines and increasing capability. You have now to consider smaller machines as part of the procurement exercise. You may decide not to do it, but you really ought to consider whether or not a small machine is capable of meeting your needs. It is your decision whether or not to use a smaller machine.

Data processing people will have to be much more concerned with the technology and the technical abilities of systems rather than with resource management. If the systems cost less, you do not necessarily need to make sure that they are kept at maximum efficiency, 24 hours a day. Technology will change the whole structure of computing because it is becoming cheaper. So decisions will have to be made at a fairly senior level. That is an indication of the way things are moving. People can no longer base their activity on established computing practice.

Let us look at process control. Process control is different from data processing because it is engineering based and is concerned with a more rigid discipline than is traditional with data processing. With process control, we are concerned with physical interfaces — pneumatics, mechanics, or electrical interfaces — rather than primarily with information, although we cover that as well. But process control is very much concerned with real-time tasks. This is where the history of micros is important because they grew out of minicomputers rather than batch processing machines, so they have an inherent ability to deal with real-time systems.

On the slide on the next page I have shown a typical block diagram of a microprocessor used in an industrial application. You will see there is the chip in the centre, connected to ROM and RAM memory. ROM memory is used for your program and RAM memory is used for storing intermediate data. The whole device is timed by a real-time clock. There is cassette storage for your programs and for storing data. There are a number of interfaces for printers, for alarms, and for serial and parallel communication. There is also the facility for analogue to digital conversion at fairly high speed, and you can multiplex your digital inputs.

The implication of microprocessors for process control is that you now perform tasks in software that you previously performed in hardware. One of the problems in process control applications is that the engineers who are processing the data need to learn about software which they have not previously been involved with. If, however, the engineers have a structured approach to hardware, then the transfer to software will not be too bad for them. However, that has not happened in many process control applications. So many people hire computer experts to write the software for them and if these experts do not know enough about the engineering process they will create all sorts of problems. The capacity for error is phenomenal.

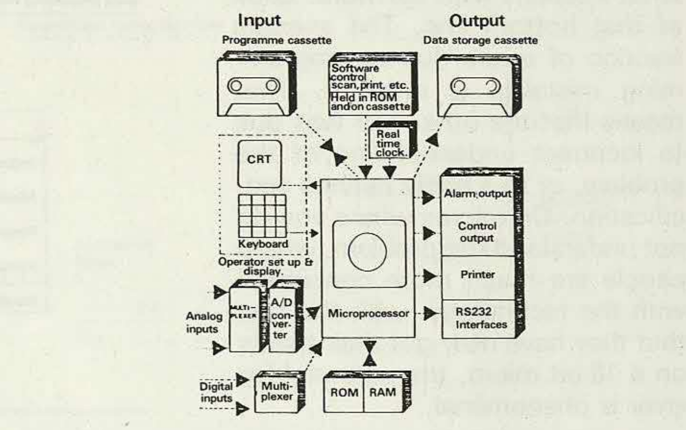

With process control applications you base a fairly critical process on something that is fairly inexpensive. If you have not engineered the device in the same way as you would engineer a much more expensive hardware solution, you can build a device that does not work. A low-cost solution is no longer low cost if it does not work. It is low cost only if it does the job that it was designed to do.

One of the problems with micros is that there are a tremendous number of fairly young engineers in the industry, without any experience of applications. They know a lot about the internal architecture of a Z8000, or an 8086 and could probably discuss the differences much better than I could. They are getting into senior positions and they do not appreciate the problems of the total reliability of the application. The situation is getting worse. You see plenty of job adverts in the technical papers for microprocessor engineers, and if you can spell the word you can get a senior engineer's position. Not only is the company losing out, but the engineer himself is losing out because he is never receiving appropriate training. When I did my training in the early 1960s, I started as a raw graduate in a fairly big industrial organisation. I made lots of mistakes, as we all do, but most of my mistakes were trapped by someone more senior than I was, who knew what I was doing, who understood what the problems were, and who shielded the customer from my problems. With micros this is not happening. Today, companies employ a micro "expert" at 23 or 24, who makes mistakes and probably leaves the company before they are found out. I know that this is a jaundiced attitude, but it is a very significant problem with respect to companies using microprocessors for the first time. If you do not understand what the experts are talking about, do not assume that they are correct.

One way of overcoming this problem is to get help in the assessment of new applications. The Department of Industry are helping here by funding the first £2,000 of a feasibility study for most microprocessor applications. But there is another problem with that scheme, because the DOI has not vetted its consultants very well, and there are a tremendous number of people who will be quite happy to take the money off you. I am pessimistic about micros so I am not going to give you a happy picture of the situation.

If you look at the trends in the data processing industry (which are shown on the slide on the next page) you will see that only programmer productivity has grown less quickly than system reliability. One thing that should be pointed out with respect to micros is that system reliability may well affect the hardware and systems software, but it certainly does not apply if you put applications software into it.

Looking now at hardware reliability, the second slide on the next page shows some phenomenal figures for mean time between failure. Micros are small, cheap, consume very little power and they are highly reliable. People assume, therefore, that they must be able to work in their application without any problems in the future. Without doubt, the hardware itself is highly reliable. But if you put the software in the hands of someone who does not know what he is doing, or does not know your application, you can forget all about reliability.

My third slide on this page illustrates the software reliability problem. This is not specific just to micros, but it is very relevant for micros. People may disagree with the figures on the slide, but they are fairly indicative of what happens with software. Look at that bottom line. The average fraction of errors due to programming mistakes is one-half. That means that the other half was due to incorrect understanding of the problem, or to a badly defined specification. On micros where you do not understand the problem, where people are much more concerned with the technology, with the fact that they have now got their hands on a 16-bit micro, the potential for error is phenomenal.

Data processing industry growth trends

	1955	1965	1975	1985
Industry	1	20	80	320
Machine performance	1	10^2	10^4	10^6
Programmer productivity	1	2.4	5.6	13.3
System reliability	1	5	24	120
Source: Art Benjamin Associates Ltd.				

The first slide on the next page shows the expenditure during the life cycle of a normal fairly large project. I have split it into four phases. The conception stage is a fairly small percentage of the cost, not because there is not much effort there but because you are not using processor time and you are not using a lot of facilities that you have to use later on.

Failure rates and mean time between failure Motorola MC6800 microprocessor

Year	Failure rate (percent/1000 hours)	MTBF, hours (years)
1974	1.27	78 000 (9)
1975	0.5	200 000 (23)
1976	0.12	833 000 (95)
1977	0.08	1 700 000 (194)
1978	0.013	7 700 000 (877)
1979	0.006	16 666 666 (1901)

During the development stage you start to use more facilities and it starts to cost a lot more. During the production stage you use a lot more facilities, and even more during the operation and support stage.

Software reliability parameters

Parameter	Characteristic Value
Average number of initial program errors	1 per 100 lines of source
Average manpower to identify each error (after failure)	2 man-hours
Average manpower to correct an error	6 man-hours
Average computer time to fix an error	2 hours
Average number of errors per new system software release	1,000
Average fraction of errors due to programming mistakes	½

If you look at the second slide on the next page you can see that the impact of the conceptual stage on the total life-cycle cost of the project is much higher than anything else. So you make decisions during that first stage which will lock-in your life cycle costs. You may be able to increase the costs but there is not much chance of being able to reduce them at subsequent stages. This comes back to the problem of micros. If you are too hasty in developing a solution to the problem, you will miss out on this first stage; you will not spend enough time in defining how you are actually going to do it.

I have an example on the third slide on this page of the different time scales, depending on whether you jump into coding your program or whether you design it before you code it. If you spend a fair amount of time designing the program before you code it, you will end up with less time being taken up by the whole project. But if you jump in without designing the program, it will cost you a lot more time to debug it and to implement it.

The first slide on the next page gives a graphical indication of the two different approaches.

How does this affect micro development? The second slide on the next page is a fairly simple graph representing the reliability of most software systems, but it can also apply to failure rates on micros. You eventually get the system debugged. Then you have to change it. Coping with change is one of the problems with micros. If you do not think about the changes first and about how you are going to bring out version 2 or version 3, it will cost a lot. It will also cost you a lot of failures. The point indicated by "A" on the slide is where the system was scrapped and transferred to a new machine, and the problems have started again.

I should like to finish by reviewing the trends in the single chip computer market between 1979 and 1984 (shown on the third slide on the next page), and relate these to a formidable force that is entering the field. The first three slides on the page after next show the dollar value, the average selling prices and the number of units shipped for 4-bit, 8-bit and 16-bit single chip microcomputers.

The 16-bit market has changed because it only started recently.

The fourth slide on the page after next shows the important thing, and that is Japanese production of electronic devices. The top line shows total production, and the line below shows imports to Japan. The bottom line shows the Japanese exports. Round about 1984, Japan will be a net exporter of electronic devices, and the States will take 60% of these devices. So the emphasis of semicon-

Expenditures during life cycle

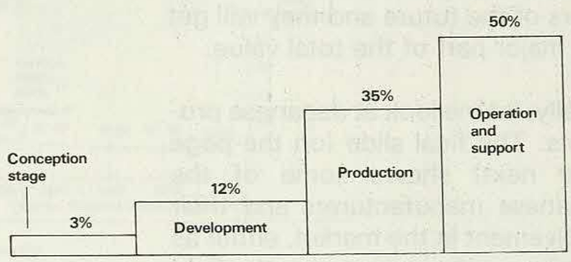

Effect of early decision on life cycle cost

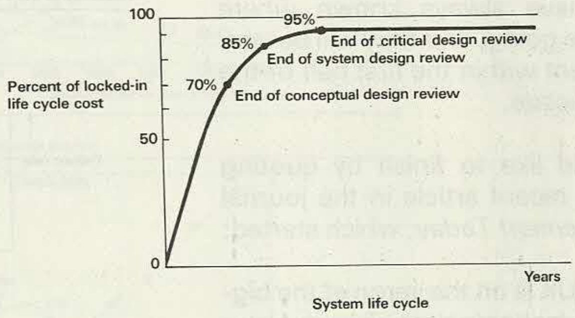

Effort applied in poor and good programs

	"Coded"	"Designed"
Program design	5%	35%
Program implementation	45%	10%
Debugging	50%	35%
Subtotal:	100%	80%
Maintenance over life	25%	5%
Modification over life	50%	10%
Total:	175%	95%

ductor manufacture is changing from the States being a net exporter to becoming a net importer.

Remember what I said at the start — that the people who manufacture devices will be the systems suppliers of the future and they will get the major part of the total value.

Finally, let me look at Japanese products. The final slide (on the page after next) shows some of the Japanese manufacturers and their involvement in the market, either as a major vendor, or entering the field with a strong capability. As you can see, the Japanese manufacturers cover a wide variety of products. That is the situation with the dominant country in the business. They have better CAD facilities than the States, and they now have set their minds on software. The Japanese have always known where they are going, and they will be very dominant within the first half of the next decade.

I should like to finish by quoting from a recent article in the journal *Management Today*, which started:

"The UK is on the verge of the biggest technological change since the Industrial Revolution, and it presents all the confidence of an old lady trying to cross the road."

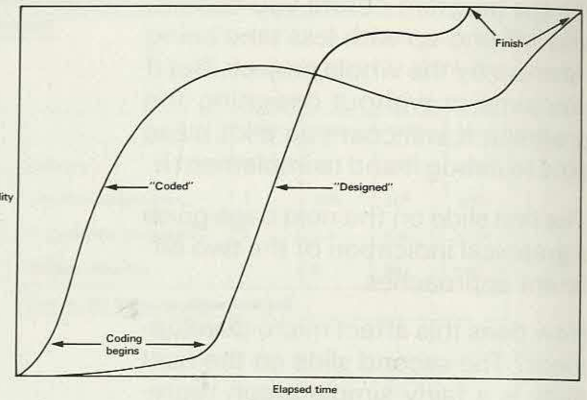

Reliability of most software systems

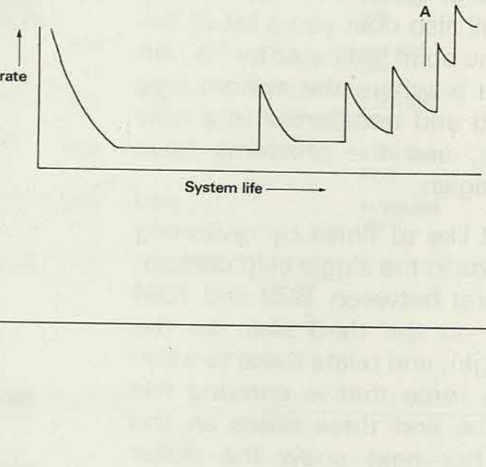

Single chip microcomputer market

Source: General Instrument Microelectronics

Worldwide 4-bit single chip microcomputer market (excludes calculators)

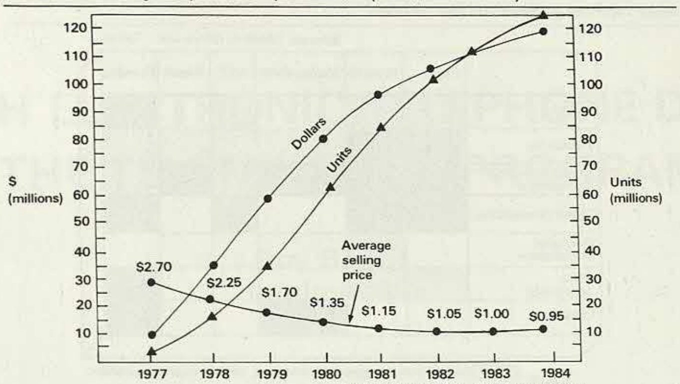

Source: Motorola Semiconductor Group

Worldwide 8-bit single chip microcomputer market

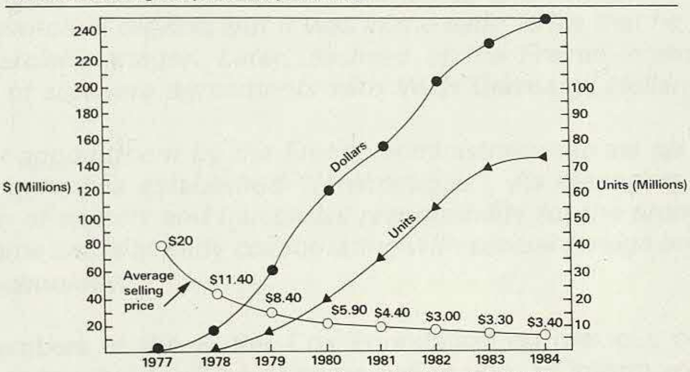

Source: Motorola Semiconductor Group

Worldwide 16-bit single chip microcomputer market

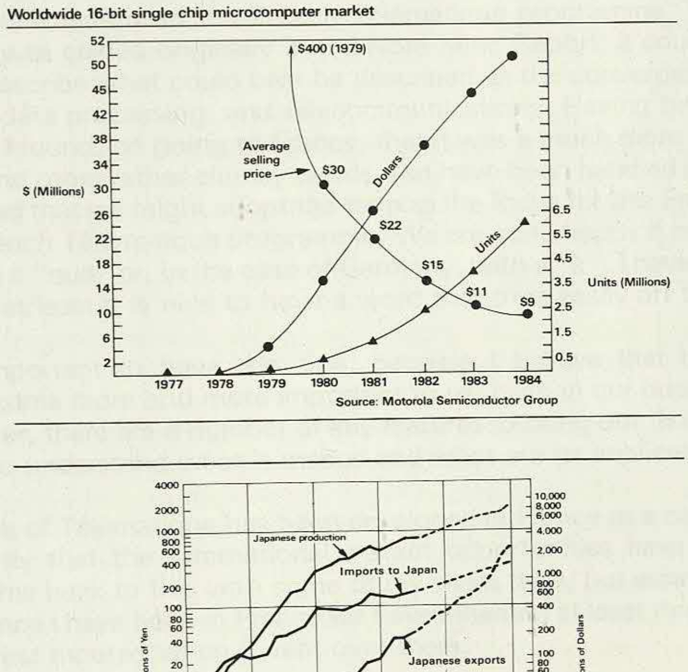

Source: Motorola Semiconductor Group

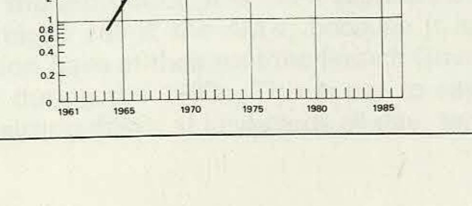

Major Japanese suppliers and their products

Source: Most & More Inc., Tokyo

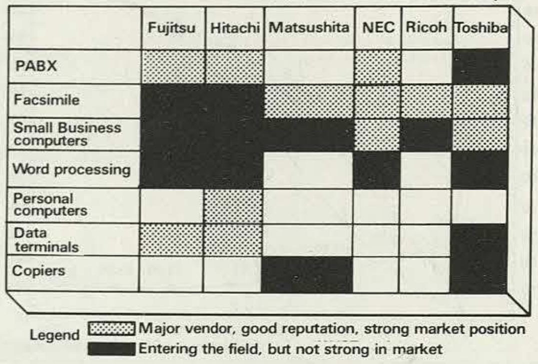

SESSION H

THE FRENCH ELECTRONIC TELEPHONE DIRECTORY AND THE TÉLÉMATIQUE PROGRAMME

Roy Bright,
Intelmatique

Roy Bright's earlier career was in the British Post Office and included the marketing of the experimental packet switched service, but it was in the early 1970s that he became the world's first viewdata commercial manager. Later, as head of the Prestel international division, he negotiated the series of software agreements with West Germany, Holland and elsewhere.

Since his more recent appointment by the French administration to set up a non-profit making commercial subsidiary, he has established "Intelmatique". As managing director, he leads a multi-disciplined team of experts and has global responsibility for the promotion of the French Télématique programme and is already collaborating with several foreign organisations planning to use this French technology.

I have addressed members of the Butler Cox Foundation on previous occasions, wearing a different hat. I hope today that I will be of some use to you, to inform you about the plans in France. The first point is that my mission is not just concerned with videotex, although that naturally is an important feature of the French Télématique programme. The programme is a generic title. Its title was coined originally in the Nora-Minc Report, a couple of years before I went to France, to describe what could best be described as the convergence of informatique, the French term for data processing, and telecommunications. Having been coined, the term tended to lapse. But I found, on going to France, that it was a much more attractive word than "compunications" and other rather clumsy words that have been bandied around over the past year or so. I suggested that we might adopt the term as the focus for this French programme, so we have now the French Télématique programme. We are quite happy if people want to spell it with a "c" instead of a "que" or, in the case of Germany, with a "k". I have asked people to feel free to do that, but at least it is nice to have a word that trips easily off the tongue.

I think that it is important to have this title, because I believe that in the coming years Télématique will become more and more important to us, both in our business lives and in our private lives. However, there are a number of key features to bring out. It is one thing to have a term, it is another to understand what it means and what are its implications.

First, the programme of Télématique has been developed in France as a national programme. It is only more recently that the international market opportunities have been identified and addressed. I will come back to this with some of my slides later, but essentially it is a dynamic programme. Ever since I have been in France we have inherited at least three products to add to those which were first mooted when I went over there.

One way of defining the Télématique programme — with apologies to a well-known American company — is the 3M's. It is multi-product, it is multi-supplier, and it is multi-application. To avoid embarrassing 3M we might call it the 4M's, because it is also multi-national. Those objectives of multiple application have at their root the French Government's ambition to create an information-based society during the 1990s. This is not to say that every home, or every person in every home will be sitting down at keyboards all day, tapping to each other. What is important to recognise is that the Télématique programme as a first objective, and yet a very vital objective, seeks by economies of scale, by mass production, and by the use of the ever-decreasing cost of LSI, to bring the price of these products and services down to a level which initially will penetrate much more into the small and single man business sector, and later even into the home. One important initiative that the French are taking — the electronic directory service — is addressing that goal specifically.

The four factors that I have identified on this first slide are an attempt to encapsulate all the objectives and the opportunities that have been described for the Télématique programme. First, a coherent R&D programme ensuring compatible products. Unlike many countries — and to some extent I include the UK in this — there has been a very serious attempt in France to coordinate the R&D work for the Télématique programme. If you look at the most obvious and striking examples of the broadcast teletext systems and the interactive videotex systems, known fondly as Antiope and Teletel respectively, they have their origins in a common research and development centre, in Rennes. Both the telecommunications authority and TDF, the French broadcast diffusion authority, jointly fund that research centre.

This coordinated R&D also takes advantage of the favourable trends of the price of technology. I suggest that, without that, many of the things that we have discussed in the past two days would not have been possible; and naturally that applies equally to the Télématique programme.

The next factor is a modular design philosophy to maximise adaptability. That is a bit of a buzz-word phrase, but what it means is that each product has been developed essentially to create the maximum opportunity, or synergy, to which these products can lend themselves in various applications. When we look at the electronic directory you will find that the technical compatibility between it and Teletel is a striking example of that goal.

The fourth factor is a commitment to achieve harmonisation with foreign needs. Those of you who follow closely the videotex discussions in CCITT might smirk a little at that, but the fact remains that a key objective in France is to create ultimately a world standard, but certainly a European standard in such matters as videotex, facsimile and so on. It also means that products designed for the French home market will not necessarily be readily adopted — or even suitable — for overseas markets. So there is an attempt here in some of the design options that are being planned into the French national system to maximise the opportunity to relate them more closely to foreign needs.

So the first of the Télématique family is the videotex system development, which is called Teletel in France. There are a number of important, fundamental points of a non-technical nature to make about this development. The first and most striking point is the objective of the French PTT not to become an information database manager. That puts it in a nutshell. Essentially what it means is that the PTT will, following the Velizy trials, back off completely from database management and that activity will be conducted by the industry itself — be it private industry, government agency, or whatever form they represent.

That has important implications, not least of which is that the industry, including the supporting industry, the software houses and the hardward manufacturers, has to be much more versatile and flexible. A centrally-managed videotex system has some strengths. For example, everybody knows which way he is going, and, provided someone gives a direction, it is easier to conform. However, in conforming, you perhaps lose some of the benefits and some of the striking features of the development of that videotex service.

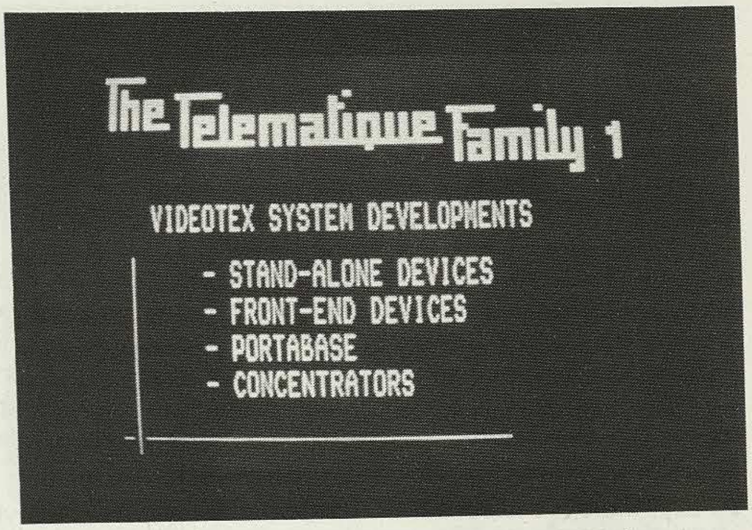

In France, on the other hand, it does mean that there is much greater concern about such matters as the type of index structure each database should use, and the type of security level that should be imposed within the database operation. That is a complication which I am sure many of you, as users of videotex and like systems, will appreciate is a key one, particularly if your ultimate aim is the mass market, and the untrained user.

The French objective does, however, bring some benefits in its wake. It means that if I am in a banking operation and I wish to make use of Teletel, I do not have to concern myself with the fact that it will take X months, or maybe years, before certain levels of security are available for me to exploit in order to allow my banking operations to be free from misuse. Because the responsibility is mine, as the operator of that system I will decide what level of security I require and whether I should invoke some sort of "smart" card.

Similarly, if I am a mail order operator, I do not have to worry about trying to transfer all my inventory, all my database material on to someone else's system. I operate my own system and I decide who I allow on to that system. I decide whether they have to have certain passwords if they are agents or the end user, and I decide how I update and index that system. Maybe a tree structure is not the best way to handle a particular application. So the French approach has strengths and it has its penalties, but one thing it does have is flexibility. I would suggest that, in many cases, the concern expressed about market confusion vis-a-vis different types of indexing structures can be overcome if the operator of the system is allowed to make his own decision. He knows better than any PTT who his customers are and what they can best comprehend when they wish to access his system. So in a way I think that it is putting responsibility where it best belongs.

Another key point that it brings in its train, having referred to non-tree-structured indexes, is that you have to assume that all terminals will have a full alphanumeric keyboard. Otherwise we will be in real trouble if we try to offer some sort of associative search technique and 80% of your users have numeric-only keypads. So again there are certain ramifications of that original decision that the PTT, in principle, should not be the database manager. These ramifications have to be taken into account, and I suggest they are worth noting when you are tempted to ask what is the difference between Teletel and Prestel. There are some fundamental differences which are not technological differences *per se* but are much more concerned with features, attitudes and policies.

To summarise on this slide, firstly I have referred to the greater responsibility on the industry, on hardware suppliers and on software houses. I have also referred to the fact that to serve the growing number of service providers planning to take part in the Velizy trial, which starts in two and a half months' time, there are different types of customer interest. For instance, a stand-alone device typically would be needed by a company which, for whatever reason, had decided to go into Teletel without any existing commitment. Maybe they do not have a computer already, or, if they do, they do not want to disturb that side of their business. Whatever the reason, they want a separate application because of the advent of Teletel. Whereas, as I mentioned earlier with the mail order example, if I have an existing database and I wish to utilise that, then I am interested in a front-end solution. The particular example that I have in mind here would be two of the major mail order companies in France. Both have IBM 370s which carry their present database. So they have been offered a solution by the industry which would involve an IBM Series-1 emulating a 3270 looking into the database, and matching the Teletel standards and protocols looking out to the network.

The Portabase is a trade name of a particular piece of hardware developed in France, but it represents a range of options, which essentially is a mini-videotex system. A number of similar developments have already occurred in the UK, because it is recognised that there are certain applications where a modest-sized, low-cost, stand-alone videotex system could be of use within an organisation, and if that organisation is geographically dispersed it does not matter. The Portabase range is an indication of those types of developments.

Finally, there is the concentrator. You will appreciate by what I have said earlier that if you have a distributed database — which is essentially what Teletel envisages — and an equally widely distributed customer base, the networking becomes a very important feature. In fact, one of the benefits in France that has enabled the PTT to take this rather more transparent role, is the existence of Transpac, which offers distance-independent tariffing. So in France you will have the appearance of local concentrators which will handle the interface between the local population of Teletel users and, if it is over a longer distance, link you in to Transpac, which will then act as a main highway to the distributed databases.

The other form of concentration should not be neglected, because the very nature of the system puts much more onus on the need for a menu — a menu of information or topics that are available somewhere in that widely distributed system. So there will be a need for some form of concentrator, but working more in a commercial sense. Indexing and maybe centralised billing and tariffing and other features can be incorporated in it, as distinct from the communications concentrator to which I referred in the previous example.

To give one angle on the approach to Teletel, this model represents a Thomson CSF device which is based on an existing TV display. Underneath the display you will observe the rather smart plinth on which it stands, which contains as you would expect the decoder, modem and so on. The lady in the photograph is holding the full alphanumeric keypad rather than the numeric only keypad to which I referred earlier.

I should add that since January 1980 the French TV manufacturers have agreed that all sets produced will incorporate a special interface for such things as video cassette recorders, but which includes an external adaptor for videotex systems. That means that anybody buying a TV since January 1980 automatically has an interface built into the set which overcomes one of the criticisms of external adaptors — namely the slightly poorer quality image that one gets if you go in through the antenna.

The second member of the Télématique family, electronic directory systems, generates more interest than any other single product in the Télématique programme. It is worth spending a few moments on it here, then we will come back to it later with some application examples. Essentially it refers to the ability, as you would expect, of any electronic directory system worth its salt to meet the needs of PTTs and carriers. That was the main reason for the French development. The intention is that there will be a supply of low-cost terminals to all telephone subscribers, enabling them to access the electronic directory databases at no cost to themselves so far as the terminal hardware is concerned.

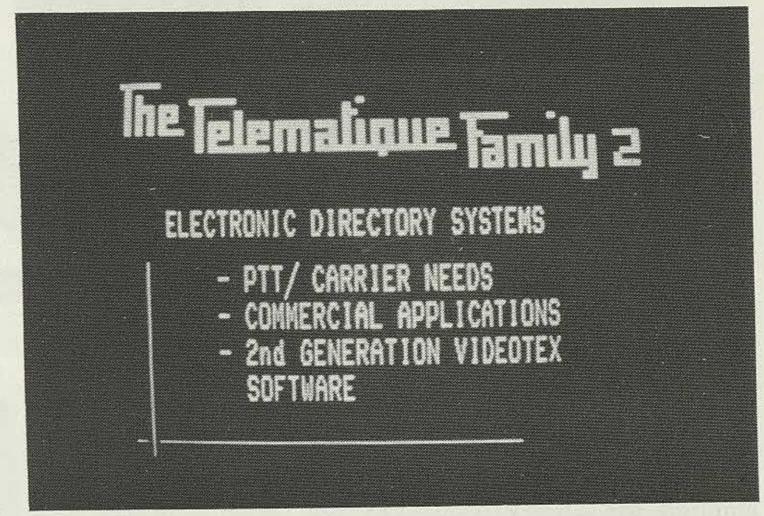

This system has already started a series of trials. One took place last June through to September, with about 50 customers involved, primarily to address the questions of man/machine dialogue — the user friendly approach that is so important if it is to succeed. Later this year, in June, another trial will start with about 1,000 terminals, and this trial will refine some of the lessons learnt from last year. In the last quarter of this year, the ½ million user trial will start. It is scheduled to take place in Ille et Vilaine. In that location the complete database of directory entries for that department of France will be mounted, and every telephone user in that department of France will be equipped with one of the electronic directory terminals.

For non-technical reasons, during the trial the terminals will not be allowed to access the Teletel databases, which by then will also have been established. But I stress that the reasons are non-technical. In other words, there are certain implications vis-a-vis newspaper attitudes to the classified ads/Yellow Page grey area. (This area is the focus of attention in many parts of the world, including the USA.) Nevertheless, I stress that the reasons are non-technical. Technically, those terminals and the network could allow that access to happen very readily.

My next heading — commercial applications — is to remind me to stress that although the system was developed for an electronic directory, there is absolutely no reason why it should not be used in non-directory areas. This is an important point because of all the publicity (and indeed the software), which is designed to emphasise the directory application. However, if you stop to consider the terminal, with a low-cost monochrome screen, with a Teletel decoder built into it — it is equally happy working with other types of information. Less obvious, but equally true, is that the software developed for this system can be exploited for many other types of commercial applications in ways that are not apparent at first sight. For example, the basic fields the software uses for the directory must include initials, name, address details, telephone number, and maybe profession if we are looking at the Yellow Page side. If you think of that as a set of four or five fields and you replace some of those headings with other topics (such as catalogue item for telephone number and item descriptor for name and address details), you can begin to visualise that there is a certain transparency about this software that lends itself to non-directory applications. It is worth considering this as second-generation videotex, because once you start to see some of the demonstrations and the capability of this software, it is very impressive. It uses full alphanumeric, it uses inverted file techniques, and it allows you to have focusing — we have taken some terms from photography like "autofocusing" and "autozoom". If you are looking up a telephone number in a strange town, one of the first things that the system will ask you if you do not have the full name and address details is, "Which district?" "Would you like a list of districts?" Answer: "Yes", and up comes a list of districts. Now you can start searching from that angle if you prefer.

All these techniques and the software to support them have tremendous value in any non-professional area that you like to envisage. Another important example would be in the problems of misspellings and phonetic problems of that nature. Again, the software has been designed to cope with those problems. So I would argue that this is a very useful example of second-generation videotex software.

Some of the hardware for the electronic directory system has been developed. The French PTT placed research and development contracts for the terminals with each of four major manufacturers: Thomson CSF, Telic, TRT (the Philips' subsidiary in France) and Matra. The PTT defined the characteristics of the service in the form of a facility specification and allowed the manufacturers to work out the technical solutions and the packaging. These four slides show an example from each of the manufacturers, although most of them prepared four or five different versions just to show the variety of treatment that they each applied to the problem.

They are very neat, and could look something like the final version that will be used in the trials at Ille et Vilaine. One of them is interesting in the sense that it has an integrated handset. This is another interesting feature, particularly when we look outside of France in terms of the use of these terminals in other countries for non-directory applications.

The next item (see first slide on next page) on the Télématique list is mass facsimile. There is nothing new in facsimile, and nothing particularly new about facsimile over the normal dial-up network. What is interesting, however, is the key word ''low-cost''. Typically, US$500 is the target cost per unit, and the production envisaged to reach that target would be a million units over a ten-year period. As you will appreciate from those numbers, we are not now talking of penetration into every home, but rather penetration

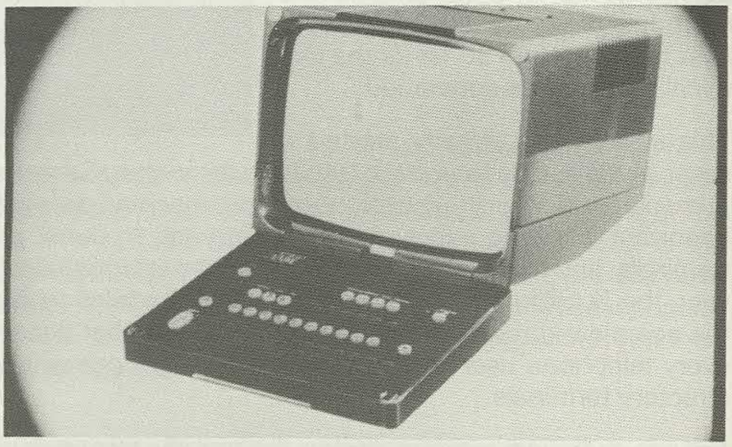

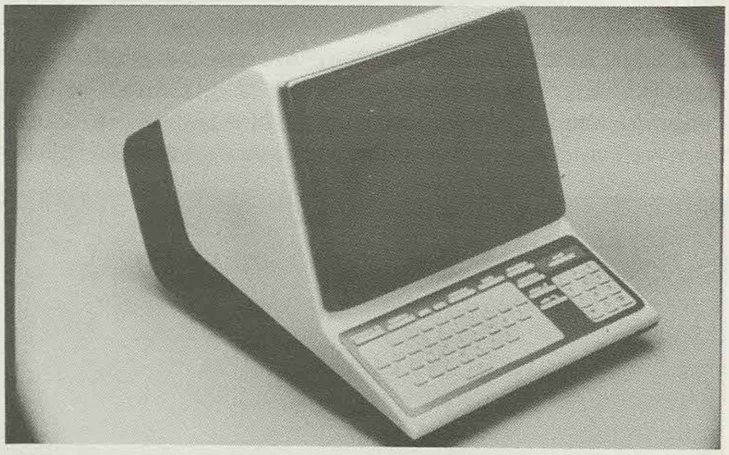

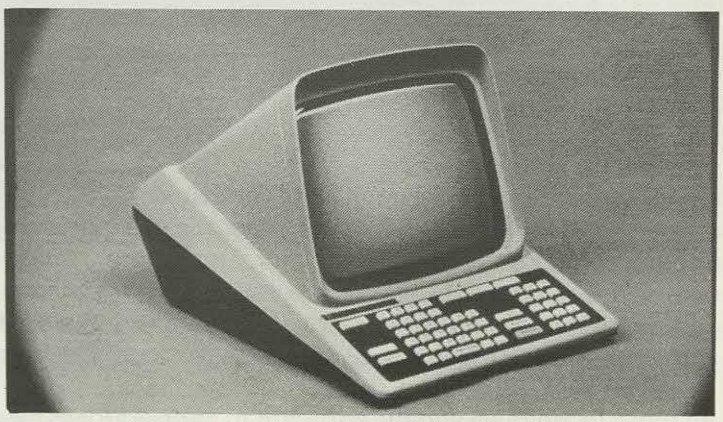

throughout the complete strata of business organisations right down to the smallest shopkeeper. Mass facsimile units could provide small businesses with the added benefits of a low-speed photocopier and hard copy videotex output, where, for example, a shopkeeper uses videotex and wants some hard copy occasionally, yet does not wish to afford the price of the dedicated hard copy device.

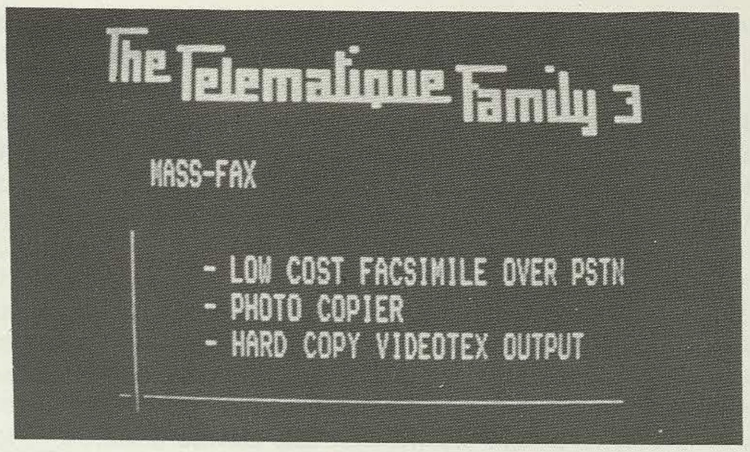

Of all the Télématique family, mass facsimile is the trickiest to give time-scales, objectives and dates on. It is not as advanced in some respects in terms of commitment from industry, in terms of definition and so on. Later this year a pre-production model should appear for evaluation, but if someone said to me today, "I'd like a thousand of those mass-fax units in the next six months," I would be embarrassed. But we would try.

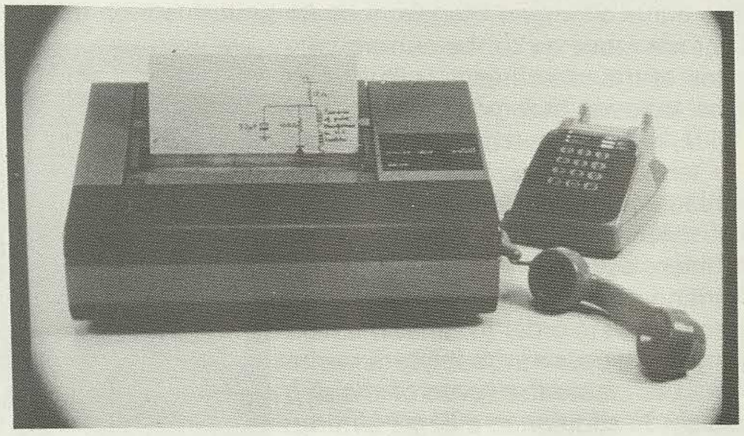

Here (second and third slides on this page) are two examples of the manufacturers' products that have been developed along these lines. The first one was Thomson, and this one is Sagem, the Safax 2000. As you can see, it is rather battered looking, because it has been carted around many places in the last few months. So mass facsimile is an objective and will be appearing in the next couple of years as a viable, low-cost product.

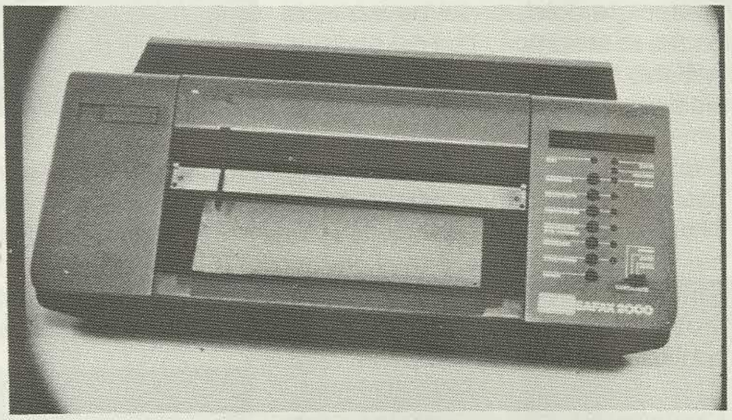

Moving now (fourth slide on this page) to Telewriter — this needs a little more explanation because it has not been around for long, although the principles and the technology are well established. It is an interactive graphics / manuscript service provided over the PSTN. The history of Telewriter is that the audiographics teleconferencing commitment in France is quite big — indeed part of the Télématique programme is to promote that development. To support it — that is to add the graphics element to the

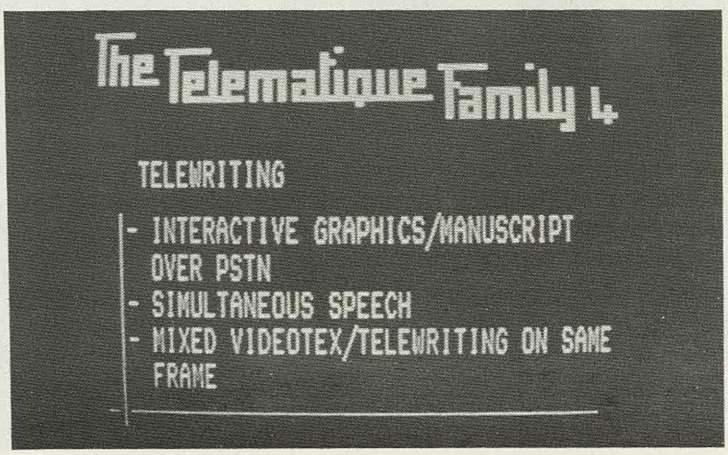

main audiographics — the Telewriter has been developed, and for that reason it is fairly high cost in terms of its comparatively low-volume production. However, I have taken the view since I came to France that Telewriter does have some interesting capabilities in its own right — on a station-to-station or person-to-person basis rather than an audio studio-to-audio studio basis over a leased line. Indeed, provided the PTT in any particular country give their blessing to the actual modem, there is no reason why simultaneous speech and telewriting should not be achieved between two people talking over a normal, dial-up circuit.

The last item on the slide refers to a second-generation development, the prototype of which appeared in 1980. The prototype has the ability to create a mixture of videotex information and telewriting graphics, such as a signature, on the same screen. This development has some software implications at the videotex level because you have to link maybe four videotex frames together to carry the mix of videotex data and the telewriting material. Nevertheless, technically, we have shown that this can be done.

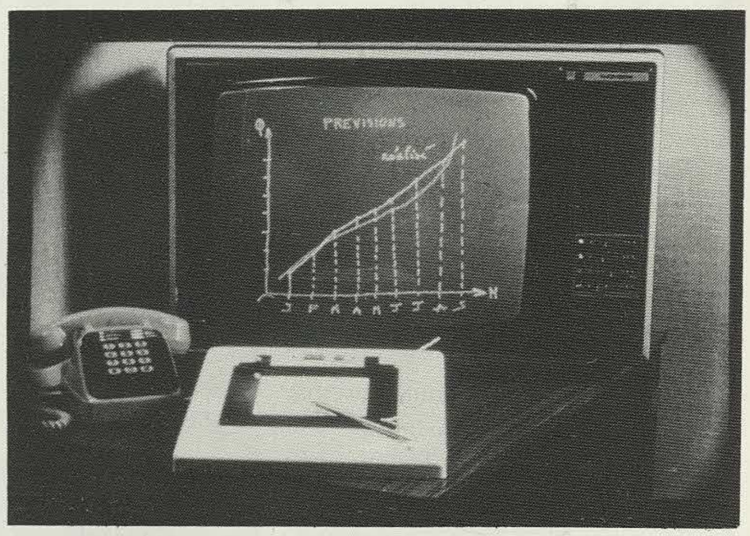

Let me move on to the current version of the telewriter. It is limited to the colours that you see on the slide, red and green, and it is limited in terms of the size of the writing area to not much larger than a sheet of toilet paper. However, the next edition does have a much wider writing area. It relies on capacitive techniques in order to sense the writing and it uses a probability curve in order to create the speed of the circuitry. In addition to a larger writing area, eight colours will be available on the next edition. The next edition will also have the ability to combine videotex and telewriting.

From my point of view I think that telewriting could be a useful development. Several American companies have their head office on the East coast and their R&D centre on the West coast. These companies have looked at Telewriter and said, "This could be just the answer to some of our needs". At present they are on the 'phone every day to their colleagues, trying to describe something they would like to see done with a circuit design, or discussing a management problem, or maybe discussing an advertising logo. Telewriter lends itself to all those sorts of applications very nicely.

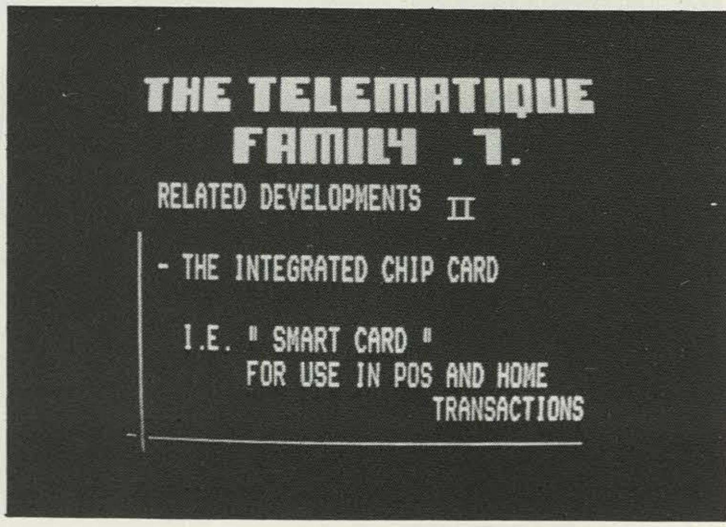

I mentioned earlier that the chip card could be conceived of in a number of applications. The first point to make is that the technology can be applied in many ways. The focus at the moment, as you will gather from this slide, is in the financial area for point-of-sale and home transactions. But the chip that appears in the plastic card has intelligence. It is powered from the reader. It has the ability to have what we might call a non-secret field and a secret field. Information in the non-secret field can be accessed by any terminal, whereas the second level of information is accessed only with the full agreement of the user, who has to key in his or her own PIN number before any access is available.

The interesting feature of the chip card is the technology of the card rather than the terminal. It is a rather thick plastic card, but otherwise it looks like a conventional credit card. (Ed.: At this point Roy Bright held up a chip card.) In this corner you can see the actual chip with the interface leads. There are eight interface leads here of which six are currently used and two are there for future possibilities. The chip card is definitely still under early development, let me stress that straight away.

The plans in France are quite interesting. There will be trials of such a card in one or more cities in France — Lyon has already been selected as one testbed — towards the end of this year or at the beginning of next year. There are four manufacturers involved, each doing their own development work on this subject — CII-Honeywell Bull, Electronic Marcel, Dassault and Flomic Schlumberger. Dassault have been concentrating on an intelligent magnetic stripe card which takes care of debiting and crediting, but now are working on a second version which will combine a magnetic stripe on one side of the card and a chip on the other. Dassault are taking this approach because for a number of years there are bound to be two types of terminal living alongside each other, and they want a card that can be used in either type of terminal.

The technology of the "smart" card not only offers greater security — hence the excitement it is arousing in the banking and money market areas — but also it is a very versatile computer. In fact the name given to it by CII-Honeywell Bull, CP8, is an abbreviation of "computer-in-the-pocket 80," and it literally is that, and it has some very interesting possibilities.

To give you one example other than banking, it can be used for medical information, where the bearer of the card has on the non-secret area of the card such statements as allergies, blood group and so on. On the protected area of the card, where only his PIN number can open the door, there would be a lot more privileged, confidential information which he must agree to disclose only to an authorised person. Similarly, chip cards could be used for security access. Another application that has been developed in France is that of TV subscription. In North America Cable TV and other people are anxious to broadcast programmes, some of which they do not mind who sees them provided that the viewer pays a monthly flat rate fee, and others of which you pay on demand or pay on use. The chip card with the appropriate reader would allow that discrimination. For example, only those people who have paid special fees for this month's box office release would be allowed to see it and they would pay by inserting their card into a reader.

The first piece of hardware to use the chip card is in the retail industry. The card allows the merchant to operate the terminal itself, but with the handheld keypad it allows the user to invoke his own PIN number. This can be used both on-line and off-line. You will see from the slide that it is a magnetic stripe card rather than a chip card. It will be on trial later this year in France, at St. Étienne.

The magnetic stripe card will eventually be replaced by the "smart" card. On the first slide on the next page is a mock-up of the configuration that might be envisaged whereby one of the electronic directory terminals could be used in conjunction with that handheld keypad. The users could not only do their browsing round the "best buys" but could pay for those best buys, or inspect their bank accounts, or check the value of that card. One of the features of the card is that it will give you a read-out of a value. It has many ways of being used financially. Not as a credit card, I stress. One of the more popular uses is that you literally buy a card worth, say, $100 and it is debited each time you use it. You can get a read-out on any terminal of the amount of money left on that card at the time you wish to check that.

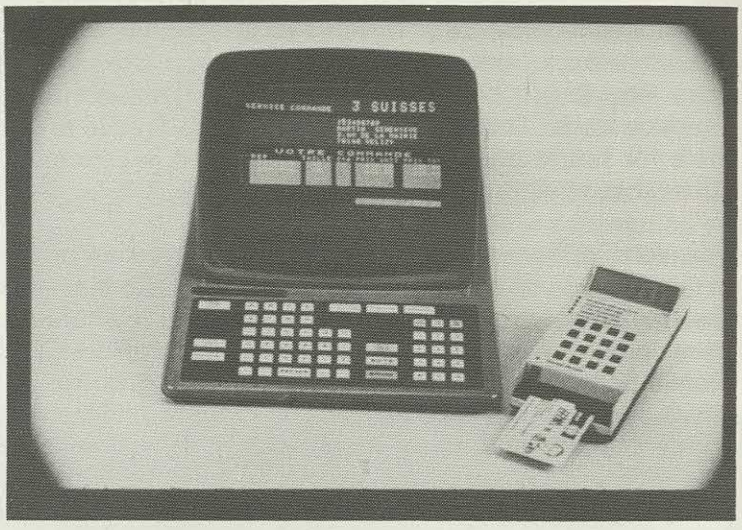

I said earlier that I should come back to the electronic directory. It is one thing to talk about terminals, but it is another to talk about what the user can do with the terminals and how meaningful the service may be to the user. What I have done on this series of slides is translate some of the typical French screens into English to illustrate the principle of the electronic directory. I go no further than that. First, it offers the user a choice either of alphabetical (i.e. White Page) listings or Yellow Page listings, or at the bottom of the screen, emergencies, doctors, fire service and so on.

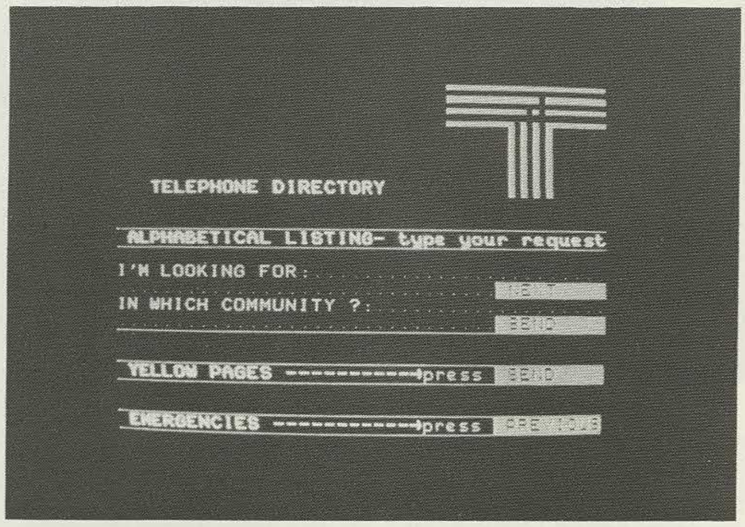

Let us assume that the user opts for a White Page listing. You will see that on the screen he is invited to key in, "I am looking for . . ." You need the name of the person and the community, with as much detail of the address as you can provide. The grey areas at the end which display the control commands NEXT, SEND, SEND AGAIN, or PREVIOUS, indicate the control buttons on the keyboard which the user would operate to invoke the next step. Let us assume that the user has, in one-finger fashion, typed in MORVAN, and for community he has typed in the name PARIS, which is a fairly open-ended search area for any computer. This is where the system gets more interesting.

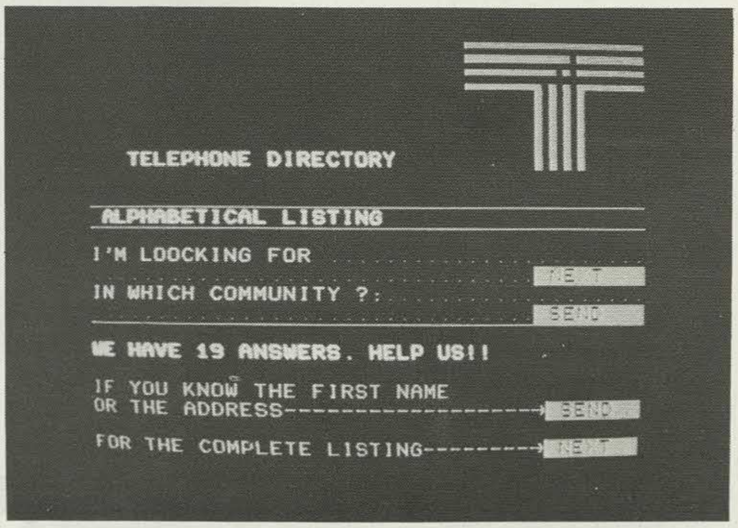

The example (third slide on this page) is rather artificial, for I am sure that there are more than 19 MORVANs in Paris, but for the purposes of illustration it is claimed that there are 19 listed with that name in Paris. If you had this in real-time on the screen, the top half stays available and the bottom half is erased, and a new entry made on the bottom half of the screen alone which says, "We have 19 answers. Help us! If you know the first name or the address press SEND," which means that you would type in whatever you know; or alternatively, if you cannot help us any more, "We can give you a complete listing" in which case press the control button marked NEXT.

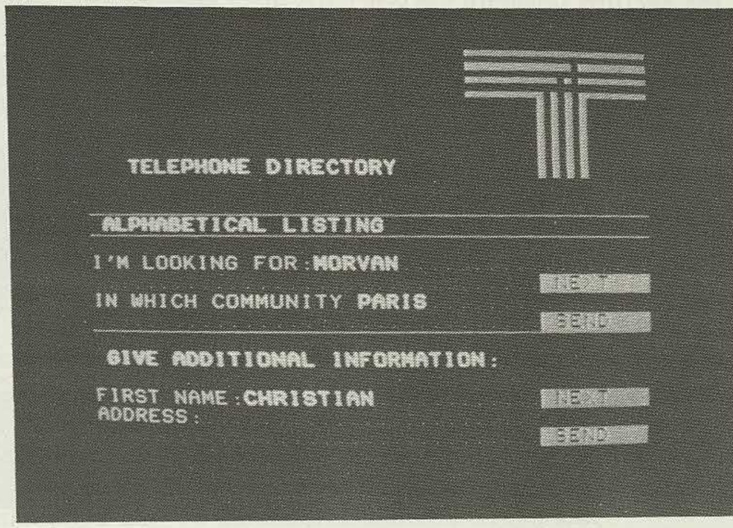

We assumed (first slide on this page) that the enquirer has been a little more helpful and entered the name CHRISTIAN, and is now able to ask the computer if it can be more specific itself.

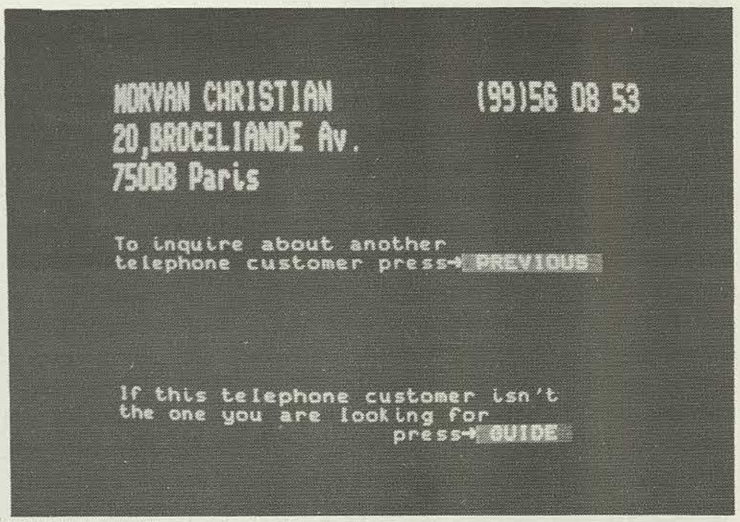

The answer is "Yes" and we are supplied with CHRISTIAN MORVAN's full name and address details and his telephone number. If you have further enquiries for other telephone numbers, you are asked to press the button marked PREVIOUS. Alternatively if you do not agree with that offering, you can press GUIDE and the system will take you to a guide section. That is important because you will appreciate two things. Some users may be very expert but not have full information all the time. Other users will be beginners and, whether or not they have full information, they just do not know how to use the system. So the system offers three levels of skill for beginners, intermediate users, and advanced users. The advanced level means that the user types in right at the beginning as much information as he can and bypasses a lot of the slower steps. On the other hand, the beginner will be very happy to go to the guide and be led by the hand through the system for the first couple of occasions. So that is one example based on White Pages.

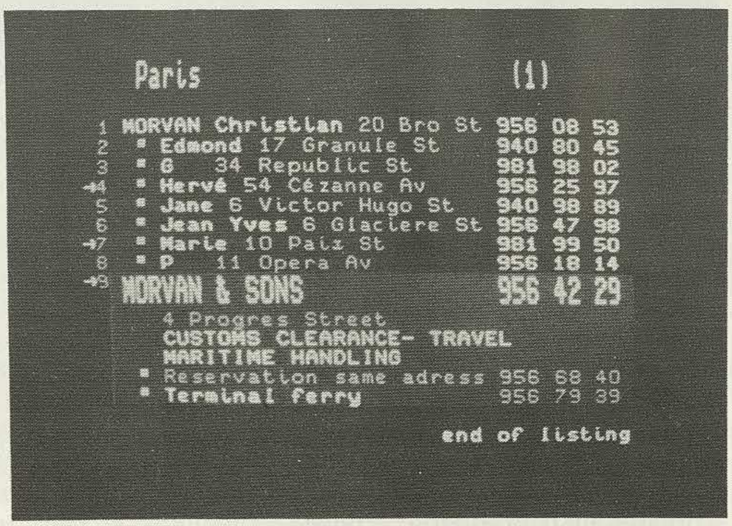

Here (third slide on this page) we have an example of Yellow Pages. It has a listing which includes this Yellow Page business card entry. If this were a live screen the little arrows on the left next to the numerics would be flashing. That indicates to the user that there is at least one more supplementary frame of information supplied by that company — and this is where we get into a more commercial environment — that the user can inspect.

Let us look at an example of a supplementary frame. This frame is inviting the user to use this particular company's car ferry to Jersey. Let me draw your attention to the use of graphics. Remember that this is typically a monochrome screen, although if you were accessing it with Teletel you would get the benefit of colour. All frames are designed for both colour and monochrome, or for both types of representation. This frame is really an advertising page.

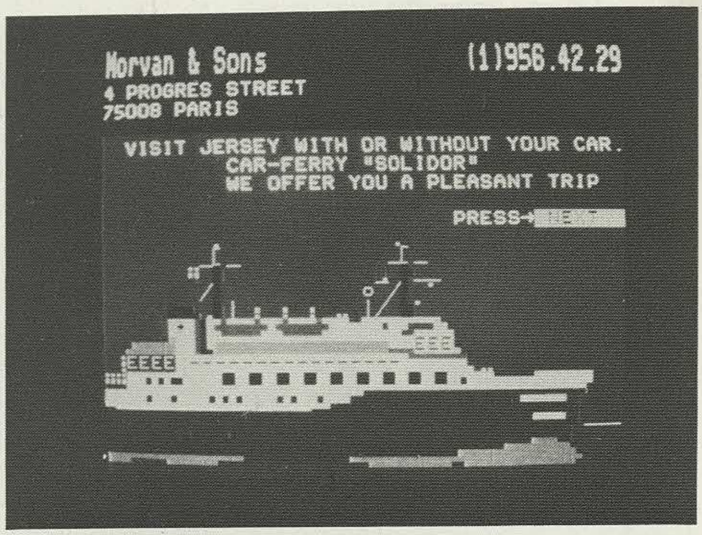

We are invited to press NEXT, so there is a second supplementary frame, and on the next frame we are given more factual information such as timetable material and the like. The important point is that we have used the electronic directory terminal. We have shown that the ability of the user to use a full alphanumeric keyboard is there. We have shown that the graphics and other features associated today with videotex are available, and there is absolutely no reason why that terminal could not be used in a variety of applications which have nothing to do with electronic directory.

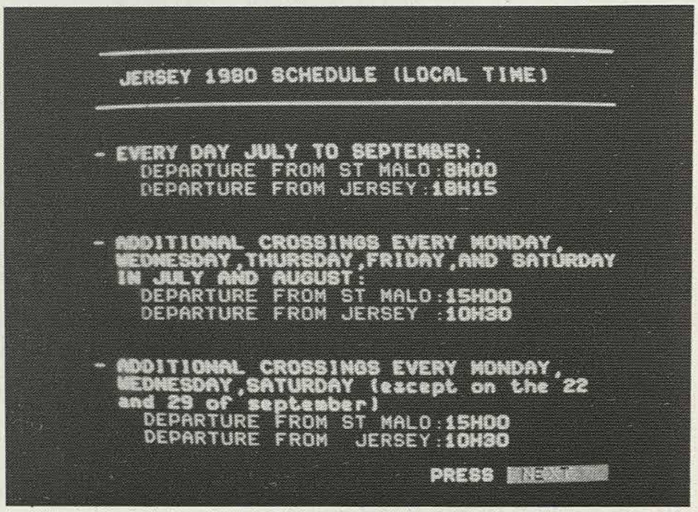

The third slide on this page shows the network for the electronic directory. Working from the top lefthand corner we have a terminal and we have used the dial-up network to access one of the concentrators that I mentioned earlier. Through that concentrator we have intercepted an interrogation centre. The first point to make about the interrogation centre is that the interplay between the user and the system takes place at that level. The next level — the orange box labelled "documentation centre" — is where the inscriptions are held and the complete database for that area is located. All that dialogue that we saw just now was the user interacting with the interrogation centre.

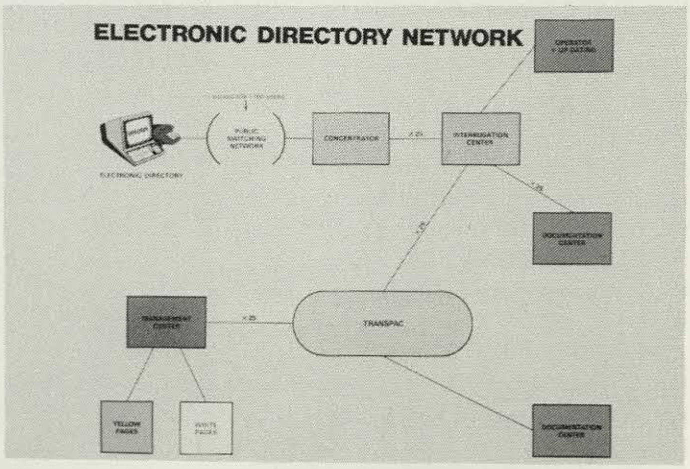

Coming down the page you will see the oblong shape that is labelled Transpac. That is to indicate that in addition to your local service directory enquiries you can get into a national environment. Each of those lines at the high level are labelled X.25 and are involved with packet switching. Off to the left you have the management centre which is responsible for both the Yellow Page and the White Page activity on the system. On the right there is another orange box to show that you could have accessed a remote database enquiry centre. The green box on the top right represents operations and updating. Naturally there has to be some sort of spooling operation and updating of entries of inscriptions between the system. In the case of the White Pages, the PTT will have that responsibility, because the PTT has the existing sales office and all the advice notes that are issued go through this office. The PTT already does this for its present directory system, which is accessed by the directory operators when they handle enquiries from the public today.

For the Yellow Pages, the contractor in France is a company called Office de Nance, which is a subsidiary of Abbas, the largest advertising agency in Europe. They have done a remarkable job. When I went to France I was surprised to see how much work and how much skill they have put into this Yellow Pages subject. I mentioned earlier the photographic terms. They have created almost a science on this subject and are becoming a very powerful influence on the design of the system, and the value of their experience can be considered outside of France.

Today I have shown slides on or referred to videotex, the electronic directory, mass-fax, telewriting, "smart" card, and audiographics teleconferencing.

So what is happening outside of these developments? Essentially Intelmatique, which was the name that I coined when I went to France for this operation, has been set up as a commercial subsidiary of the French telephone administration, of the Directorate General's Office of Telecommunications. That is, we are not a department within the PTT but rather an arms-length operation, but nevertheless wholly owned by that organisation. Our mission is summarised on the slide. It is to promote the Télématique programme internationally; to stimulate customer interest worldworld; to assist customer experimentation; to advise the French administration on harmonisation — the point I made earlier about matching the needs of foreign markets to the products; and finally, to liaise with the French Télématique industry. I have invented yet another phrase, the French Télématique industry. It is a multi-supplier situation in France. You have heard me mention some of the larger companies, such as Thomson CSF or Matra; software houses like Stéria, and CAP-Sogeti, and others, which are very small, but have found a little niche to develop editing terminals or specialist modems. We in Intelmatique are representing their interests worldwide in terms of the Télématique operation. From what I have said today you will appreciate that it is a tall order for any one company to have a perspective of all these products and to know where they are going, to know what is happening and which are the leading edge activities. It would also be naive to expect one company to talk about other companies which may be their rivals in certain product areas. In a way that is what Intelmatique has been created to do: to promote and create an awareness, up to the moment when a client organisation wishes to get down to commercial negotiations. Even then we might be involved in the discussions. We then hand over to the end supplier. Eventually, a contract is drawn up between the customer and his supplier or suppliers if we are into a system environment.

I think that will illustrate to you the seriousness with which the French Government regard their Télématique efforts. They are pouring a lot of money into it; US$27 billion went into the telephone system, including all this work, in 1975 in the seventh five-year economic plan. Recently, just before Christmas, the eighth plan was announced, and US$23 billion has been devoted to such matters as telecommunications, microelectronics and computing. That is big money, even in today's terms. So I think that is an important point to register with you.

CONFERENCE CONCLUSIONS

David Butler,
Butler Cox & Partners Limited

The title of our conference was "New applications, new technologies" and I believe that during the past two days what we have heard from the speakers has been a series of extremely interesting, well-prepared, challenging and stimulating presentations, and that they have amply fulfilled the mission which we entrusted to them.

The thing that sometimes puzzles me, going away from a conference such as this, is the sheer multi-dimensionality of the whole problem. You can look at it down the economic axis and think of it in terms of trying to extract profit from business operations or to improve efficiency or effectiveness. You can look at it along the technological axis. You can look at it along the human factors' axis. You can look at it along the social axis. There are so many different axes along which you can look, and you see the same picture from a totally different viewpoint. What puzzles me is how we can reach a stage where we can get some rounded view of the whole area at which we are looking; some rounded view of the whole area of information technology and its impact on our way of work, our way of earning a living, and our way of life.

For two or three years we have been struggling to move towards such an understanding of the areas that we all discuss together, but I believe that in the last year or so we have made signal progress towards such an understanding. I believe that from this meeting, and from others like it in the last year or so, have come at least the beginnings of an understanding of that global picture which it is so important to understand.

Some of the inputs that have been most valuable to me came from a meeting which many of you attended, the one that was addressed by Professor Stonier in London, when he talked about the way that as technology advances it has the effect of turning garbage into wealth. The oil, the coal and other garbage were in the ground long before we knew how to extract them and use them. It is technology which turns that garbage into wealth. But along the road it also means that more and more people get shaken out of traditional occupations, as technology advances. I am sure that those of you who were there will remember the very graphic term that Professor Stonier used to describe that process. He called it the "ratchet" effect, because the technology turns in one direction, but you find that when the cyclical boom comes along the ratchet does not turn back more than one or two notches in the opposite direction, so that, of the many thousands of people who fell out of occupation as the technology moved, only a few thousand move back into occupation as demand rises.

I have been looking recently at the work of a speaker who we have not yet had to a Foundation conference, but who I should very much much like to have at some stage in the future, whose work similarly contributes to such an overall understanding of the problems that we are trying to resolve. I refer to Professor Daniel Bell, the author of *The Post-Industrial Society*. Professor Bell looks at things from a slightly different viewpoint than Professor Stonier, and sees three principal characteristics of the kind of society which all of us are trying to build.

The first is the change from the goods-producing society to a service-oriented society. We have all seen the figures that, for example, we published in a Foundation report as long ago as 1977 showing how people had migrated from manufacturing industries into service industries. Professor Bell emphasises that you do not necessarily have a society which is less wealthy in goods, only that it requires very much fewer people to produce a larger volume of goods. But what Professor Bell calls "the axial principle" of his post-industrial society is the codification of theoretical knowledge for innovation. He contrasts the way in which modern product development gets carried out with the hit-and-miss methodologies of the past. He points out that, for example, Alexander Graham Bell was a speech therapist who was simply trying to find a better way of communicating with deaf people, and happened to stumble on the telephone; and that Edison, perhaps the greatest inventor of all time, was a mathematical ignoramus who knew very little about the theoretical basis of the things which he invented.

Professor Bell is basically a technological determinist. He sees all this knowledge becoming codified, mobilised in a highly systematic way, and thereby transforming our ability to produce new goods and products, and to turn much, much more of Professor Stonier's garbage into much, much more wealth. He calls that the creation of the new, intellectual technology.

A third speaker that we had at another meeting was Gordon Scarrott. He talked about turning from the computing slave which was the computer of the past to the knowledgeable servant, the system that can help us to do things that we could not do before. He said that if Linnaeus had been classifying mankind today he would not have called him *homo sapiens,* he would have called him *simia ordinans*, the organising ape. Gordon Scarrott sees technology as being fundamentally linked to man's ability to run a civilised society. For those of you who were not at that meeting, he divided the history of mankind up into huge blocks and said: "this is when we learned to write; this is when we learned to write on paper; this is when we learned to print". He sees the advent of information technology as a similar major turning point in the whole history of civilisation.

He makes the fundamental point that all those advances in information technology — learning to write, learning to write on paper, and then learning to print — were changes in the method of transmitting information not changes in the method of processing information. Therefore, since processing power — the human brain — remained distinctly finite and expensive, it was centralised. Now we have methods of processing information which are cheap and plentiful, Dr. Scarrott says, ergo, the politics of the last hundred thousand years are hereby declared obsolete.

All of these differing viewpoints offer different perspectives on the technology about which we have been talking for the last few days — communication satellites, microprocessors, voice systems, spatial databases and so on. I think that it will be a very long time indeed before we can say that we have anything like the overall grasp of the nature of this subject that we are all striving towards, but I do believe that some overall picture is emerging. At least we are beginning to get some idea of the scale of the revolution in which we are all involved.

I should like to close this conference, perhaps appropriately in view of our last speaker and the presence of a number of very welcome guests from France, by quoting from the famous report *L'informatisation de la Société*, by Simon, Nora and Alain Minc. They said:

> "If France does not respond effectively to the serious new challenges she faces, her internal tensions will deprive her of the ability to control her fate."

That is a very important and challenging statement with which the report opens, because it is not about technology, it is not about economics, and it is not about working in offices. It is about the very fabric of the society within which we live. I believe that is a valid perception of the scale of the problems with which we are all trying to wrestle.

I should like on your behalf to thank all the speakers who have contributed to the conference this week. I should also like to thank you, our members, for participating so actively in the discussions, for your challenging questions and your intelligent comments. I look forward to seeing you all this evening at dinner and to successful workshops tomorrow. I hope that you will all thereafter have a safe and speedy journey to your homes.

The Butler Cox Foundation

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

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

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

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

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

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