Transcript

Management Conference

Stratford-upon-Avon
November 15–16, 1977

The Butler Cox Foundation

Butler Cox Foundation

MANAGEMENT CONFERENCE TRANSCRIPT

Stratford-upon-Avon, November 15-16, 1977

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

CONTENTS

LIST OF PARTICIPANTS

BOC	Mr K Austin Mr G Owen Mr R Tester
Burmah Oil	Mr P H Groves Mr I G Handyside
Grandmet	Mr G Hawker
Plessey	Mr G Cochrane Mr D Holloway
Post Office	Mr B Cartwright Mr C Hartshorne Mr H K Khalil
Reckitt & Colman	Mr R Ferris Mr G Gregory Mr P Lancaster
RHM	Mr D Batts Mr P Goble Mr A Wright
Spillers	Mr R E Brook Mr B J Gladwin Mr G G Weaver
The Thomson Organisation	Mr M Adler Mr G Consterdine
Turner & Newall	Mr G Calverley Mr R Gleed Mr J Hill
Unilever	Mr B Maudsley
United Biscuits	Mr V Hope Mr R L James Mr R Talbot
Vickers	Mr G Proctor
Speakers	Mr B W Manley Mr D H Roberts Dr J Evans Mr E Fauvre Mr T W Hart Mr R D Bright Mr H Donaldson Mr A d'Agapeyeff Dr H C Zedlitz Mr L Elstein
Butler Cox & Partners Staff and Associates	Mr David Butler Mr George Cox Miss J Haffenden Mr J Kinnear Mr K Kozarsky Miss B Lacey Mr P Landauer Mr R Malik Mr R Woolfe

THE SYSTEMS INDUSTRY OF TOMORROW A SURVIVAL KIT

D. Butler

BUTLER: Gentlemen, good morning. Welcome. I never interfere in conference arrangements because I have found, from long and bitter experience, that whenever I differ from Roger Woolfe on that subject it is because he is right and I am wrong. But I really do feel obliged this morning to express to you the horror which I feel in welcoming you to Stratford-upon-Avon, for a reason which I feel constrained to explain.

Just a little over 400 years ago a young man, almost a boy — I think that he was 14 or 15, exceptionally young even at that time — went up to Cambridge, and over the next 10 years or so proceeded to carve out for himself one of the most remarkable careers that this country has ever seen. He became the leading dramatic playwrite of his day. He also became a notable figure around Cambridge, a dashing, talented linguist and drinker. His name was Christopher Marlowe. But behind the Marlowe that we know, the poet and man about town, was another Marlowe about whom we knew absolutely nothing until the 1920s; because in addition to being the Christopher Fry and the Aga Khan of his day, Marlowe was also the James Bond of his day. He became a spy.

Those of you who doubt that thesis might care to read Calvin Hoffman's book on the subject, together with my article in *Esquire* for September 1956.

My aims this morning are twofold. The first is to give you a very brief progress report on the work of the Foundation up to now. The second is to set the scene for the rest of the two days that we are going to spend here together.

Normally, it would be my practice in opening one of these regular management conferences to offer a particular welcome to those companies which are new members of the Foundation since we last held a management conference. But on this occasion I think that it would be inopportune to do that because all the companies here represented are in fact new members since our last management conference in May, when we first discussed in public the launch of the Foundation. If you look on the back page of the Agenda which has been circulated to you, you will see a list of the companies which are represented here today and which have become the founder members of the Foundation. I am also pleased to be able to announce that since we went to press on that, two further companies have indicated that they will be joining in the work of the Foundation. The Weir Group from Scotland are joining straightaway, and ITT will be joining us as members in January. So we will have an opportunity to welcome them to the next conference in April.

As far as we are concerned, we feel that the membership of the Foundation is satisfactory from two points of view: first, that we have enough members to hold interesting and useful meetings; and secondly, looking down the list, we feel particularly pleased about the quality of organisations which are participating in the work of the Foundation. They are all organisations with a lot to offer and a lot to contribute in the way of experience and expertise in discussions such as we will have over the next couple of days.

Looking briefly at the activities which have gone on so far within the Foundation and what is going to happen in the next few months, we have published the reports on Data Networks, and Display Word Processors,, and there will be reports coming out in the next few weeks on the subjects of Terminal Compatibility, and Office Automation Technologies. Other subjects on which we will be producing reports early next year include Private Automatic Branch Exchanges and Network Traffic Planning, and also a basic report on the Convergence of Technologies, which seeks to tie together a lot of the detail about which we have been talking.

Up to now we have held four of the regular professional and technical seminars. I think that most people in this room have attended one or more of them. We have held them on the subject of Word Processing; Data Networks; one on Computerised Branch Exchanges and the evaluation thereof; and one on the Recommendations of the Carter Committee on the Future of the Post Office, which was held earlier this month. There is also a forward schedule of those meetings and, if you are not familiar with it, Roger Woolfe can certainly let you have further copies.

The fourth point that I should like to touch on in this brief progress report concerns the management of the Foundation. As I think you all know, it has been our intention since we started to have a management board which would consist of some representatives of Butler Cox & Partners and some representatives of the member organisations, so that we could sit around a table and discuss as freely as possible the way that things should be going in the future. In normal circumstances, we would expect the representatives of the members to have been elected by the members as a whole; but this year we felt that the members did not really know each other well enough to make an election meaningful. So we simply co-opted on to the management board three people whom we were confident would be helpful to us. They include Brian Maudsley from Unilever; Geoff Dale from the Post Office; and Roger Tomlin from the Thomson Organisation.

In the near future, we shall be seeking to fulfil the other commitment which we made several months ago, that when the UK group of members of the Foundation had reached what we regarded as a critical mass, when there were enough members to be able to hold meetings like this and have discussion with a varied range of interests, we would then begin to internationalise the Foundation by establishing members' groups in overseas countries. We shall be proceeding with that in the New Year, when we open our first European office.

The last but one point that I should like to mention is concerned with input from the United States. As I think you all know, one of our partners, Karl Kozarsky, is resident in the USA; and I should like to put on record just how useful it has been to have Karl there, looking out for things that might be important for the future. Obviously, we do not expect to have our thinking too much influenced or dominated by what goes on in the United States, but we need to know about it and take it into account; and that is one of Karl's roles.

The final point that I should like to make about this and future conferences is that we are preparing, and it will always be our policy to prepare, a transcript of the conference which will be circulated to the members after the conference. The transcript will be limited in its distribution to the members and, aside from that, these conferences are off the record in the sense that we have not invited, and we do not propose to invite, journalists or magazine writers to attend these conferences; so you can speak without fear that what you say will appear in *Computer Weekly* next week. I should explain that my friend, Rex Malik, is here today not as a journalist and not as a writer, he is here because he and I have agreed that he will act as a sort of catalyst and challenger to many of the things which we jointly agree. When we all agree on something, I think that it is useful to have somebody in the room who is highly likely to disagree, simply because we all agree. I think that Rex will fulfil that role admirably. But it is important to recognise that these sessions are private working sessions and they will not be reported in the Press.

The conference over the two days as we have organised it is divided into three modules, and each one has its own distinct purpose. First, we will look at some of the technologies which are emerging right now; some of the technologies which promise so much for the future in terms of cheaper, more effective and more reliable systems. We will concentrate on some of the ones which seem likely to have the biggest impact on the cost effectiveness of systems over the next few years. But, of course, having the technology available is only part of the story; another major determinant is how that technology will be brought to the market; how it will be packaged and put in front of the buyer. The final section is concerned with what users are currently doing to try to take advantage of that technology. We hope that by the end of the conference we will have formed a logical chain of what is likely to emerge in the area of technology; how it is likely to be presented to the user; and what users are doing right now to move in the directions of using these advanced technologies. We will try to preserve a logical link throughout those three modules, and at the end of the conference, my colleague, George Cox, will have the unenviable task of trying to sum the whole thing up to see what lessons have been learned out of the two days.

What I should like to do in the time that remains is to try to set the scene for the speeches which will follow. I should like to do this by looking at a number of different problems. First, what are some of the problem areas which are now arising in the field of data processing and data communications? Secondly, what are the trends in some other areas of what one might term the "electronics" industry? I think that if we look carefully at what has happened in the world in the past few years in other branches of the industry, we can get some pointers to what is likely to be important in the future. Thirdly, again looking to problems which may well afflict us in the next few years, some of the problems of definition of data processing and data communications in the USA. We all talk a great deal about the convergence of technologies, about the fact that boundaries between data processing and data communications are becoming increasingly fragile and artificial; but this poses certain problems, particularly in the area of who does what, which are likely to be of importance to users of communication services in the future.

THE SYSTEMS INDUSTRY of TOMORROW

- **Problems arising in dp/dc**
- **Trends in other areas of electronics business**
- **Problems of dp/dc definition in the U.S.A.**
- **Lessons for Europe : a survival kit.**

Finally, drawing on the three analyses that I have done above, what are some of the lessons for Europe and what sort of survival kit do we need to help us to get through? I may say *en passant* that I think that some of the points which we will discuss in the next few days may touch upon the problems of structural change within both the computer business and the communications business. These are subjects which have been discussed frequently in the past and doubtless will be in the future. But in the United Kingdom certainly, and to some extent in Europe, there is in my view now a greater willingness to discuss some of these structural changes, changes for example in the role of common carriers, than there has been in the past. Certainly if one looks at the situation in the United Kingdom, now having a new Chairman at the Post Office and a new Managing Director of Post Office Telecommunications, having the Carter Report in the air, and having a new Chief Executive for ICL, there seem to be more prospects of change in the wind now than there have been for some time.

One way of looking at the problems, and one to which I will come back in a moment, is to think briefly about the data processing business and about the data communications business. I guess that most of the subjects that we will hear about in the next two days will touch on either data processing, data communications, or both. We have already talked at these meetings and others sufficiently about the basic phenomenon of the convergence of technologies, the way that this is happening at the technical level, the way that it is happening at the market level, and the problems which it is creating as far as the user is concerned. But perhaps one could be a little more precise about the nature of those difficulties and how they have arisen elsewhere, and how they are likely to arise here; because I think that the important thing to recognise is that although the industries are converging, they do have characteristics which in many ways are importantly different.

INDUSTRY CHARACTERISTICS

DATA PROCESSING	COMMUNICATIONS
Fragmented	Monolithic
Competitive	Monopolistic
Sales-oriented	Reactive
Short write-off	Long write-off
Laissez-faire	Regulated

If one looks at the data processing industry, for example, in general it is fragmented. Although there is one near monopoly operator, the rest of the industry is certainly fragmented; and it is an industry in which we have seen small companies go from nowhere to a rather important position in the market very quickly. It is in general competitive. Most companies which buy data processing equipment take tenders from a number of different suppliers and can seriously consider the possibility of changing.

It is sales oriented. Most of the progress in fact in the computer world has been achieved by salesmen knocking on the doors of customers and convincing them to buy the product rather than customers battering on the doors of the computer manufacturers saying, "I insist on having one of your machines." It is characterised from a financial point of view by relatively short write-off periods. It is a *laissez-faire* business. If you want to design and market a new computer, unless you want to sell something which could be used for producing improved nuclear weapons to China or the Soviet Union, by and large you are allowed to do it.

The data communications business, on the other hand, is relatively monolithic and monopolistic. If you want to use data communication facilities, you do not have to spend very long choosing your common carrier, it just depends on which country in the world you are in. It is reactive in the sense that, in general communications' administrations will seek rather to respond to proven public demand than to speculate that they might be able, by launching a service, to create a public demand. It is certainly characterised by long write-offs, both in the public and private sector. The amortisation period for most telecommunications equipment in our experience is roughly double that of most data processing equipment. It is a regulated environment in the sense that if you want to connect something to the public network you have to seek the approval of the necessary regulatory body in whichever country you happen to find yourself.

These are fairly formidable differences between the two market places. If we are saying that increasingly the activities of suppliers will bridge the gap between those two traditional industries, then we should think pretty carefully about how that gap will be bridged, about who does what, and about how we give the user a reasonable chance to get his hands on the facilities that he needs, and not be held up because there are enormous conflicts going on between interested parties who want to extend themselves in one way or another.

Later I will examine some of the evidence for the way that this particular conflict is taking place between AT&T and IBM in the United States, because it seems to me that there may be some lessons that we in Europe can learn from the exercises which are going on in the States right now. It may seem, on the face of it, that the differences between these two sides of the systems business are largely philosophical and academic, but I think that when we come to look at the US experience, we will find that the difficulty of reconciling these two sets of characteristics actually lies at the heart of some of the problems for the user, who after all is not much concerned with who provides the service as long as he can get the service that he wants at a price that he considers is worth paying for it. So I believe that the differences which I have mentioned are fundamental and that they very much affect the environment within which the user can buy services. I think that perhaps the difficulty that we now see in buying computerised branch exchanges, when it is really rather difficult to determine what sort of amortisation period one should go for, is perhaps the tip of an iceberg which we need to examine much more carefully.

Secondly, may I look briefly at what has been going on in the past few years at some other branches of the electronics industry, other than the computer business and the telecommunications business; because I think that there are some important lessons to be learned there both by the suppliers and the users of such services.

What I have noticed, and I think that it shows itself in the evolution of the electronics business over the past few years, is that we seem quite recently to have come to something of a watershed in the electronics business. A few years ago, I think that it was true to say that what distinguished a competent supplier in these fields was some pretty low-level skills: the skill to manage batch fabrication of assemblies with a reasonable degree of reliability and a reasonable yield; the ability to hire labour at costs which did not crucify both the supplier and the potential customer. I think that in some ways the most difficult and important claim that a computer manufacturer or a telecomms manufacturer could put forward two or three years ago was, "We have a product which by and large works for most of the time." I think that was a reflection of both the technology and our understanding of the problems of production management in those areas of the technology.

Some of you may have read in the Press that our company recently had an opportunity, as a result of some work which we were doing for the National Economic Development Office, to go round and talk to all the computer manufacturers in the United Kingdom in a relatively short space of time. What struck me, talking to each of them and finding out their current problems, is that we do actually seem to have got beyond the point at which the basic problems of fabrication and production control in electronics are the major dominant problem. The impression that I had, talking to the computer manufacturers in particular, was that they were now pretty confident that they could manufacture their product and that it would work pretty well; and the problems of bad yield, bad reliability and crucifying costs seemed largely to have gone away and to have been replaced with other problems, problems concerned with packaging, product presentation, marketing, investment in software and things of that nature.

Of course, there are some obvious reasons why this should be so, and we will be hearing about them later in the conference, mostly concerned with the declining cost of components and therefore the declining significance of labour cost as an element within the total factory price.

If we are moving away from a situation where the key factors were the ability to buy the cheapest labour, and manage production, and get a reasonable yield from fabrication and assembly processes, what are the factors that will make for success in the systems industry of tomorrow? Can we get any pointers on what they may be? What are the important skills that the manufacturers will have to develop and that the users will have to look out for if they are going to pick a reliable supplier for the future?

I want to look at two other branches of the electronics industry: the TV industry and the pocket calculator industry; because I think that they have some useful markers for the future. Let me make it absolutely clear that this is not to say that I am automatically assuming that the computer industry and the communications industry will go in the same direction as the TV industry or the calculator industry, because there are obvious differences. It is certainly not to say that I am assuming that the computer industry and the telecomms industries will become consumer industries as such. I am not arguing that, either. What I am saying is that in a field where one is, to a limited extent anyway, coping with the same structure of problem, the same problem of putting components into a presentable package and marketing it in a way that people will be interested in buying, there may be some useful pointers from the TV industry and the calculator industry.

In the TV industry, we have witnessed over the last few years a battle between the US suppliers, the European suppliers and the Japanese, which makes anything going on in our neck of the woods of systems, computers and telecommunications look rather genteel and well-bred. Let me remind you of some of the latest moves in that battle to the death.

When the US Government decided that it was time that it started to attempt to reflate its economy fairly gently a couple of years ago, there were consumption tax reductions applied to television sets in order to attempt to stimulate demand. But these were across the board and they applied not only to sets manufactured in the United States, but also to television sets which were manufactured overseas, and particularly in Japan.

The US Treasury was approached by the indigenous television industry in the United States and asked whether it would impose countervailing duties, at least not to leave the Japanese suppliers any better off than they had been before the consumption tax credits; but the Treasury refused. However, the United States regulatory mechanism being what it is — if you cannot find one body to do what you want you can normally find another — the Customs Court in New York overruled the Treasury finding on a petition from a US manufacturer.

At this stage, a further regulatory body, the International Trade Commission, recommended that it was not right to penalise the Japanese *vis-a-vis* the rest of the suppliers, but there should be some direct intervention to help what it described as the "injured" US TV industry, which became known as the escape clause.

However, in May of this year, all of those negotiations and legal processes were swept aside when President Carter intervened decisively to set up a so-called voluntary agreement with the Japanese television industry. It was agreed that the Japanese companies would import into the United States no more than 1.56 million sets a year throughout the period of the agreement, plus just under 200,000 assembleable sets, so that the total input to the USA from Japan would be less than 1.75 million sets a year.

If one looks at the reaction in the United States to this intervention by President Carter, I think that one can begin to get one clue on what a European survival kit should look like in the future; because what has happened in fact is that the US importers of overseas television sets have immediately increased their prices in order to cope with the problem of lower volume; and they have increased their prices yet again because they suspect that later on they may be facing retrospective duty additions for which they have to have some money in the jam jar right now. So as has been pointed out by Gene Gregory, whose work I should have acknowledged earlier in producing these figures, it is in fact the US consumer who suffers most from this protectionist policy, because over the next few years something like 4 million American buyers of television sets will pay more than they would otherwise have done for their sets.

It is not surprising of course that America should respond in this way. Let me not give the impression that it was a rash or hasty action, it was not; by the time it was taken 11 of the 18 US TV manufacturers were still in business, the other seven having gone bust; but eight of them were running at a loss. Japanese imports in September of last year had already increased to 30% of the market. But the important thing is that the protectionist angle which the States has been ultimately bound to pursue in this particular industry is being financed out of the pockets of their own consumers. I think that is an interesting and revealing look at one particular industry.

I will come back to the television industry in a moment because I think that there are reasons for believing that the Americans are going to cope with this situation in a rather more creative way than they have up to now, and there may be some lessons for us in that as well. What the television boys are doing now is trying to learn some lessons from a situation where the United States appeared to have got itself blown out of the water, and then made a comeback in a very big way — and that, of course, is the calculator industry.

JAPANESE SHARE OF U.S. CALCULATOR MARKET

YEAR	MARKET SIZE ($m)	JAPANESE SHARE BY VOLUME	JAPANESE SHARE BY VALUE
1966	N/A	Nil	Nil
1970	224	N/A	40%
1971	N/A	60%	45%
1974	750	N/A	21%

Let us look at the Japanese share of the US calculator industry. I do not have the figure for the market size in 1966 but I do know the Japanese share of it — they were not in business in the States at that time. By 1970, just four years later, the market size for calculators had gone to $224 million and the Japanese importers had secured an unbelievable 40% by value of that market. In the following year the Japanese had secured 60% of the market in volume terms and 45% by value. Now at that stage I suggest that any rational man in the calculator industry in the United States would have sold his business and got into something else. If he had done so he would have made a catastrophic error because, by 1974, the size of the market had grown to $750 million, but the Japanese share of the market by value had been cut in half in that three-year period. Somebody described it as "the most significant comeback since Lazarus" and I think that one can see from the figures what they mean.

THE U.S. CALCULATOR INDUSTRY: SECRETS OF THE GREAT ESCAPE

MOS price/performance
Use of price/demand elasticity
Declining labour costs
The entry of the vertically integrated
- Texas
- Rockwell
- NS
'learning curve' production and pricing policies.

I should like to spend a few minutes just talking about how this transformation was achieved, because I think that it may have some lessons for us. First, MOS technology and the price performance associated with it. I remember once in the United States meeting the President of National Semiconductor, who at that time struck me as being about 18 years old and being a zillionaire. He said to me, "There are two things you've got to recognise about our industry. The first is that everybody's research department works for everybody else. The second is that I spend my entire life in aeroplanes looking for labour that is .001% cheaper than the labour that I'm buying now. That is the secret of my business." I asked, "Where are you off to next week?" and he said, "Britain." He had done the Caribbean.

Of course, this question of the ability to hire labour at cheap rates favoured the Asian manufacturers enormously, but as the price of the technology declined relative to the product as a whole so marketing, research and development costs and, in particular, transport costs became more and more significant. So the advantage of the Asian manufacturers and their pool of cheap labour became less significant.

Secondly, a lesson that they had learned from the Japanese, the use of price/demand elasticity in this particular market. I cannot put this any more clearly than to quote my colleague George Cox at our last conference, when he said, "When calculators cost £50 a time, in my house we had none of them; today, my wife doesn't even understand that you can change the batteries." The Japanese mastered the business in the calculator market of forecasting accurately how fast demand would accelerate, given a certain rate of reduction in price; and the Americans certainly learned from that.

Thirdly, declining labour costs, which I have already mentioned; and fourthly, and probably most important, the entry of the vertically integrated manufacturers such as Texas, Rockwell and National Semiconductor, who came into the business in 1972 with a degree of vertical integration which obviously gave them components at a price which other manufacturers could not match. Finally, what have been called the learning curve production and pricing policies which may also be extremely significant for the computer and communications industries for the future. What the American calculator manufacturers realised was that, learning from the example of Boeing and other well-known cases, every time you double production volume you improve your production performance by a percentage which is not only real but also predictable; that you can actually map on a graph the extent to which your production skills will improve and your production costs reduce each time you double your volume of production.

If you believe that to be true and if you can establish what that degree of improvement is — and they have — then you reach the situation where you can plan your price movements very, very carefully, ahead of time, and exploit the price/demand elasticity to be where the new market is going to be, faster than anybody else. It is that linkage of production efficiency to price planning to market planning which has led to the renaissance of the US calculator business.

Switching back for a moment to the television industry, of course all the American television manufacturers are now looking at the calculator boys and saying, "How did you do it? Can we learn anything from you in fighting off the Asian invasion?" The answer is probably they can. In the 1960s, the business of making television sets was very simple and highly labour intensive. If that were so today, the chairmen of the television companies would be doing exactly the same thing as the President of National Semiconductor whom I mentioned a moment ago, hunting the world for labour that was very marginally cheaper than that of his competitors.

But it really is not like that now. Solid state TV is as different from the old-fashioned TV as these calculators are from the ones of the 1960s. If I could offer you one figure, the rate at which a human being can insert components into a PCB in a television factory is approximately 300 components per hour; the rate at which a numerically controlled machine tool can do the same job is 72,000 components per hour, which is work for 240 people. So we are getting a better yield; better maintenance; more accessibility in the product. It now seems, if you look at the structure of the costs of the TV industry, that because of this change the Americans can probably compete with anything that the Japanese can now put forward in the area of labour cost and capital intensitivity. They also have a big advantage that, as energy costs rise and transportation costs rise, *vis-a-vis* their own market they have an edge there.

You will probably say, "If that's true, why did I read in the paper only this year that Grundig have set up their latest television manufacturing plant in Taiwan?" The reason for that, which would have been unimaginable two or three years ago, is not to get the advantage of cheap labour rates in Taiwan but to save transport costs when they attack the Asian market.

I do not think that any of us in this room should think of ourselves as slavish admirers of US technology and management, but I think that it has to be said that in these two industries the United States has pulled off a remarkable recovery in one, and seems about to do so in another. But it is not so much that about which I want to talk as identifying why they were successful in moving out of very perilous situations into ones which are slightly better.

One can see three characteristics which have been fully mobilised in these recovery programmes. One is dynamism, the speed at which the American companies brought the new technology to the market, and learned how to manage it and how to market it. One is integration. I doubt if any of the companies which did not have vertical integration could really have made the transformation in the calculator business that the integrated boys did. Finally, there is what one might call a global perspective, a willingness to put production and to seek markets in almost any corner of the world.

I will come back to this in a moment because it seems to me that there may be lessons here for European suppliers and purchasers, and at the end of my talk I will try to put forward a survival kit for European companies of some of the things which we should be thinking about if we want to learn to get out of our difficulties in the way that these two industries are apparently doing.

Before I do that, however, I should like to look at the industry characteristics that I mentioned near the beginning: data processing being fragmented, competitive, sales oriented, given to short write-off periods, and a rather *laissez-faire* market, as compared with communications which is more monolithic, monopolistic, reactive, given to long write-off times, and a regulated environment. To put it at its bluntest I suppose that the question is: how do companies operating in one area or the other of those two industries decide just how far they want to go in the opposite direction and just what range of services do they want to provide? That seems to me also to be a problem that we need to resolve in Europe if we are not to waste an enormous amount of time on regulatory hassles between computer companies and telecommunications companies.

In the USA, the task of drawing that line and somehow rendering it defensible rests on the shoulders of the Federal Communications Commission. For some years past, one of the major objectives of AT&T has been to persuade the Federal Communications Commission that it should be allowed to offer both a data communications and a limited data processing service. The reactions of the interested parties are predictable. IBM, through the medium of Wallace Doud, said, "If that were to happen, a large segment of potential data processing applications would be relegated to monopoly carriers." The choice of words is exquisite, isn't it? — "relegated to monopoly carriers." Do you know what is the IBM word for an installation that has somebody else's equipment on it? It is "contaminated".

The riposte from AT&T's Paul Villiers, the Assistant Vice-President for Network Operations was, "It would be a blow if the Bell System couldn't offer its users a complete data communication package." In the United States there is an edict of the Federal Communications Commission which states quite categorically that it is really very simple: common carriers are not allowed to offer data processing services. But I think that if one wanted a vivid example of the process of the convergence of technologies and the convergence of markets, one need look no further than this particular situation; because then somebody turns round and says, "Well, what is data processing? You'll have to give us a definition." The FCC said, "Data processing is the electronically automated processing of information where the output information constitutes a programmed response to input information." Somebody pointed out that that is an almost perfect description of the telephone network. So perhaps the telephone network is data processing and Bell should not be allowed to be in that either. I don't think that the FCC quite meant that.

At that point, the situation became even more confused because the Justice Department, choosing its language with magisterial dignity, accused the FCC of "goofy rule-making". Fundamentally, what is going on is that AT&T are pressing for a definition which bundles up data processing and data communication as much as possible and fuzzes the boundary as much as possible. What is IBM's view? Actually it is a surprise. What IBM is saying is not "Keep AT&T out of our patch," but "Please let them into our patch through a separately accounting subsidiary." The reasons for that, of course, are that it would make it easier to convince the Justice Department and the rest of the world that IBM was not an unchallenged monopoly in the United States if AT&T had what IBM doubtless expects to be an inefficient but very, very large data processing department.

EUROPEAN SURVIVAL KIT

A strategy opposed to protectionism

Dynamic use of technology

World-wide business horizons

Vertical integration (where required)

A creative regulatory framework with the user's interests at heart.

So what lessons can we learn from this in Europe? I think basically five, and I hope that during the rest of the conference we shall see our speakers addressing themselves to at least some of these requirements for the future. I think that we need a strategy which is opposed to protectionism pure and simple because the evidence suggests — and I don't think that it should surprise us — that if we just draw lines around markets and say, "Thou shalt not enter," in the end it is the consumer who pays the extra price for the product, as in the case of the Japanese television sets.

I think that we need dynamic use of technology. Perhaps it will become clear during the course of the rest of the conference just whether we feel that we are taking up the technology fast enough. We find we are asked, when discussing these matters particularly with managers who are perhaps not directly exposed to the technology, "That's all terribly interesting, Mr. Butler, but when will that technology be available?" Time after time one finds oneself saying, "Well, actually it's available today. What we've been describing is something you could do today." Are we taking up the technology fast enough?

Worldwide business horizons. I think that European organisations have not been anywhere near as adventurous as American ones, either from the point of view of where they incur their cost or where they seek to establish their markets.

Fourthly, there seems some evidence at least that in some markets vertical integration may be a *sine qua non* to effective competition. We have to say that in Europe the degree of vertical integration among companies concerned with the systems industries is very, very small indeed. What is more, I think that it is also true that there is very little in most European countries of a national industrial strategy to help companies to move in that direction.

Finally, coming back to the question of AT&T and IBM, a creative regulatory framework with the users' interests at heart. I think that far too often the regulatory frameworks in Europe seem tó be prejudiced in favour of the *status quo* and to take account of all kinds of what in my view should be subsidiary political issues, such as employment prospects in certain areas of the country or whatever. It must be remembered that the main purpose of regulatory mechanism is to serve the interests of the consumer, the user. I think that too often we tend to forget this.

If you say to me, "All the points you've made are ones which I am inclined to agree with. But is this the Europe that we know and love? — the Europe whose political, economic and technical motto might be 'Vive la difference!' " then I have to say to you, "No, I don't think it is." If, on the other hand, we have to think about some fairly fundamental changes in the way that we structure our industries, the way we manage them, and the way that we think about national regulatory policies, and particularly international regulatory policies, it seems to me that the time to start thinking those fundamental thoughts is now; and this, perhaps, is as good a place as any to begin it.

THE IMPACT OF NEW TECHNOLOGIES IN OFFICE COMMUNICATIONS

B.W. Manley

BUTLER: I am particularly grateful to Brian for agreeing to come along and speak to us today because I happen to know that he is in the rather challenging position of doing two jobs at the present time, and it was very good of him indeed to find the time to come and speak to us. His current job is Managing Director of Pye Business Communications Limited and he is due to become Managing Director of Pye TMC in about 10 days from now. So with no more ado, I'll hand over to Brian.

MANLEY: Thank you very much, David. It does seem as though I have two jobs at the moment; although strictly speaking I am resting between engagements, because I finished as Managing Director of that first company on the last day of last month and I start as the new one on the first day of next month. That is the explanation why the overhead slides that I show you had to be done by my own fair hand, because I have nobody else to work for me at the moment.

Just as a small digression, I was interested in what David was saying earlier about the component business and the television business. Up to a couple of years ago, I was with the Mullard Company; and if you are reading your newspapers at the moment, you can see the politicking that is going on between Mullard and Hitachi over the building of the new Hitachi factory in County Durham. I recommend it as good reading on how to conduct a protectionist racket, but I have a particular angle on it. I hope that we can see a little bit of protectionism for the component business in this country, because without it I think that it can succumb, not just in this country but in Europe, to what is a very carefully planned and strategic attack on the electronic industry of Europe by the Japanese. It begins with the component business; and without a component business there can really be little electronic industry in total.

That was a digression, although it does lead into the point that, in talking about the impact of new technologies in office communication, it is important to see that the opportunities that have been given us in that sector arise from the component industries, from new components and new materials.

Having chosen my subject, I found it extraordinarily difficult— and maybe that difficulty will emerge as I go on and you will see it — to talk about technology and to separate it from techniques and applications, and not to wander about too much between one and the other.

I also found it a problem to approach it in a logical way, to see where the technology would impact upon the present situation via techniques and applications. It is a very broad spectrum. What is most difficult to discern is the speed with which some things will happen. It is very easy to say that in the next decade we will see enormous changes in the area of office communications and business communications. But at what speed? I find that very difficult to answer.

In order to understand why that is, one can see that many of the factors that will affect the introduction of new technologies are environmental and social, not technological at all. To get that into perspective, we have to look a little bit at the background of industry and commerce today. If we look at the manufacturing industries, we can see that investment decisions are tolerably well-planned and the techniques for evaluating them are pretty well known and, broadly speaking, this encourages an objective approach to putting money into the improvement in manufacturing technology.

The situation is not the same in the office. To quantify improvements in the business environment outisde the direct manufacturing sector is far more difficult, and indeed is not really tackled at all, I would say — certainly not in the UK industry sector.

	1974	1979
MANUFACTURE	7.7.	? 7.0
OFFICES	8.5	10.0

EMPLOYMENT IN THE U.K.
(MILLIONS)

Whether this is cause or effect I do not know, but if you look at this projection you see some interesting trends. The number of people in 1974 employed in manufacturing industry, in the productive industries was 7.7 million; it is certainly dropping at the moment. For constant output anyway, which unfortunately is what we have largely seen, it certainly needs to drop; in fact I think that it is only slowly becoming apparent within the UK economy just how far and how much efficiency in production needs to increase and therefore the number of people involved must drop. Whether it will reach 7 million by 1979 is a question of what kind of efficiency improvements one wishes to put into the equation and what kind of expansion in output one expects to see, but probably it will be around 7 million.

If you look at the numbers in the office environment, the figure is 8.5 million and it is trending up; it will probably reach about 10 million in 1979. You can talk a lot about the manufacturing sector and the factors that influence it, but not very much about that second line; and I do not really think that we understand what is happening and why it is happening. It is certainly true to say that it is an environmental and social problem rather than a quantified and carefully calculated process.

I did a little bit of a market survey the other day, to try to see what people knew about this second element in relation to business communications. I asked one organisation — which happens to be my parent company in the UK — how much it spent per year on computers and on the whole area of data processing. Within seconds, I had a very detailed answer which separated this among the different elements of cost, and I was referred to at least half a dozen committees and experts who could give me even greater detail than I was already given.

Then I tried to find out how much was spent on communications in the company. It is a multi-site operation of between 20 and 40 sites. And nobody knew. I do not think that this particular company is unusual in that respect. It is treated as a totally fragmented problem. I had in the final analysis to get a lot of budgets out, and then do some quick calculations on what we spend on paper, mail and postage; what we spend on telephony, telex and so on. I came to the rough answer that we spent about £6 million in 1975 on computers and all their aspects and about £12 million on communications. Yet we did not really know.

I think that there is an aspect here which affects almost everything we will do in the coming decade about the introduction of technology. If we do not understand the problem, there is not very much chance that we will understand what happens as we approach new technologies in that sector.

What are the social and environmental pressures that we are discerning and which will determine the adoption of new technologies anyway? I am happy to see that we have some representatives here from the British Post Office, and that must be the biggest impediment to the introduction of new technology that we have in the country. First, let me say that I am sure that the UK is not the only country which finds this to be the case. David said earlier that what was important was user satisfaction. Unfortunately, I think that we are going into an era — and we will be in it for a long time — in which one of the primary elements of that equation will be the preservation of employment in the Post Office. This means that the pressure on the *status quo* is enormous. The regulatory aspect of what is done in communications and the regulatory aspect by the British Post Office really will be a great determinant to what happens.

I do not wish to cast stones at just the Post Office, because the second major impediment is industry itself, the telecommunications industry. We are in the middle of the most appallingly difficult period for the communications business that it has ever been in, I think. That is not simply in anticipation of my new job next week, but I think that it is really a problem. If one looks at what is happening in employment patterns, in demand and so on, if one looks at the American pattern and the number of people employed in the industry, it has fallen enormously fast and is doing so also in this country. I think that, two years ago, there were some 80,000 people broadly in the telephone and communications business in this country among the major companies involved. I think that figure today is about 50,000, and within five years will probably be no more than 10,000 or perhaps less than that.

How does a company act and react in those conditions? Clearly in a very self-protective way. Therefore, it is not necessarily looking for the ways in which it can most quickly diminish its own productive base and its own added value. I believe that David put his finger on a most important point, that of vertical integration. We see the borderline between what happens in the electronic component area and what happens in the equipment assembly area moving steadily in favour of the component; and the point at which design of the system comes in is also moving. It is a whole area which one could well discuss at great length. The fact is that the structure of the industry will change enormously. It is quite impossible to foresee in a few years time the number of companies that are involved today, even in the UK. There must be liaisons, alliances, combinations and disasters in the industry which will change its total structure and the number of customers that are involved. So we see the second element in this equation determining the rate of adoption of new technologies as being the protective aspect of the industry itself.

The third point, which I think could be the most positive and strongest of all, is the nature of the office itself. Because it is very difficult in many respects to be objective about it and measure it, and to quantify what is happening, one is bound to accept the fact that to a great extent it is a social aspect of life. What happens in the office is half about the social animal and the way that he behaves. If you offer him the opportunity of changing his secretary for a mechanical robot or even an electronic one, I do not think that he will do it if he has any way of avoiding it. So I think that you have this other factor which is very strong, namely that the environment of communications within an office business is a major determinant on the speed of adoption of technology.

But even if there were no such impediment, then one must still say that the determinants are means and available technology. There is no point in having a new technology simply because it is new. To stand any chance at all, the new technology must give some discernible advantages even if it is not objectively measurable. We can prepare a check list of what

- LOWER COSTS
- INCREASE SPEED
- ENHANCE ANALYSIS
- IMPROVE RELIABILITY
- EASE USE

they have to be. It must do one of those things. It must lower the cost of the system. It may lower the cost simply by having lower equipment costs; reducing the space that is required; increasing reliability; requiring fewer people although, as I said before, that does not seem to be a factor which is really a very powerful one; or consume less power and so on

Increasing the speed. Clearly by widening bandwidths of systems; giving faster access wherever one wants the information and so on. These are important factors.

Enhancing the analysis. That can mean a whole variety of things: ensuring that information is available in the right form, at the right place; more flexible handling.

Reliability. It goes without saying, although I will come back to that later, the need for few systems failures; fewer subscriber failures, or however one measures it.

Easier use. Clearly you need to have more systems intelligence in order to do more things without increasing user skills. It can also include more ready accessibility and such things as that. That really has to be our check list when we look at the way in which new technologies can impact the communications business.

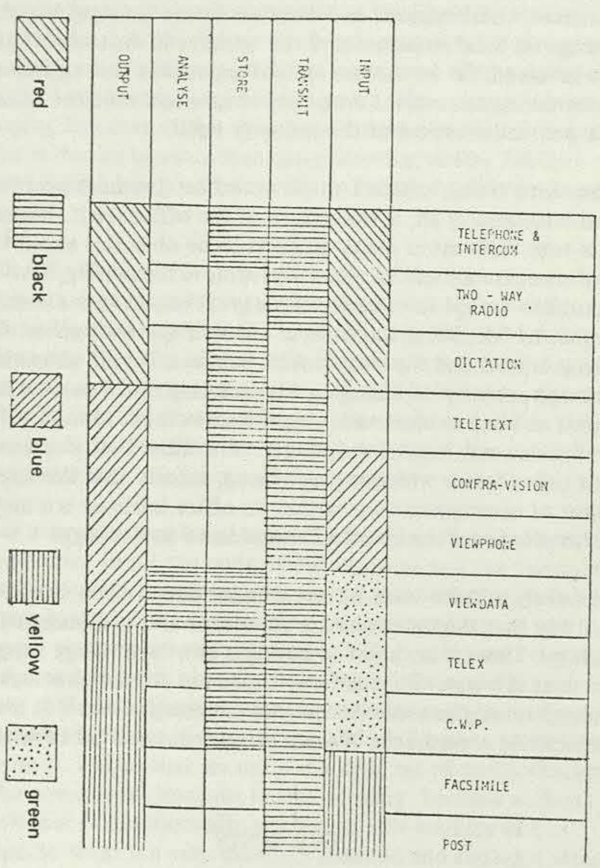

But what are the techniques around which we are going to examine these criteria? This is one way of looking at the problem, where we can look at this business of flow of information between and among groups. These are the functions that we have to fulfil: input; transmit; store; analyse; and then output. The techniques: the telephone and the intercomm; two-way radio, which can include paging and personnel location; dictation; Teletext; Confravision; Viewphone; Viewdata; Telex; communicating word processors; facsimile; post. They all have an input and they all have an output; and all, apart from dictation, although that is a somewhat arbitrary choice, do not actually have a transmission process.

I have tried to analyse these techniques in terms of the input and output methods. Here, in the pinky-red one, we are talking about sound input and output; the telephone is the obvious example. Confravision, Viewphone for video input. The green one is data keyboard input; and the yellow one is the hard copy. So we can classify by sound, vision, data or hard copy; and in the same way on the output.

Then one has to make one's choice. Of course, the information can be of a transient nature, as with a telephone call; or it can be hard copy. Generally, the information will either be in the head of the transmitter or from some stored bank of information, so you can classify it in a number of ways. But we now have to make a choice to see which of these are susceptible to new technologies and which are the areas in which we can see the impact or the potential impact of new technologies.

	LOWER COSTS	INCREASE SPEED	IMPROVE RELIABILITY	EASIER USE	GREATER CAPACITY
VIDEO INPUT			•		
COPY READERS	•	•	•	•	
TRANSMISSION	•	•	•		•
STORAGE	•		•	•	•
DISPLAYS	•		•		

I have made my choice and you may make yours, but I think that we must look at storage. Clearly, that is a technology which will impact many of these techniques and others. Displays are very important indeed. Transmission technologies. Video input devices. Hard copy readers. Let us just think about those. There may be other functions which will be impacted by technology, but let us look at those and then think how, from that previous analysis, we can identify the significant benefits which technology can offer.

Again, to some extent these are judgments about what is important. If we look down the righthand side we can see those criteria. I have put red blobs where I believe that we can see the key things where technology can impact those particular functions of video input, readers, transmission, storage and display. There are two lines there that are clearly dominant in the sense that they occur in most cases, that is lower costs and improved reliability.

I suppose that lower costs is self-evident. It is self-evidently important in almost everything that we do. But what about reliability? It is probably true that in most of the things that we do today in terms of our communications package, we are reliable enough. As an example, if you take the telephone system, we are fairly tolerant in telephony, but the fact is, looking at some Post Office figures in the UK, the number of failed calls is about 2%. It does not seem like that on my telephone at times, but that is what the Post Office tell us so it must be true. The most likely reason for failure is that the chap at the other end is not there so you cannot make the call. It is certainly arguable that if you are going to spend money on the system, you would spend it on personnel location and not on improving and diminishing the failure rate. Better spend it on a paging system, so to speak, or offering that kind of facility; or improving the access to a chap who has moved away from his desk and is at some other place rather than diminishing the fault rate. That is one example where, with a bit of tolerance from the user — and I think we are very tolerant — we have a system which is about reliable enough.

But now let us look a little further forward and see what is happening. As we add complexity via technology, quite clearly the potential for failure becomes greater. As we put more control into the system, the potential for catastrophic failure also becomes greater; your whole system goes off the air or you lose your memory in the electronics sense, and then you are in grave trouble; so the penalties for unreliability are higher.

But there is a further aspect, which David touched on earlier, and that is the question of service costs. Service is essentially a high labour cost area. We see the two things going hand in hand and relentlessly increasing the cost of applying service to any system whether it is domestic, whether it is consumer, whether it is in the office, whether it is a national system like the telephone network or whatever. So that in time we see that the costs of ownership of the system will become the dominant cost. Ask your wife what happens if the washing machine goes wrong and she discovers, when the chap comes to repair it, that just coming and opening the back of it will cost her £12 or some such figure. Already we are at the point now where we can see almost the impracticality in the future of having the kind of service operations that conventionally exist to service electronic goods today.

We have to find another solution. I would say that within the next five or so years that solution has to come through a far wider use of self-diagnostic systems, and indeed of self-healing systems. One can think of ways of doing that electronically, if not in the washing machine certainly in the telephone system and in the communications sector.

Again, an example today on an SPC PABX, I believe that all of them have self-diagnosis as an essential part of the system. Gone are the days when the telephone engineer could walk in and, by listening to the Strowger things going round, he could tell exactly what was wrong. He cannot do that any more, so he needs to have a diagnostic routine which prints out exactly what the fault is when it occurs. Clearly, it is then a small step to have the output of that diagnostic routine directly reported via the network to the telephone service operation. The self-healing aspects, of course, are already done in the sense of having two processors checking each other and switching in as appropriate. Self-healing in that respect really only means duplication, and when hardware costs become as cheap as we expect then the duplication is a fairly straightforward one.

The key technology in that whole area is, of course, large scale integration of circuits, and VLSI is the next step so that one can put more intelligence into systems. So reliability and lower costs are key factors.

I think that we should have a look at each of these in turn to see what is happening in the area of technology. I start with storage partly because that is the sequence in which I have put these pieces of paper, but also because it is perhaps a central aspect of what is going to happen in the future. If one looks at what is happening in dollars to the storage costs under different headings, one sees from 1968 to 1977 the trend of reducing costs for the minicomputer, the on-line memory, the off-line memory. We are talking of a cost fall in about a decade of an order of magnitude and more. That process is going on continuously. Indeed, if we look now at the area of CCD and

	MINI COMPUTER	ON-LINE MEMORY	OFF-LINE MEMORY
1968	*PDP 8* 13,000	*MAIN FRAME DISC* 25,000	*MAIN FRAME TAPE* 20,000
1971	*PDP 11* 8,000	*CARTRIDGE* 10,000	*MINI TAPE* 9,000
1974	*NAKED MINI* 3,000	*FLOPPY DISC* 3,000	*CASSETTE* 2,000 *CARTRIDGES* 3,000
1977	*MICRO* 1,000	*BUBBLE CCD* 75	*MINI FLOPPY* 1,000

bubble memories where $75 now is the kind of price about which we are talking, whereas previously a mainframe disc in that area was $25,000 a decade ago, we are approaching a very interesting point. I think that the bubble memory situation is intriguing and I have no doubt that we will hear more about that in the course of the next couple of days. Commercially, one can obtain a 9200 bit memory on a small unit. Certainly, currently existing are 256K bit stores; and within a couple of years we will certainly be talking of megabit stores.

This produces the interesting situation that storage itself has reached the point where it ceases to be a cost problem in terms of the applications about which we are talking. I guess that the average office needs some few tens of megabits of information, if one chose to store it in the office. I think that it is arguable the extent to which one would and the extent to which it can be stored electronically anyway, but clearly at these kinds of price levels now we have reached the point where cost is no longer the impediment.

What is the problem now, if one thinks in terms of the office situation, is the problem of inputting data into this store. Let us look at that problem a little. If we think of the material that flows through and into an office, we each have a particular point of view and a particular view of that problem; our difficulty then is that our view differs according to the business that we are in and according to our own aspect of that business.

I think that it is true to say that we do not know much about it. It is another of these areas where we really have not sufficiently analysed the nature of the flow of information through a business chain. There are a couple of factors which are constant. Within a major, large organisation, about 70% of information — it can be 80%, it can be 60%, but it is certainly more than 50% — will be within the organisation information; probably only 20% or 30% will actually flow outside the organisation. A proportion of it — I don't know what that proportion is and I suspect that it varies enormously — is material which readily lends itself to handling via keyboard into store or on to display.

I was talking over coffee about this, and I think our friends from the Post Office might take the view that perhaps 70% or 80% of information could be handled in that way; in which case, of course, one can then see the impact of storage on the office and in communications within the business environment as enormous. We see it either on a central basis, or more likely on an office by office basis: banks of storage available for at least 80% of the data flow through the company. Looking more generally at the businesses with which I have been involved, I think that the answer will be much less dramatic than that and much slower coming. A very large proportion of material that we handle at the moment physically in hard copy will stay in hard copy; it will be preserved in its natural state. I think that letters and memoranda and things of that kind lend themselves to electronic handling, but only a proportion, maybe 50% or less, will fall into that category.

If that is the case, you still have all the normal systems that you have at present, with an overlay, or an add-on, or an adjunct which is your desk store or your central electronic store; and I think that the impact will be much less dramatic than one reads about in some of the journals.

The difficulty about this other kind of material that you cannot simply punch up through your keyboard is that it comes in all kinds of shapes and sizes, and it is all so difficult. Sometimes even the media is the message itself. If you get a nice vellum letter from the Queen, you do not want to put it on your CWP; you would be very upset if it arrived that way. So I think that we have to accept that sometimes the actual nature of the piece of paper or whatever is important; it does not have to be from the Queen, it can be from your bank. So there is a whole variety of material that falls into this area, where handling is the biggest problem; that is why hard copy reading is perhaps an area of technology where least at this moment is done, and which perhaps could have the biggest long-term impact on what we do in terms of storing information in the office.

It is not an electronic problem. It is very easy to think of a whole variety of electronic means of reading information. The problem is a mechanical one of handling it. That is not so amenable to technological price fall; it is not even simply amenable to solution. The consequence is several fold. I think that that in itself will be the reason why facsimile, for example, remains a rather interesting but not very dramatic area; it will not grow significantly. The reason that it will not is partly because one can do things with CWPs, or Telex, or whatever, but more particularly because the sheer difficulty of handling the material for it will be the determinant and the problem that will stop its being used. So we come to this rather undramatic conclusion about storage, in my view, that its impact will be more particularly in the conventional data processing area, that it will impact the areas of conventional letters, memoranda, written material, reports and the like selectively, but that perhaps the bulk of the present hard copy within a business will remain in hard copy form and will be stored in its natural pristine state. Disappointing. But I am sure that many of you will totally disagree with that.

Let us have a look at video input. What do we need and where do we see technology going there? Sometimes I think that although I suppose I am a technologist, way back, I tend to be a pessimist about the speed with which things will happen. Let us look at the problem of video input. Here we really mean the camera tube. What is happening there? A few years ago, we saw the CCD being promoted and great forecasts made about the way in which it would take over from the normal videcon type of primary sensor. Some of us spent some of our early years working on other kinds of camera tubes and remember that it took 14 years to bring to the point of production the camera tube which is currently used in all your colour broadcasts, the plumbicon; and I well remember that they were 14 very difficult years. Having seen that, and even with the greater speed with which technology advances now, one could forecast that it will be a very long time before one can do for video input in the solid state what one can do in the vaccuum tube.

The problem — the area relating it to that chart where we see the advantages coming — is not in cost particularly, but in reliability. The conventional camera tube does not last long enough, it is 10,000 hours; and that does inhibit it in many of the areas of application in business, not just in the simple communications sector that we have been talking about, but in security aspects, monitoring what happens in a buidling and so on. If one could go not to 10,000 but to 100,000 or 10^6 hours, then we would see a much wider use.

CCD can do that. It will start with only a few hundred element picture size, but in order to impact the areas of communications that we see here one needs to have a capability in terms of resolution and colour performance which is as great as the current camera tube. I think that is many, many years away.

At the other end of the video chain there is the problem of display. It is very interesting. In that sector the CRT has dominated the situation from the beginning. It is flexible. It has the great advantage of employing an electron beam as a switch, and it is unequalled as a switching technique. It is not so much that it is an efficient way of displaying a picture on a cathode ray tube, but it is the switchability of the CRT that really scores. The disadvantages that it does have increase with size and resolution. So if you want a very high resolution display, say 120 points per inch in something bigger than a 12-inch picture, especially if colour is required there really is nothing on the horizon that will touch the cathode ray tube. This is the kind of performance that is required if one wants to have the electronic equivalent of print.

Then one has to ask, "Are we really concerned about that? Is that really the criterion? Are we too constrained in our view of what we want to display electronically?" Most people would agree that some kind of display on a word processor is necessary, but does it really need to be a full page? Probably not; a few lines may be sufficient. Viewdata in the office, which no doubt will come — at least I hope it will — really does not, in my view, require to have a full size CRT display. We happen to have got ourselves geared to that because we started with Teletext and that was the domestic set, and so we got ourselves hooked up on a CRT display without really meaning to. If we can re-think that, then perhaps we can re-think the display problem. Once we come down to more modest requirements on the display, then other kinds of system have some advantages.

	MAX SIZE	RES- N ●/IN	LIFE HR.	POWER	EASE OF USE
C. R. T.	26"	120	10^4	HIGH	MOD.
G A S	12"	40	10^4	V. HIGH	EASY
E. L.	STRIP	40	?	HIGH ?	EASY
L. C.	STRIP	40	?	V. LOW	EASY

If we look at the options of cathode ray tube, gas displays, electroluminescence for displays and liquid crystal displays, we can see this rather interesting point that the maximum size of a CRT is, let us say, 26 inches; of a gas display probably about 12 inches. When we come to electroluminescent and liquid crystal displays we are essentially talking of strip displays because of the multiplexing problem. Resolution somewhat lower, well below the CRT display. So we cannot approach the detailed, high quality print kind of criterion, but we can get about 40 points per inch. We are in the 10,000 hours life bracket for most of these kinds of gas or vacuum displays. We do not really know about electroluminescence or liquid crystal; indeed, the problem is that we do not really know what technology we are talking about. It can be thin film electroluminescence or powder electroluminescence. When it comes to liquid crystal you have the big problem that, in order to make a display, you need some kind of non-linear effect or some sort of switch effect, and I do not think that anybody has really discovered that for liquid crystal. So you need a hybrid affair probably with thin film transistors behind the liquid crystal. So we do not really know too much about it, but we can already see the advantages in the different sectors.

When we look at power consumption, we have rather high powers for CRT; electroluminscent displays may also need high power; liquid crystal, very low. That is enormously important in something like a desk display, especially if it happens to be linked to a telephone system where one may think in terms of powering it off the system and not having a separate power supply into it. So this looks very attractive.

Ease of use. Most of them are pretty easy to use and, if they are not easy, one can put that in by adding intelligence into the system. So when one looks at that picture, I would say that we will probably reach the point where for large displays with these extreme requirements of quality, the CRT will be with us until we have long retired. But in other sectors — and I think that those other sectors have to be thought out — I believe that we will see liquid crystal becoming the dominant means of display in the next decade.

Let us talk a little about transmission; but very little because shortly we will be hearing from Dr. Evans who will tell us about perhaps the most interesting development in the transmission sector. It is the common theme among all the systems that we are talking around. But it is not a virgin desert and I think that is probably the biggest problem of all. I think that the British Post Office's fixed assets are about £5,000 million, and a lot of that is in copper wire, sunk in the ground; and in main exchanges which are related to the transmission problem. So one needs to see that that value is retained. Therefore, I think that one must look at the way in which new technologies — particularly fibre techniques for transmission — will impact the situation.

If one divides the problem between the local network, that is the telephone instrument, the wires and the exchange, the investment in copper is so great that I do not think that we will see new techniques being used actually in the transmission medium. Better use may be made of it by VLSI, and perhaps going digital in some respects; but it will be making better use of the existing means of transmission. In the trunk network we see the possibilities there for the introduction of the new technologies of fibre. I have no doubt that we will see gigabit per second data rates, with tens of kilometres between repeater stations; but I am sure that we will hear more about that shortly. However, I should just like to highlight two problems there that we should think about. One is a technical problem, which again we may hear a little about, and that is the life of the laser; reliability; the repeater itself. Ten thousand hours, which I think is current technology, is far too short when one has frequent repeaters, and one needs to have a much higher reliability there.

The second one is this business of the industry and the need for its own survival. I saw somewhere a calculation that one cubic metre of glass would satisfy all the transmission requirements for ever, everywhere. That is interesting, but when one relates it to the fact that this has to be shared out among all the companies in the world in the cable and telecommunications business, then you have the phenomenon that not only are the cable companies making glass fibres, but the glass companies are coming into the act, which I suppose is not surprising. But then you really have a self-destructive situation in the industry. I do not know what is going to happen, but I am quite sure that it is no business to be in. No doubt we will get some other views on that.

In summary, the overriding technological development is integrated circuit technology, coming with increasing impact mainly to simplify what happens; in other words, to make things easier and to increase the range and extent to which systems can be used; doing more complex things without adding to the skills of the user; reducing costs — almost to vanishing. I will hazard a guess that there will be a cost reduction even from today in integrated circuit technology per function of two orders of magnitude within 10 years. You can see the factors there that will bring it about. It may be electron beam lithography, which immediately means that you can place many more functions per unit area; or the learning curve processes that we have heard about. So that is an overriding technological aspect.

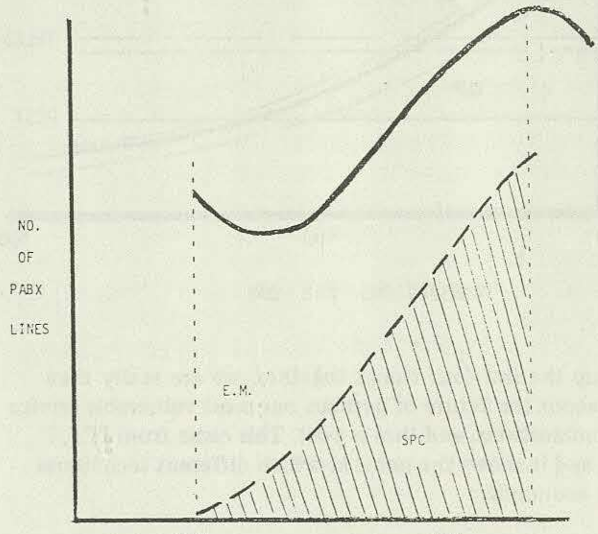

Now let us go back to functions. What is going to happen in the voice area? It is an interesting picture. If we look at the number of PABX lines in the UK, by 1980 we calculate that about 80% of new PABX lines installed in this country above 100 lines initial, if I can exclude the Post Office bit of it, will have stored program control. This is a development from 1972 up to 1980. This has come about largely through the adoption of new technology into the business in the LSI area.

We will see a much wider use of direct speech systems, which gives a number of advantages that we will come to later. Two-way radio. That is a whole area of technology, but I think that in terms of applications we are largely talking about personnel location. I think there will be a great stimulus in terms of the use of paging for personnel location; high call-back costs and the ability of VLSI to provide a message service, and in effect stored messages, will increase that enormously.

Dictation. One day, the truly voice organised system will come, but I think that it is so far away that we will not see dictation systems changing from what they are today. Confravision. Limited. I do not think that will change. I think that what will increase considerably is audio-conferencing. I do not think that it will change the way in which we handle our business or that the airlines will suffer enormously as a result of more people being prepared to conduct their business from the office rather than travel. I think that there will, however, be quite an upsurge in the use of audio-conferencing as we improve our telephone system and our intercom direct speech system. That will become a major growth area because it fulfils all the requirements that we want to see with new technology.

Viewdata: I have great hopes for Viewdata. I think that we are all waiting with bated breath to hear Mr. Bright talking about that and the Post Office trials that are about to start. It is a great new area for advance. It depends entirely for its success on harnessing the information providers. I hope that will go successfully. Without that, what one has is a rather sterile system. But I believe that is an area where we will see great advances in the future.

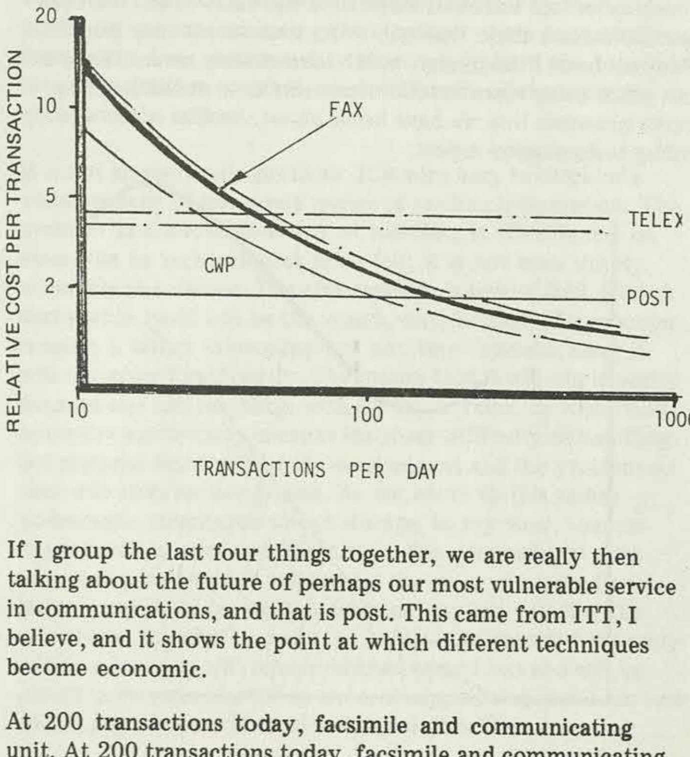

If I group the last four things together, we are really then talking about the future of perhaps our most vulnerable service in communications, and that is post. This came from ITT, I believe, and it shows the point at which different techniques become economic.

At 200 transactions today, facsimile and communicating unit. At 200 transactions today, facsimile and communicating word processors become cheaper than post. The difficulty is that you do not have the degree of standardisation that we require genuinely to take advantage of that, and we will not have that for many years. So although that is what could happen, I think that it will be delayed by a very considerable time scale, because solving technological problems is easy, but solving standardisation problems is another thing altogether.

In the long term, however, I think that we will see distribution of our traffic amongst these methods. Obviously, post costs will increase; it is a labour service, 80% or more of its costs being in labour. I am afraid that its reliability is bound to decrease further, and in the end we will see it limited only to bulky items.

POST	COSTS INCREASE. RELIABILITY DECREASING. USE LIMITED TO BULKY ITEMS
CWP	RAPID INCREASE "INHOUSE" SLOR ADOPTION EXTERNALLY
TELEX	INCREASING USE AT EXPENSE OF 'PHONE'
FAX	SLOW GROWTH LIMITED TO SPECIALIST USES

Communicating word processors; a rapid increase, of course; but largely for in-house use where one can achieve the standardisation of format and so on which is necessary; but rather slow adoption externally, for all the same reasons.

Telex. I think it is a much under-used service at the moment and I think that it will increase, especially at the expense of the telephone. Properly marketed by the Post Office I think that could happen, although perhaps they have a vested interest in not marketing it in that way. But it is a very convenient and relatively cheap way of communicating.

Facsimile. As I said before, I think that is limited by other factors and I do not think that we will see very much happening in that sector, outside the area of specialist use.

"A SYSTEM NOW COSTING $ ½ M

MAY COST ONLY $5000 BY THE

MID 1980's"

THE ECONOMIST

29.1.1977.

One last point which bears upon all these systems. I do not know whether this is true or not, but it made an impression on me when I read it, and I want to leave you with this last thought. It relates to this business of costs. If that is true, and I think that in many respects it self-evidently has some truth in it, then that could be one of the main determinants of the adoption of new technologies; but it is not the only one. I would leave you with the thought that the principal determinants are environmental and not technological.

BUTLER: Thank you, Brian. We have a few minutes in which we can pose questions to Brian and which I am sure he will be happy to handle.

QUESTION: You have emphasised the difficulty of handling in office copying. Surely the solution to this is intelligent copiers, already being described in the US.

* *Questions were not recorded verbatim (though answers were), so they appear in abbreviated form here.*

MANLEY: Of course, that is the easy bit, if I can put it that way. That is the actual process of reading or copying or whatever. Once one has a signal, one can do any variety of things; it is putting it there which is the problem. If it is a single sheet of paper written on one side, and of a preferred size, then you have no problems. If it happens not to be any one of those things, then it does get to be very difficult. It is a handling problem, not a copying or reading problem.

QUESTION: No, I'm sorry, what I am saying is that an enormous number of documents are handled now . . .

MANLEY: It is all relative. Of course you can solve the problem. One can have a whole Heath Robinson system or do it how you like to handle anything. You do it in factories every day and, as you say, you do it in a certain way in offices every day. What I am saying is that it is an area where it poses you a number of problems, which are about standardisation and hardware which is not susceptible to cost reduction and so on in the same way as the electronic side of the house is; and it still leaves you in many cases with the problem of needing an operator. The time that you want to use those things is usually at unsociable hours, outside the time scale of the normal office period when you want to use your bandwidth for sending this kind of material. I think that it is the most difficult problem and probably the least susceptible to technological advance.

QUESTION: I am confused between the speaker and the questioner, because Brian did seem to make a point about the continuing and increasing costs of the labour element of keyboard input; yet the questioner seemed to be making the point that you do not need a keyboard, because once something has been keyboarded it will be copied almost inevitably and one can generate the signal for transmission from there. With respect, Brian, I don't think you are taking the questioner's point in this respect.

MANLEY: I hope that I am not ducking the question, nor do I want to over-emphasise it. The point that I am trying to make is that when we look at technologies, one can see all kinds of things happening, especially in the electronic area. In the mechanical handling area, I think you have problems. They are problems which I do not think can be solved with the same ease as in the electronic area. Therefore, things that require mechanical handling in some way or other will impose limitations on their adoption.

Now I agree that in there lies an important point. I rather glossed over the fact that if you do keyboard input, that is a costly process because it is a handling problem. If you have a standard thing which you can scan electronically, that is not a handling problem. You can use this solution if you have standard material which is readily processed and can be handled mechanically, there is no problem. It is really the difficulty of knowing what proportion of material that flows through an office lends itself to that. I made the point earlier that I do not think we know enough about the nature of the material that flows through an office environment. We need to do far more studies on that.

BUTLER: May I on your behalf thank Brian for a very stimulating and challenging talk. One thing that is clear in his new role: whatever his organisation may lack, it will not lack forceful and strongly-held views from the top; and I think that will be a great advantage to them. Thank you, Brian.

DEVELOPMENTS IN MEMORY TECHNOLOGY

D.H. Roberts

COX: I think that everyone here is well aware of the exciting developments which are taking place in information technology. I think that what will become clear to you over the next two days is that these are not just problems that we face in interpreting this technology from the viewpoint of the user, but that it also poses a number of questions for the supplier. I think that we, as users, often sit there thinking that the supplier has all his cards there and, if only we knew what was in his hand and how he was going to present it, we could work out rather better how we could use it. That is not strictly the case. The manufacturers and suppliers not only have to solve problems of "can the technology be developed?" and "what is its capability?" and "can it then be manufactured economically?" but they also have to consider what their major competitors are going to do, the major market movers. They have to consider questions of legislation, and the moves of the Post Office and the common carriers. Thirdly, they have to interpret what we, as users, will actually buy and put into application, with all the questions of human behaviour that many of these new technologies will pose for us.

So there we have the manufacturers. They sit there, knowing what they have got, and often unclear of what to back. I can think of technologies in the past which have just been there and have failed to be taken up and used. One company which is in this position, along with all the others of a similar size, is Plessey; a company which we think of as being a telecommunications organisation, but which of course also has a great deal to offer in areas such as computing, microprocessors and memories. On this last point, we have invited Derek Roberts, who is the Managing Director of Plessey Microsystems with a number of different memory technologies that can be applied, to consider what the market wants, how we might use the technologies, and therefore which should be presented to us. Therefore, I can think of no better person to talk to us about developments in this area. Derek Roberts.

ROBERTS: Thank you, George. Gentlemen, I assume that I am here to demonstrate just how few cards suppliers have in their hands and how badly they are in need of reshuffling. I think that it is inevitable that I am going to repeat some of the general observations made by Mr. Manley. I make no apologies because I think that many of the points that he made, and I hope that one or two of the points that I will make, are sufficiently important to be worth saying more than once.

I will concentrate primarily on memory. First, let me say a few words about why I think that memory is of sufficient importance. A rather trite way of demonstrating this is to look at a typical computing system and see how much of it represents some utilisation of memory as a function. Obviously a wide variety of technology, even in old-fashioned terms of disc, tape, high speed scratch pad, core memory and the working memory and so on. Nevertheless, even in the context of old-fashioned and orthodox computing systems memory was a fairly pervasive technology.

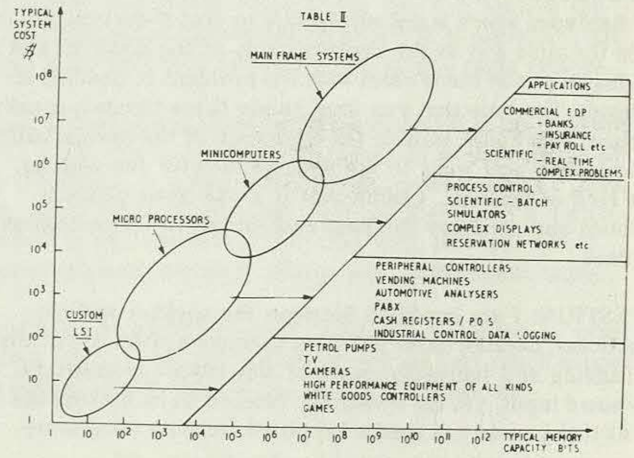

This also sets out to demonstrate another way in which I think that memory is a very significant technology. This is endeavouring to show, in a fairly simple-minded way, the relationships between such things as mainframe computing systems, minis, microprocessors, and custom LSI; which in a sense are all different ways of implementing programmable or pre-programmed projects. The vertical axis is typical system costs, going up into the 100 megabyte region and down in the $10 region for custom LSI. The horizontal axis shows the typical memory capacity, going down from no bits up to the order of 10^{12} bits of storage associated with the larger systems. Then on the righthand side we can see the various applications listed.

This illustrates again that to a large extent, when talking about differences, for example, between minicomputers and microcomputers, one of the first order differences between them in terms of the class of application that they can serve is the size of memory. They will be physically small and cheap in both cases, but in so far as there are distinctions — and they are getting fewer — the memory capacity which is associated with the machine is a significant parameter.

I will be talking this morning in such a way as to try to deal with three questions. The first is why memory is important; then I want to talk a little bit about some of the things that are happening in memory technology; and finally, we will concern ourselves with some of the future implications of these newer memory technologies. So maybe I should first define what I mean by "old" and "new" technology.

Memory Types

'OLD' TECHNOLOGY		'NEW' TECHNOLOGY
magnetic drum		semiconductor
magnetic core		CCD
magnetic tape	– serial	magnetic bubble
	– transverse	optical
magnetic disc	– fixed head	
	– moving head	
	– floppy	

By old technology I am thinking essentially of magnetically dominated technology; magnetic drums; core; tape; disc; fixed head, moving head and so on. The new technology: semiconductor; charge coupled devices; magnetic bubbles; and optical. They will be some of the things that I will be talking about. But before doing that, there is another general aspect of memory about which I should like to say a few words.

Impact Of Technology On Memory Economics

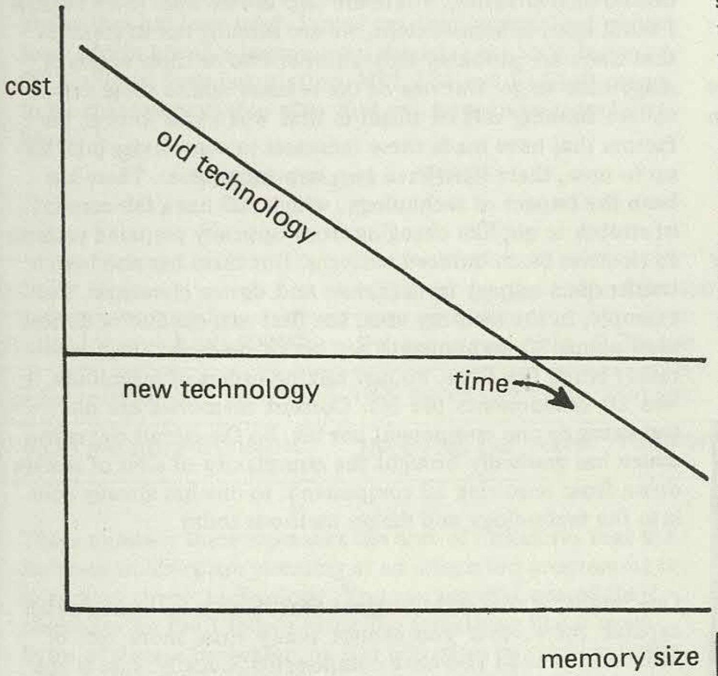

One of the interesting characteristics of old technology as compared with new technology is that if you look at the cost per bit as a function of the memory size, it is a fairly common characteristic that with the old technology the cost per bit falls dramatically as the size is increased, because there is a very large fixed or overhead element in the cost. That is a very significantly shaped curve because that, to a large extent, is one of the prime reasons why over the past 20 years computing systems have become fairly large, complex and centralised; because only by having one very large memory could you get the cost per bit down in size.

New technology, whether it be semiconductor, bubble, or CCD, has a much flatter curve, relating cost against memory size. This has several interesting implications. First, it means that you can economically start to break your total memory to divide a particular system up into smaller lumps, so that you can adopt the principle of distributed memory which is at the foundation of distributed processing and distributed intelligence. So it is the flatness of the new technology cost curve which has made distributed processing a reality rather than just something which would be very exciting if you could do it.

There is another significant thing which stems from the comparison of those two curves. First, it suggests that when you are introducing new technology you should start by feeding the new technology to applications that do not require very large memories; in other words, do not go in at the right-hand half of that diagram where from day 1 you are losing money, but start at the lefthand area where from day 1 even the new technology, before you are very far down the learning curve, is able to be cost competitive with the well established technology. Then as time goes on, the cross-over point there will gradually move to the right, and down.

That same comparison on those two curves, incidentally, supports the point that Mr. Manley made when he was comparing display technologies, because CRT shows the old technology curve; it is very cheap if you want a thousand characters but a bit expensive for six. On the other hand, the new technology, including LED, the variety of plasma panels, liquid crystal and so on, exhibits a much flatter cost relationship; hence the observations previously that the new technology is relevant for smaller displays and that the old technology will take a lot of beating for big ones.

I think that this is a very important characteristic of the introduction of new memory technologies. It tells you on the one hand how you should introduce new technology. But it does something else as well: it also suggests that frequently new technology will not simply penetrate existing markets but also create new applications. That is why, to some extent, in looking at new market opportunities for memory, I am not too put off by Mr. Manley's comments, with which I agreed, that electronic memory will not automatically supersede paper. I will not. Luckily, it will supersede some paper; but what is more important, it will create a whole area of new applications and new demands for memory in its own right. To a very large extent, the new markets will be created by the lefthand half of that curve in the first instance, in my view.

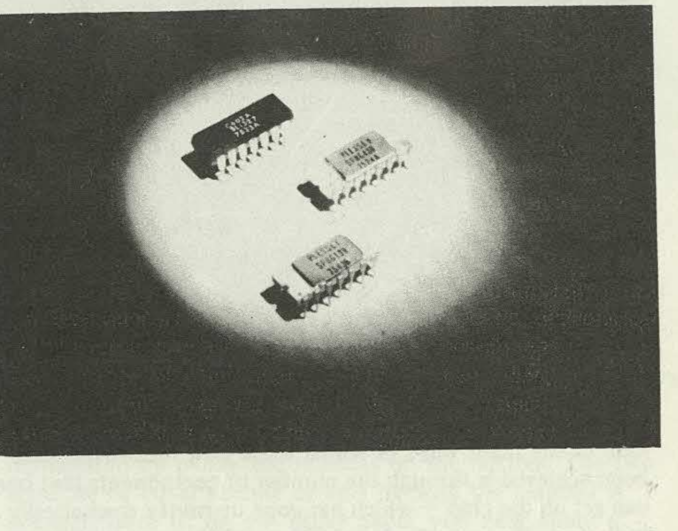

As has been said before, a lot of things that are happening are made feasible by the silicon integrated circuit industry, the industry of making nasty little things like that. If any of you have been lucky enough not to see inside one, that is what

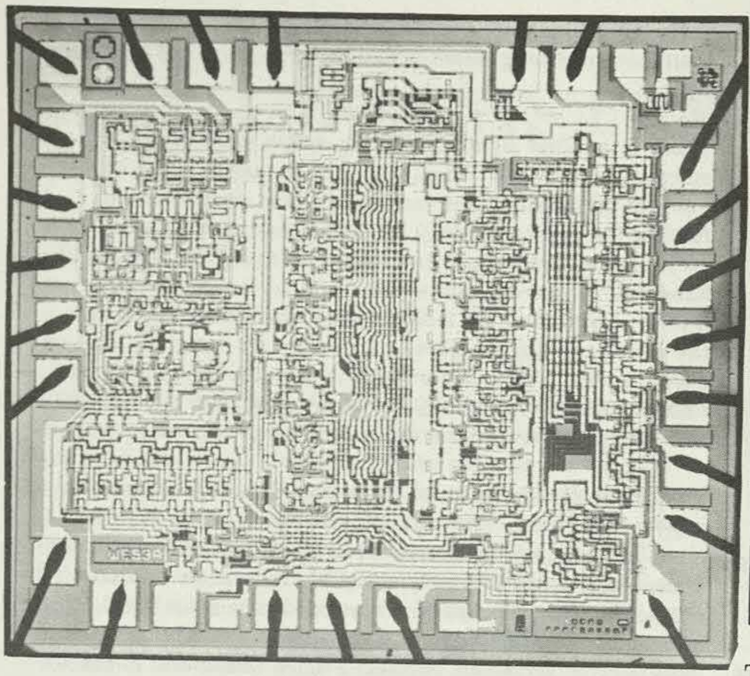

they look like. It is all based on the strength of batch fabrication or on slices of silicon which not many years ago were one inch diameter, and now we are talking about four and five inches. That has led to the ludicrous situation, as seen through the eyes of somebody who is in the semiconductor business, that, for example, the price per gate — using a logic gate as an indicator — has dropped by something like five orders of magnitude over the past 20 years. Again, as has been suggested, there is no reason to believe that that will not carry on for the next few years.

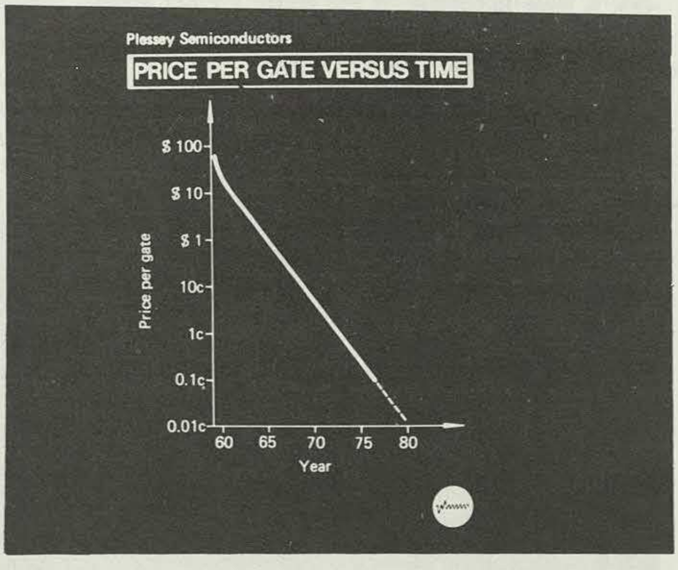

One of the main ways in which these cost reductions have been achieved is through the number of components that one can get on the chip — which has gone up pretty dramatically with time. If you go back to 1959 or 1960, one transistor per chip was what everybody could make; and then simple, multi-input gates came on the scene. Now we are at the stage where 64K bit memories, CCD memories, are available. Again that is a trend which will inevitably continue.

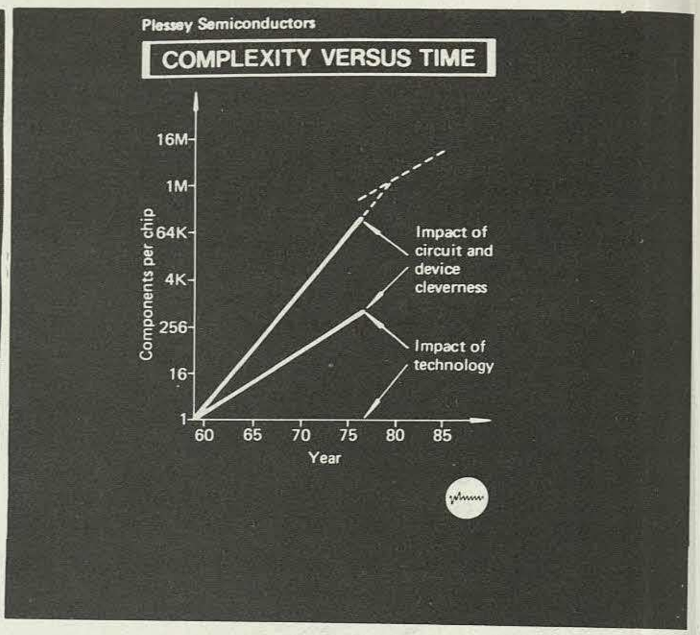

The upper curve suggests what has been going on, and the dotted extrapolation. There are two dotted lines there because I think that, to some extent, we are running out of steam in that there are probably only another two or three orders of magnitude to go. But one of the reasons that to some extent we are running out of steam is that when you look at the factors that have made these increases in complexity possible up to now, there have been two important areas. There has been the impact of technology, which still has a fair amount of stretch to go; like changing from optically prepared patterns to electron beam induced patterns. But there has also been a tremendous impact from circuit and device cleverness. For example, in the memory area, the first semiconductor devices used about 10 components per bit of memory; they were rather crude flip flops. So just talking orders of magnitude, it was 10 components per bit. Current memories are one transistor or one component per bit. So the circuit cleverness which has gradually brought the complexity of a bit of storage down from requiring 10 components to one has already gone into the technology and design methods today.

One might, at first sight, suggest that there is no more stretch capable there, that you cannot really store more bits of information than you have components. Luckily, that is not true. There are indeed ideas around, although I think that they are still for the future, of multi-level storage whereby you introduce a degree of analogue thinking into the storage so that you can indeed store several levels of logic on a single component — a single capacitor or a single MOS transistor. But having said that, I think that a lot of the circuit and system ingenuity has already happened and I suspect that there will not be quite as much scope in future. So I think that inevitably there will be some degree of tailing off. But the level at which it is tailing off is at the level of the order of one to 10 million components per chip.

The other number that I should like to put out at this time is that I have a theory that it does not matter what the complexity of the semiconductor device is, once it is made in reasonable volume it homes on $2 as a selling price. It does not matter what it is. It was true with 256 bit memories; it is true of 1K and 4K; and it will be just as true — it is only a question of time — of 256 Kilobit memories and 10 megabit memories; they will all come down to $2. That is the cost of the package, the gold wire and so on; the silicon chip tends to zero.

SEMICONDUCTOR MEMORY PROGRESS

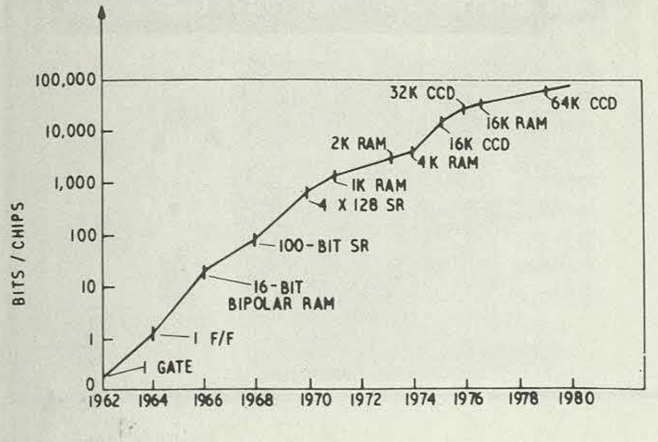

Just to illustrate that same tendency by putting some of these examples in, again you can see that it is not just a theoretical curve. That shows what has happened in terms of the introduction of new products with time. So already one is quite close to 100,000 bits per chip, and this is going to change. One of the ways in which it will change is demonstrated by the Japanese programme on VLSI. (Again this is a phrase that has been used. I must say that personally I object to all of this horrible jargon, even though I use it, in terms of SSI for Small Scale Integration, MSI, LSI and VLSI. It seems to me that the next step after that will have to be something like FLSI!)

VSLSI targets – 1980

bipolar logic	100K gates per chip	10n sec
bipolar logic – high speed	2K gates per chip	0.3n sec
bipolar memory	16K gates per chip	10n sec
N-MOS memory-dynamic	2M bytes per wafer	150n sec

Those numbers there represent the sort of objectives that the Japanese industry are pursuing as an integrated programme on integrated circuit technology. You can see that one of their objectives for the 1980s is achieving something like 2 megabytes of storage per wafer, or 100,000 gates per chip in terms of logic. But I think that there are two other things that are worth saying which I think are rather frightening about this programme. First, I think that they will probably achieve it. Secondly, the customers of the semiconductor industry also believe that their semiconductor industry will achieve it, and so they are already developing their system thinking now so that it will be able to incorporate these improved levels of technology when they are available.

It is also interesting to see why the Japanese are doing it. There has been a lot of adverse comment recently in the USA from the US semiconductor industry, which even though it has been so very successful, high-growth, virile and all the rest of it, is now getting very protectionist because they are very concerned about the Japanese competition. In response, the general manager of the US branch of one of the Japanese semiconductor companies said, in a rather defensive way, "I don't understand what all the fuss is about on this VLSI programme. We're not pursuing this enormous research programme . . ." — incidentally they are spending about $1,000 million over the next four years — "in order to make life difficult for the US semiconductor industry, we're doing it to screw IBM."

I think that this is what it is all about. Here we have an example where this technology is being invested in at a very high level, and it is being done because it is recognised that system economics and system design is being done at the level of the silicon chip.

I was interested in an observation that David Butler made earlier. He implied that the changes in technology and the fact that all this silicon is getting so much cheaper have led to the fact that when you look around the electronic manufacturing industry, at the level of building equipment, problems of yield and reliability in production costs have disappeared. They have; they are all replaced by the same problems in the hands of the semiconductor industry. I think that the problems are still there; it is just that there is a shift in the balance. The people who make these silicon chips now have the quality problems, the reliability problems, the yield and the cost problems. They have not disappeared. But I think that the fact that those problems are now predominantly within the semiconductor manufacturer's house rather than the user's is just symptomatic of the fact that increasingly design and system performance is being determined at the level of the silicon chip. I think that this is a very important sign of things to come; and it is one of the reasons why there has to be concern about the Japanese programme; because incidentally I do not share David Butler's optimism about the buoyancy of the American calculator companies in fighting back. I think that they fought back in market share by buying market share at the expense of horrendous losses, and I do not think that is the basis for a healthy future.

There are other implications of this Japanese VLSI programme. I think that a peripheral thing that will emerge from this will be CCD camera tubes of the kind to which Mr. Manley referred, not as rather esoteric replacements of existing plumbicons, but cheap and nasty enough to make life difficult for Kodak in the domestic market for home movies. That will just be regarded as a spin-off from this programme.

DEMAGNETISED CRYSTAL

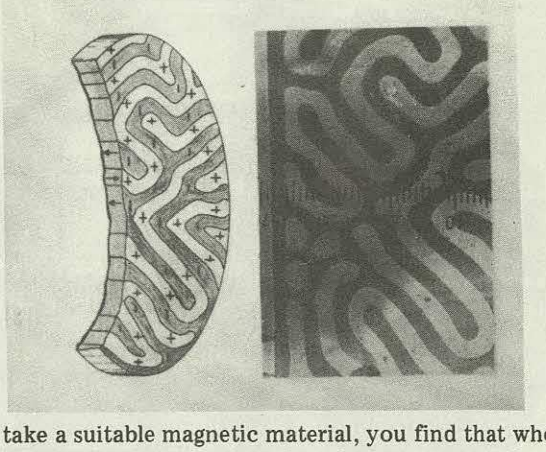

If you take a suitable magnetic material, you find that when it is in a demagnetised state, when there is no net magnetisation, if you find a way of looking at it with polarised light, you can identify the fact that adjacent regions of the material are magnetised with the North Pole on the top or the South Pole on the top, and when it is demagnetised North and South are represented by pluses and minuses on the bottom. The two areas are of equal volume or area because the material is demagnetised. If, on the other hand, you now apply a permanent magnetic bias to that, you make one of these regions grow at the expense of the other.

The beauty about these particular magnetic materials is, first, that there is a significantly wide range of applied magnetic field where these quite small, cylindrical magnetic domains are the stable configuration, and also in which these cylindrical domains will move freely if you apply a magnetic field within the plane of the material. If you have stray fields around, they just move around like soap bubbles; hence the concept of magnetic bubbles.

They are very interesting in that form, but they are not particularly useful. What you have to do is to find a way of impressing information or data on to such a pattern. That is done by applying to the surface of your magnetic material what was originally thought of as a T-bar pattern. You have a magnetic metal film, Permalloy or thin film; and you have a rotating magnetic field so that you can induce changes in the magnetisation of that Permalloy so that a North Pole, for example, will gradually work its way through the pattern. So as the North Pole winds its way through the pattern, a magnetic bubble will follow it through. You can actually form a bubble, make it go into the pattern, rupture it and start again, so that you can actually control the generation of bubbles into the T-bar pattern.

Essentially, the way to think of these magnetic bubble devices is that they are a little bit like magnetic tape memories; but instead of moving the tape relative to the write and read station you do not move anything around except the state of magnetisation in the material. The magnetic vector is moved in the material, but there are no mechanical moving parts. The sort of realisation that one has to achieve in practice is that you make these very fine patterns, and again these bars and T-shapes in a Permalloy film which has been evaporated on to the magnetic garnets. It has to be photo-engraved, using a lot of technology which is common to that of the silicon integrated circuit business, which is one reason why bubble technology has made reasonably rapid progress over the last two or three years, because it has not had to invent every technique for itself.

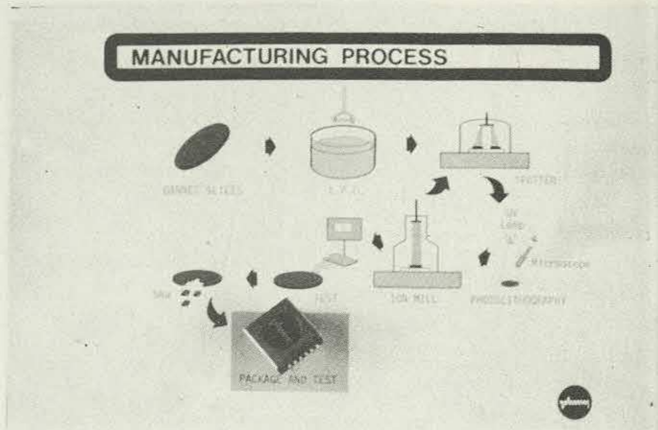

This is just to illustrate the fact that the same sort of batch fabrication methods are used with these: two, and soon three inch slices of garnet in which you can fabricate a fair number of these memory devices side by side on a batch fabrication basis (and then you chop them up and throw most of them away later).

The manufacturing process is very simple when you draw it like that. You start with garnet slices. You go through a phase of so-called liquid phase epitaxi. You spot a metal on to it. You go through a photolithography stage. You machine that, using ion milling. You test the slice; you break it up; and you mount the chips in the package at the bottom. Several of the processes that are used there are very similar to those using silicon. Some of them are uniquely different, but overall I think that there is enough in common with many other aspects of semiconductor technology for this to be not too great a burden to get out of the laboratory into a manufacturing situation.

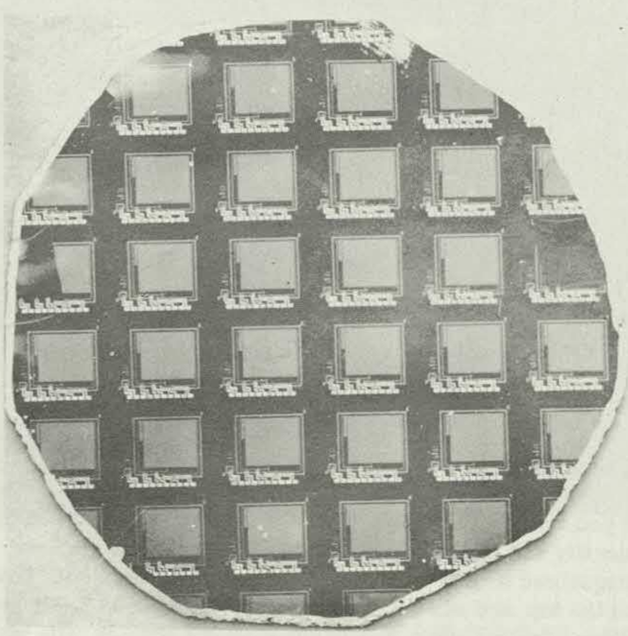

That is a singularly useless picture, but it is a problem in resolving the full structure. That is a chip with 64,000 bits of serial, non-volatile, bubble shift register on it.

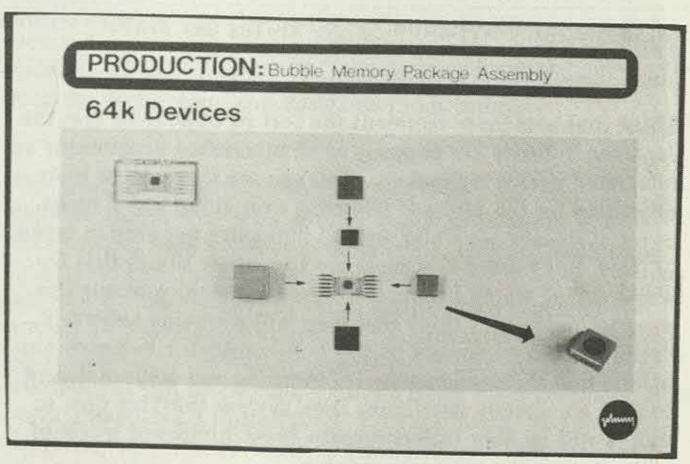

That shows how a device is assembled. The bit in the middle there is the bubble chip. It is mounted on a frame just like a silicon chip. That is then assembled in a dualled in-line package. It has the addition of one or two extra bits. The black shapes are the permanent magnets that provide the magnetic bias field to make sure that the cylindrical bubbles are the stable configuration, and there is a pair of orthogonal drive coils to provide the rotating magnetic vector, so that when you switch power on to those two coils you actually move your data through the bubble; and when you switch power off the data stays there, but it stays fixed in space wherever it happens to be at the time.

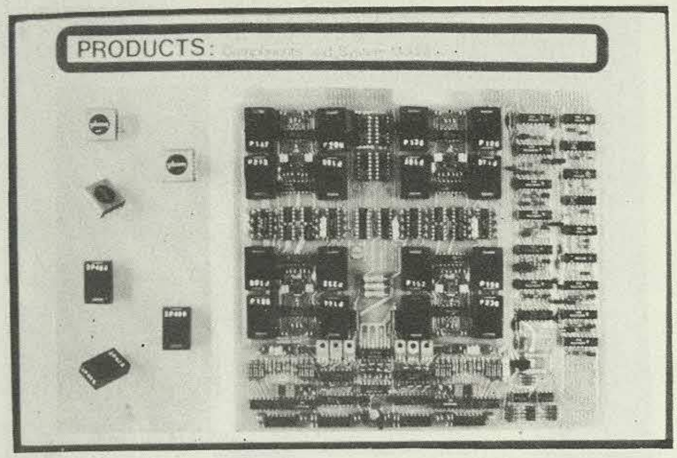

Those devices can then be put together on a normal printed circuit card, together with ICs for the provisioning of power supplies, the drive signals, the write signals, sense amplifiers, and construction of standard logic levels to go in and out. So, as far as the user is concerned, if he has a system card of that kind, all he needs to know is if, for example, it has 16 packages, each of 64K bits, he has a megabit that can be organised in a variety of ways of non-volatile storage.

I think that the easiest way of seeing the attraction of this kind of thing in new applications is not to look at the cost per bit; the cost per bit, certainly in the early days of this technology, will be rather higher than floppy discs and fixed head discs, but it is ideal for the kind of application that does not want even a floppy. There is no way that you can buy 10% of a floppy, so the cost per bit goes up by a factor of 10 if you are using only 10% of the capacity. The point about this technology is that you can tailor the capacity and hence the cost to the particular problem that you have in mind. So with today's technology, something of the order of 50 to 100,000 bits in one package assembled on a system card to provide a megabit represents the state of the art that many people round the world can now do. Again, as Mr. Manley mentioned, it is clear that within two or three years a million bits in one of those packages will definitely be feasible and commercially available.

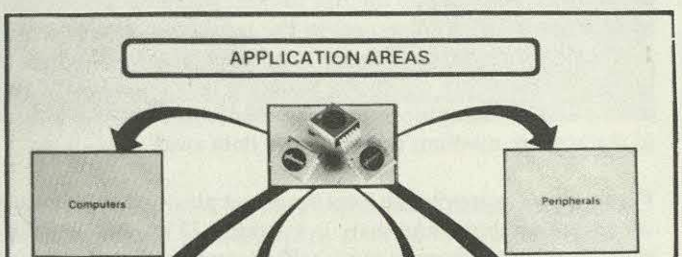

This shows the sort of application areas for magnetic bubbles — you just list everything that you can think of. I think that these are fairly justifiable.

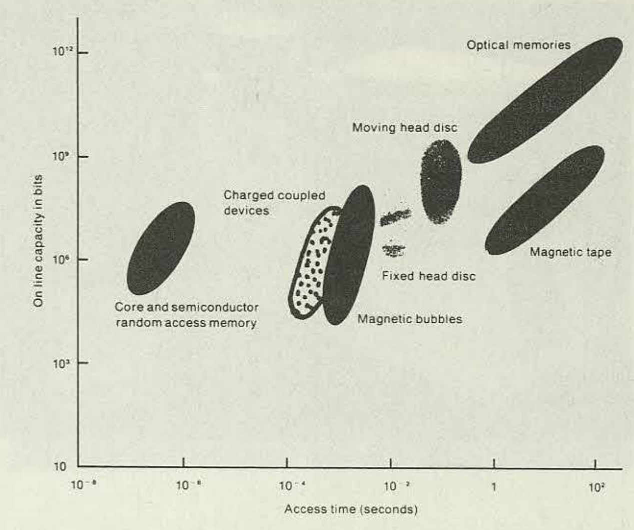

The impact of new technology concentrates on the lefthand side of that diagram up to now, where this is looking at the capacity versus access time for core and semiconductor random access memories, for magnetic bubbles and charge coupled devices coming in between the random access memory and the fixed head discs in terms of size and access time, but they do very little for the large memory long access time technology because they are too expensive.

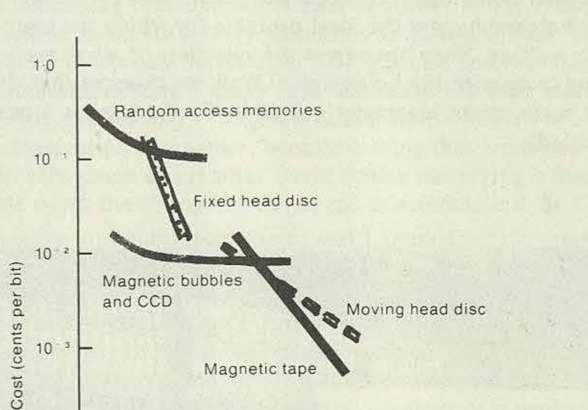

This shows the cost in cents per bit vertically against the on-line capacity. You can see that we have been talking really about technology up to now which is at the 10^{-3}, 10^{-2}, 10^{-1} cents per bit level. But the next technology that I am going to say a few words about is that of optical storage where we believe that it is feasible to get down in cost by another two or three orders of magnitude, but only if you are doing it at the level of, say, 10^{12} bits. There is no way that you can get one bit of optical memory for 10^{-5} of a cent, unfortunately.

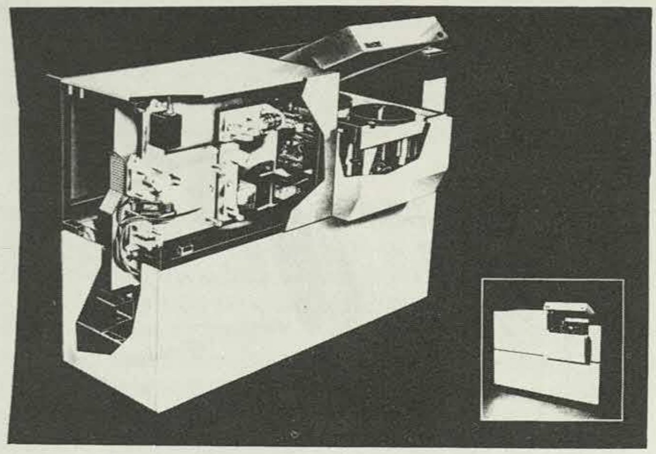

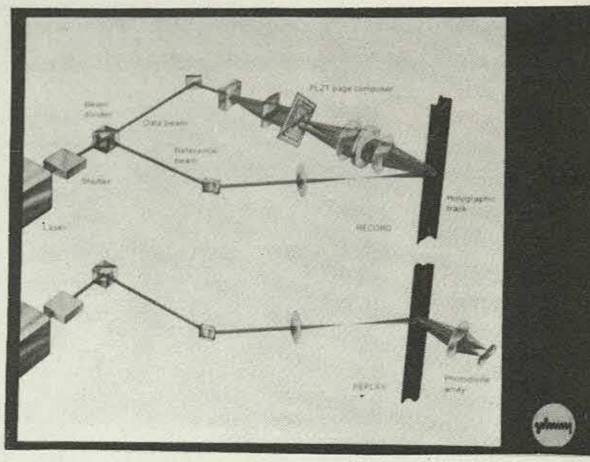

Optical memory. First, just to show you the hardware; it is roughly a metre on a side, just to give you a feel. The bit on the right is an orthodox disc drive, except that instead of using magnetic tape it uses optically sensitive tape, either photographic film if you want to make a read only memory, or a photochromic material that enables one to erase and re-write the information.

There is a laser around the back. In fact, that is really why optical memories have happened. For years and years, the laser has been the solution for which nobody had a problem; and then holography came on the scene and it was obvious that holography was the ideal problem for which the laser had been waiting. Then there was the question of what are you going to do with the holography? Well, we thought that we had better make memories out of it. So the laser is now justified!

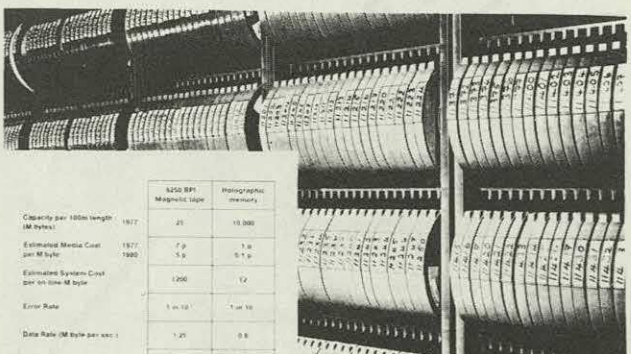

		6250 BPI Magnetic tape	Holographic memory
Capacity per 100m length (M bytes)	1977	25	10,000
Estimated Media Cost per M byte	1977 1980	7 p 5 p	1 p 0.1 p
Estimated System Cost per on-line M byte		£200	12
Error Rate		1 in 10 [illegible]	1 in 10 [illegible]
Data Rate (M byte per sec.)		1.25	0.8
Copying		On-line process	Off-line contact printing
Read Tape Speed (inch per sec.)		200	3
Volume per M byte (cu. ft) (cu. m)		3 [illegible] 8 [illegible]	3 [illegible] 8 [illegible]

Just to give you a feel for the comparison, if you compare the capacity for a hundred metre length of mag tape and for the optical holographic memory, you can see that it wins by a factor of about 400. The media cost comes down. All the favourable parameters go up and the unfavourable ones come down, inevitably. It does look like a potentially very interesting storage technology, mainly because one can make use of the very high optical packing density.

Let us see how it is done. On the righthand side you have a moving film. You take the optical signal from your laser. You divide the beam into two. You impress data on to one beam, using an optical modulator array or a page composer. You have a reference beam so that you mix the two, and you form the hologram and expose that on a moving film. When you want to reconstruct it, you move the film through a rather simpler piece of apparatus in which much the same reference beam is transmitted through the film and it reconstructs the holographic information into binary information on the photodiode array.

The special optical tape that has been developed using organic photochromic materials enables you to get an acceptable sensitivity so that you do not need a megawatt laser to write information in; but neverthless the information, once written, appears to be stable for the order of a hundred years at room temperature and you can erase and re-write the information something like a few million times before there is any detectable fatigue.

The advantage of holography as opposed to purely optical storage is that basically it makes it easier to do; and if it was not easy to do we would not be able to do it. You take out a lot of the mechanical tolerances. Even though you are putting in information on a one micron packing density, you do not require one micron tolerances in the mechanics. Likewise in the optics, you do not require very expensive optics. But particularly you are not in the situation which you would be in in a non-holographic mode — of dust, scratches and defects in the storage medium taking all the data away.

For example, suppose we have a normal photographic image of the initial machine, and there is a scratch. If you reconstruct that image in the normal way, half of your machine is missing. But if the same picture of the original equipment was stored in holographic form and again with the same scratch out of it, you do not see the scratch in the reconstruction. You lose a little bit in terms of signal to noise ratio, but you do not actually lose any data.

You can take it a stage further. You can take most of the information away. Say there is only 10% of the data initially left by virtue of its having been damaged. You can still reconstruct a rather fuzzy image. You have degraded signal to noise ratio, but you have not actually taken out complete chunks of the data. We consider that to be a very significant factor for very large, secure information.

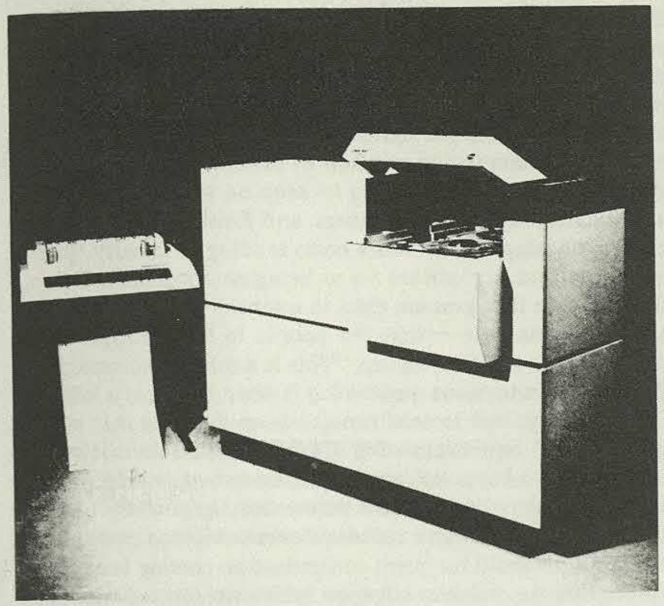

Finally, that shows a view of the general machine.

I should like to say a few words in drawing one or two of these things together. One of the first reasons why memory is important is that it was the development in silicon memory that led to the development of the microprocessor. There were two reasons. The first one was that the fabrication skills that were necessary to make complex memories made it technically feasible to make microprocessor chips. But I think that the other driving force was that the semiconductor industry needed microprocessors to help them to sell memories. One should think of microprocessors, at one level at least, as a marketing aid for semiconductor memories.

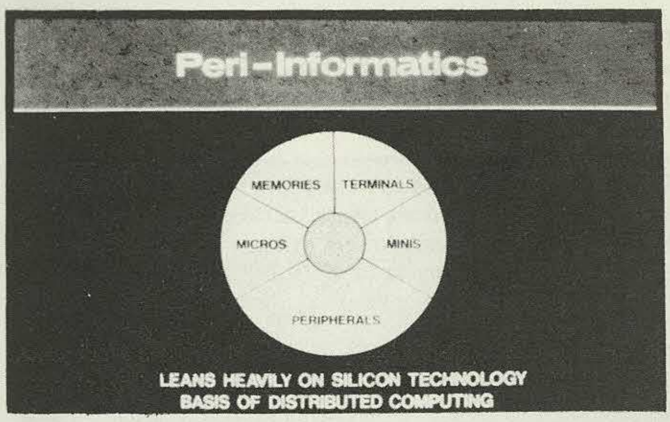

The French and the EEC have had some interesting thoughts, coming up with terrible words such as 'peri-informatics' describing that area of digital system technology that puts together the microprocessor, the minicomputer, memories, terminals and peripherals as a single competitive market place. I think that this is a very important area which in total is very much sitting at the centre of several of the things that will be talked about in the next two days. The whole of that leans very heavily on silicon technology, and it is that area of peri-informatics which lies at the centre of distributed computing.

Another interesting thing about this field of peri-informatics arises when you look at the competitive situation, and again I was interested particularly in one of the things that David Butler said this morning when he was comparing the characteristics of the data processing and data communication industries. The interesting thing is that when you look in this

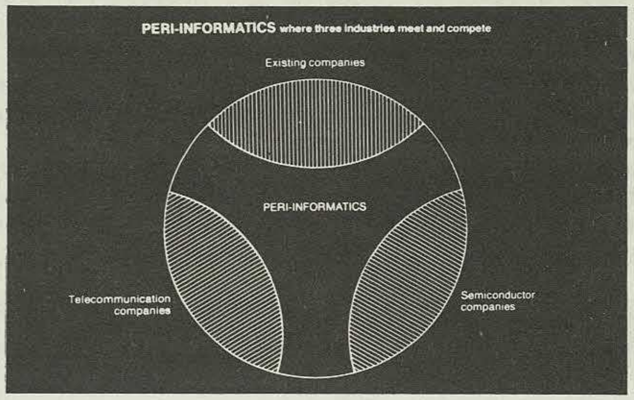

area of peri-informatics, it is an area where three major industries meet and will compete. There are the telecomms companies or the data communication companies; what I call the existing companies, by which I mean the existing companies in the definition of peri-informatics, the existing data processing companies of all sorts and sizes. And then there are the semiconductor companies because the same drive that has led several of the US semiconductor companies into vertical integration to make watches and calculators (and to lose money that way!) is also leading them into vertical integration in this area of peri-informatics.

It is interesting that companies such as Mostek, Zilog, National, Texas and so on are all very ambitious in moving up into this fringe area between the classical EDP business and the classical telecommunications. I think that will make life even more difficult, particularly — to give a biased view — for the telecommunication companies; because a thing that interested me again very much about what David Butler was saying is that he wrote down the characteristics of the telecomms and the EDP industries and contrasted them, and I think that the semiconductor company characteristics line up very much with the EDP companies, only more so. The other point which to my mind makes life difficult for the telecommunications companies *per se* is that the changes in technology that we are talking about impacting on the market — changing the classical methods of structuring the industry, creating new market opportunities and so on — the time scale in which these new technologies are emerging and need to be faced up to is obviously consistent with the time scale of the semiconductor companies, because that is where it comes from; it is not too inconsistent with the time scales that the EDP industry has become accustomed to; but it is totally inconsistent with the time scale that the telecommunications industry, up to now, has ever been accustomed to. So I think that these things, coming together, will create some major problems.

I am sure that it is true, as Mr. Manley was saying, that one can get so excited with technology that one loses sight of the fact that there can be other factors that restrain the development and creation of a market. There is more to life than just technology. I am sure that is true and that many of the other aspects that he was talking about will indeed inhibit some of the potential application of this technology. But it certainly will not inhibit all of it, and I think that many of the social problems and the problems of change relate to where new technology is coming in to supersede existing technology, but they do not relate quite so much once the new technology is coming in to create completely new market opportunities. I think that new markets can be created more rapidly than you can replace old technology and existing markets, because there is not quite such a well-established old guard to fight it off.

Let me just say another word about Japan because I feel that we have to recognise that the Japanese threat is a real one. I think that there are two things that one should say about the Japanese situation. One is that we certainly should not under-estimate their power, their intention, their dedication and their overall strength in terms of technology, investment and everything; but at the same time, it is a battle that we cannot afford to lose. In the eyes of many people the battle has already been lost. I think that we have to resist that attitude.

COX: Question time, gentlemen. As they used to say on the radio when I was a little younger, "Five pounds here for anyone who can stop the memory man."

QUESTION: Could you give us any idea of likely costs of holographic memories?

ROBERTS: I am very anxious not to give you a single hard number, because it is a fairly complex price. It is not just a simple component and it depends so much on what goes with it. If I can just put it in perspective by saying that the kind of box that I was showing has a capacity rather like the IBM 3850, where they have this mechanical monster running up and down the room pulling out cassettes on a random basis. I think that at around 10^{12} or 10^{13} bits, typical costs would be, say, $1 million or $2 million, talking orders of magnitude. For the raw hardware that we are talking about here, with very little in the way of frills, software support or control, but as a raw peripheral on an OEM sale basis, we could see our way on this, in reasonable volume, substantially below £100,000.

QUESTION: What, if any, is the difference between reader and writer costs?

ROBERTS: It is the same device. The diagram, just to make it simple, split it; but the hardware is the same piece of hardware. It is one of the options that is available for the kind of user who may wish to have, say, one facility, writing information and creating data, and then producing copies of it. One of the things that I forgot to mention is that one of the virtues of this holographic recording is that you can cheaply produce error-free copies of data by just using normal photo-graphic contact printing. So having gone through the expensive machine to write data on to your 100 metres, say, of film, you can then run that through a very simple machine and produce a hundred copies. Suppose it is a fingerprint file. You may well have one file where you create it, and then you might send 50 copies out into 50 localised centres in the UK; and you only want those centres to be able to read the information, you do not want them to be able to change it or input information. The read station then could be significantly cheaper than the sort of cost about which I was talking. The cost that I was talking about was for a complete read/write and modify unit. In volume, a read station could probably be £5,000.

QUESTION: Distributed microprocessors and memories such as you have described could lead to fundamental changes in the way that control is exercised within a business. But in reality the suppliers of this equipment will encourage central-isation. What are your comments?

ROBERTS: You may be right, but it is a very depressing prospect if so; because it seems to me that there are several implications of this technology that we have been talking about, some of which we have touched on already, in terms of the fact they they do, at least in principle, make distributed processing economically feasible. The other thing that this does, which again has been referred to earlier, is that it makes hardware costs tend to zero. I agree that these things cannot happen overnight, but unless we are going to re-think our overall approach to the solution of problems at the level of system design, if we are going to keep on with the old habits and the old ways of doing things, and finish up with hardware costing nothing and software costs tending to infinity, there is no way that that strikes me as being an optimum solution. So I am sure that you are right in a sense — and I also see the same characteristic — that the people in the microprocessor business who start by saying, "This is a microprocessor, It's $15. Distribute your processing," then go on to say, "Incidentally, in 10 years' time, also on the same chip we'll give you the equivalent of a 3750," and you are back in centralisation but smaller hardware and cheaper. I think that there is a logical disconnect there, and I personally believe that there needs to be a more systematic view to new system thinking, to avoid incipient centralisation coming back in, to reconsider the balance between hardware and software; to look at it in terms of more not just distributed processing, but more dedicated hardware and to be less intent on ongoing programmability. It seems to me that programmability when you first make it, and dedication thereafter, is a more intelligent way of using the transient technology.

If we do not, but simply use the technology to make a large, central processor that costs nothing, and continue to incur all of the other costs of running it, operating it and developing complex programs for us, that may well happen but I am convinced that that is not the right way to use the technology.

QUESTION: Another serious problem is not technology but accounting. There is a lot to be said for getting the high technology items out of the capital account.

ROBERTS: Yes, and certainly get it out of the capital account, and out of the hands of the people who have built up careers in the last 20 years doing centralisation. I think that those two things go together.

QUESTION: How are costs to be apportioned, above the line or below? Capital items are treated differently to operational expenses. It could be argued that the new economics of com-puting should lead to much of it not being treated as a capital expense. But how do we make the transition?

ROBERTS: It will happen. The people who will make it happen will be the semiconductor companies. It was the semiconductor companies that made the digital watch happen. Whether or not you want it does not matter: it is there. Likewise the pocket calculator; and likewise this area of specialised programmable but dedicated hardware as providing a 20th century, or maybe a 21st century alternative to time sharing on a large, complex machine. That will be made to happen by the semiconductor industry, and they will create new markets; and penetrating and attacking the old-established ones is not the prime concern.

COX: You almost won a fiver then, Rex! On that note, we will break for lunch. I should like to close this morning by thanking Derek for that very entertaining and provocative session.

NOVEL OPTICAL COMPONENTS AND SYSTEMS

J. Evans

BUTLER: Inevitably, every two-day conference has two sessions that are known in the business as the "graveyard" sessions — the ones immediately following lunch. Part of the trick of conference organisation is to pick speakers for those sessions who are guaranteed to stave off the effects of imminent repose. You have to do that rather carefully.

A few weeks ago, I was lucky enough to hear our next speaker in action for the first time and decided that it was not only appropriate to ask him to speak at this conference, but also to give him the dubious privilege of occupying one of the two graveyard sessions. Joe Evans is the Director of the Materials and Components Laboratory at STL. He is going to talk to us about some of the fascinating things that are happening in the area of optics.

EVANS: I have a problem because people have been saying all sorts of nasty things about the telecommunications industry and, as they pay some of my fixed costs but not the variables that I am incurring today, I have to reply to that. We are a multi-national company which, in the case of STC, means that its top management is all Scottish, the middle management is all Welsh, and the factories are run equally by English and Pakistanis.

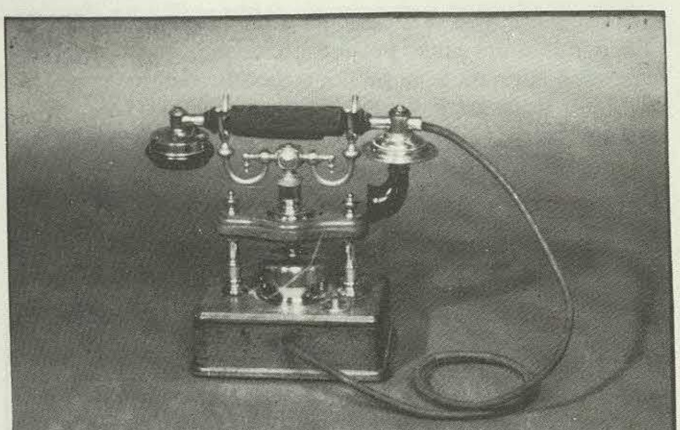

You may wonder why I am showing you all this. The reason is that Derek and I have been involved in the middle of this, and so has Brian Manley, and we sometimes forget the tremendous changes that have occurred in technology. Let us look at one or two reasons why this happened, to see what might happen in the future.

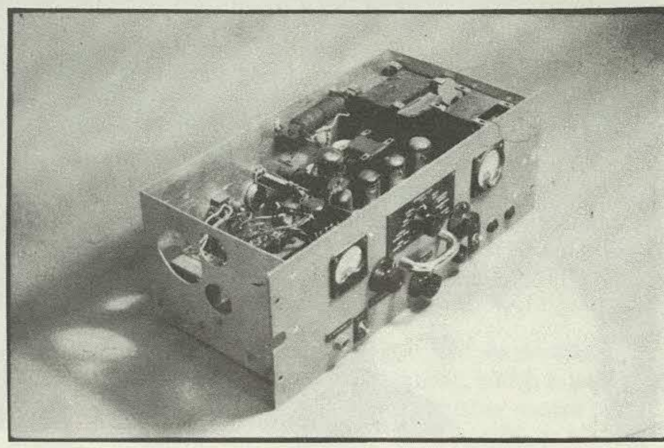

Some of you may remember thermionic valves. That is a chassis there with thermionic valves in. These things are heated, so they have to dissipate heat and therefore need a lot of space. They have power supplies, huge transformers and capacitors to drive them. Because they have a limited life, since they get hot and things wear out, they have to be plugged in and out, so they have to have sockets. So you end up with a tremendously heavy material-intensive technology to make tube circuits. The reason why we did not put electronics into telecommunications before now is because of the limitations of that tube. In fact the telephone in the slide above is exactly the same as the telephone set that most of us use today. The carbon microphone, the moving diaphram earpiece, and the electronics inside which is not electronics but electrics, have hardly changed at all over about 50 years until the transistor came along.

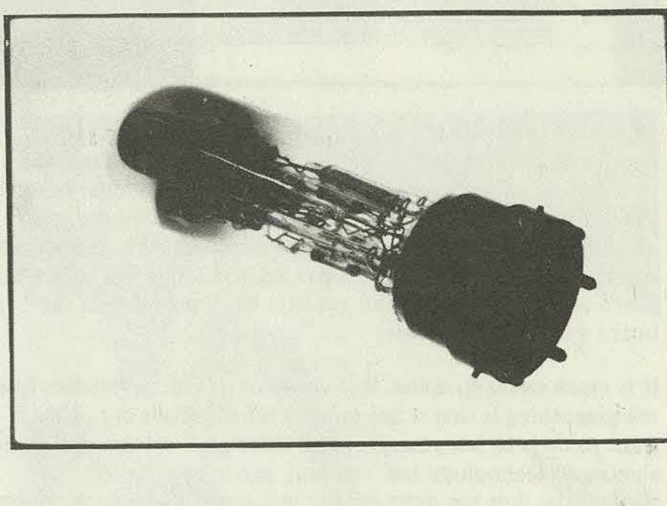

What happened was that people said, "We must get away from the tyranny of these valves. Let's try to integrate many functions into one valve." The idea of an integrated circuit is not new. This thing, like a lot of other important things in the

UK, was designed in 1925. I was designed in 1925; many of you were. That contains several thermionic active devices, and some resistors and capacitors inside one tube. It failed, not because the idea was not good but because the individual failure rates of the individual components in there were so high that most of the time something went wrong and you had to throw the whole tube away. So the idea of integrated circuits is not new.

What did we do about it? We came to the transistor, and then the integrated circuit; and you have seen some of that happen. What effect has this had on us? The effects are dramatic. This

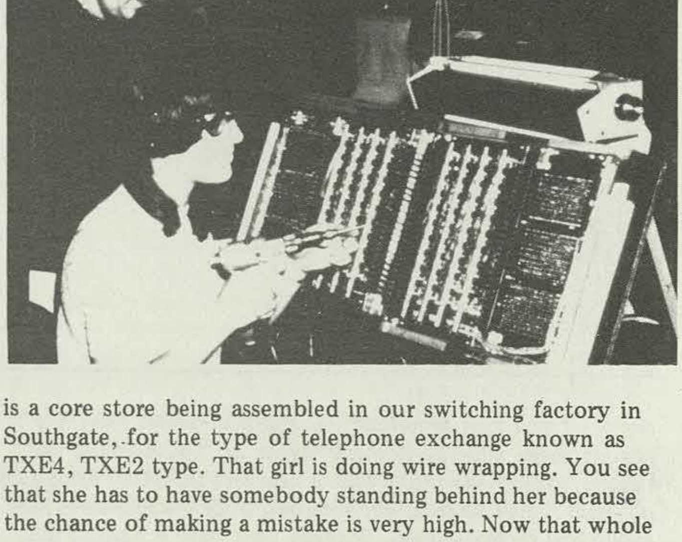

is a core store being assembled in our switching factory in Southgate, for the type of telephone exchange known as TXE4, TXE2 type. That girl is doing wire wrapping. You see that she has to have somebody standing behind her because the chance of making a mistake is very high. Now that whole

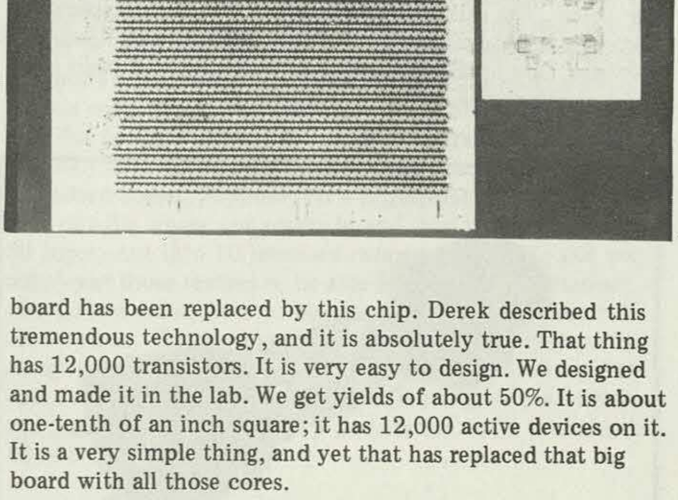

board has been replaced by this chip. Derek described this tremendous technology, and it is absolutely true. That thing has 12,000 transistors. It is very easy to design. We designed and made it in the lab. We get yields of about 50%. It is about one-tenth of an inch square; it has 12,000 active devices on it. It is a very simple thing, and yet that has replaced that big board with all those cores.

It is much easier to make. It is cheaper; it is more reliable. But the great thing is that it has taken a lot of labour out. What I want to do is to tell you, as we go through, how this change in electronic technology has not only given you various new possibilities, but has dramatically influenced the way in which we do business in telecommunications. It has taken the labour out.

We have big squabbles in ITT as to whether the component man should do something or the system man should do something, and these miss the whole point; which is that the semiconductor revolution has not just transferred the added value or the labour from one place to another, it has just taken it out. Therefore, the need for manpower is considerably reduced. That is a structural thing, not a temporary balance of payments or terms of trade problem, it is a permanent structural change.

What I am going to suggest today is that just as we have seen this terrific change in solid state technology due to the invention of the transistor and the integrated circuit, a similar thing might be happening in optics. I should like to take you through a few ideas which we have been developing to see how this might happen. The combination of optics plus micro-electronics will make a tremendous number which may influence your business.

First, we look at the optical spectrum. The window that we see, the visible range, is very limited. There is a tremendous spectrum of electromagnetic radiation. We have used some of the radio waves which are familiar to you. Optical gives us very high frequencies, and there is a big stretch of infra-red frequencies which we will use in future.

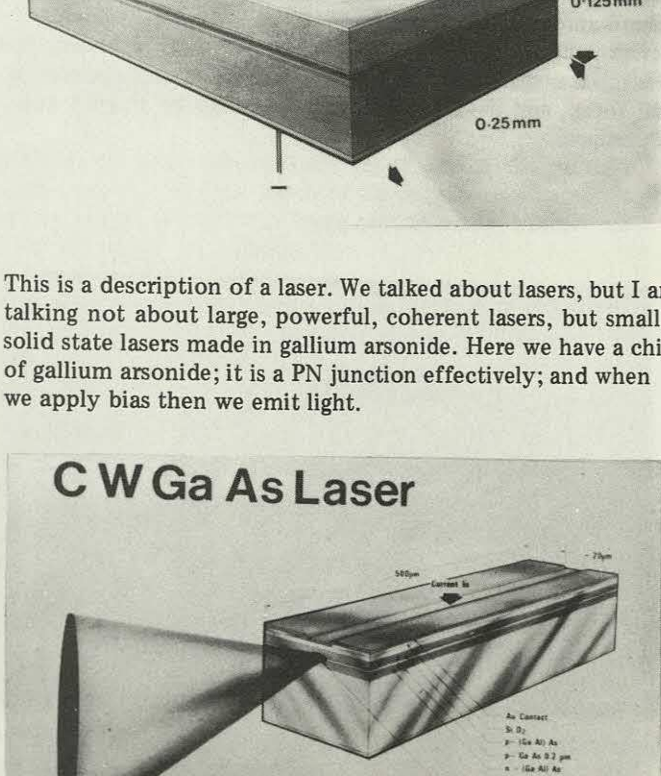

This is a description of a laser. We talked about lasers, but I am talking not about large, powerful, coherent lasers, but small solid state lasers made in gallium arsonide. Here we have a chip of gallium arsonide; it is a PN junction effectively; and when we apply bias then we emit light.

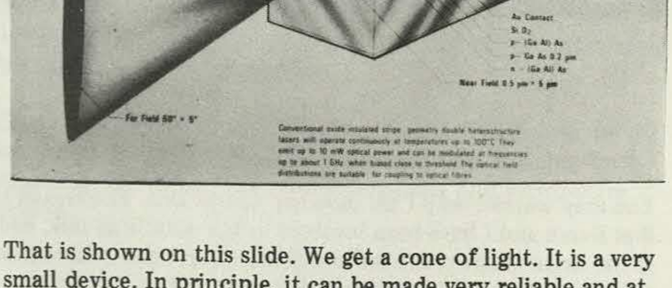

That is shown on this slide. We get a cone of light. It is a very small device. In principle, it can be made very reliable and at very low cost; and so we have a source of infra-red radiation which can be turned on and off by an electrical signal. I think that will lead to all sorts of new possibilities which we have only just begun to see.

Until recently, we could not make these things reliably, cheaply and consistently; but now these problems are being overcome, taking advantage of some of the developments in silicon technology.

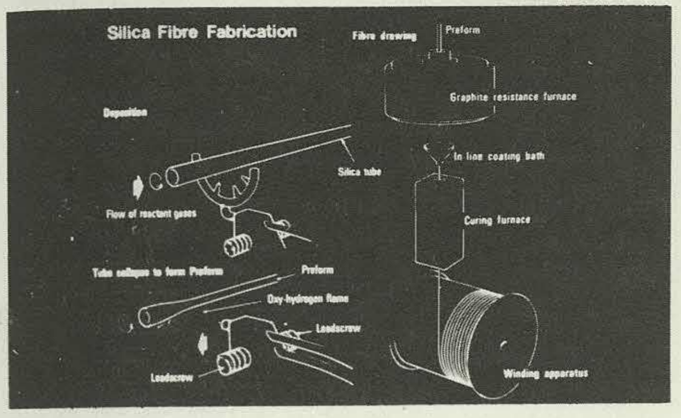

But to go with that, of course, we need a transmission medium; we need the fibre. We have been working, as have many other people, on two kinds: either a silica or a glass fibre. Pure silica seems to be the preferred method at the moment. We need to make a wave guide structure. We need to have some way of confining the light in a fibre; so we need to have a core of one type of refractive index glass, surrounded by a cladding with a lower refractive index so that we get total internal reflection. You will see that described in a film that I will show later.

So we have our tube which is got by means of gases passing through it. We have deposited the various layers that we want. We then collapse that tube into a rod. We put the rod, called a pre-form, in the top of the righthand diagram and we melt it and pull out a fibre which we wind on a drum. So we start with a metre or two metres of rod, and we end up with several kilometers of fine fibre which will carry light with a very low attenuation.

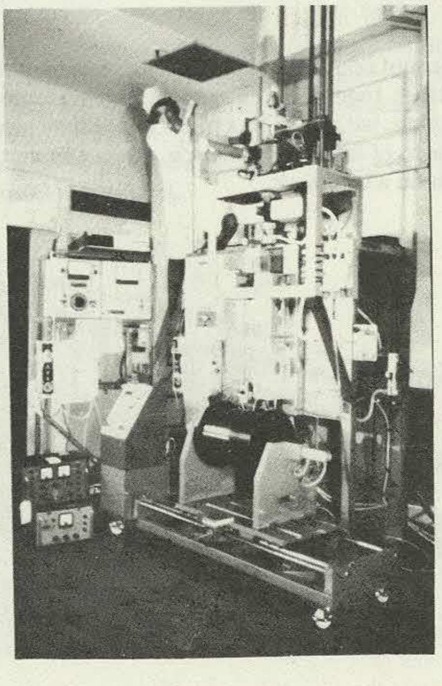

This shows the actual equipment. There is a girl standing at the top, which gives you the scale. She has a silica rod there. It is going into the furnace. It is being pulled; it is being coated with plastic immediately it falls in order to protect the surface, and then being wound on this drum. That drum, made of plastic so that it does not expand and put strain on the fibre, will have a couple of kilometres of low loss fibre.

Then we have to turn that into a cable; and we do that by passing it through a conventional extruder, which is traditional cable technology. Then, having got that fibre coated with plastic, we want several of them to make a cable, so we take

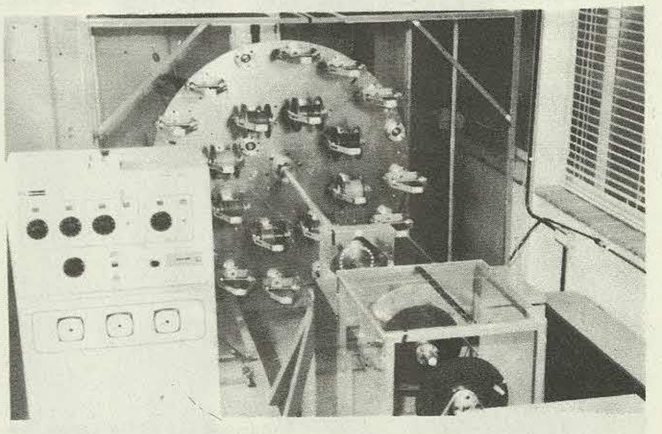

these bobbins. Each of these bobbins now has a plastic-coated silica fibre in a reel. We then twist those together in this machine, and we produce a cable.

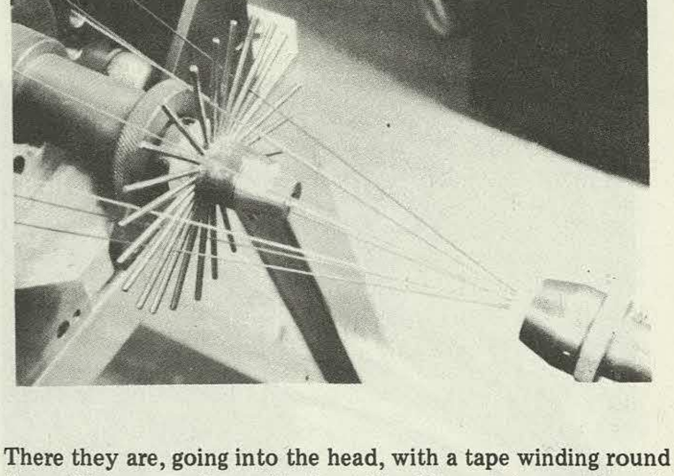

There they are, going into the head, with a tape winding round them. Then we end up with the cable. I have some samples of

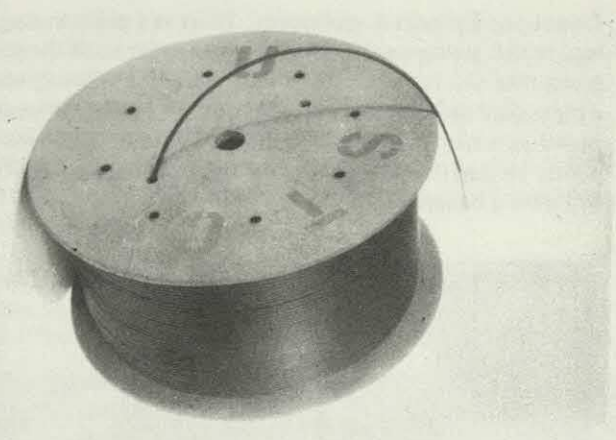

that which you might like to look at. What I am trying to get across here is that this cable has to be shown to be a reliable, rugged thing before people will start to use it. We have to establish credibility in the components that we will be using for optical systems before anybody takes any interest. So we have made several miles of this stuff and subjected it to all sorts of tortured treatment, some of which you will see in the movie.

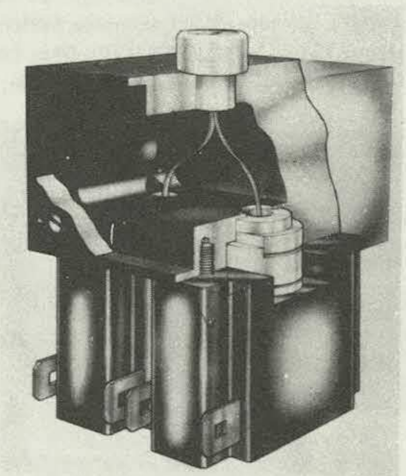

What had happened to the laser meanwhile? Of course, we really wanted to develop a laser for military applications, high power, infra-red beams for doing all sorts of ranging and illumination techniques. This shows how, if we want to increase the power, we can take two lasers and put a fibre on the face of each one, and combine the fibres at the top so that we have twice the power coming out.

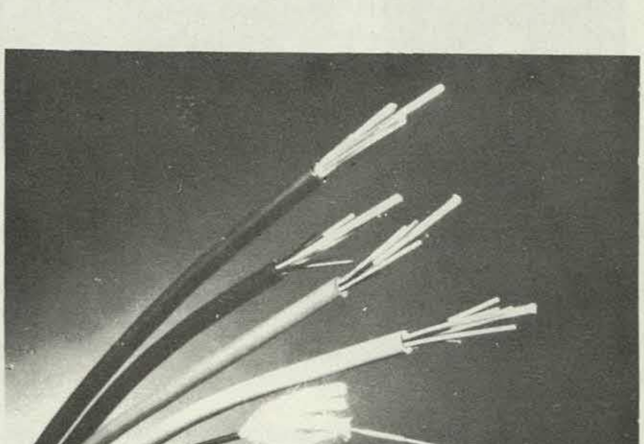

This is a selection of different types of optical cable, sometimes with only one fibre; sometimes with a number of fibres; sometimes with fibres of silica and copper wires; sometimes entirely plastic so that it cannot be detected, there is no metal in it; sometimes with a steel strength member; sometimes with a plastic strength member. There are all types of cables which one can use.

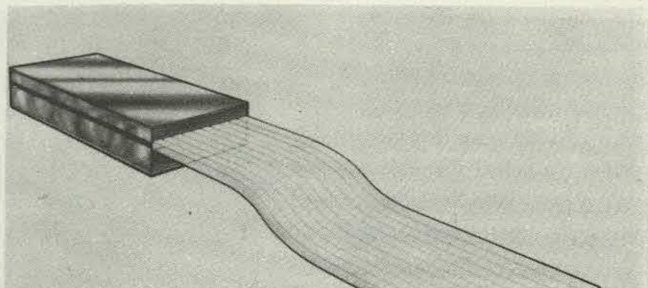

A technique has been developed in which a fibre optic tape is attached to the output facet of an unpackaged laser chip, accurately aligned with the active region.

Typically 90% of the laser output is emitted from the free end of a 2" length of fibre optic tape.

I was talking about combining the outputs of a number of lasers. The next thing to do is to get the whole face of that laser chip to emit and then, if we provide a number of square fibres, we can collect all that energy. We are beginning to use the fact that we can conduct light through fibres to do all sorts of things; we are going to combine those outputs.

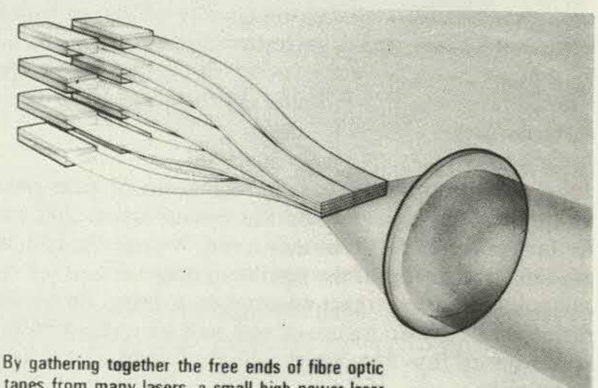

By gathering together the free ends of fibre optic tapes from many lasers, a small high power laser source can be produced suitable for optical collimation

Then we could take several of those lasers and all these fibres, and put them together, and make a sort of square emitting surface which, with suitable optics, could then generate a very intense beam of infra-red radiation. We can mount that sort of thing then on a heat sink.

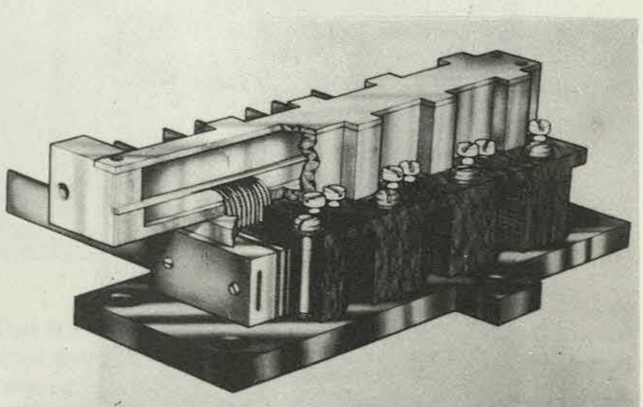

This is a traditional heat sink. The lasers are all in here, with individual fibres coming out; their output is all combined in that highway. We apply the bias through the junctions in series, take the line out in parallel, and we have ourselves quite a powerful source of infra-red radiation, which would burn a

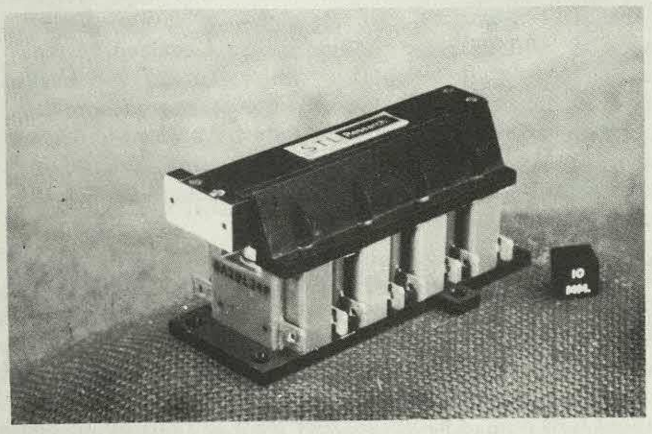

hole straight through your retina if you were so unfortunate as to look at it.

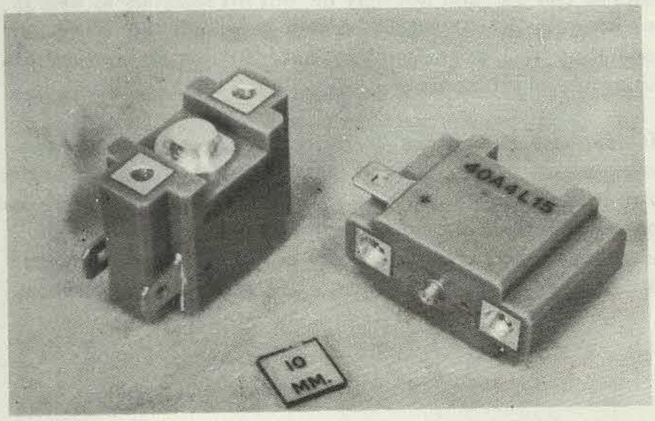

Here are two small modules showing how one would actually encapsulate this into a practical form. You do not see all these delicate fibres waving about, it looks like a solid, rugged thing to you.

We start with the single laser, then we build up modules of lasers. We have talked about having a long piece of fibre to carry laser light. But most of you are not concerned with very long distance transmission, you are interested in short range things. You want to put data into a network, you want to take it out, and do things with it.

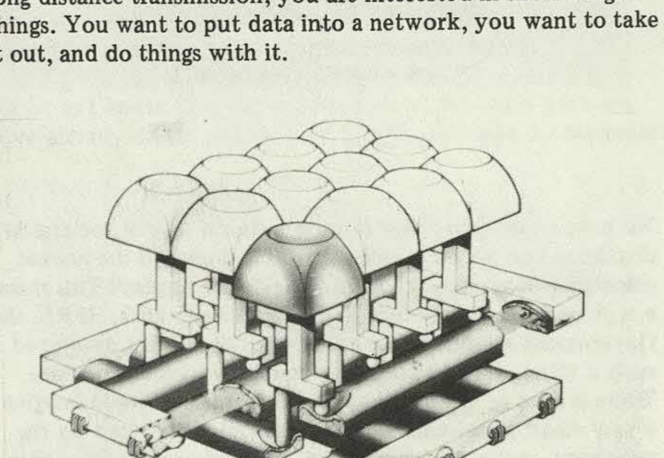

We thought that you might like to do that in the optical mode. This is a very simple idea. This keyboard has a bank of keys; and when you press a key what you do is not open contacts or squeeze a bit of PZT, or anything crude like that; you squash a rubber pipe through which light happens to be passing. Light is very squashable, unlike water; and when you squash it, it just stops.

What we have is a number of emitters, light emitting diodes which people like Derek have made very cheap, reliable and available. At the other end we have a photodetector; and so we have a matrix, if you like, of light paths. When you press the key, it will squash two of them; and then some clever electronics, which in this case is a microprocessor, will tell you which key you have pressed and, hey presto, you have a key block using optics.

What is the advantage of that? First, it has very long life because these things made of neoprene and so on are very strong. It has solid state emitters and detectors. It has microprocessors for the circuit work. There is nothing really to wear out or go wrong. But one other thing is that the signal comes out in optical form; and you can either transfer that back into electronics; or you could keep it in optical form if you wanted to, and send it along a piece of cable to a remote device. So you could keep your signal in optical form, avoid interference, pick up all sorts of problems. We are just beginning to find some of the advantages of that.

That is an early model, and the later one is so good and so clever, so ingenious and so simple, that they would not let me show you a slide of it.

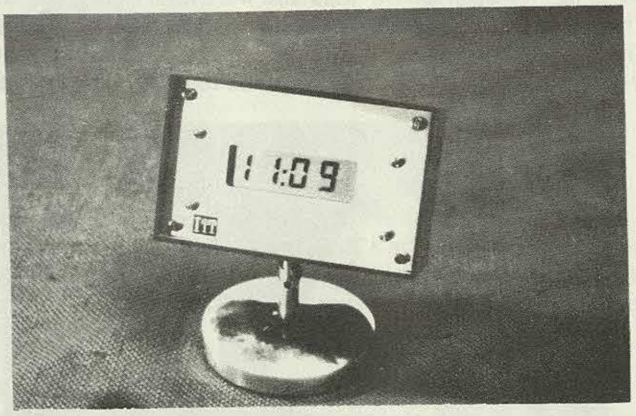

Let us go on to something else. This is a liquid crystal display, and it has the advantage that it is completely passive. This merely reflects light which is available, ambient light, like printing — the so-called Caxton display that some of you may know. You can also send light from behind and make it transparent.

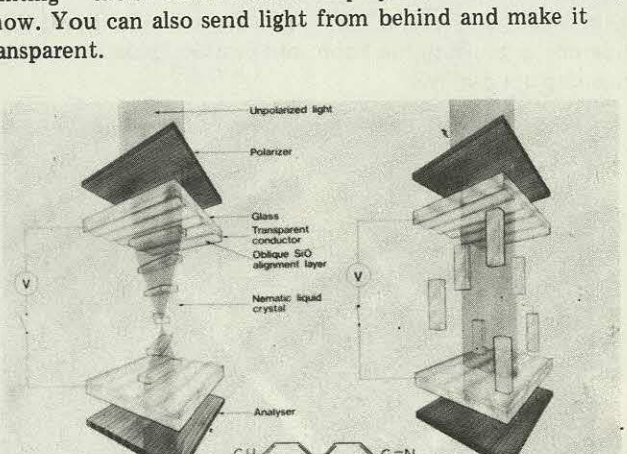

How does it work? This is a liquid crystal; and we arrange to send light which is unpolarised through a polariser. These liquid crystals — the particular twisted pneumatic type that we are using here — have the property that the molecules form this helical arrangement which rotates the plane of polarisation so that when it gets to an analyser which is set at right angles to the polariser, the light goes through. If I apply an electric field, just a field which takes very little current because it is a high impedance material, then I can align all those molecules; and the light, instead of now being rotated so that it can go through the analyser, is blocked. So I can either let the light go through or stop it, according to whether I do or do not apply a small electric field. This, of course, gives you the opportunity, with very little dissipation, to effect a display.

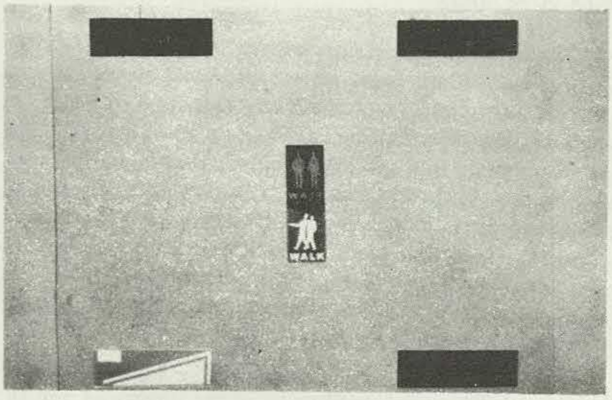

You can then light it from behind, and you can do all sorts of things. You can have a 'walk and wait' kind of symbol. If this is a fuel tank, you can make the thing chunter up here as the fuel fills the tank; and there are various other things which you can do. You are not only confined to characters, you can

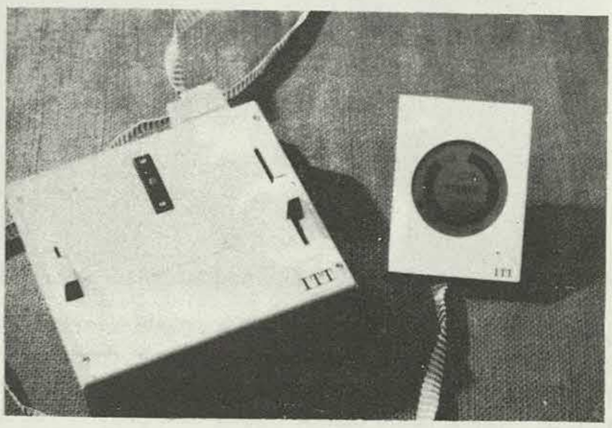

make analogue displays. This is actually a circular thing; it just unwinds as you turn the knob; and you can make some very nice displays like this.

This is an LCD Sharp pocket calculator which uses a liquid crystal display, very thin, recently designed in Japan. But it has another very interesting feature which I commend to you. They wanted to make an extremely thin calculator which you could put into a wallet, and of course the keys are one of the biggest problems. The battery has to be small, and you get that by having a liquid crystal display and low power circuitry. But the keys themselves are a problem. They always have a finite height; they are mechanical; they have to have a feel and so on. They decided to get over this problem by having a keyboard as a completely flat sheet with pressure operated keys. Another difficulty is to know whether you have actually pressed the key or not, whether it has actually taken the signal.

To check that, they have a little bleep. A musical note shows that. Every time you press a key it just gives a little bleep. That is a very nice bit of ergonomic design. When we go from semiconductor integrated circuits to things like display and other features, we have subjective influences. We must think about how human beings actually react to a display, to a push button when they push it. These are things which are very difficult to define objectively. So more and more we have to keep checking with the customer that this is the kind of thing that he can use; and not only he, but his secretary as well. If we are going to have lots of remote terminals with secretaries operating them as though they were typewriters, we have got to avoid the problems.

Incidentally, the optical keyboard is interesting because you can press two keys together, but you can arrange for it to give you the output of both keys, the first one first and the second one second, even though they were separated only by milli-seconds. So you can avoid some of the practical ergonomic problems of mechanical displays by using optical techniques.

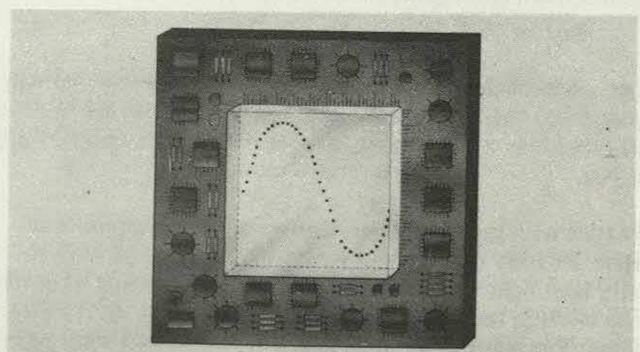

LIQUID CRYSTAL OSCILLOSCOPE DISPLAY WITH ON BOARD ELECTRONICS

We have shown that you can use a liquid crystal for digital display — you are all familiar with the clock and the pocket calculator. What about having an analogue display? This shows a type of CRT display that one would like to have. RSRE, the Government establishment at Malvern, has in fact developed such a thing, about two inches square, which is very nice. There is a lot of work still to be done. But you could imagine a very small panel with all the electronics mounted on the periphery, and a portable pocket type of oscilloscope, with limitations on performance which are being improved all the time. But that, to us, is a very attractive thing to envisage.

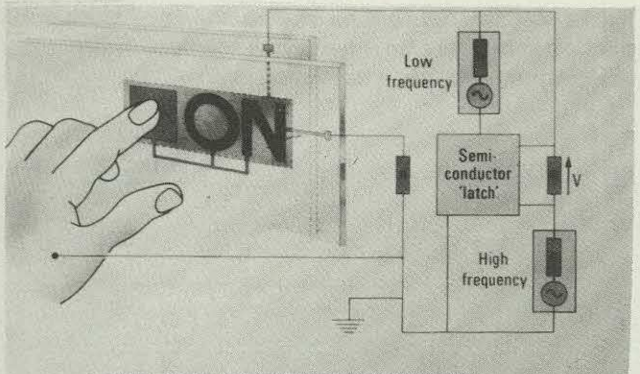

LIQUID CRYSTAL REACTIVE KEYBOARD

You could then go a stage further and say, "This liquid crystal panel that you have is a high impedance device. Couldn't we actually use the capacitance of a finger to switch it?" You can arrange that by having two circuits, a high frequency and a low frequency circuit, one to operate the display and the other to detect that you have touched it with your finger. You can make a reactive keyboard with liquid crystal so that, touching it, it will light up with the symbol that you have touched; and if you have made a mistake, you can immediately see that. We have begun to explore ways in which you could do that.

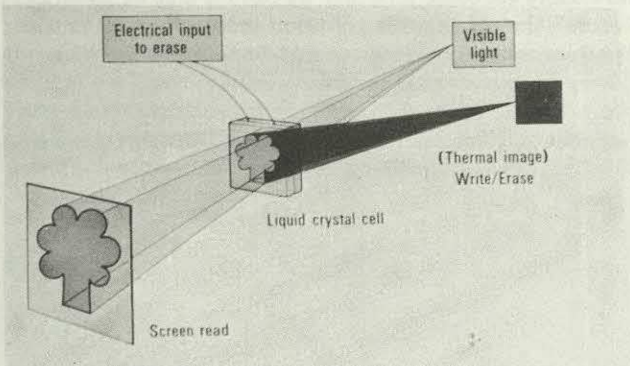

THERMAL IMAGE FOR NIGHT VISION

Another possibility is to say, "We've got liquid crystals, we've got lasers, can we combine them?" One idea that occurred to us is that we might take an infra-red image — often it is easy to get an infra-red image but difficult to translate it into visible light — and to take an ordinary liquid crystal cell, say two inches square, to project visible light through it, as I am doing here through two-inch square slides on to a screen; and then to arrange the transparency of the cell to be altered by the arrival of a thermal image which might be a picture of a tree, say. So where the thermal image falls in the cell, it heats up the local areas of the liquid crystal, and changes the scattering properties so as to give you an image on the screen.

As far as I know, that has not been done; but what has been done already is this: people have already used this for facsimile.

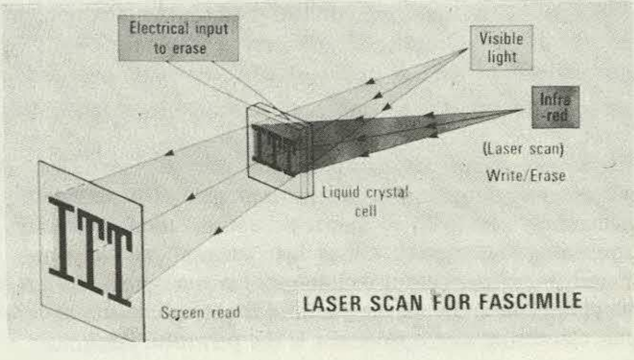

LASER SCAN FOR FASCIMILE

Western Electric and IBM have both made a machine like this. You take a liquid crystal cell. You project light through it on to a screen. You write in, by means of a scanning laser beam, electromechanically scanned with mirrors and so on, information which you then put on. You can then project it on to a big screen; you can make photographic images; and they have done art work for printed circuit boards and so on in this way.

That suggests all sorts of things to you. I suggests a large display in an office, using a projection system like that, with only a small active cell and various methods of putting in information. it is still at a very early stage, but to me it is a highly ingenious idea.

We have got some other things to do first, of course. We cannot get very far with optics until we provide people with a whole family of components. One of the commonest components in electronics is the solder joint, where you have lots of wires coming together and they are soldered; and current going in through one goes in through all the others.

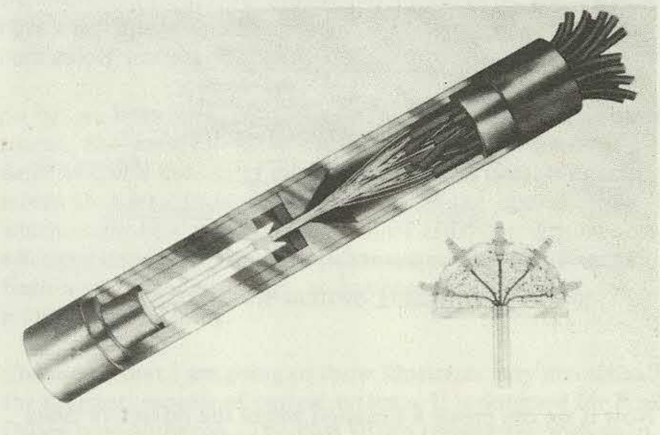

This is the optical equivalent of a solder joint. We have 19 fibres — 19 because they pack in a circular symmetry — coming down here. We have stripped the cladding off. The light goes down and is reflected off a mirror and back out. So light on any one goes off on any one of the others. That is an attempt to start a catalogue of optical components which can be made available.

Another nice idea would be to have a 3-D display — not for broadcast television, because you would have to change the whole system of broadcasting, and it is hard enough to change the licence fee, let alone the technical system. What this

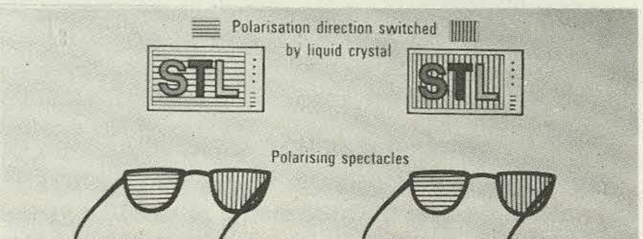

FRAME SEQUENTIAL STEREO PAIR

demonstrates is that if you look at a scene with two eyes, right and left, through polarised spectacles, then you can arrange for the polarisation to be alternately switched between two types, vertical and horizontal, with a liquid crystal panel in front. So you can get two views — a lefthand and a righthand view — from two cameras in the studio. That might be interesting for air traffic control or for inside a building where you want a 3-D display. That is technically possible, although there is a lot of work to be done in making large liquid crystal panels.

Finally, I should like to talk a little bit about switching. We had a great chap at the Labs, called Alex Reid, who invented PCM, OBOE and a few other things; and he always thought that optical transmission was obvious and trivial and that the real difficulty was optical switching. He envisaged that one day an exchange would consist of a whole assembly of light rays, flashing back and forth across an empty space, and connecting customers to one another by means of light signals. There would be no moving parts at all, just the light beams whizzing around. So we would like to find a way of switching light.

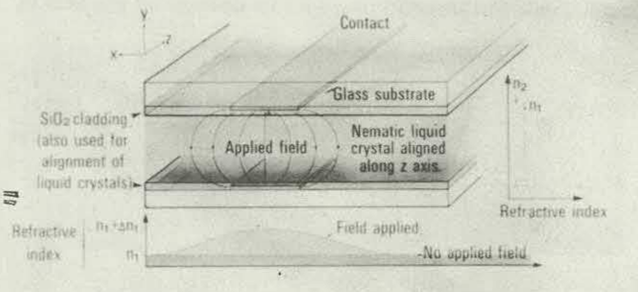

Forms waveguide for light propagating in z direction polarised parallel to y axis

SWITCHABLE OPTICAL WAVEGUIDE

Now if we can create a situation where the refractive index is high in a certain channel and low outside it, we can confine light to that channel by total internal reflection. How can we do that? The liquid crystal material has a refractive index which is electrically changeable, so we might be able to make a switch.

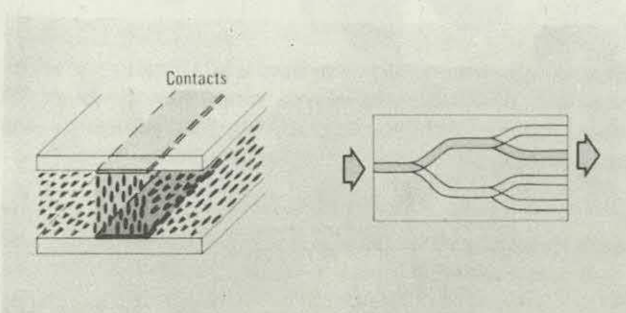

SWITCHABLE OPTICAL WAVEGUIDE

This shows an artist's realisation, which is always so much more elegant than the thing that you see in the laboratory. It has two plates, with liquid crystal material in between. We put on a metallic network on the top and, by applying a field, we can steer our optical signal through various channels like a railway junction. That still has to be tried because the losses may be too high in it.

What I really wanted to show you is that one has a whole family of optical components coming along, which might mean that you could do some things more cleverly.

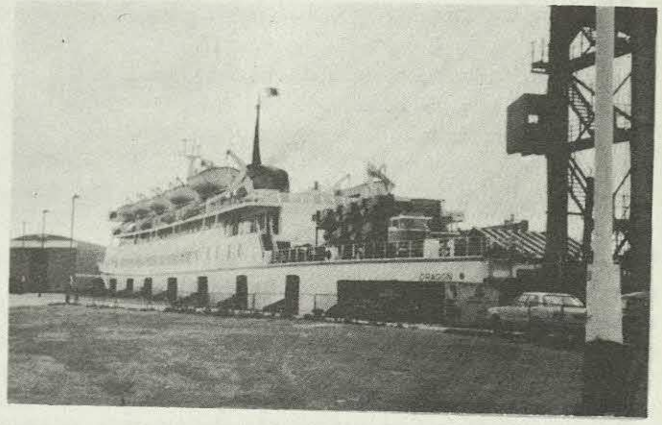

One last example. We have our light sources and our fibres: now there are some jobs that we can do that we could not do before? One of them is pollution monitoring, oil in water. We were asked whether we could find a method which would detect oil in water at the parts per million level for monitoring pollution near tankers. Every ship is required to have this equipment, if one can make equipment that can stand life on a ship. We have just installed one and I will describe it briefly to you.

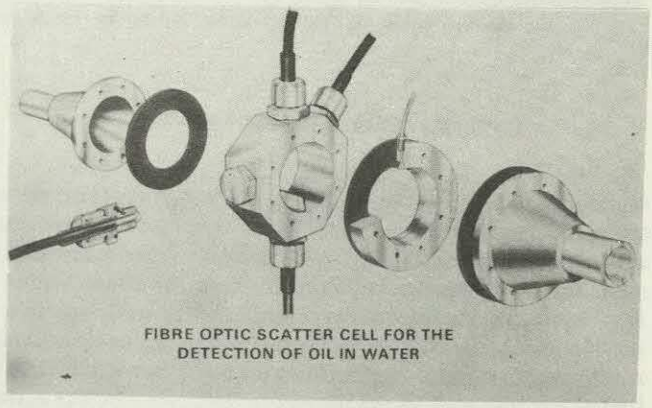

FIBRE OPTIC SCATTER CELL FOR THE DETECTION OF OIL IN WATER

The principle is very simple. We have a scatter cell. Just as the light beam from that projector is being scattered by the dust particles in the air here, so if you send a laser light through a thing like this, if there is oil in the water, the oil particles which you have created by homogenising in a separate agitator will scatter light; and then you can pick up the direct signal and the scattered signal. So you have a laser and you have photodetectors; that is all you have. The water goes in at one end and out at the other; and in that cell you get a signal.

This is very simple to do from that point of view. The complicated thing is processing that signal. This is where electronics comes in — microelectronics, microprocessors; processing that signal; taking into account all the non-linearities in the system; building in various equations that you want to. That is the clever part. So the combination of optical plus electronic microprocessor is the powerful thing.

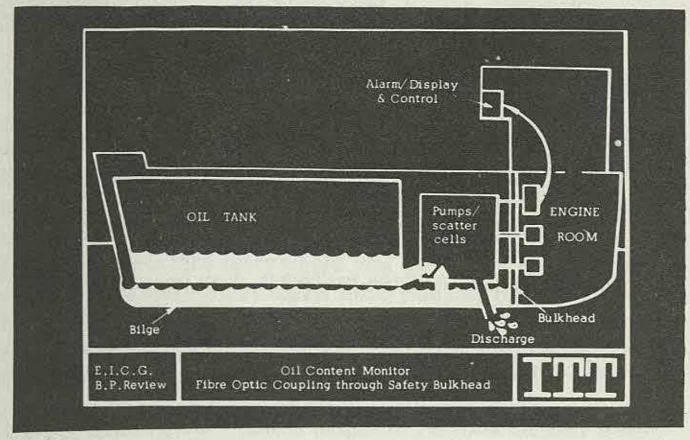

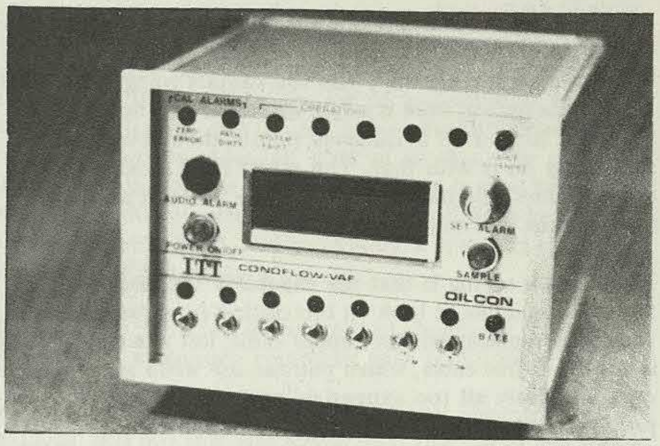

We discovered a terrific bonus which we did not realise. There is a big safety problem with oil tankers; they do not want any electronics, any electrics in the pump room. The pumps in fact have a shaft which goes through so that all motors are at the other side. When we wanted to pull a fibre through this bulkhead, there were no objections because the fibre obviously does not carry electricity; but the fibre carries the signal. So we have our scatter cell on one side, and the electronics on the other; and that was a tremendous selling feature which we did not foresee.

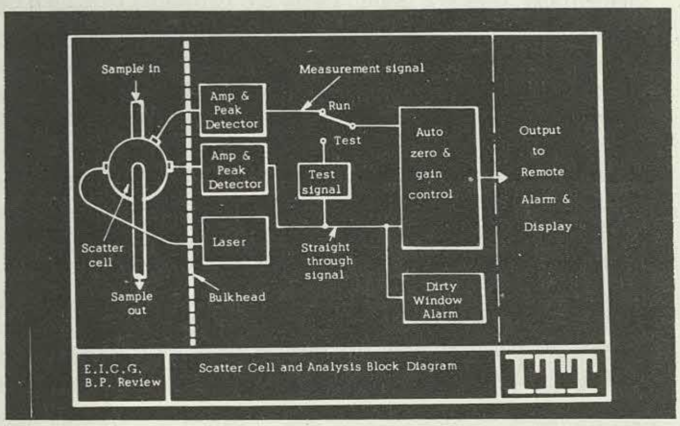

This demonstrates how the whole thing is arranged, showing the scatter cell with a sample going in and coming out; it shows the fibre leading the light from the laser up to the detector, a straight through signal, and then all the electronics.

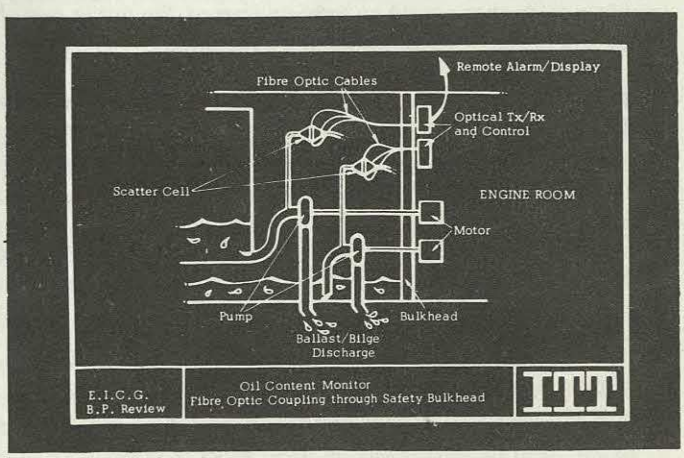

The whole thing is very portable and can be mounted in a ship. You can check the bulkhead, the ballast, and the bilge through the bulkhead.

They have two separate problems. This equipment now can be led to a very simple display. This is an LED display, with various alarms; when it reaches a certain level it creates an alarm. You have various sytems for flushing the cell through, checking that it is clean and so on. This has just received Department of Industry approval as a prototype, and we now have a certificate so we can put this on to ships, and it is a completely unexpected application.

So far, we have tried to show that the presence of a solid state source, the presence of low-loss fibres, the existence of detectors, and the use of microelectronics to process signals, means that we can begin to do things in the optical mode which might be simpler, cheaper, more reliable and more cost effective than doing them in other ways. This could range from a whole lot of things, to keyboards, to displays, to pollution monitoring.

The movie that I am going to show illustrates very dramatically the practical aspects of optical systems. It is designed for Post Office type audiences. The Post Office cooperated with us in providing the ducts and so on where we put in the cable. I do not want you to get the wrong impression. We are not saying that the application of fibre optics is only to long-haul, Post Office type communications. If it were, we would close business tomorrow, because we could very quickly replace all the existing co-ax with a few fibres and that would be the end of the business; just as in transistors they did not only replace tubes, they made possible a whole lot of things which were not even conceivable in the days of tubes. So we are hoping that fibre optics will make possible a whole lot of things that are not even conceivable with co-axial cables, multi-pairs and so on.

(Film of the 9 km long, 140 Mbit/sec optical fibre telephone link between Hitchin and Stevenage)

BUTLER: Gentlemen, we have a few minutes left for questions and discussion.

QUESTION: You said that a few of these cables had the capacity probably to replace the GPO network. Could you expand on that please.

EVANS: The thing is, of course, that they have a tremendous capacity. I have some samples here which will illustrate that. I have a typical multi-wire copper cable, with maybe 2,000 pairs, the equivalent co-ax with about 18 tubes; and also the optical cable which you saw in the film.

Now this is not a very fair comparison, because one of the things that you can do with multi-wire copper cable is to take each wire off to a different point, and similarly you can break this co-ax down into a number of channels; whereas we have a lot of channels on one fibre. So if you have a very broad bandwidth highway, with lots of data going from point A to B, the optical fibre is a very attractive solution. But if you want to feed off at different points, then you would not necessarily do it with that. You might have a cable with a lot of fibres, and merely take just one fibre off at point A, the next one half way along and so on.

What we are saying is that the capability for frequency is there. I would like to look on this optical fibre not as a very wide-band replacement for copper cable, but as a strong, light, interference-free cable, which you can use when space is short, or alternatively all too expensive, or there are interference problems and so on. In other words, I do not want to think of this optical fibre as just a replacement for copper cable because there is no future in that. (We would work for a short time, replacing the copper cables, and that would be the end of the business). I want to think of ways in which you can do jobs now with this optical fibre that you could not do before; and that is really the whole aim of it.

The Post Office has sponsored the optical telephone link because it saw it as an important part of its network of available systems, but I see it as much more versatile. Just as the transistor did not just replace tubes but made things possible that you could not do at all with tubes, and could not even think of doing, I would like to see ways in which we could use optical fibres where you cannot use copper wires.

QUESTION: What sort of work have you been doing to identify applications now?

EVANS: What I discovered very early on is that most people are like you: they are not interested in the fact that it is new and different, they want to know how it helps them and what it does for them in economic terms. So we are really looking for people who have a problem that they cannot solve by existing techniques, or which are better solved this way.

An example is diving. We had a diver. They wanted to send an umbilical cord to a diver and monitor physiological data; they wanted to take his temperature, his pulse rate, his blood pressure and so on. They wanted to do this in a very light and non-restraining cable which did not interfere with all the chains, pipes and so on that he has already. We provided a link like this — at least the factory did. Then it turned out that it was very interesting to the diving people, because they saw a way of getting data up and control signals down to a submerged vessel — in this case a diving helmet, but you could also think of it as a diving vehicle — which was very much easier to handle than a great big, thick metal conductor type cable. So that is one area.

Now they are not just asking for a piece of cable, they want you to solve the problem. They say, "Look, we have this situation. We're prospecting on the ocean bed," or "we're mending oil rigs," or "We're exploring the underneath of Soviet trawlers. We want some way of feeding data down."

The Electricity Board and other people want to have data in power stations, between pylons and so on; and they have the problem of interference, of pick up. They have the problem of voltage levels between different parts of their equipment. This is an insulating cable; so you could go from a high voltage to a low voltage with a cable like this, without worrying about earthing problems and stand-off voltages. So we are really looking for people who have a problem that they cannot solve at the moment, but might be able to with this optical fibre. Unfortunately, before you find those people you have to do the work and make the cable. You have to make 20 miles of cable before anybody will believe that you can make cable at all. So it is unlike Derek's thing, where he makes six circuits and shows them around, and if they want them he makes a million, and if they do not he forgets it. We have to make 40 miles before we start because people say, "Optical cable? Surely that will break when you pull it."

So that has been one of the problems. You have to develop the technology, make some samples, get people to look at it and play with it, people who are not used to using this kind of technique at all.

Motor cars are another example. Inside a motor car you might have a ring main controlling all your appliances, which is an optical multiplex system. People have looked at it. Of course, you do not need silica for that, you can use plastic; it is the short-range, high-loss system.

BUTLER: May I suggest one other application, Joe? You could sell them to the captains of Soviet trawlers who need detector equipment to find out if anybody is looking at the bottom of their ships.

EVANS: Ah yes. He knows that you know that he knows . . . Incidentally, that is a point. There are two slides that I did not show you: one because the Ministry of Defence asked me not to, and the other because my own commercial people asked me not to. All that suggests to me is that there are problems around where we have suddenly hit on something which they think, "Ah, maybe this can solve it." That, to me, is interesting. Something may come of it and something may not. But we are really at the stage where we want people to exercise imagination now, on the problems they have always had lying around. Here is a chap with imagination.

QUESTION: I've been wondering about the security aspects. It seems to me that with the optical fibre there's no way of breaking in and tapping information. Am I right?

EVANS: You would be able to know that people had done it. Yes. One of our applications was the police. We had two applications for the police. There was a lightning strike in Bournemouth, and it put out of action the link between their central computer and their out-station; they had a list of stolen cars and things like this. This lightning storm took out all the data that went along their cable, and they asked us to put in a cable that would be lightning-proof. We put in one of these and it has worked beautifully ever since.

In the other application, the police asked for a length of cable and two transmitters, but they did not tell us what they wanted to do with it. So I lift my carpet in the office occasionally and have a look!

BUTLER: Joe, could you give us any idea what cable mileage of optical link is now in service in this country?

EVANS: Very little. The Post Office have put in a number of systems. They asked us and they also asked Plessey and BICC to supply an experimental cable, which the Post Office are now studying. My worry is that, for the next two years, the Post Office will have all their work cut out testing these, sending data along, measuring error rates and so on, and they really do not want to buy any more. They've got them in the ground and they are playing with them. We really cannot wait for that kind of market to come. We are looking for things which are not telecomms, not Post Office, not long-range.

QUESTION: Can you send power down optical fibres?

EVANS: In the old days we used to have a copper wire with an enamel coating; I want to change this round and have the glass in the middle and the metal on the outside, and just see what you could do with it. You can send power, yes, but again the amounts of power that you can send are trivial by comparison with metal. We would like to put down submarine optical cables. I was in Japan recently, and the thing that frightened me is that everywhere you went they were working on this, as on every other damn thing you can think of. They had application areas: optical submarine cables. This worried me; and every time I left a lab I noticed that they took that down and said, "Ah so, got him worried." They then put it in a taxi and sent it to the next place that I was visiting!

When Charlie Cale first thought of this, the losses were 10,000 DB per kilometre; now it is typically 4 DB per kilometre. So one of our chaps said, "What about .01 DB per kilometre?" and everybody said, "That's crazy," and he said, "Think of Charlie Cale." It is like the song, "They thought Marconi was mad". If you have very low attenuation, you could have a link across water. There are lots of stretches of water which are 50 kilometres between islands, and even along rivers. So if you could get a low loss, you might be able to have optical underwater. Then you have the repeaters across the ocean. We thought very hard about how we could power these repeaters without sending the power along the wire, because having to put in metal spoils the beauty of the design; if it is all optical you can make it very small and strong, and very difficult for Russian trawlers to pick up. They pull the submarine cables up regularly and cut them. They are also being adapted to catch fish now, in the newer models! If you could have that, this would be very attractive. But some way of sending power from A to B which is not along wires, we have thought a lot about it.

Nuclear source is one idea, and we went to Harwell. Just talking to those chaps — the cost per minute that they charge you just to talk to them. I would never dream of giving them any work to do. So nuclear powered sources with and without wire would be attractive. There are all sorts of possibilities like that.

QUESTION: Have they been used as data storage devices — like delay lines?

EVANS: Not the fibre itself, as far as I know. It's ever so quick, that's the trouble: it's 186,000 miles a second, isn't it?

QUESTION: How cost effective are they compared with copper cables?

EVANS: It is too early to say, of course. These things are all made by hand, by PhDs; and the cost of advertising films has to be loaded on! But Derek was talking about parallel processing, about one slice of silicon four inches in diameter; and so many hundred circuits, and the cost of the hardware tending to zero; well this has glass which is extremely cheap (SIO_2) and it has plastic which is also very cheap. When you look at it, intuitively you feel that this must be a cheaper thing. Although at the moment it is all handmade and lovingly constructed and tested, the cost of the materials is very low and it must be cheaper in the end than almost any other cable; especially co-axial cable with all the copper in that. We of course have moved from copper to aluminium in the Post Office network, and unfortunately, soon after we did that, the price of copper fell and the price of aluminium started to go up because of energy costs; but long-term, aluminium is still cheaper than copper. But when you think of this optical fibre, it is cheaper still. It is similar to our experience with semiconductors: a small amount of material, very carefully prepared, in the end does the job of vast amounts of iron-mongery which we used to use. Here a very thin silica fibre, carefully prepared, does the work of all that. One of these six-foot spools holds several kilometres; so it must be cheap, although at the moment the costs are all artificial because they are loaded with R&D, and everybody is spending enormous amounts of R&D on it.

QUESTION: So in fact we are talking about cost reductions of orders of magnitude compared with traditional cables?

EVANS: Oh yes. At the moment it is not cheaper to do it this way, it is just that you would do it this way if you could not do it any other way. Eventually, it will become cheaper; it must do. Already, those stranding machines that you saw which were all designed and built at the Laboratory are being re-thought; people will build them much larger and much faster. When we started we thought that glass was a very delicate material and that it needed very precise and careful handling, but it is not true, it is extremely strong. We are looking on this now as a light, strong cable. You can imagine a soldier winding 200 or 300 metres of single core cable round his waist. He goes to the field and he pays it out, and he has a link which is strong and undetectable; he leaves it behind when he has finished. It really is an exercise in imagination on applications from now on. I think that the costs will come down as the usage goes up.

BUTLER: Gentlemen, let us close the session at this point. It is also the end of the module of the conference devoted to new technologies. After the tea break, we will turn to the module of the conference which is concerned with how those technologies are likely to be brought into the market place. But before we break, let me ask you to express your thanks for what I think has been a characteristically witty and provocative presentation. Joe, thank you very much.

THE SHAPE OF FUTURE PRODUCTS

E. Fauvre

BUTLER: We now move on to the second of the three main parts of our agenda, which is about how the new technologies will be packaged and brought to the market place in the form of new products which the users have to decide whether or not they wish to buy. Our first speaker is Ed Fauvre, who is the Group Manager of Commercial Engineering with Digital in the USA. An impressive list of companies with which Ed has gained experience is in the agenda, and I should like to tell you now that Ed has kindly agreed to be here throughout the conference so that you will have the opportunity to talk to him, get to know him a little, and ask him any questions which arise which we may not have time to deal with today. Ed has taken as his subject a small and modest topic: the shape of future products.

FAUVRE: Good afternoon. I will take questions during the course of my talk, because I am covering a number of topics. I have tried to think of what you would be interested in and to keep it free of particular product and commercial content. But please interrupt if I can clarify anything or if you would like to lead me on to a particular path, and I will leave it to the chairman to turn it off if it gets too long. Also, in terms of anything specific, I'd be very happy to meet with all of you, or any of you, during the next day and a half.

As you know from the introduction I'm Ed Fauvre, Chief Engineer of the Commercial Group of Digital Equipment Corporation, whose corporate headquarters are in Maynard, Massachusetts, USA. The Commercial Group is perhaps one half of Digital's products at present.

Digital may be viewed as the world's largest minicomputer supplier — and these are the only statistics that I will give you — with a billion-plus dollar revenue this past year; this first quarter we did an increase of 48% over the last year's first quarter; we have over 36,000 people; and our income/share is up now from 43 cents in the first quarter last year to 66 cents this year. The truth of the matter is somewhat more obscure and complicated than these facts would seem to indicate. For us better to understand one another, I owe you a more intimate, revealing view of Digital in general, and my groups in particular.

Our view of ourselves is that, although the financial metrics are obvious and exciting, they are so primarily because we have a corporate orientation to what we call intelligent, dedicated people trying to do the right thing. In our case, doing the right thing has been supplying dynamic system solutions to the problems of the people and companies who relied on interactive data processing when we started, and who rely on information processing now; and who will rely on image processing in time to come.

During the past 20 years, since our inception, Digital has been a leader in the developments associated with the transition from interactive data processing to information processing. In the future, we want to be leaders in the transition from information processing to image processing.

It has been very challenging, and it has been rewarding; because overcoming resistance to technological innovation is not a new development. The early restrictive and nuisance laws on horseless carriages versus the horsedrawn variety was one case in point. However, the pace of technological innovation is, of course, of concern to us. We are not in the mode of being technologically driven — either pushed by the technology or pulled by the market *per se*. But we are in the mode of applying leading edge technologies — hardware, software and systems — to user needs and requirements.

That is enough of the background and Digital's perspective. We are all here, I think, to talk about where we are going, having assayed where we have been. This morning and this afternoon we have heard about the systems industry of tomorrow, the importance of new technologies in office communications, developments in memory technology, and novel optical components and systems.

My thoughts, addressed to "The Shape of Future Products", will be along the lines of how we tend to look at the world in view of its dynamic history, and what that history has been from both philosophical and technological perspectives, where we see the world going, and how we of DEC will continue to help the world get there.

First, philosophically viewed, the history of data and information has been relatively simplistic until quite recently. But, as I will demonstrate, there are some significant parallels between what has happened with data and information and what has happened in technological developments.

Before Guttenburg, data and information were conveyed by couriers and scribes in local languages, and monks, who wrote, usually for posterity, in the universal language of Latin. Then, with the invention of the printing press, distribution of information became more widespread. But, considering the levels of literacy, relatively few people could use data or information in meaningful ways. Over time, and with the availability of information, more and more people, who were forced by necessity and by their drives to better themselves, learned how to use the information. Besides, more and more people became literate, interested and involved.

Production printing broadened, making information more affordable. Newspapers in local languages conveyed data. Specialised periodicals in narrower, specialised languages began to proliferate. Books of esoteric value became available. More recently, electronic media — TV — and paperbacks in the printed world made all sorts and kinds of information and data — both trivial and non-trivial — available to almost everybody. Data and information availability has gotten to the point of having been characterised as "the information explosion".

About 20 years ago, one of the early pioneers said that four computers would take care of all the world's needs. I am sure that if that same man had been alive in the 1400s, he would have said, "We don't need printing presses, there are enough scribes to write all the books that we'll ever need." I think that we are at a similar point in evolution today with our computer in terms of the kinds of technologies that are beginning to be available and the pressure on what will eventually come out the other end.

In the technological world of computers and communications, in the days of ENIAC, computers were gigantic in size, relatively unreliable, and usable by only small numbers of scientists who wrote in machine language. I wrote in machine language. I thought zeros and ones were great.

The computer has been used as a pivot on which to treat movements in our physical, economic and social worlds. It — the computer — is both model and metaphor. It has infected and altered all parts of science, technology and, more recently, business. It is the atom-smasher, microscope, telescope, agent of control, modeller, simulator, filing system of the world. Digital has been one of the leaders in the domestication of this inanimate device, which both blends into the background and, at the same time, usurps the foreground of the world's interests.

But, again, let us look at some history in the business world. First, there were manual systems with complete reporting. Then manual systems with exception reporting. Then there were electronic accounting machines. Next, there were computers, affordable by limited numbers of wealthy companies whose resources permitted them these new luxuries of financial controls, *post facto* reporting and so on. But we all know about these evolutionary trends. Similar trends have occurred in the scientific world's use of computers.

Strangely, however, there have been no really fundamental changes in the nature and/or organisation of the computer since its conception and inception. I remember, back in 1960, when I was at Douglas Aircraft Company programming a 7090, when a scientist came into the computer area and said, "Here's a formula. I hear you have a computer. Please run it through." This man was looked upon as a lunatic and eventually was guided to a bunch of FORTRAN programmers who taught him the mysteries of life. But he was right: why should there have to be an army of FORTRAN programmers to convert his formula? Why couldn't it just run through the machine? The population of these stored program, electronic machines has increased dramatically over a relatively short time frame.

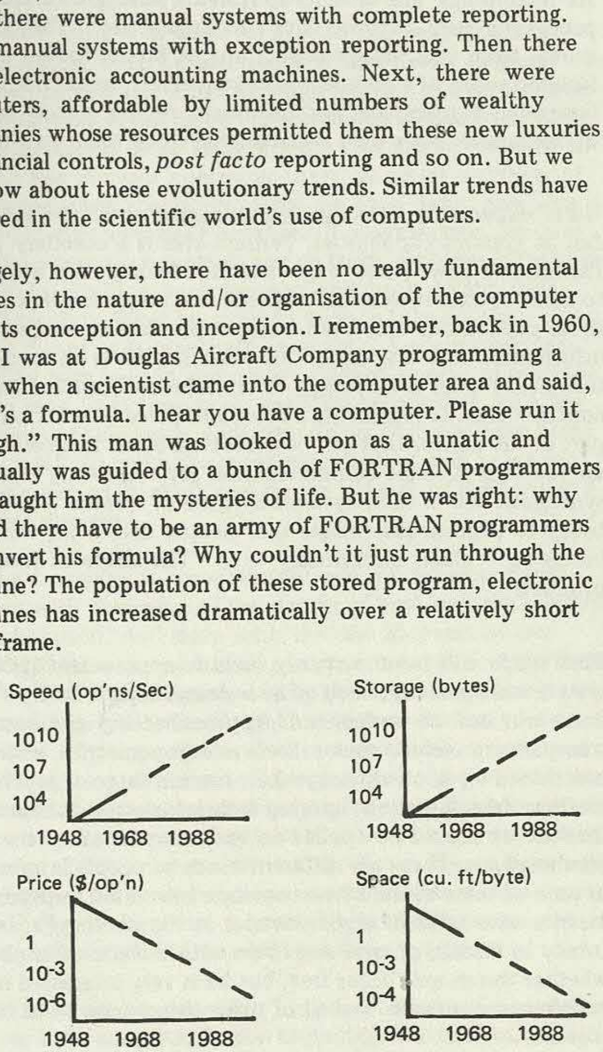

E. David, in his "Some Thoughts about Production of Large Software Systems", from the NATO Conference Report on Software Engineering, has given us these historical and projected trends associated with computer hardware developments. I believe that these tie in similarly with what we have seen today. This chart simply says that mainframe speeds are increasing; the price is decreasing; the cost per byte of mass storage is dropping; and the space per byte is decreasing as well. All of the hardware and what we call "real estate" is decreasing rapidly.

With the hardware transitions from discrete components to integrated circuits to medium scale integration to large scale integration, and now to very large scale integration, we have

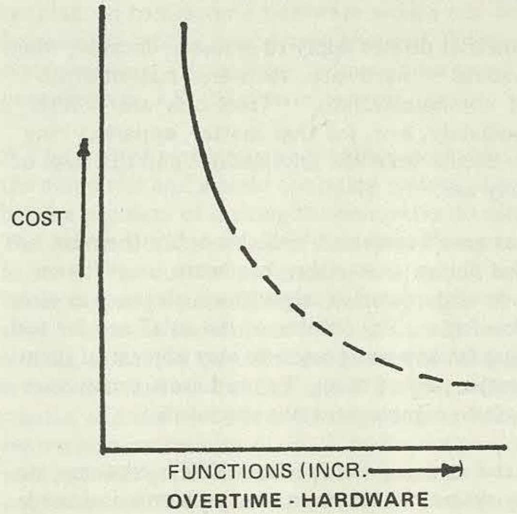

witnessed a progression like this. We at DEC do not know exactly where we are on this curve, but we do not think that we are very far along it, since packing densities on a chip will in all likelihood get heavier and perhaps approach molecular densities and even surpass that.

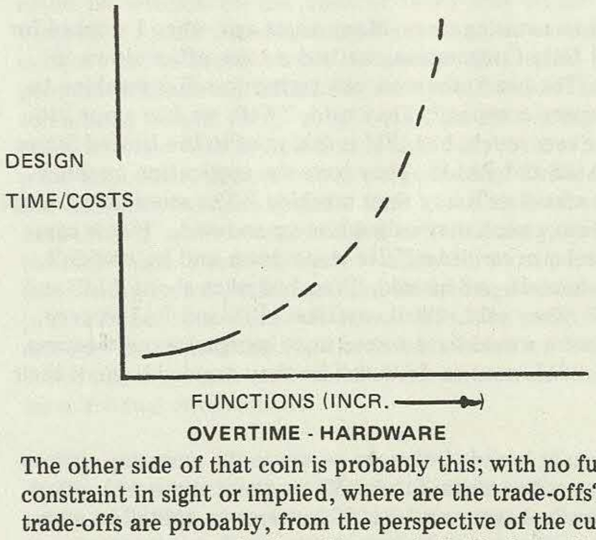

The other side of that coin is probably this; with no function constraint in sight or implied, where are the trade-offs? The trade-offs are probably, from the perspective of the customer and the user, not to be found in any individual simplistic view of an individual hardware or software technology or cost.

Candidly, our pragmatic view of these trends suggests that technology breakthroughs — that is, cost-effective solid state replacements for electromechanical devices — will happen when they are badly enough needed for us to push hard for them. Certainly, we have every reason to believe that, at least in this regard, history will repeat itself. We also maintain that similar algorithms apply to sotware engineering, when they are viewed singularly, superficially and simplistically. Breakthroughs will occur when they are badly enough needed for us to push hard for them.

I am sure that all of you have seen *Byte* magazine, the computer hobby shows, and the consumer computer. It reminds me of back in the early '50s when I was making home radios and hi-fi, winding my own relays, building cones and amplifiers; it is at that stage that we are in. In a way it is a kind of hobby. But the fact that we have solid state devices today, and equipment selling in the low hundreds of dollars, is creating an appetite among not only business people but amateur users who program and create their own libraries of subroutines. I believe that that pressure — not economic but social — will in fact push the industry into doing things quicker than it would have done ordinarily.

But the same metrics do not apply to systems; because, when we analyse systems — hardware, software, man-machine interfaces and communications — trade-offs are neither apparent immediately, nor, for that matter, apparent long range, until we decide what the implications and the roles of the system really are.

But, first, what aren't systems? Well, probably they are not tightly-bounded things that either hardware or software engineers can do with materials, algorithms, elegance or state-of-the-art technologies. The "states of the arts" are far too rapidly changing for any one person to stay abreast of them all, or even a major part of them. We need more renaissance or generalist people to complement our specialists.

To repeat what I said previously about DEC's business, we want to supply system solutions to user problems and needs that are generated by users' dynamic environments and the concurrent requirements that they do things in their own businesses better, smarter and more productively; but use the machines in the way that they want to use them, not the way in which the computer manufacturer chooses for them to use them.

There is an amusing story. Many years ago, when I worked for Control Data Corporation, we had a sales office down in Mexico. The head salesman was trying to sell a machine to an insurance company. They said, "Well, we like your little machine very much, but IBM is taking us to the United States to see ALIS and PALIS. They have the application packages, and I'm afraid we'll buy their machine." The salesman gave up; and the next week they called him up and said, "Please come down and sign an order." He went down and he couldn't contain himself, and he said, "Yes, but what about ALIS and PALIS?" They said, "Well, we like ALIS and PALIS very much, but it would have forced us to reorganise our business, and we didn't want to do that." So they decided to do it their way.

What then in broad strokes do we see as the emerging environment, the office of the future? Throughout the world, where we presently have over 100,000 systems in operation, our commercial data systems, typesetting systems, text management and word processing systems are well known. Little known is that, within DEC, we have started an electronic mail system for our own use, just so that we can understand all of the segments of "The office of the future" and to make our own office people even more productive. We are certainly not replacing our secretaries, but they like using word processing systems because it gives them the time to do other things which are more important to them.

Additionally, our planners and specialists are trying to understand and share their understanding of the technological convergences necessary to complete that picture. Where and when will facsimile, word processing, video, teleconferencing, typesetting and so forth converge? What kind of systems can we provide our users which still better complement and best support both piped and unpiped communications networks? Are electronic point-of-sale systems (EPOS) and electronic funds transfer systems (EFTS) inextricably intertwined? What do we need to make them most useful to our customers? How can we help our customers exploit these systems?

What can we do to provide our users with the most highly secure systems, taking into account not only the various levels and kinds of security, but also those of privacy? How do we build impervious systems? I find the move to the need for secure systems somewhat analogous to the experience that we had in the United States with the automobile. First, we invented the automobile. Then we built a lot of roads that it could run on. Nobody thought that, out of that, the suburbs would emerge and be a threat to the cities. In a similar kind of way, once we started interactive processing and the control of the data base went out of the central computer room, it was only then that the data processing manager started to think that now his data was in jeopardy. Now we have to think very seriously, because of interactive computing, about how to deal with the problem of security and privacy. What kind of systems do we provide our customers that are indigenously high availability? What is high availability? Representing the highest Digital priority, what kinds of systems will fulfil our need to be socially responsible to ourselves, our customers, and the world?

As technology has advanced, systems have grown from processing pure numeric data and English text. They have grown from stand-alone collections of boxes, locked in an isolated room of a company's headquarters, to distributed and interrelated processors and terminals, serving a wide variety of users with their own departments or at their own desks.

Users' expectations have grown at least as much and at least as fast as systems capabilities. Perhaps this is a corollary to Parkinson's Law: Users' expectations of computers will grow to meet or exceed systems' capabilities, no matter how great those capabilities are. Users and people deal in images, not information and not data. This presentation is an image. A technical report is an image. It contains information and it contains data. It represents a complete piece of work on the part of the author that can be evaluated by its recipients. We are the biggest single communication gap separating users from synergistic use of their systems because of systems' inability today to process the images that people deal with, which brings us — albeit circuitously — to "The shape of future products".

That shape will most certainly include approachability and system transparency; both quite serious subjects that we are doing our best to understand. Approachability and system transparency include such subsets as ergonometrics, which was mentioned by a previous speaker; human factors; psychometrics; data depiction; imaging technologies; and the metrics that are or should be applied to each, any or all of the combinations. There are different kinds of people interested in each of these areas. I have someone interested in psychometrics who talks of ergonometrics as "knobology". He is totally in disdain of anything to do with the size of knobs and whether things are flicker free, but he is very interested in how people communicate. But all of those things have to be tied together.

All kinds of esoterics have been discussed in such contexts. Forward thinkers have based their postulates for the future on analyses of the past. The common denominator of all the postulates would seem to be the eternal or ultimate goal of man: image and thought transference; first, artificially; then naturally.

Thus far, we — the technologists — have been particularly unsuccessful in this context, probably because our technological implementations from the onset have tended to build functionally bigger, better and faster machines, and adapt or force-fit humans to machine "ways of thinking". I was talking to a gentleman during the break about an article written by someone from Xerox describing an experiment done with children where they developed a language called "Small Talk" to help in graphics. After reading the article, I felt that rather than trying to adapt the machine to the children, they were teaching another community to learn one more language, now called "Small Talk". That is the thing that we have to get away from over time.

In the process, we have brought generations of specialists into the gaps of man-machine communications. We have called them programmers and systems analysts. Too many systems analysts are in reality systems designers, who are adapting human processes to the computer, instead of analysing how the computer may be adapted to solve the problems in a manner that people want to solve. Initially, the programmers and systems analysts were the scientists who used the machines directly; and, by nature and necessity, they communicated with the machines in their machines' only language — machine language — binary ones and zeros.

There were too few of these people. So our next step was to provide a more widely understandable and usable language — Assembly language — which could be learned and understood by more people. Then came interpreters and compilers. Each step or increment being an additional abstraction designed to broaden and increase computer usage by larger numbers of not-so-specialised specialists. What we have done, now, has resulted in ever increasing numbers of programmers, systems analysts, computer operators and so forth, who are in fact user to man to machine back to man and back to the user communicators. What we need to do is to take another, closer look at "what evils man hath wrought." Let's reverse the trend. We have been in the mode of adapting man to machine for too long. It is time now, partially because we have the beginnings of this technology, the things that you have heard about today, to make the machines adapt, as well as we can, to man's way of thinking and tip the balance of the scale in a different direction.

But the foregoing is a conclusion; and my purpose here is to share with you some of the details of how we have reached this conclusion. As I have said, the last 25 years of our industry have seen a continuing improvement in cost/performance. Similar improvements can be predicted for future years. These advances in technology have repeatedly forced a redefinition of our products. Each year has seen major, new, cost-effective applications of systems. As a speaker said this morning, for the normal cost curves, for every doubling of our output of microcircuits, I believe that we have a reduction in the order of 28% in terms of cost; or, if you like, engineers like straight lines. So if we don't have a straight line we put enough graph paper underneath it until it eventually comes out looking like a straight line, and then we're happy.

"The total cost of ownership" is one of the useful models that we have adopted. In the beginning, the cost of the computer hardware was a dominating concern. In the future, the costs of using a system will dominate. In order to demonstrate that fact, consider a standard Digital system of perhaps five years from now. For $N one will be able to purchase a system comparable to what is $10N today.

Assume that the system is used in a small enterprise, and perhaps 10 individuals use the system regularly, as a necessary part of their day-to-day work. Most the the major factors in cost of ownership are as follows, along with possible areas of technological impact:

1. The cost of the basic system. Our hypothetical system will be complex, and support a complex operating system, suited to the hardware configuration. Today's operating systems running on tomorrow's hardware would not be adequate, because the setting would have changed. Whereas a $1 million system can sensibly require dedicated operators and systems programmers, a $50K system cannot.

2. The cost of making an application work. Having purchased the computer and a basic operating system, the customer still has the problem of making the computer do some useful work. The fundamental problem is the cost of programming as we know it. The cost of programmers is bound to rise unless better means of training high-quality programmers are found, simply because the number of programmable computers in use is rising at a rapid rate. The problem with hiring a programmer is twofold: the cost of the initial programming will be substantial and the cost of continuing program maintenance may represent a continuous drain. A dedicated programmer is well paid and will get better paid as the years go by. The capital equivalent of five years' amortisation period can be a very expensive number, probably higher than the cost of the computer itself. If complex systems are to find large markets, some way of eliminating the traditional application development costs must be found.

3. The cost of training. When computers were rare, operators could be selected on the basis of being able to use the computer. As the use of systems becomes more common, the system must be totally usable by average, that is,non-computer trained individuals. In our hypothetical example we can assume that of the 10 people using the system regularly there are none with any particular aptitude for machines or, more important, any interest in programming. Those of us in the computer systems business have come to assume some capability in systems analysis and understanding, and may tolerate — albeit unconsciously — many forms of poor human engineering in computing systems. In the large volume markets poor human engineering will impact sales. The smart customer will buy the system that has the fewest elements of computers for computing's sake, and which is the most natural to use in his individual environment.

4. The cost of use. Assume that the 10 users have an average salary of $10K a year, which is very conservative for five years hence. If this is true, the salary cost of the users far exceeds the cost of the computer system. Historically, we have spent much effort in the optimisation of the computer resource. The problem in the future is optimising the usability. Part of the problem was discussed previously, namely, eliminating unnecessary extra people which, in the purest sense, includes operators and programmers. A second part is making the system most productive for the necessary end or application users. Requirements vary according to application: a supermarket point-of-sale system that delays the checkout clerk will rapidly become counterproductive. In the case of point-of-sales the needs are evident. But in other cases, the usability problems may be much more sophisticated.

5. The cost of failure. Systems have traditionally failed. When a component fails, some, or all, of the system function may disappear. In the sense that the system price is a function of the number of parts (or weight), the reliability of a $50K system may remain more or less constant, given the increasing number of gates on a chip. But unless explicit steps are taken, the system will become less repairable because it becomes more functionally complex.

The problem of repairability is compounded by the decline in quality of service personnel as systems proliferate. This problem of failure is compounded further because the systems become more valuable. Expensive systems have typically been used to perform selected functions that comprise only part of our day-to-day activities. If the systems failed, these activities were delayed, but others went on. As the costs of systems power decline, more and more functions will be performed by them, and our users will become more and more dependent on them for their day-to-day activities, until at last they are as necessary as electricity and telephones — which, in fact, they will control. If our hypothetical system were to fail for a week, it would have a high direct labour impact of many dollars. More likely, it would cripple an enterprise for a week, and even longer in terms of recovery. Many applications cannot survive being crippled for an hour, not to speak of a week. If a steel mill does not put out a schedule, it can never recapture that revenue. If a toy distributor does not put out a bill of loading for its warehouse, it does not ship that day; it has forever lost that money, especially during a particular season.

Does that mean that necessary systems will not be built? Only if high availability systems cannot be built economically. Given the rate of component price decline, doubling the numbered parts of a system does not seem very important, especially if the more expensive system can deliver highly reliable service at a small fraction of the labour intensive field service cost.

6. The cost of evolution. No matter how well we plan our systems, they will become obsolete over time. We assume that the sophisticated customer will realise that, and treat the costs of evolution as a cost of ownership.

We will continue to explore the leading edge of hardware technology at DEC. We are into customised LSI chips and Very Large Scale Integration. We are working with bubble memories and CCDs; we are using them as part of discs and for other purposes. We are interested in laser technology and in the silicon that we talked about today. We will continue to expand and integrate our systems — interactive data systems, word processing, networks, transaction processing, distributed processing, and more and more availability and security which we feel is very important, into more and more humanly approachable systems, with intelligent terminals and stations, sometimes generalised, often specialised application terminals, for different needs of people, for different industries and different applications.

BUTLER: Thank you, Ed, for that very clear analysis. I particularly appreciated the analysis of the future cost, which is something of interest to everybody in this room. I'd like to invite you to launch the first question.

QUESTION: You described your view of the trend to distributed processing, and you said that your definition of an image was one man's view of reality. Could you elaborate on this?

FAUVRE: I can explain it in several different ways. We don't have the solutions. I would certainly not like to give you the impression that we have solved all these problems. But let me give you three views of an image. First, when we do applications and we talk with people, people in different industries have a different vocabulary. Certainly, people in the computer industry very often attempt to speak in straight language and find out that people understand only half of what they have said, because they talk in certain jargon. It is very difficult, if you are in the business, to avoid that. People in the typesetting world talk in terms of that world, and they would really like to communicate with the computer in that way. People in the medical world or among lawyers are talking another set of terminology. People in the trucking industry talk about things that they are doing in a different manner. I feel that the computer should be able to talk to them in that kind of language. At a very elementary cut, there is COBOL and FORTRAN, but it is very primitive; and business people really don't say, "Perform A. GO TO B." The only person who does that is Grace Hopper. She said that when she went to Europe, since she was the founder of COBOL, the only time that she could sometimes communicate was when they would talk COBOL at the Standards meetings. If they did not speak the same language, at least they could communicate in COBOL. But that is terribly primitive.

Scientists can talk in DO loops, not really because they want to but because FORTRAN is something close to them.

At the same time, almost all of our systems generally come out in English, and the translation to a different alphabet is difficult. Sometimes algorithms, for example; in a hyphenation justifaction algorithm I found that in one of my products the entire algorithm is based on a 26-character alphabet. The person who did that did it with the purest of spirit, but he did not realise that I would be selling a system to Spain and we would need a different alphabet; and now we must re-write the hyphenation and justifaction routine.

When I did a COBOL compiler, 10 years ago, the first ANSI compiler, we did the reserve word list by simply setting up a table, external to the compiler itself, which could be changed to any language; and therefore anybody could write COBOL in whatever language suited him. That did not take any particular talent, but the awareness is what is important. I believe that as the computer companies and the software designers come to have an awareness that they are dealing with a world that is complex and different in its needs, we can start to approach those kinds of problems.

At Control Data, years ago when the Master system was developed, they decided that they wanted to know how to put in the date. So they sent a questionnaire to all their customers, and they finally put in a certain date. The date standard was the Australian standard, because a customer in Australia was the only customer who answered! So we had the Australian date standard. That is one cut at the universe in terms of images. It is just the concept of getting something close to what people are doing.

But when you look beyond where people are today in terms of using data processing, people have problems that they want to solve. Problems in astronomy. They are using graphics. Now graphics are fairly expensive, and coloured graphics are expensive, and the programming of them is expensive. But if an astronomer wants to use the system, or somebody wants to do automated design, he does not want to do a complex amount of programming. So as these areas become more and more inexpensive, and we get more and more modular building blocks, which become possible with solid state, very cheap memories, so that the cost of memory is not the issue, then we can have those things sitting out there at a minimum of expense, and you can get these very sophisticated usages of the machine done at a nominal price. What the user has to do to the building blocks to get to his solution, he is not even talking in a computer language, he is dealing with another level of abstraction. Also, he will be able to use his senses. We will have colour. He will be able to use sight and sound, and voice.

I know people who are managers who choose not to type. After I sent my secretary to a word processing course, she came back and I said, "Now do you know what computers are?" and she said, "Yes." That was her concept, because typing was very natural to her and she felt good about it. But when I was at MIT with a friend from the Telephone Company of Pennsylvania, I sat him down to try to run a regression analysis at a terminal, and he froze; he did not know how to use the keyboard. We have to get rid of the keyboards; we have to get rid of these things. We have to get a natural man/machine interface between people and how they want to talk to a computer, the same as they want to talk to a secretary, or to an accountant, or to their doctor. That is what I mean by the beginning of the concept of imagery. Does that answer the question?

QUESTION: I was very pleased to learn of your awareness of the total cost of system ownership. However, I am worried that although you were making a case for justifying hardware more easily in the future, you did not comment on software. When can I as a commercial user expect to see systems and programming costs come down?

FAUVRE: I was trying deliberately to stay away from products *per se*, but do you mind if I go into products a little? Today, the natural or the most used method of doing commercial languages are either through an RPG or COBOL, or in Digital terms something called DIBOL, which is a small business oriented language. Well within the state of the art in products that we are delivering today are things like DECFORM, which is a forms entry terminal. The systems designer puts on what is to be in the form, and this is a black box that talks to the terminal entry clerk. She can fill in the blanks and this automatically goes into the machine. This is transaction processing. It also adds to the protection of the system because the people who use the terminals then cannot get to the data bases, they cannot contaminate the system; they can only deal with a particular transaction that they are used to. In doing that kind of thing, we have made it easier for the programmer; instead of having to think about back-up, fail safe recovery, journalising and so forth, all these things start to be done for him, so that the actual programming module begins to shrink in terms of what functions have to be put into it. He has to worry less and less about his file system. I think of a computer system as a terminal talking to a data base; so we start to put all of the validation into pre-canned programs, and we take all the file handling and put it there; and therefore the person who writes his transaction task has to do very, very little. So that takes down the cost of programming.

We then have products which are query and report writer. We have one, recently announced, called DATATRIEVE. That is a query and report writer that takes all files, all data types from all languages, to our record management system. He can deal with those; pull off a file; sort it; select on it; and do a number of things. We have it set so that, for a totally unsophisticated user, a non-programmer, he can use most of the functionality of that DATATRIEVE system, with no training; he can train himself.

Then we have another level of complexity that can be used by the programmer or the sophisticated systems analyst. What I expect to see over the next very few years are actual programming environments where people like financial analysts, managers, people who are just not programmers at all will be able to create algorithm solutions through a combination of ultra high level languages, metalanguages; through the use of on-line teaching aids by example; and on-line, soft copy documentation. I believe that this is very possible, feasible, and well within the state of the art. This is the path that we are pursuing. We want to solve the problem of the cost of programming; we want to reduce it significantly; the cost of maintenance; the cost of documentation; and not only the cost of the programmer but the fact that he is unavailable.

I was with a bank, two weeks ago, which uses PL/1 as a language. They do not mind spending the money for the PL/1 programmer, but they cannot find him. So for their casual reports and their casual look-ups and so forth, they want to use something like DATATRIEVE, so that they do not need the skilled programmer to do these one-time reports, one-time updates, one-time extractions. So this is the general development. We are doing research in the area. We have models and so forth. We have products. I, for one, am very much interested in pursuing this to its final end, which is, first, so that people who understand the problem that they are trying to solve can have a tool that they can use very easily.

We have some products today; and we have models. We like to put out products that really work, and work properly. Word processing, which we have in the field now today, we tested out with our secretaries; we kept it in house; the same thing with our electronic mail system. We use our time sharing for testing our own computers in the factory. I have models of things that my own financial analyst uses. I have an SPR, which is a software problem report. This is a totally untrained individual. I have computer operators, and I give them models of things which are advanced languages that I hope to have out in the next two or three years.

I find two things; that totally untrained people, who never programmed, never will and never want to, can, within the space of literally minutes, get things on the air. Then I find that something else interesting happens: that these people like financial analysts will then begin to use the computer to create the kinds of reports that they really wanted, not the kind that they could get.

I have on the one hand researchers doing the work, and people who can use the models of the products on the other. Interestingly enough, I find that the researchers come to me and say, "Your financial analyst is not using this system the way I designed it." Because the people who think up the tool have a different kind of imagination from the people who want to solve the problem. So I have a feeling now that what we need is slightly simpler tools, not to throw in all the elegance, not to throw in all kinds of fancy constructions, but leave a little bit to the imagination and ingenuity of the individual to use it. He starts to use it in very innovative ways, rather than if you give him a highly circumscribed system that he cannot use very loosely.

One thing that we have at DEC is this synergy going between the research, the advance development and the product development; and then a whole body of people who will use our products. So we have a good time debugging them very early. But that is the direction that we are getting towards.

BUTLER: I am going to exploit the chairman's privilege to follow that line of argument a little further, if I may Ed, because I think that it is a very interesting one and perhaps lies very close to the heart of what you have been telling us. Perhaps I could ask you to cast your minds back to the time when we were making the transition from the so-called second to third generation of processors, and much the same arguments were put forward at that time. The phrase that was used then was "the brute force solution to the problem"; since hardware is growing cheaper, you just throw more hardware at it and don't worry if the programs are rather inefficient, because that is more economic. But my recollection of that period is that we discovered that, however fast the cost of the hardware declines, the ability of inefficient software to consume hardware grows even faster. So perhaps what you've given us is a little bit of a do-it-yourself, home assembly, bad news kit; and that in spite of the reduction in the cost of hardware, our total budget will get bigger and bigger and bigger. Please tell me I'm wrong, Ed.

FAUVRE: You're wrong. When I was talking about do-it-yourself, it was not in terms of dealing with anything like operating systems, or even languages. But, if you can, picture a model of elements whereby a normal human being, a non-programmer, non-algorithmic thinker, can get on the air in the course of an hour, and can put out fairly complex-looking reports in two or three hours. Actually, I go to my financial analyst regularly, and he will bring me a report and I'll say, "How long did it take you to do that?" It would be a programming task, and he'd say, "As long as it took me to type it in."

BUTLER: I don't want to press you further than you wish to be pressed, but is your conviction that the trade-off between the improved productivity of people and deliberately wasteful use of hardware resources will in fact be in favour of the user, the result of the research done by your people or just a gut feel that it must be so?

FAUVRE: No. In terms of the kinds of things that I have talked about, in terms of these kinds of things which are not image processing but fairly conventional things, I believe that based upon hard experience, these kinds of systems can be made to work in small configurations, with floppy discs and so forth. Nothing fancy; no huge discs; no huge mountains of memory and so forth. No, I find that absolutely practical.

QUESTION: Well, I'm not convinced. I'm a commercial user and my computer costs keep going up and up!

FAUVRE: Well, they do go up. There is an interesting part of the curve which shows computer hardware going down and software going up, because part of that software is in fact commercially-supplied software or software application tools; but other things have got to do with the sophistication of some of the applications today. Now some of the applications that I wrote years ago were simplistic compared to some of the demands that we have today, in terms of validation, data verification, upper and lower case, editing. Think back 10 years ago to the kind of systems that we had then, with 7010s, 1410s, CDC 3000 Series, Honeywell 200s, and compare them with the things that we can do today.

As you were talking, I was thinking about the fact that I have a brother-in-law who used to work for Univac. They were trying to put in some 1108s with Exec-8 multiprocessing, at White Sands missile range. I talked to him and asked him, "How's that going?" and he said, "Well, the mean time between failure is five minutes — as long as we don't try to initiate a job." It was five minutes mean time between failure in the idle loop, it was really terrible. But that really has changed. I think that part of the business of the economical computer is the fact that you now can have a convenient single user system today, or a small, multi-user system.

Interestingly, within my programming group, we have a number of machines — we have many computers, probably 40 — of different kinds. I calculated that I could give each of the programmers an 11/34 in their office, cheaper than their using some of the systems in multi-use. You want to use these in multi-use because you do want to share data base, communications and so forth; but just in terms of doing a very basic job, it would be cheaper for me to give each programmer a computer. As I do, I give them all a terminal, and then they can hook up the various kinds of computers; and then the cost comes somewhat higher, because of the communication cost, line cost, disc cost and different kinds of peripherals. But you can do a very sophisticated job today, and very quickly do the programming, with these kinds of systems.

I will tell you a story in terms of some prototype systems. Somebody mentioned ADAM the other night, which is one of the emerging systems. There is a system on the West Coast of the United States, called GENESYS, which is copyright and owned by a software house. There was a contract that was let by the Los Angeles Fire Department for an application; and they had bids out to several software houses. Four or five software houses all gave bids very close to one another, for about a year and a half calendar time, and about five man-years of work, to install the systems. This little company came in a low bidder, by far, and they were awarded the contract. They got it up in 28 days, using their system. That presumed a productivity increase of between 10 and 20 times over COBOL, and I find that these systems are very practical.

Some of them are in the field. Some of them are very little known, because they are very small companies. Some of them are bigger. But that is the way we are going, and it is eminently practical. There are problems where you get into difficulties. I have friends at Harvard who have done research on the difference in terms of how people think algorithmically. This is one of the problems with the software design, because the software designer, the researcher and the programmer do not realise that users think in a different frame of reference; and that they have to design programs so that people can use them, not so that other programmers can use them. People like to design things that are elegant and interesting for themselves, and put in all sorts of fascinating functionality which is of no use at all to any outsider who would use the system. They clutter up the system, they make it overly complicated; they generate bugs and all the ancillary problems that come about that way. But these systems are on the way, there is no question about that.

With the mass memories, as you saw some of the performance curves on memory systems, solid state memories as they begin to take over, electromechanical devices — when that happens, which I believe will be at least, if not more, significant than the microprocessor itself, we can begin to throw away our file systems, the things which make the programmer type in and say, "This is my name. This is my password. This is my account number. This is my job number." "You didn't answer fast enough so I'm going to throw you off the system anyway." We will get rid of that.

Today, we have file systems and, as we get these cheaper sets of memories, we are going to front end them with a man-machine interface which will make it transparent to him, so he has to say, "My name is Joe, and I'd like to look at my file. Maybe I'd like to run a job." Then you can converse and talk to the computer back and forth, without going through idiotic sign-in procedures and so on.

These things are coming, and I think much faster than people realise.

QUESTION: You've been talking about software. On the hardware front we also have problems of compatibility, problems associated with putting a variety of devices onto a network. How easy is this going to be in the future? Secondly, do you foresee problems of compatibility as between DECNET and HDLC, SDLC etc?

FAUVRE: Let me answer those questions in order. First, yes, we find networks very important; and our philosophy in DECNET is to be able to interconnect our various machines together, so that they can talk on three different levels; one on a line protocol to what we call DDCMP, which is not too different from SDLC except that it is our own, and we have always built systems that do the kinds of things that we want them to do. Then we have NSP, a Network Services Protocol, and something called DAP, which is Device Access Protocol, which will allow us to transparently have someone work from a computer program in his machine and access a file in somebody else's machine, without the programmer having to be aware of that. So we are working on that issue.

However, we do support things like 3271, 2780; we are very interested in X.25; we are doing packet switching X.25 both in Canada and on the Continent today, and that is very important to us. We are very interested in, first, the interconnection of our equipment; and secondly, the interconnection with other hosts, very obviously IBM because they are the largest, but we also connect with CDC and others.

QUESTION: On a more global front, can you comment on future peripherals — typewriters and so on. Will it be possible easily to plug them into a network?

FAUVRE: That is very important to us. First, we want the network to be able to talk. We want the users and the programs and the data bases to be able to talk across those networks as transparently as possible; which means that in our file systems we are beginning to design in the things whereby, if a program asks for a file, the record management system itself will know where that file is, go find it on another system and bring it in; the user will be unaware that it is on a different system.

Now there are two aspects of this whole networking thing that have to be dealt with. One is the whole business of integrity, diagnosability and repairability. This year, we are spending almost $1 million just on system diagnostics. I was Manager of Diagnostic Engineering at DEC for the past three years, and we are spending almost $1 million in network diagnostics; not network programming, but network diagnostics. The other thing that we have to be able to do is that in networks, first, we have to know that whatever we give, we have the integrity that we will get the job done and analyse it. So while we have routing, we do not want to support random topology, which means that any node talks to any node, until such time as we know that we can backtrack and in the end case condition find out what is wrong. That is very important to us. It is also very important that we deal with the issue of network accounting, otherwise you may buy a network and put in an application, then find that the random, casual user will use up the time, and create very bad responses, unpredictable responses, and that is intolerable.

So whatever we do in the way of features, the integrity of this network is most important and the measurability of the network is very important.

You asked about intelligent terminals and so forth. Two things: we are putting in intelligence; we have intelligence. We have LSI 8s and LSI 11s in our terminals; we are making more and more specially oriented terminals for particular applications. We are putting intelligence into our communication multiplexors; we are putting intelligence now into our disc products, intelligent discs or distributed data bases, that is an area that we are getting into now. It is very important to us. All of the characteristics, both hard and soft copy terminals, that make it useful in a production environment to people. It is all right to sit back 10 feet and look at a television set flickering, but we are very concerned about people who have to use word processing. We are concerned about people who use data processing for eight hours per day.

BUTLER: At this stage we must move on to the next scheduled item on our agenda. But thank you Ed for an interesting talk, and a very informative discussion period.

THE COMPUTER AND THE OFFICE WORKSTATION

T.W. Hart

COX: You can already see, in certain areas, how the concept of the multi-function workstation is beginning to be offered on the market. That is one of the reasons that we have invited Terry Hart, the Managing Director of Jacquard Systems, to give us the final session of the day, and to put forward another view on what a manufacturer sees as a market opportunity and the products that might be offered.

HART: I am going to break my talk into three sections. First, I will explain the scene with regard to developments in the computing industry, and particularly developments towards the multi-function workstation, as Jacquard see it at the moment and as I personally see it. Secondly, I will then explain what we at Jacquard are doing to take advantage of this scene; in other words, how we have aimed our product at this particular market and at the trends which are exhibited by this market. Finally, I will end with a short summing up, and we can then go on to question time.

THE WORK REVOLUTION

Material Processing

Data Processing

Information Processing

JACQUARD SYSTEMS LIMITED

This, very briefly, summarises what we are talking about. We have so far been through two major revolutions in work. One is the material processing revolution, which we all know as the Industrial Revolution. Then, more recently, in the last 20 or 30 years, we have been through the data processing revolution. Now we are entering into the information processing stage. This was referred to by Ed Fauvre, who included image processing in this stage. We do not define any differences between information processing and image processing as a company. We see it as a total trend, a total development in the industry which we are in.

Why is this trend taking place? What is happening at this moment in data processing, in information processing and image processing, which is causing this trend in our businesses?

TREND TOWARDS THE OFFICE WORKSTATION

rising cost of staff
diminishing cost of hardware
diminishing cost of transmission
under utilisation of existing equipment
distribution of data processing

made possible by

increasing skills
micro processors
emerging technologies

JACQUARD SYSTEMS LIMITED

First, there is the rising cost of staff. The cost of clerical staff is perhaps one of the fastest rising elements in any business. Secondly, there is the diminishing cost of hardware; I am sure that we all understand the speed at which cost of computer hardware and computer technology is falling.

Thirdly, there is the diminishing cost of data communication, the cost of using lines, the cost of developing networks; from a hardware point of view, this is also an area of rapidly-falling cost. Fourthly, there is the under-utilisation of existing equipment. This does not only apply to computer equipment, but to office equipment in general. It does not matter whether you are talking about a photocopier, a facsimile transmission machine, a word processing system, a computer system, a badge reader — the total scene is of under-utilisation of equipment. So the logical trend is to try to combine these different areas of hardware into smaller, lower cost and more highly effective systems; because the more that one can combine the different office functions in one piece of hardware, then the higher utilisation that one will get out of it; and consequently, the more cost-effective that system will be in your business.

Finally, there is the distribution of data processing. This is another trend — a trend at a tangent if you like — towards taking processing away from the central computer installation and putting it on to small, low cost computers, away from the central computer; in other words, the decentralisation of data processing and the decentralisation of information processing.

This is already happening. The convergence of the different technologies is not something that the users are in a position to install; the convergence of technologies is still relatively new, a phenomenon of just the last two or three years. The distribution of data processing has been going on for the last seven or eight years, and is something of which I am sure everybody in this room has experience, either directly or indirectly.

It is because of these trends that we, as a manufacturer, and you as a user, are being encouraged towards a certain course of action. This course of action is to combine as many functions as possible in the one hardware and software system.

Now what is happening in the industry which helps us to achieve this, and within a relatively short space of time? First, there are the increasing skills not only of the people who are manufacturing the equipment, but the people who are installing it. There is more and more knowledge of communications, of small computers, of the use of simple programming languages. The general level of ability of people concerned with these systems is rising all the time.

Secondly, there is the emergence over the last three, four to five years of microprocessors. Microprocessors are at the heart of this new technology. We could not do it without microprocessors, because the processing part of the system necessarily must be small, and it necessarily must be highly reliable. I think that these points were also brought out by Ed Fauvre in his session. The microprocessor is the answer to our requirements because it does exhibit these facilities.

Finally, there are the other emerging technologies in data transmission, in word transmission, voice transmission, video transmission; the new technologies which also enable us to hang on to one microprocessor a number of different office functions, not just data processing functions but image processing functions and word processing functions. In other words, the whole of information processing.

What I am talking about has been happening for the last couple of years. If I can use our own Jacquard system as an example of this, until three or four years ago, the three areas of the data processing business, the hardware business, were fairly clearly defined. If you talked about a mainframe computer, everybody knew exactly what you meant. If you talked about a minicomputer, most people thought they knew — even if they did not. If you talked about an intelligent terminal, likewise people reckoned that they understood what it was. There was certainly a very clear break-point between an intelligent terminal and a minicomputer. You had one or the other. The break between the minicomputer and the mainframe was less clear. Companies like Digital and Data General did their best to make this division even less clear by introducing large, powerful minicomputer-based systems which, in terms of power and ability, were in fact mainframe computers.

As far as our own product is concerned, we have introduced a system which starts at about £8,500 end user price, and grows with one compatible range of hardware and software through to a medium sized minicomputer system at about £80,000 to £90,000. You will see that that also covers the lower ranges of the traditionally-based mainframe computers. I am giving you as an example our own particular product, but this is also a trend in the industry. There are a number of companies which are blurring the images in this way, and it is based on the use

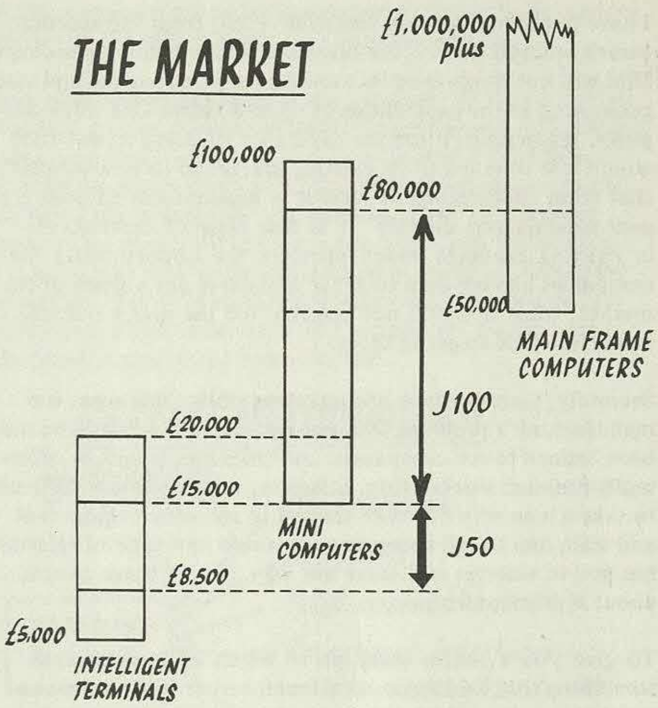

of microprocessors; the fact that one can put together a small, intelligent, single user system at about £8,500, which has a microprocessor CPU, can carry out a number of functions, including communication functions, and can be enhanced and developed until it becomes a very large, stand-alone minicomputer or small mainframe computer.

This is what is happening in the computer business itself. As far as the office multi-function workstation is concerned, this same system has also been designed to take advantage of the developments in the other areas of communications and information processing.

Before I do that, I want to propose to you a few words of caution, because the multi-function workstation will not happen overnight. Even though the technologies, particularly the hardware technologies, are in most cases here, the actual implementation of these technologies, both from the point of view of the manufacturer and the point of view of the user, is quite another kettle of fish. What we *can* do is quite different from what we are actually doing in practice.

CONSTRAINTS ON PROGRESS

Investment in existing products based on "old" technologies.

Lack of marketing skills

Lack of user skills

Lack of coherent policies towards Information Processing

Lack of knowledge

Cost justification

JACQUARD SYSTEMS LIMITED

I have put the reasons on this slide. First, from the manufacturer's point of view, is the investment in existing technologies. IBM will not throw over its existing ranges overnight and start promoting all the capabilities of its new Series One minicomputer. It can do it; it has the capability of doing it; but IBM would lose so much of its existing market to its new market that from IBM's point of view it is uneconomic to push its new product too quickly. It is this type of investment in existing products which provides the opportunities for companies like my own to get in and carve out a share of the market, while it is still not possible for the major manufacturers like IBM to get in there.

Secondly, there is a lack of marketing skills. This again is a manufacturer's problem. It is not easy to take a man who has been trained to sell computers, and train him to sell an office multi-function workstation. Likewise, it is even more difficult to take a man who has been trained to sell office equipment and train him to sell computers. A whole new type of salesman has got to emerge; and there are very few of these people about at the moment.

To give you a simple example — which again relates to something that Ed Fauvre mentioned earlier — the subject of word processing, which is basic to the multi-function station, requires the ability to converse in a whole new language. In my office we have just prepared a list of word processing or typesetting terminology which runs to seven or eight pages. If you are going to sell a word processing system effectively, you have got to know and understand the terminology. Not only has the salesman got to know and understand it, but so has the user; otherwise you have complete non-communication between the two sides.

Thirdly, there is a lack of user skills; in other words, no user is going to embark wholeheartedly on a new development before he feels confident about being able to obtain the benefits. So what we find in the business is that companies are just getting their feet wet by installing one system.

They may need hundreds of them, but the trend at the moment is to install one, try it, see how it works, get some experience, and then develop from there. So this again is slowing down the whole process of the implementation of this new technology.

Fourthly, as far as the user is concerned, very few companies have a coherent policy towards information processing. In most large organisations there is the data processing side on one side of the business, the man in charge of telecommunications on another side of the business, the office manager looking after typewriters and so on in another part of the business. It is very rare for the three of them to sit down with perhaps the O & M manager, or the management services manager, and work out a coherent policy about where they should be going. It is extremely rare; and that applies not only to this country, but to Europe and to the United States as well. I do not think that we shall see a really sharp take-off in this type of equipment until the large organisations determine their policies towards this type of equipment. This certainly is not happening yet, other than on a few rare occasions.

Then there is lack of knowledge on the part of the users about what products are available. There are so many new products coming on to the market at the moment in the computer business — microprocessor-based products — that it is almost impossible for a user to keep abreast of the latest developments. So this is a major problem: how does the user find out what is available, what are the new trends and the new products, and what can be done with them?

Finally, there is the constraint of cost justification. There is a very interesting point here: that as far as our own particular system is concerned, we have two quite distinct marketing policies between the United States and Europe. In the United States, we sell our system first and foremost for word processing, but with a data processing capability. In Europe, we sell it the other way round: as a data processing system, with a word processing capability. This is purely because of cost justification. In the United States, it is cost justified to install a small, single user microprocessor-based computer, simply as a word processing station. In this country, it certainly is not. About the only country in Europe where this is beginning to happen is Germany. So there is this element of justification. How will we obtain the maximum cost benefits from this new technology?

All I am saying here is that, although we have the technology, there are a lot of things that we can do at the moment which we will not really see taking off, as far as the user is concerned, for another two, or maybe three years, because of these inherent constraints in the application of the new technology from the manufacturers' point of view and the users' point of view.

That is the scene as we understand it at the moment. What I should like to explain to you now is what we at Jacquard Systems are doing to take advantage of this scene. We are one of the small, California-based companies that Ed Fauvre was talking about earlier. We are very unlike DEC in that we do not have 30,000 or 40,000 installations. In the September edition of *Datamation*, we were listed as having just over 600 installations throughout the world. That is a slight exaggeration, but it is not too bad an exaggeration. We have been in operation since 1969. We started life as a systems house producing customised systems based on other people's minicomputers for specialist applications. In fact the people who started the company all came from TRW and, prior to starting Jacquard Systems, were working on the NASA Space Project. So when they started the company they had very extensive experience of micro-miniaturisation techniques, communications, and of developing small, highly-efficient operating systems. It is these three areas of knowledge and experience which have been pooled in creating the Jacquard product.

The ideas for the product were generated in about 1973/1974; and the first prototype system was produced around late 1975. We have been developing it, selling it, and enhancing it since. The big advantage which we did have, and which other companies like us have, is that we did not have an inherent investment in other equipment; we were able to start from scratch. The philosophy which we adopted in 1973/74 was to design a system which would be small, low cost, microprocessor-based, hgihly flexible, and would follow the concept of a multi-function workstation. What I shall explain to you now is what we have done and how far along that path we have actually been able to develop.

Our design criteria as far as the product was concerned are set out on this slide. We wanted to design a system which would carry out with equal facility data processing, word processing and communications. The system had to carry out these three functions simultaneously, using any combination of peripherals. It is a screen-based system. The main method of entering data and processing data is that you enter data through a screen and keyboard, and the principle was that

SYSTEM DESIGN CRITERIA

- Carry out with equal facility
  data processing
  word processing
  communications simultaneously
- Standard peripheral interfaces
- Use existing products
- Interface to planned products
- Easy to install
- Simple to use
- Fit ordinary office environment
- Small
- Extremely reliable
- Low cost

JACQUARD SYSTEMS LIMITED

any screen should be able to do data processing, word processing, or initiate communications, in any combination with any other screen; and it had to be simultaneous.

The system had to use standard peripheral interfaces. What I mean by this is that we were not interested in designing our own discs, our own tapes, our own card readers; what we decided to do was to produce peripheral interfaces which were standardised on the industry leader in any particular type of media. So that our system could accept discs manufactured by X manufacturer, a card reader by somebody else, tapes by somebody else, an OCR reader by another company — it did not matter; what we wanted to do was to design a CPU which would take advantage of each new development in information processing as it came along. So we had to have these standard interfaces. We decided that our best policy was to standardise on the industry leaders.

This brought us to the conclusion that we should design a system which would use existing products as far as peripherals were concerned. We also decided to use existing products wherever we could in the rest of the system as well; which means, for example, that we do not build our own microprocessor. The CPU itself uses a microprocessor produced by another company. In no way did we wish to re-invent the wheel, we wanted to take advantage of the current latest products. So it was flexibility and standardisation which were the keys to the system.

In addition, we wanted to be able to provide an interface to new products as they came along. I will explain later how we have attempted to do that.

The final part of the list is pretty common to any system which is going to be successful in an office environment. It must be easy to install. It must be simple to use; in other words, it is designed so that the average person doing a clerical task in the average office can use it. It has to fit an ordinary office environment; it must not require any special air conditioning. The system plugs into a 13-amp plug; the CPU runs from a 13-amp plug; and all the peripherals plug into 13-amp plugs. It has to be small. It has to look like a piece of office equipment, not like a computer. It has to be extremely reliable in that one cannot afford to have it serviced by an engineer on a regular basis. It has to be placed in an office and run like a typewriter, and if it goes wrong you call the service engineer in. Finally, it has to be a low-cost system.

Those were our design criteria and certainly we have achieved them to a greater or lesser extent — I think mostly to a greater extent, and I will explain to you exactly how we did this. As far as the hardware methodology is concerned, as I explained just now, we decided to use existing items of hardware wherever we could. We did not want to get into the business of making our own peripherals because there are very good companies, mainly in the United States but with one or two exceptions, which make excellent peripherals. So we were not interested in producing items like this.

SYSTEM METHODOLOGY

HARDWARE

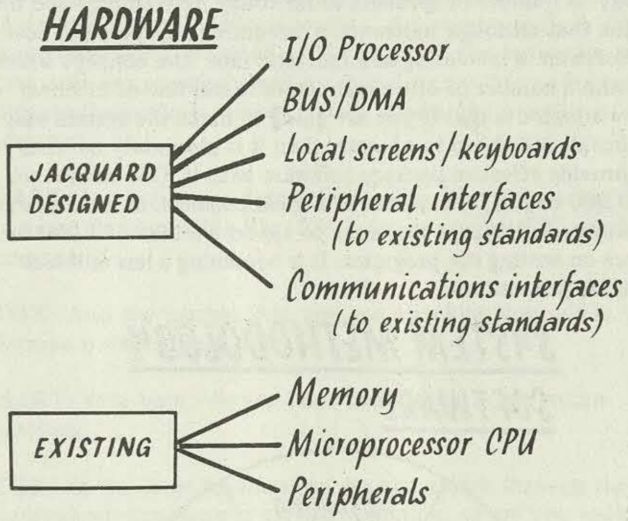

JACQUARD SYSTEMS LIMITED

Jacquard-designed items in the system include first the I/O processor. What I mean by this is that between the microprocessor CPU, which is a bought-in part, and the standard interfaces which control the peripherals, we had to design our own I/O processor to enable the two parts of the system, the I/O part and the CPU, to interface most efficiently and effectively. So that is a Jacquard-designed part of the CPU.

The system is based on a typical minicomputer bus, a DMA channel, which is again a Jacquard-designed part of the system. The local screens and keyboards for the whole system had to be able to provide as fast and efficient response as possible; it is a screen-driven, transaction-driven system, and we cannot afford to have a deterioration at the screen simply because there is a job running in the background which is hogging the CPU, for example. So we designed our own screens, our own keyboards, and our own controllers for the screens and keyboards, to ensure that, from a hardware point of view, this type of contention did not happen.

We designed our own peripheral interfaces and our own communications interfaces. Here again, we took advantage of the industry standards in communications in designing interfaces to meet the two standard types of industry interface, which is simple asynchronous and synchronous/bisynchronous (which is the IBM standard for intelligent communications).

As far as bought-in items are concerned, there is the memory; the microprocessor CPU; and, finally, the peripherals. Apart from the screen and keyboard, we do not manufacture any peripherals, and we do not mind which peripherals a user hangs on to the system. We say to our users, "OK, as long as it has an interface which is one of the standard interfaces we provide, you go along and plug it in." Nine times out of ten it will work; on the tenth occasion, when it does not work, we come along and sort it out. As a matter of principle this is a great idea, as a matter of practice, it does work. We have not yet found a peripheral which does not interface with one of our standard interfaces. If a peripheral does not have an industry standard interface, then it is another matter; then it is a whole new job of designing something separate, and we do not like to get into that if we can avoid it.

The methodology as far as software is concerned was also carefully thought out from the beginning to provide this same flexibility and standardisation. If anything, the software methodology is more important than the hardware methodology. A number of speakers so far today have emphasised the point that although hardware is becoming cheaper, the cost of software is increasing at a fantastic rate. The concept which we and a number of other companies in our line of business have adopted is that if you are going to make the system easy to install and cheap to operate, then it is absolutely necessary to provide effective package software with it. One cannot sell a £20,000 or $50,000 system to a small business, or even a large business, and expect that user to spend another $50,000 or more on writing the programs. It is becoming a less and less feasible proposition.

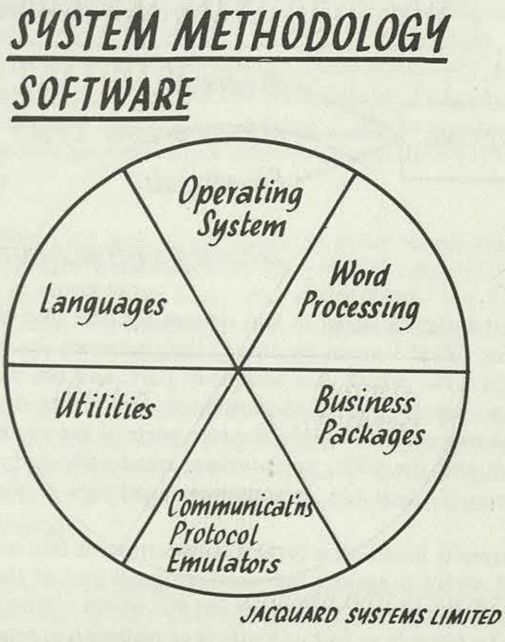

So with the system we developed and implemented a real time, multi-tasking, multi-function operating system; two programming languages; a range of commercial utilities; a range of communications protocol emulators to enable the system to communicate with an IBM computer, for example, or an ICL computer, or a Univac computer; and a range of business applications packages. I think that it is worth saying here that in fact to do this we had to develop two lots of packages. We developed one set of business packages which were suitable for the United States market; and another, rather more flexible, set of packages for the European market. It was quite impossible to achieve acceptance in this country, and even worse in places like Germany, of the constrained, 100% defined, rather inflexible packages which are quite acceptable in the United States. So we had to develop two sets of business applications packages. Finally, word processing. The word processing package is teated by the operating system just as a business applications package or an application program. One calls word processing in just as one calls any other program in, and off you go.

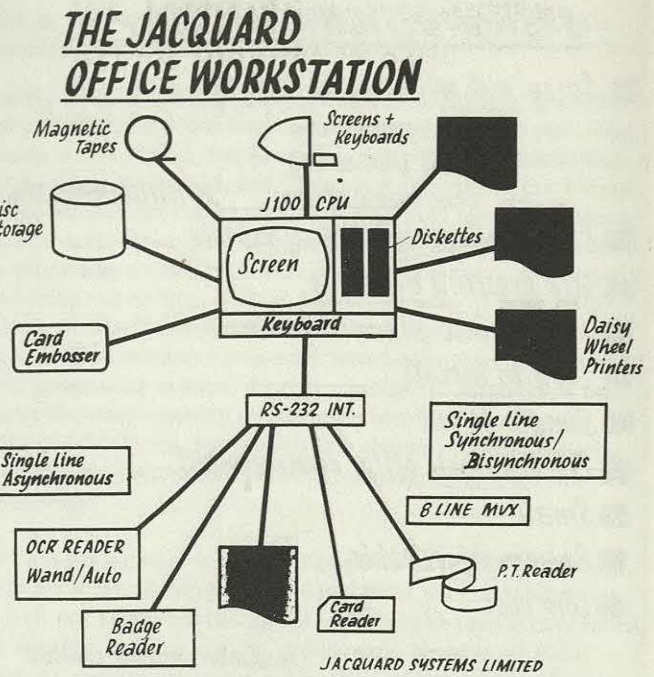

So what we have ended up with is something which looks rather like this: a central desk top computer which consists of a CPU, memory, screen and keyboard, a couple of floppy discs, and a couple of flexible discs. And it can all be run as a single user multi-function workstation. Into this system one can plug, at the moment, a range of peripherals. In addition to the standard minicomputer peripherals of disc storage, tape storage, line printers, matrix printers, and the screens and keyboards, one can also plug into it word processing or daisy-wheel printers as they are called, to produce a typewriter quality word processing print out.

Then through the standard RS 232 interface (the standard teletype interface) one has this whole range of plug-compatible devices which are appearing on the market at a fantastic rate at the moment. There seems to be appearing, certainly in the United States, a trend towards standardisation on the RS 232 interface for the small, cheap, plug-compatible, office piece of equipment. At the moment, we have things like OCR readers, either hand-held or automatic; graph plotters; badge readers; paper tape readers; card readers; and various other devices which I did not have room to put on this slide.

In addition, we have the standard communications interfaces to provide communications not only to a mainframe, but also to another Jacquard on a network basis; one Jacquard to another Jacquard to a Teletype to another remote intelligent terminal to a mainframe. One can mix and interchange the different elements within the system with a great amount of facility. The reason for this apart from the hardware aspect is again a point which was mentioned by the previous speaker, which is that the operating system treats any peripheral as a file. The application program goes to that file, picks up whatever data is in it, and then processes that data, depending upon the application which is being run.

That has taken you quickly through the design concepts or the philosophies which we adopted in developing this system, and what we have been able to produce in practice. Everything that I have explained to you and shown in these slides is here and now; it is either installed in Europe or in the United States. I have not talked to you at all so far about what we intend to do in the future. The reason is that as long as our design philosophies are correct, then we do not have to develop our own facsimile transmission device, for example, or our own voice transmission and receiving device, because these are being developed far better than we could develop them by other companies.

To give you an example, we know of a company in the United States at the moment which is developing a voice input and output device with a standard RS 232 interface. As soon as that device is available to us we will plug it into the Jacquard, and we will then have that facility. I am simplifying the thing a lot because we do not know — and I am not sure whether anybody else knows — exactly how successful this particular product will be. But as long as the principles which we have adopted are proved to be correct in practice — and they have been proved to be correct so far — then we do not have to worry ourselves too much about developing these other areas of information processing capability, because they will be developed for us. What we have to do is to make sure that as soon as that item comes on to the market we get one, plug it into our system, and make sure that our software is capable of handling it.

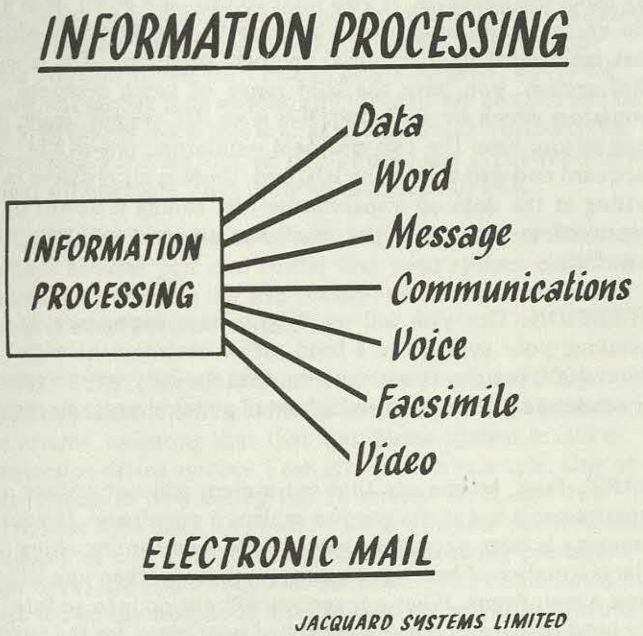

To sum up, the subject of information processing really covers the processing of these seven areas: data processing; word processing; message processing; and the subject of communications in general, consisting of the four areas we have covered on the Jacquard system so far. We have the systems installed and actually operating in these various modes. The areas of voice, facsimile and video we cannot do yet; but on the other hand, I do not know of any other company that can combine these with an office multi-function station.

I would not pretend, and I hope that I have not given you the idea, that our system is the only one which has been designed with this philosophy, and the only one that has reached this present state of development, because there are certainly a number of other systems coming out of the United States at the moment which exhibit these same characteristics. But as I said earlier, it is a question of getting to know about them, being able to see them, play with them and use them.

The final point on this slide: the subject of electronic mail. I must admit that I don't really understand what electronic mail is, because I think that, at the moment, apart from being able to transmit a signature on a letter, we can transmit the rest of the letter. Whether one needs to transmit a signature as well as the whole letter, to really have electronic mail in all its terms, I am not too sure. Electronic mail is becoming a little bit of a catch phrase, and is another term which is open to misunderstanding and requires definition. The major point is that as far as the office multi-function workstation is concerned it is not a catch phrase, it is something which we are producing as an industry; and is certainly, we believe, where the future of the industry lies.

Gentlemen, thank you very much. That is the end of my formal presentation; I shall be pleased to answer any questions you may have.

COX: Terry, perhaps I may throw in the first question, in trying to think of where a company such as your own and its product offering fits into the market. You mentioned something like 600 installations to date. Are these mainly with companies which are small users going into computing for the first time or upgrading small machines, or are they with large organisations which already have a substantial investment in systems?

HART: It is a mixture. I don't know the exact figures, but it is probably something like 40% large companies/60% small companies.

COX: And the market that you are attacking yourself in Europe is which, primarily?

HART: It is both. We see both those areas as our major markets.

COX: In the large organisation do you attack through the centralised management services function, or are you selling direct to the user?

HART: Generally, the only way we can attack is through the centralised management services function, because our *entree* in Europe has to be, at the moment, as a piece of distributed data processing equipment, because word processing at present does not turn anybody on very much. But what we find happens is that when we talk to the director of management services or the management services manager, then the word processing capabilities, and the other capabilities to hang on the other types of office peripherals, suddenly capture the imagination; and we find that, in many cases, we have the head of management services doing our selling for us, in selling our system to his users within the company. "Look what it can do. Not only can it do your accounts payable and receivable, order processing and stock control, but you can put your secretary on it as well." So it is a rather oblique sell.

In the United States, it is quite different. We sell probably the majority of our installations in the United States on their word processing capabilities; and data processing and the ability to put the system into a branch and have it communicate with the mainframe on an order entry stock control basis is secondary.

QUESTION: Earlier, Ed Fauvre was saying how in DP it has taken us 15 years to begin to understand the needs of users. In the office have we really analysed the needs, or are we just jumping on a bandwagon in order to sell more hardware?

HART: It is very complex. I doubt very much whether anybody has analysed the office environment 100% effectively. Even if it were completely analysed, I doubt whether it would have much effect on the trends about which I have been talking, because everybody wants to jump on the bandwagon. I think that probably the main basic trend is this one of the cost of labour in the office. Not so much the low productivity of labour in the office, although that is a significant factor, but the sheer cost of labour; and maybe combined with that, the lowering of standards. Again we find, talking about the United States, that a machine which will produce clean letters every time, even though a girl uses it slowly, is taken to be a fantastic thing. In Europe, this is not so; we can still get secretaries who work to a high standard of performance. But the pattern over the last 20 to 25 years in the United States has happened in Europe a few years later, maybe only a couple of years later; and there is no reason to believe that this will not happen in this instance. But I cannot say to you that it will follow a sort of logical, responsible pattern, because I don't think it will; it will just be getting on the bandwagon.

COX: You can't blame the manufacturers. You must remember that the manufacturer's real motivation, and his only fair one, is to sell kit.

HART: Perhaps that is an area which companies like Butler Cox should be devoting their efforts.

QUESTION: It seems to me that for office devices there comes a price break point above which the devices have to be used intensively to be justified, and below which their use can be less than intensive. In the US the price break is about \$2-3000; for the UK it seems to be about £1,000. Now I am interested in when this price break point might be achieved in the UK for a simple workstation with a CRT and printer. When might we see a simple device like that selling for less than £1,000?

HART: I think your question makes a lot of sense. Not only does it make sense, but I think that already microprocessor-based computers are falling rapidly in price. In fact, what is happening is that the peripherals tend to stay at about the same price, but it is the microprocessors which are coming down to rock-bottom prices. The cost, as far as the manufacturer is concerned, is going to be the software all the time. Once one has covered one's software development costs, then the manufacturer will be able to afford to sell the actual CPU at maybe just a few thousand dollars. So it is quite feasible that what you are saying will happen in four or five years' time.

QUESTION: You mentioned your emphasis on cutting user software costs by supplying packages. This is a very important point. Packages can be all right, but they are constraining. Are you implying that the package philosophy which you have adopted allows the user more flexibility to develop his own system variations than has been the case with mainframe packages?

HART: Yes. It is the philosophy in the design of the European packages — not the US packages — that the packages themselves are much more complex than one would normally expect; but because they are more complex, they are much simpler to implement and much more flexible. In fact the basic design philosophy is this: there is a screen formatting package which enables one to set up the screen formats without programming. In other words, it is completely parameterised, so you or I or anybody working in an office can sit and, once having learnt a few rules, set up the screen formats. There is another module which does the same thing for all the printing. So there is no programming involved in any of the printing, it is completely parameterised. There is yet another module which looks after the file structures. So you have three standard packages — an input, an output, and a file creation. Into those you slot your applications modules, whether they be order entry, invoicing, stock control, accounts payable, accounts receivable and so on.

This principle we have installed in about half a dozen installations in the UK. They are installed in Germany, in Holland, and in Denmark at the moment. The changes required by the different tax structures, the different VAT structures and the different government regulations have been implemented in the packages in a matter of a couple of weeks. So the design philosophy in practice is working in the way that we expected.

QUESTION: Can you use this package philosophy in conjunction with a mainframe in order to allow the use of large centralised files while still retaining local flexibility?

HART: Well, yes you can, because you do your local processing using the packages. If you want to call data down from a file on the mainframe, then you have to communicate with that mainframe by a standard protocol emulator. If it is an IBM system, you have the IBM range of batch protocol emulators which we also offer; if it is an ICL system, again, as long as you have the two protocol emulators, one in the Jacquard end and one in the ICL end, there is no problem in getting at the data on a mainframe file, calling it down, or transmitting data from the small system end into the mainframe.

QUESTION: Can you tell me if you have experience of installing your systems in a head office environment with about 1000 people. It seems to me that the only way to gain an economic justification would entail a major reorganisation.

HART: First, let me say that our system will not replace a mainframe; it is not designed to replace a mainframe. If your company is large enough to have a large, number crunching or a large number of batch processing operations, then you still need a mainframe. What our system will not do is to go into the head office and cover a variety of operations for the head office. You still need the mainframe to do the large, batch processing runs, and then maybe you can take a system like ours and put it into the purchasing department for it to do purchasing, and the typing and the clerical work in that purchasing area, but to communicate with the mainframe for the main purchasing accounting routines for the company. So it is a rather different philosophy or different concept of use from the mainframe concept.

QUESTION: Do you have any installations where a substantial number of your workstations communicate regularly with each other, and if so do they do it via a mainframe?

HART: The best example that I can give you, without quoting any names, is a user in London that has a CPU with four local screens. The screens are not in the same office as the CPU, but they are within a hundred yards, in the central building. That CPU is communicating with an IBM mainframe in another building about seven miles away. Then there are three other Jacquard CPUs in other offices, in other parts of London, communication over the telephone network with the central Jacquard CPU.

We have a system similar to that in Holland, which is in a firm of motor distributors, also running in much the same way, with a link from IBM to our system, and screens which are distributed both on a local and remote basis. Does that answer your question?

QUESTION: You are saying you don't have any systems acting independently as communicating word processors without the support of a mainframe?

HART: We have experience of Jacquard CPUs (say six or eight) in different locations, communicating with one another, without any one being defined as the central one. I cannot think of any installations like that in this country, but I certainly know of some in the States.

QUESTION: Could you describe the nature of the application in this US example?

HART: It is mainly a word processing system; in other words, each one of the computers is designed to do local word processing, but they are also sending messages, letters, if you like, from one station to the next. So you type messages in on one station, it goes over the telephone network, and is printed out on another one. I might add that we are also using the same principle in our own offices in Los Angeles, where we threw out all our typewriters; and the girls are not allowed to use typewriters. Now this is a bit of a gimmick obviously for advertising purposes, but we do not have any typewriters at all in our offices there, they were all sent back to IBM.

QUESTION: You seem to be avoiding knocking mainframes, perhaps because you said earlier that your system could be extended upwards to the size (nearly) of a mainframe. Could you give us an example of this?

HART: Well, it certainly cannot be built up into a large mainframe, but it can be built up into the power of a small mainframe, assuming that that mainframe system is still a transaction-driven system. I can give you an example. One of our systems in New Jersey has 10 screens on it, and two 80 million byte discs, and two 300-line a minute printers, and a couple of high quality Diablo word processing printers; a tape deck; and also a couple of Jacquard CPUs communicating with it over the telephone line. The system works beautifully; in a transaction processing mode you can update the file on disc, and it is a very large file with 250,000 records. You can have all the screens going at maximum rate and you get no degradation of the screens. But on the other hand, where it will not work very satisfactorily — this does not apply only to our system, but is a common factor with many computers — is if you put a batch processing job (say a complex sort) to run in the background, then it will run very slowly because the system is not designed to do long batch runs. Just as mainframes tend to be very efficient at batch processing, and rather inefficient when it comes to transaction-oriented processing, so minis are the other way round.

QUESTION: I'd like to know how you deal with the problem of security while working in the mainframe mode — which CPU handles recovery in the event of communications channel failure; and in Jacquard system to Jacquard system mode — in the event of network channel failure.

HART: The answer is a systems answer. It is not an answer I can give you because we do not build into the system any fixed way of doing this. What we say to the user is, "If this is your problem, let's sit down and work out the best way of solving it." It is solved purely from a systems point of view. The only recovery procedure that is inherent in the CPU is a power fail auto re-start procedure, which means that if it does go down you can start again automatically from the point at which you left off. But if you have MOS memory in the system rather than core memory, you will lose everything in the memory in any case, unless you have a battery pack to support it (which is a facility you can have if you want). But these are very rudimentary security procedures, and in my experience the answer to your question is different with every installation.

So our philosophy is to provide the minimum, but to help you build something more complex if you want to use it.

COX: Whenever a manufacturer or supplier is on the stand, you can only fairly expect him to put forward his own products and their particular appeal, and I think that is quite legitimate. But I think that what Terry has been describing, the philosophy of his company and the philosophy underlying his kind of product has rather wider implications. I am very grateful to him for a very clear explanation, and for answering some quite difficult questions. Thank you very much, Terry.

THE INTERNATIONAL PROSPECTS FOR VIEWDATA

R.D. Bright

BUTLER: We regard Viewdata — as I am sure most informed people do — as one of the most important experiments which is being carried out in public information and communications systems anywhere in the world right now. I am pleased to say that today, for the first time, we are going to hear about the international prospects for Viewdata. We are going to hear about it straight from the man who is responsible for developing those international prospects and turning them into reality: Roy Bright.

BRIGHT: One of the problems with Viewdata is that it is so diverse in its facets and interests that it is always difficult to know where to start. International, incidentally, cannot be disentangled from what is going on within the UK, and therefore, although many of you were present at the Butler Cox Seminar in September, it might be useful if we get on the same wavelength by running a film, which I promise does not last more than five or six minutes. It neatly encapsulates the many facets of Viewdata, and for those of you who have not seen the service demonstrated or presented on previous occasions — and incidentally I feel rather like a conductor without his baton, without a terminal here to perform with — I thought that this film might set the scene nicely to get us into the right mode before I go into the international implications.

FILM SHOW

Those of you who have seen Viewdata more recently will appreciate that some of that film is even now out of date, although it was only made about a year ago. The pages and various other features have moved on quite usefully since then. I guess that the feature which the film underlines is that Viewdata has many facilities which are capable of being exploited in a number of ways; and I do not think that we have seen the end of one's ingenuity to further enhance those features.

Two other points that I should just mention before we move into the international scene proper is the position in our plans for the market trial, and also a brief reference to Teletext.

Some of you may still be a bit confused about Teletext, what it is and where it stands in relation to Viewdata, so I have brought along a couple of slides to bring you up to date on that.

WHAT IS TELETEXT?

Teletext is the generic title given to the proposed broadcast information services of the

BBC – known as CEEFAX

and

IBA – known as ORACLE.

We take the view that Teletext is a good thing. Some people might feel that is a rather odd situation when it seems to be competing with us. But basically, if you think of Teletext as a low cost, and indeed low volume, source of information, that is really the main difference between it and Viewdata. Of course, Viewdata has many interactive features which Teletext cannot achieve. If we look at a couple of items here, I have tried to bring out some of those differences.

WHAT ARE THE MAIN DIFFERENCES BETWEEN TELETEXT AND VIEWDATA?

VIEWDATA has a technically UNLIMITED database.

It is interactive

It is a chargeable service

With Viewdata, as you may know, we do not have any real technical limit to the size of the data base, whereas with Teletext the technical limit is down to 800 pages per channel and in practice, because of people's impatience with waiting for their pages, it tends to nearer 50 or 100. As I mentioned just now we are interactive, whereas Teletext is broadcast and therefore plays really to a passive audience. But on the other hand Teletext is free once you have a acquired the set; whereas we will be charging for Viewdata, both at the telephone call rate and, in certain cases, the information itself will be charged for.

HAVE TELETEXT AND VIEWDATA ANYTHING IN COMMON?

Teletext and Viewdata are technically compatible using the same decoder circuitry and the same character coding (ISO 7)

They are certainly more complementary than competitive.

Because, as the final statement on that slide indicates, we regard them as complementary rather than competitive, it made good sense for us to work together to achieve common standards on the de-coder. So in future — for example in next year's market trial — all the sets that we are having provided by the TV industry will be capable of receiving Teletext, by definition. So one is able to say that those sets will become combined Viewdata/Teletext sets. We have adopted the same standards in the de-coder, and this same character coding which conforms with ISO 7. Those of you in the computer game will appreciate the importance of these international standards in future developments, which of course brings me on to the question of the international scene.

WHAT IS THE PROPOSED TIMESCALE FOR INTRODUCTION OF VIEWDATA?

1st January 1976	Pilot Trial Commenced (To define 'Product')
March 1978	Market Trial Commences (To test market the product)
1978/9	Public Service Commences (If trials satisfactory)

Our international plans are related to our plans for the UK development of the service. As you can see from this slide, we started a pilot trial back in January 1976; and we plan to open the market trial in June 1978; that is to say, the first end users of the service will become active around June 1978. Given that the results are judged successful, both by ourselves and the TV industry and the information providers, then we would expect to see some form of limited public service commence during 1979.

We have a world lead of the order of two to three years. The reaction is often one of surprise that we have this world lead over what one regards as the power houses of customer technological innovation, namely America and Japan. I think that the reasons are that the USA, for example, is hamstrung with common carrier regulations, particularly AT&T with its consent decree of 1954. It is almost ironical that in the land of free enterprise such a situation exists. Nevertheless, as things stand at the moment, AT&T could not even operate an arms-length subsidiary to handle a Viewdata type operation. But in addition to that, other factors should be considered; the penetration of cable TV in the States, and indeed in Canada, coupled with the rather more fragmented television broadcast situation over there, has led to a situation where currently they do not have the equipment of a Teletext service. Let us recognise the importance of Teletext in this country in helping to create the right atmosphere for Viewdata.

On the other hand, in Japan the differences are rather less apparent at first sight. Certainly one would not claim that there are necessarily political or commercial obstacles, but it is more a question of the language problem. As you will appreciate, the Japanese language does not consist of an alphabet as such, but of things that do not lend themselves readily to narrow-band transmission on which Viewdata is based. They have talked for the past decade about their plans for a wired city, and I think that this might be the explanation of the Japanese situation.

The UK Post Office obviously has one or two axes to grind when it comes to exploiting Viewdata internationally, the most obvious being one of creating some revenue opportunities based on our sales of the know-how software abroad. But there is more to it than that. Those of you who follow the Press will recognise that nationalised industries — and the Post Office is not outside this criticism — are sometimes accused of not encouraging British exports. But we think that with Viewdata we have an opportunity to redress that balance by making openings for the computer industry in this country, TV manufacturers, and indeed for software houses — all of which could capitalise on this world lead along with ourselves.

The other important area that we are very anxious to consolidate is international harmonisation — such obvious points as adopting ISO 7. Clearly, international harmonisation makes good sense. It is not just a question of directing one's thoughts to the ability to set up calls across national boundaries, although that of course is an important ingredient; but if you think about the technology that was shown in the film, the large-scale integration, the components industry and the semiconductor industry are very conscious of their markets being measured in multi-national terms rather than just national terms. Anything that one can do to encourage more countries to take on the same standards, the greater the opportunities for bringing down the cost of those chips or making chips available to set manufacturers so that the total price of those sets with Viewdata Teletext becomes more acceptable to the general public.

So those are really the three keys to our international objectives: create revenue opportunities for ourselves; encourage exports; and encourage harmonisation. They translate into a strategy which I think I could sum up in a few words, namely, we are already, with our German Bundespost sale, pushing to sell this know-how and software to other PTTs, particularly around Western Europe. That will encourage harmonisation on the one hand, create greater markets for the semiconductor industry and so on. So that achieves all the earlier points that I have just listed.

Another less apparent strand to our strategy to which I would like to return a little later is the ability to exploit Viewdata on a private business system basis. So far, most of the publicity and most of the thinking, publicly anyway, has been directed at Viewdata as a public service. But already there are signs both in this country and in others that people are beginning to see the value of looking at Viewdata as an in-house tool with the ability to handle internal communications at perhaps a price which is much lower than hitherto, and that makes good commercial sense.

Having mentioned private business systems, I should make reference to another feature of the public service which sometimes confuses people. If you consider the ability to identify the user as well as the terminal, you can now start to make arrangements for certain parts of the data base on the public service to be restricted to only those people with certain passwords. This is a facility that we commonly refer to as the Closed User Group. A couple of examples illustrate the ramifications. First, a company can afford to put its own information on the system which only its employees should be able to access, and there are many examples of that. Another variation is when a company offering a service to clients might wish to put information on there that only its clients can access. That again makes good sense from everyone's point of view, because those hitherto specialised services can now ride on the back of the public service with the relevant economies of scale.

One other strand to our international strategy must be considered, particularly in the international context — and that is access to third party data bases. Many of you know of the trend in recent years to set up specialised networks — even now the forthcoming EURONET is yet another example of this — enabling terminals in one country to access data bases in another, possibly using as a communication vehicle someone else's communications network. Obviously with Viewdata we are in a position to capitalise on that type of application. As the system grows in the UK, we would expect there to be occasions when our access to a data base is not a case of creating the duplicate of that data base on the Viewdata computer. Rather we would use our computer as a gateway into someone else's data base. Clearly one can extend that beyond national boundaries and look at this as a means of accessing foreign data bases.

That brings me to one important point, which I thought I might dwell on for a moment or two: the political implications of information transferred across national boundaries. I do not profess to be an expert on this, and indeed I am sure that many of you here have addressed the problem more closely than I have. But it does occur to me that this trend towards what is sometimes called "transnational data regulation" is an important influence on our thinking for the future.

I recently took note of a couple of quotes. One was from an OECD Symposium in Vienna, earlier this year:

> "A current and concrete issue of international interest is the possible effect on multinational communications of emerging national privacy laws, which may establish different rules governing the export of personal files."

Another one which I think was from *The Times* of London:

> "The absence of an international agreement governing computer networks and the transmission of data across national frontiers affects many computer users. Not only the large companies and firms offering international computer services are involved, but also smaller companies, some of which are having their data processed abroad without being aware of it."

This is very true.

These are points which Viewdata will bring even more into focus, and it is something that we have to address in our plans when we talk about "international Viewdata" or "international transfer of information".

To give you an example of how it is already impacting our national plans, let alone international plans, our recent dealings with Germany have highlighted a problem that they have which results from legislation. This lays down that the Bundespost, the German Post Office, is able (subject to government control) to handle the Federal Communications System of that country. However, the responsibility for information, or what they loosely call "cultural aspects" which concerns information, is the responsibility of the State Governments in Germany. So straightaway they are in a dilemma with Viewdata, because the Bundespost as a communications authority, the carrier, can handle Viewdata in the way that we are handling it over here, but as soon as they start looking at the information content involved they suddenly realise that this legislation creates a barrier to the sort of information that we are currently planning.

They can get round this to some extent, because this particular regulation or legislation is not so tight as to govern their use of Viewdata for specialised business applications. So already they are making plans for business information activities, but the general public information content that we plan to involve will be the subject of further legal wrangles in Germany over the coming months.

Sweden is another interesting example. The cause is different but the effect is the same as in the German example. Their legislation lays down that the act of broadcasting information is vested in the State broadcasting authorities. Unlike our Broadcasting Act, which has carefully included the word "wireless" broadcast, theirs does not. Even in the past, the Swedish PTT have had to obtain permission or special dispensation to put out information services over the telephone like TIM and WEATHER; so they have a problem to consider.

I am just making the point that Viewdata raises more than the conventional communications issues. It raises a lot more political and legal issues as soon as you start talking about the information *per se.* Those are some of the legal or political problems. Then there are the more practical problems that face us with Viewdata activities abroad. An obvious one is the difference in languages, as I mentioned in the Japanese case just now. But if you stop and think about Arabic and Cyrillic languages, you will appreciate that they take up a lot more horizontal space — they spread along rather like shorthand does — and often, just to confuse matters further, they read from right to left, which does not help.

Then there are TV standards which differ in different parts of the world. You will be familiar with our own standard, PAL, which operates throughout most of Western Europe. I say "most" because the obvious exception is France, with SECAM, which in turn is also used in Eastern Europe and Russia. Then if we go to North America and Japan, they have their variation called NTSC, which is a 525-line version, whereas our is 625. One of the first things that we discover when we look at a 525 standard is that the number of lines of information that we can accommodate, which is 24, now suddenly compresses down to nearer 20. There is even doubt whether we could achieve the full 40 characters across the row; it may drop down to 36 or so. So there are problems there.

Another problem with which you will be familiar from your other hardware involvements is the difference in power supply around the world. That is not a major hurdle, but it needs to be borne in mind in your plans.

Having mentioned language differences, there are even fairly modest differences between ourselves and other European languages, typically the umlauts in Germany and Sweden, and the accents in France and other countries. However, fortunately the Viewdata character set, with its use of escapes, can accommodate those, and we are working on a version which will be available some time next year to overcome those particular problems.

Turning to North America, the Chairman did suggest that I take advantage of my recent visits to Canada and Atlanta, for the Intelcomm 77, to give you a feedback. I warned him that we were so busy on our stand and in the conference that we did not have time to do much exploration of other people's activities; but I think that it is fair to say that the reaction that we were getting on Viewdata, both in Canada and later at Atlanta, was that people had not realised that we had gone so far. Most of them, if they knew much about Viewdata at all, assumed that we were still very much at the laboratory stage; and the ability to show that we had some 10,000 pages already on, and over 100 information providers active (the commercial end of the operation as distinct from the technical) I think really impressed them.

As I mentioned earlier, they immediately looked inwards and asked, "Well, what are our chances of Viewdata in the USA." If one were to make some sort of forecast about the possible penetration of Viewdata in North America, one might see this as a much more fragmented type of operation than we are expecting in Western Europe; various companies acting as entrepreneurs perhaps setting up their own Viewdata operations, not necessarily local but certainly limited to various parts of the USA, rather than the broad national plan that we are designing for UK purposes.

Another feature in America is the penetration of cable TV. It is interesting to look at Canada and discover that 60% of their television customers are on cable TV; and if you go into places like British Columbia, that figure jumps to 80%. So cable TV is a big thing out there. The only trouble is that the current generation is uni-directional; that is, although they have the bandwidth, it can only flow in one direction. To overcome this, they would have to spend a lot of money putting in bi-directional repeaters, which will take time and money. So we think that cable TV, in the early days, might well provide a souped-up version of Teletext, that is, a much enlarged data base capacity but still a broadcast-only type of activity.

If we can turn briefly to the interactive features of Viewdata, you saw on the film some of the message capabilities demonstrated; mention was made even of the Telex capability. Consider a terminal like a business terminal rather than a glossy residential colour set, with its normal telephone access via the PSTN (Public Switched Telephone Network); and, by means of the PSTN, access to one's local Viewdata computer centre.

Firstly, we have a public data base, with residential magazines and business magazines. Then we have the ability to put up partitions to allow private information or restricted information to go into these Closed User Groups, which is still accessed over the same path and via the same computer but with a different disc set up. If we look at the adoption of Viewdata by other computer industry interests or product business interests there is the possibility that using the flexibility of the dial-up telephone, one could make separate calls to one's own private computer which is Viewdata-compatible. Hence you can still use the same terminal. And, because you are using the dial, it matters not whether you have dialled the public service, and via that the public data base or the Closed User Group, or, alternatively, a private computer.

The first addition that one could add to that, based on more recent developments, is that we should recognise that via the computer we should also be able to link out to another terminal with the message service. That is a store and forward message service. If the recipient is there, who needs Viewdata? You just ring him up and speak to him on the 'phone; but if he is not there, as you saw on the film, then store and forward is an important ingredient of Viewdata.

Then Telex. What we have done is to define an interface between ourselves and the Telex service which, for example, has to accommodate such things as the differences in alphabet; ISO 7, IA 2 is the Telex alphabet. We can make calls via a connection through to a Telex subscriber. But what is more interesting is that Telex itself is an existing international service, if Telex can provide access to 80 other countries for example, Viewdata now has a ready-made link to an international message service called Telex. This makes good, interesting sense.

So not only are we generating traffic on an inland basis, because Telex hitherto was regarded as a businessman's service (you did not find many Telex machines in the home), but in future one can see Viewdata penetrating the paths of communication and between the business community with their Telex machines, and residential or other users not so equipped with Telex using Viewdata.

I mentioned also the third party data bases. Almost certainly economics will come into play eventually and we will decide that we should have links out to other third party data bases, and thereby create an even bigger pool of information.

Now I should like to bring your attention to one fact that is common throughout all these statements: we are using the same terminal. Those of you who are familiar with VDU prices will appreciate that anything under £1,000 is getting right down to rock-bottom prices for VDUs. Already, the prototype business terminal that we have had developed and some of you may have seen glimpses of recently — it happens to have been developed by GEC but other companies are working down the same path — could come out, even in year 1, at something under £400. That includes the built-in modem, and an alphanumeric keyboard. If you then go ahead a couple of years and talk in terms of a growth of demand for this, then that figure will drop even more dramatically. Again, I stress, compare that to your so-called low cost VDUs.

Having made the point that the terminal itself is low cost and therefore businesses can afford to spread them more liberally around their organisations rather than having them in specialist offices only, we now say that this same terminal, armed with a conventional dial-up facility, can accommodate all these different operations. Such encouraging news does that make that I think it would be nice for a moment to play a little game and put up a supposed "Day with Viewdata in the Office", which I sketched out roughly this morning.

09.10 – Call out previous night's USA incoming message.
09.30 – Retrieve latest market prices and RPI.
10.00 – Check flights to Brussels – enter reservation.
10.15 – Send message to colleague – out at meeting.
11.00 – Access CUG to retrieve agency MR info.
11.30 – Retrieve yesterday's sales figures from private database

LUNCH

14.15 – Display urgent message stored during lunch absence.
14.20 – Hold conversation with Prod.Mngr – hard copy of figures.
15.00 – Return lunch-time call on Message service.
16.00 – Check Business Card files on tendering company.
16.20 – Retrieve latest Stock Market quote on tendering company.
16.45 – Check BR Travel Flash – 17.35 cancelled – catch 17.15.

Leave power on to store incoming overnight messages from USA/Far East!!

I have tried to illustrate here the various applications, with perhaps a little bit of artistic licence, but nevertheless I think that it brings home the point.
09.10: call out of the store and forward facility with last night's incoming USA messages. You happen to have left quite early, at about seven o'clock, and something came in even after that.
09.30: use Viewdata to call into the business magazines; for example, check latest market prices and perhaps the Retail Price Index figures which were published yesterday and have now been updated on to the system overnight.
10.00: you discover that you are going to have to get to Brussels and you want to sort out some timetable information; and, if you so wish, you can now place a reservation for a seat on a given flight.
10.15: you try to get in touch with a colleague on the 'phone, find that he is at a meeting, and you say, "Right, I'll put a message into Viewdata and it will be waiting for him when he returns from his meeting."
11.00: we access the Closed User Group facility. Let us assume that our company subscribes to a market research agency which puts its information on to Viewdata. We are equipped with the appropriate password; we can get that market research information out of the appropriate Closed User Group.
11.30: retrieve yesterday's sales figures from your private data base. That might be on your own computer or another Closed User Group within the public service, it does not matter. The terminal can get access to either source.
We have a break for lunch and, at
14.15: we come back only to discover that an urgent message has come in which was brought to our attention by the 'message waiting' facility on the terminal during our absence; which leads us to have an urgent conversation with the production manager and, because the information is relevant to a decision, we decide that we had better use a hard copy attachment on the terminal in order to achieve some record of what we agreed.

This result enables us to return the lunch-time call only to discover that the person in Italy is himself now missing at a meeting, so we leave a message for him on the message service. By four o'clock we are looking at various company tenders and we decide that we had better call into the card file, operated by such people as Exchange Telegraph, to have a quick profile of a company which is tendering for our contract. A few minutes later, we decide that we will have a look at that same company's Stock Market quote, at closing prices. Finally, at
16.45: we have a quick look at the travel situation only to discover that our normal train is cancelled, so we push off a little early and catch the train that has not been cancelled — if there are any — leaving the power on overnight to collect further messages from abroad.

That is a little lighthearted, but I think that there is more than a grain of truth in what it tells us. Indeed today in my office I can dial up the Bundespost computer in Darmstadt; they have given me a user number and I can retrieve information from their data base. And they only started playing with the service back in August. It just shows how quickly this thing is moving on.

My final comment is that I could have put on that slide "How did we manage before Viewdata came along?" but I will leave you to draw that conclusion. Thank you.

BUTLER: Thank you very much, Roy. Now who would like to begin the discussion?

QUESTION: On the subject of message switching, will you have an interface to Euronet or an alternative packet switched network? If so where would the packet network interface be — in the TV set or the local exchange?

BRIGHT: One has to make allowances over what time frame we are addressing in your question. But let us talk about the shorter term. Certainly we already have a study going on in our research department which is looking at the implications of packet switching related to Viewdata. The obvious first objective is to use packet switching as a means for inter-communicating at the high level of the network; that is, between the computer centres which are scattered round the country. One is constantly sending packets of information up and down those main network links in order to pass information say from Scotland down to London. So from that point of view, it makes good sense to exploit packet switching as soon as possible; and with X.25 hopefully being resolved and implemented in a subsequent Post Office packet switched service in place of EPSS, one would like to see that as being one of the prime applications.

Additionally, many of our information providers will be what we would call bulk suppliers. While many will be content to put in a modest number of updates or fresh pages every day or once a week, others — let's say the Central Stats Office — may well be putting in hundreds, if not thousands, of pages at certain times in the month. Anything that one can do to improve the ability to bulk update makes good sense. So we are currently working in collaboration with one or two of these bulk information providers to define technical interfaces such that, if you like, file to file transfers could be achieved more readily between their central computer operation and wherever the host Viewdata centre is that they are feeding the information to. And again, packet switching would be an ideal added feature of that type of interaction.

However, I think that the possibility of packet switching coming right down to the end user in a universal sense is well outside my scope of forecasting. What I would suggest is that, if there are end users who themselves become quite heavy users — rather like the bulk information provider — then one would look at the merits of a packet switching terminal being hooked in; and treating them if you like, as some form of information provider. At least in terms of the technical protocols involved we could do that. So I do not rule that out, but it would be the exception rather than the rule as far as the end users are concerned in the foreseeable future.

QUESTION: If you were to develop a packet network interface in the TV set, could you not have a fast fax service accessible from within the home?

BRIGHT: I suggest that once we get over the initial impact of the market trial, these sort of refinements and other possibilities will certainly come to the surface in order of priority. What I am saying is that already the order of priority recognises the impact of packet switching at the bulk information level. I think that your question is going beyond that. It is simply a question of how far we can project ahead in terms of what is being done and what might be considered necessary at a later stage. But I do not rule it out.

QUESTION: Roy, may I ask a question — a highly technical one? When you retrieve data from the German Viewdata centre, whose responsibility is it to ensure that you comply with German laws on data security and privacy? Is it the responsibility of the Bundespost, your responsibility, or is it a responsibility that you both have to satisfy?

BRIGHT: I think that the short answer is that we don't know. At the moment we can get away with it because the information content on the German data base is virtually nil, with the exception of a few pages of education and one or two other bits and pieces. Basically it is a duplicate of our indexing data base. For the very reason that I explained they are treading very warily down this minefield of legal constraints that presently exist in Germany in so far as the information content is concerned. So the problem has not arisen yet because there is no effective end information involved.

I think that what we have to recognise — and this is true whether it be UK, Germany or any other pair of countries — is that information has already an existing legal connotation over this question of copyright. In tackling that, I am sure that we will find that there will be some important questions to be resolved before we can say that the situation is stabilised. I guess that the point that you are making is that that legal refinement will have to be picked up at the same time as the whole question of copyright and other things are examined for the international exchange of data.

The point that that leads me on to, if I may elaborate a little further, is the question of how will people wish or be able to access Viewdata internationally. I was explaining the problem to someone in the bar last night. Let us take just the UK as a closed unit. If you look at the problems of the accounting and billing system that will be necessary to handle Viewdata for the UK population of users, that billing operation is a pretty horrendous problem in itself. The way that we are tackling that is to put users on their appropriate host local centres which will have a look-up table, which obviously will have to be kept right up to date, in order to check whether that terminal is a registered terminal and can be allowed into the system and so on. So you spread the load round all your local centres.

Consider the international situation in which any one of several hundred thousand users in Belgium, France, Holland or wherever, could dial anywhere in the UK over the ISD, which technically, of course they can do (when I dial Darmstadt I am simply dialling in over ISD). Think of the billing implications of handling that one, with Monsieur Dupre of Paris happening to call in the London centre and trying to get access. He would have to be a registered user on that London computer. The whole thing becomes too difficult to contemplate.

Furthermore, just another important legal aside, whereas we have some redress against bad debts within our own country, what redress would we have against bad debts from Monsieur Dupre in Paris if he failed to pay his Viewdata bill? So for all those reasons, I think that what we are moving towards is an assumption that terminals wishing to access information proper to a foreign Viewdata service will do so through their own national gateway; in other words, the terminal in Paris would dial the French gateway, which itself would then call up the information from London, or on some prearranged knock for knock basis.

But again, I make the point that if we get into that sort of arrangement the copyright of the information that is being shipped around between foreign administrations now, not just between an odd user in Paris calling in, must be seriously looked at and resolved. There are two ways of skinning that particular cat. On the one hand, you can say, "Spasmodically interesting information can be stored on its national centre and we can have this remote access arrangement via the foreign national centre". Alternatively, if it proves to be a popular piece of information in another country, why not have an arrangement whereby it is shipped in bulk and stored as part of that country's own data base, with appropriate pricing and some form of remuneration for the provider of that information? So there are a number of ways of looking at the information content; but I think that the practical ones of billing and accounting point to this solution of a national centre connected across to the other national centre rather than terminal direct.

QUESTION: I suggest an open magazine to add to your list of present magazines, called "Current Bad Debtors".

BRIGHT: That's an interesting one because we have a situation already on the pilot trial, where each user has his own "stats page" which keeps a running record of use and hence expenditure on Viewdata from a given date. From a practical point of view — I don't know whether we could get away with it legally — there is quite a lot to be said for using that as the bill itself; in other words, have an arrangement with the viewer which could go on his page as a reminder at a certain point in time in the month, saying "Would you please within the next 24 hours have a look at your stats page and make out a cheque to the amount shown". Now if you do not get that within a few days, you start sending him the conventional printed letter. But at least it would cut down some of the costs and the overheads of a billing and accounting operation. As I say, whether we can do it legally remains to be seen.

Notice that what users do within their block of rented space on the data base is largely their own concern. So if within a particular area of the data base Joe Bloggs Limited had several pages, some of which might contain normal information, but a couple of which might be devoted to a billing arrangement, then I think we would be quite happy.

QUESTION: Isn't there a danger of the whole system being clogged up by the billing procedures?

BRIGHT: If you mean by that, is the cost of administering the system going to be greater than the value that the user perceives of obtaining the information? I don't think that is a real danger. Then again you get this other problem with billing as I am sure you are aware: public utilities, but in particular the Post Office, are constantly harangued for not giving their users more detailed billing information. Our answer, which is quite a legitimate one, is, "Well, if you're prepared to pay more to cover the overheads that creates then fine, we could do that technically." Fortunately, with Viewdata we can go a little bit further down that path; we can not only keep this running record that I mentioned on Magazine 79, but also at the end of each call we have a facility which shows how much you spent on that last call, both in terms of handling charge and the information pages that you accessed. So I think that we are making progress on that front so far as Viewdata is concerned. All this should help us — which gets back to your original question — to keep down the administrative overheads in a way that, for example, the telephone service could not, without a considerable increase in its charges to customers. So I am hopeful on that point.

QUESTION: Do you anticipate that all the services you have shown will be available via a 'short distance' modem? Will the £400 business terminal be able to access international centres?

BRIGHT: I think that the answer to your question is that it is early days to say; but we are encouraged to think that our short distance modems would certainly have tolerances in them enabling them to work beyond the conventional terminal to local group switching centre, which was the design target. It looks as though there is sufficient tolerance to take a more ambitious view on that. That does not mean to say that any user, anywhere in the country, will never need — let's call it — a sophisticated modem. There are black spots, as you well know, in our network, and on particular routes from A to B you do encounter problems, such as from London Airport across to Barking or somewhere like that, you get quite a few transmission problems.

Given that our philosophy of local catchment areas is adopted, typically users would not need to go beyond that short distance simply because that is the availability of their nearest centre which happens to be a short distance away. But I do not think that we will find the modem restrictions such a big problem as we at first envisaged, certainly while we stay at 1200. If we start talking about 2400, which is a possibility as a later generation transmission speed for Viewdata, that situation may change; but on the present evidence, the 1200 rate is quite encouraging.

BUTLER: Ladies and gentlemen, I am sorry that we will have to leave it at that. I think that this discussion will doubtless go on enthusiastically over coffee. Clearly there are extremely important points which remain to be sorted out not only for the longer term future, but questions that are quite near term are still subject to debate and discussion. But every time I hear more about Viewdata, I receive an overwhelming impression that the project as a whole is now receiving and developing such momentum — not only from the point of view of what it is but from the point of view of how it is perceived by the world — that every day it seems to me the odds against the public service failing to materialise and to become an important part of our national and international communications facilities are beginning to reduce quite dramatically. In my opinion, the main credit for that rests on the shoulders of Roy and his team who have done such a brilliant job in explaining the product and the service, particularly to people who are not accustomed to thinking in such terms.

The second point that I should like to make is that those of you who have not been exposed in depth to Viewdata before I suspect, if you are like the rest of us, have now been implanted with something which we call the "Viewdata virus", which means that you walk around the whole time, looking at what is going on around you and thinking to yourself, "This is a good Viewdata application; that's a good Viewdata application." It becomes rather obsessive after a while.

My own personal favourite Viewdata application is one which I shall launch as soon as I can afford to become an information provider. It is called "Telepenitence". It is for people who are too lazy to go to confession. You have a list of all known human sins and you simply go through and check off the ones for which you wish to receive absolution, and the system tells you what penance you must make. It also has another advantage, that it puts forward a menu of interesting alternatives. Roy, thank you very much indeed.

GETTING RESULTS WITHOUT RISKS

H. Donaldson

COX: Gentlemen, the bulk of these two days is concerned with presenting developments in technology to you and exploring some of their implications. Because of the nature of the conference, most of our speakers are suppliers or manufacturers, or people with a vested interest in that technology. One would therefore expect them to speak about it knowledgeably, but with a fair degree of enthusiasm and not a little bias. To redress the balance somewhat, we thought that we should invite at least one or two users of this technology along, to give what they would consider to be a well-reasoned response and assessment.

The first of these speakers is a long-standing friend — up until this talk — Hamish Donaldson. You will notice that Hamish has already thrown a note of controversy into the proceedings by insisting on rearranging the podium to give a sense of polarisation. Hamish Donaldson is currently Director of Management Services at Hill Samuel.

DONALDSON: Thank you, George. Gentlemen, in our industry we are dealing with a technology which has a hypnotic fascination. George has, in the past, accused me of total technical fascination, of getting involved in detail and being carried away with enthusiasm; and that certainly was true until I became a user. Your attitudes change when you have to deliver next day! I will not be able to talk to you about the latest white-hot technology, because we do not actually use it; all I can talk about are the ones which we currently use.

The main point I want to make is that we need to be on our guard when looking at technological improvements. Things are not always what they seem. If you remember, numerically-controlled machine tools were going to revolutionise the engineering industry a few years ago; but when they were put into jobbing shops, everything went slower and was much more expensive. The greater degree of automation did not result in progress; it put it back.

Having said this, my talk will not be entirely destructive. I do not expect you to agree with everything that I say, but I hope that I can leave you with a framework which you can then use to make your own assessment of my subject areas.

When I was trying to get a theme for this talk I plumbed the depths of my literary knowledge and worked out that it is all a bit like "Alice Through the Looking Glass". Lewis Carroll, you see, was one of the early computer thinkers; what he was talking about in "Alice Through the Looking Glass" was "as you went forwards you often went back; as you went backwards you often went forward".

That is the depth of my literary knowledge, so I then went to my Dictionary of Quotations to go further. I found this:

"Beware the Jabberwock, my son!
The jaws that bite, the claws that catch!
Beware the Jubjub bird, and shun
The frumious Bandersnatch!"

The programming Jabberwock; the data Bandersnatch; and the Jubjub bird of word processing; these are my three subject areas. Three areas which are currently, if I read the course brochures correctly, the most popular ones to send your staff on. Improving Programmer Productivity is the first; the second is Data Base Management Systems; the third is Word Processing; and I want to round up with some sort of attack on Technical Complexity.

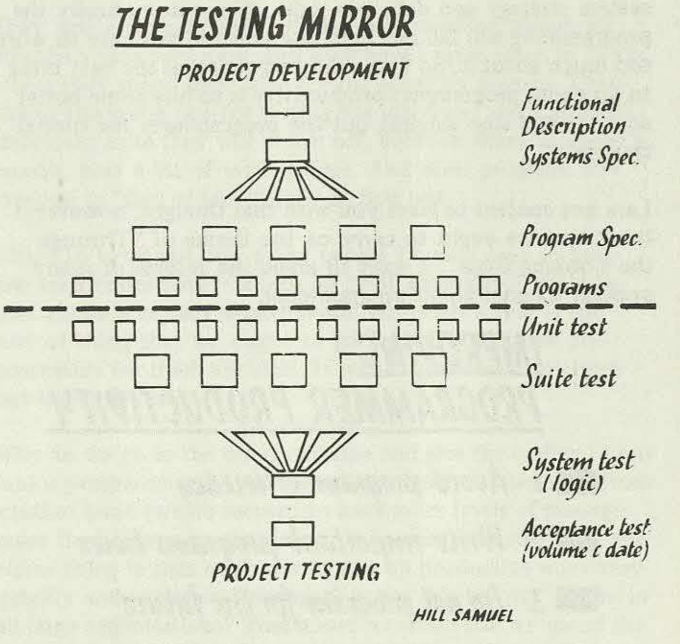

Let us start by talking about Improving Programmer Productivity. Just to continue with the "Through the Looking Glass" theme, let me show you a slide that we use when we are talking to our staff about testing. Our projects start off with the functional description, we then go on to system specs, and the problem has been broken down now into procedures and is getting bigger. Then we move on to our program specs, and we have quite a lot of programs now. Finally, we are on to individual programs.

Then we start testing, we unit test; we suite test; we system test; and acceptance test. (What I am saying, by the way, is not particularly profound, it is known to all of you already; but bear with me.) Any errors that you find in your programs are discovered in unit testing; any errors that you find in program specs are discovered in suite testing; and errors that you find in systems specs are discovered in system testing; and errors in the functional description are discovered in acceptance testing, parallel running, or actually being live.

Now who is the member of staff who is 'holding the baby' when any of these problems are discovered? It is always the programmer, isn't it, because it is the programs that are wrong. Therefore, a myth has built up that it is the programmer who is the cause of system problems. Lewis Carroll mentioned it:

> "The Queen was in a furious passion and went stamping about and shouting, 'Off with his head!' or 'Off with her head!' about once a minute."

The person who is holding the problem when it is discovered is the person we instinctively want to try to attack. We attack the programmer.

My argument is that they are the *last* people to worry about. My contribution to improved programmer productivity in my own company is to ask the programmers to program the job *once.* It gives perhaps the most dramatic improvement of all. I think that it was Hedley Voysey who once wrote in an article, "Why is it that we never have time to program it right, so we always have to program it wrong three times?" If you start off by thinking that you can improve programmer productivity by starting with programming, I believe you are on the wrong track altogether. Structured programming is not what we want; first we need structured problem solving. If you start the thing right and solve the right problem, get the system strategy and data sets right, then automatically the programming will fall out right and you will not have to worry too much about it. So I am really arguing that the best thing to do about programmer productivity is to hire some better analysts and stop singling out the programmers for special treatment.

I am not content to leave you with that thought, however. I think that we ought to carry on the theme of "Through the Looking Glass". I want to argue the reverse of many popular theories about programming.

IMPROVING PROGRAMMER PRODUCTIVITY

1. Avoid program efficiency
2. Write important programs twice
3. Do not program for the future
4. Do not maintain documentation
5. Forget program testing
6. Learn to structure three programs
7. Cut out pre-processors.

HILL SAMUEL

If you want to improve programmer productivity the first thing to *avoid* is *efficiency*. 'Efficient' programs are almost unmaintainable. They are complex; and usually they are efficient in the wrong direction. After all, we all know that only 10% of the code will be used often; and if the programs will be used only once a month, efficiency probably does not even matter. And what do we mean by efficiency anyway? Surely not program efficiency in isolation; it is *business* efficiency for which we are trying to strive. Getting the job working is a lot more important than saving the odd micro-second.

Secondly, it seems to me that you can improve your programmer productivity by *writing* your *important programs twice.* The important point here is that it is all very well structuring a procedure in the abstract, but until you have actually got into it, what you think is difficult because it is intellectually difficult is probably done in five lines of code; what you think is pretty straightforward because it is intellectually straightforward is five pages of coding.

Think of writing a report. It is all very well to write the contents list down at the beginning, but until you have begun to write the report you can not really get the balance right. This is true also in programming. So if programs matter, you should be prepared to scrap the first attempt and re-write it.

My third maxim is *do not program for the future* — it is time wasted. If you structure the data files for the future, the system will be capable of running in the future, but not otherwise. So encourage good programming practice (for example, table driven programs) but do not let programmers waste a lot of time thinking about what might happen in the future; get them to concentrate on getting the programs working and the system live.

Next point: *do not maintain documentation.* Let me explain why. Programmers in practice will not maintain documentation and the most practical solution to the problem is not to have any. We strive for good system specs and we like our programmers to structure their programs. If they want to draw flow charts they can do all of that, so long as they tear them up once their programs are working; if they structure their programs right, the code is a better guide to what is going on than all of these, probably out of date, charts. It is about all we can do to get system specs maintained; to get everything else changed as well is almost impossible and also unnecessary. Our program documentation consists of the system specs, the program structure and the code. It works much better.

The next point; *do not spend time on program testing.* Why not? It turns out that all the programmer will test are the things that he built into his program anyway. If he built a check digit verification in, for example, there is no real point in his testing it, because it is bound to work (unless he has made a coding slip, which is almost irrelevant). What you ought to test are the things that he did *not* do. What is needed is to link test the program with minimum data and then get on with the system testing. That is when you will discover what the programmer has missed out and what he has forgotten about. So do not waste too much time doing program testing.

Now for structured programming. I get the feeling that many of these improved programming techniques are designed to help IBM get over their lamentable productivity in developing system software. But how relevant are they really to the great bulk of business users? There are only three classes of program that we ever write. So why not learn to structure these three — (the vet, the update and the report). You do not have to go on any more courses.

There is room for creativity in programming without doing regular things wrong every time. Programmers, for example, are incapable of solving the page printing problem first time. Page printing is a problem because we have to allow room for

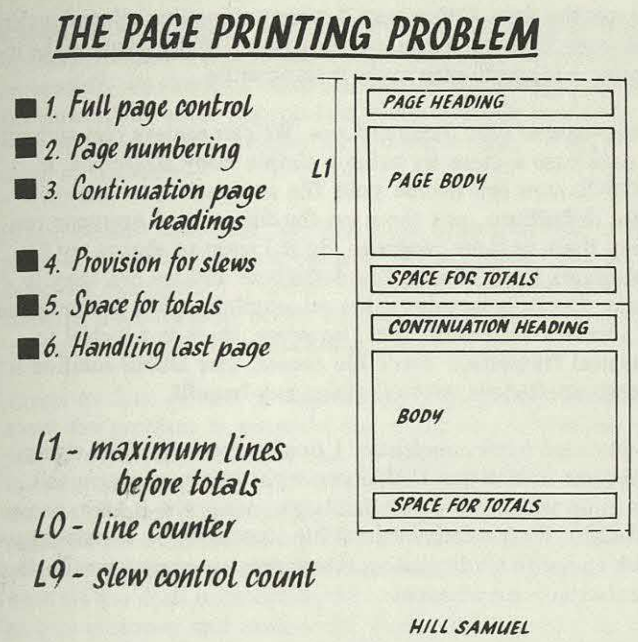

page headings, a page body, and space for totals (because we do not want the totals to go on the next page with no lines there, so there has to be at least one line for totals to make the document look sensible). The continuation heading may not be the same size as the main heading. (If you are submitting statements to a customer then the continuation statement will not be the same as the first statement). We have to carry out page numbering and continuation page numbering. We have to have provision for slews without changing control loops all the time. We also need to handle first and last page conditions; in particular, we do not want to print page headings if there is nothing actually to be printed (quite a common requirement with exception reporting).

M700 – PRINT

```
A700 IF LØ<L1 GO TO C700

B700 MOVE SPACES TO RP1, ADVANCE TO HOF, MOVE 1 TO L9
     ADD 1 TO L3 (page counter), move to edited field
     WRITE RP1 FROM HEADING -1 BEFORE ADVANCING
     WRITE RP1 FROM HEADING -2 BEFORE ADVANCING m LINES
     MOVE (m + n) to LØ.

C700 WRITE RP1 FROM WP1 BEFORE ADVANCING L9   LINES.
     ADD L9 TO LØ
     MOVE 1 TO L9
     MOVE SPACES TO WP1

EXIT 700 EXIT.
```

It is intellectually a difficult problem and very boring when it is wrong. So why not solve it right once. This is the code for the case when the continuation heading is the same as the normal heading; you replace the middle block by a different block if you want continuation headings. This is not a sub-routine, it is a module in the program (because subroutines make the thing unnecessarily complex). It solves the problem in a dozen instructions, is simple, reliable and comprehensive.

It is not really worth asking programmers to get this sort of thing wrong every time. Similarly with updates. Why should we always get the first and last record wrong every time? We ought to have a way of getting the next record, which goes and gets it and does all the matching for us.

What I am really arguing for is this: work out the things which you are likely to want to do often, and make sure that you do them very well.

Finally, *cut out preprocessors.* We used a good preprocessor once called COBAL, which generated COBOL code; it was a shorthand for generating COBOL. Another organisation with which I am involved had a preprocessor for BASIC PLUS on the PDP-11, called MPG; it generated vast blocks of code. One of my contributions to productivity has been to remove both of these preprocessors. What it comes down to is that if you do nothing but code they are very useful; but if you do other jobs as well, every time you go back to them there is a re-familiarisation process.

In certain situations they are very useful. Let me tell you about Stan and Dave, two Americans whom we employed once to program our peak jobs. Their productivity in coding was high, up to a thousand statements a day. They did nothing but code, which was their job. They were not interested in analysis work or anything else, they worked from specs. Again their emphasis was on the three classes of program: vet, update and report. They coded them always in the same way. They coded them extremely fast. You can see why a shorthand method is helpful to people who do nothing else, and they developed their own shorthand method; this is why they created COBAL.

They are extraordinarily productive, working about six months a year very hard; and they have a large ranch in Texas where they spend their time for the other six months. They charge the going rate for getting the job done, and a 30,000 statement suite they will polish off, between them, in about a month, plus a bit of testing time. And their programs run through to 'End of job' usually at first test.

I tried to work out how many people I would need to keep one such coder busy — about 50 or 60 analysts! It could be that above a certain departmental size such a skill is the sort of thing that we ought to go for; in which case pre-processors for them are ideal. It is their profession to churn out the code fast.

Why do we go to the other extreme and give the coding to our junior programmers? I was listneing recently to someone from another bank (which seemed to have more levels of management than I have programmers). It became very clear that the clever thing in that bank was to give up productive work very quickly and become a 'manager'. Is not this the risk we run in all large organisations? That above a certain size we spend the time on political advancement and fail to recognise the worth of the doers. Coding should be recognised as an important function with status for those who do it well.

So much for the programming Jabberwock. Now we come on to the data Bandersnatch.

I wrote down all the data base objectives that I could think of, as objectively as I could, and you see them on the next slide. I suggest in fact that the headings are the objectives of any good file structure, but they are commonly used as arguments for using a data base management system (DBMS). I define data base management systems (as opposed to a conventional file structure like ISAM) as inherently centralised (such as a Central Customer file) with a chained structure or similar. The first argument in favour of a data base is flexibility for change. If there is business change or reorganisation, then perhaps you can change the data base more easily than if you have a conventional file structure.

DATA BASE OBJECTIVES

FLEXIBILITY FOR CHANGE

program/data independence

progressive implementation

RELIABILITY AND SECURITY

error detection and recovery

ACCESS TO DATA

availability

currency and consistency

EFFICIENT USE OF COMPUTER

storage, processing and access

EASE OF USE

maintainability

portability

HILL SAMUEL

You can keep program and data independent and thus change the file structures without changing all the programs, you just change the dictionary. You can achieve a progressive implementation, because you know what you are going to do. You can achieve reliability and security because it is all controlled by one person. Error detection and recovery should be built in. Access to data is achieved by having it centralised. It is available because we can answer almost any pattern of question. There is currency and consistency; currency because you have one set of data and therefore it is current; and consistency because it is one set. You get efficient use of the computer because by holding data only once you reduce your storage requirements. It is conceptually easy to use.

The attraction surely is that you do not have to analyse the problem too much to know it is right. Also you can probably achieve better control of data within your organisation, because you have a central grip on it. Obviously flexibility is achieved as well because you can be flexible about not knowing what you want to do. Another attraction is that it enforces a discipline on the users.

Having enumerated these supposed advantages, I should say that my difficulty is that I have never yet found any actual situation where *any* of them really apply. Against any of these headings, it is difficult to find an instance where a data base is superior to a conventional file structure. Obviously, I have not seen the applications that everyone else has seen. But I would ask you to stay with me as I go down the list again to see how a DBMS scores against a conventional file structure.

Flexibility for business change. By business change, presumably we are not talking about a totally new application because that would be the same for either approach. Is it that we are opening a new division and we want to reorganise the sales data into this new division? I do not see how that is easier in a data base than it is in the conventional file structure. In practice it is likely to be quite a problem. If a customer is taken over by another (maybe it was an independent before and now it is a multiple) what do you want to do with its history? Does its history want to be as an independent or does its history want to be as a multiple? I think that the answer to that depends on what you want to do and how you are going to use the data. Either way it means going through and making selective changes. I cannot see that a DBMS would make it easier — it might even make it impossible.

Program and data independence. We can achieve that without a data base system by using a simple copy statement. In COBOL you can decide your file structure, you can write data definitions, put them on the library and everyone can copy them in their programs. So if I want to change my file structures I can change the definition and we can copy it again. The only penalty is the recompilation of a few programs. If I have a data dictionary, however, there is a logical to physical translation every file access. The DBMS solution is grossly inefficient without giving any benefit.

Progressive implementation. I don't agree with that either; it is clear that with a DBMS you must know what you will finish up with before you can begin, otherwise it never comes out right. With a conventional file structure you certainly have look ahead to study system interfaces, but progressive implementation is much easier.

Reliability and security. Security is not achieved best by denying access to information but by denying knowledge that it exists. An example will illustrate the point.

We have a common name and address system, which was originally designed to be a common name and address system for all our customers bankwide. A customer who uses more than one part of Hill Samuel will appear once — although he may wish to use many addresses (including his broker, accountant and so on). This central approach runs the risk that different customers can get confused across two parts of the group.

Such a problem occurred last week. A certain Margaret Joan Bridlington (the name is disguised) has been a banking customer for quite a long time. Recently, we got a new customer in our investment company who was called Margaret Joyce Bridlington, and was entirely different. So the investment company rang up the central people who allocate numbers (because we have a common name and address system) and ended up with the same customer number as the other Margaret. Because the investment company is a different processing centre a different sub-account number was used. So the entries were posted independently. But of course the system needs the customer's address when sending out statements. In this case address one was referenced by both accounts. It was not long before the banking customer started complaining about all these unexpected statements that she was getting in addition to her own. Meanwhile, the investment customer started wondering why she was not getting any statements. So the people in the investment management company looked at the file and realised that they had the address wrong; *so they changed the address* and she started getting all the statements.

Now if we were running the names and addresses in separate branches (and to improve security that is what we shall do) we just would not have had the problem. The data base approach has not proved to be secure enough in practice. Nor does it assist privacy of information.

Error detection and recovery. I was talking to a friend in the train the other day, who works for a large organisation that is using IMS. I said to him, "Look, I am trying to be objective about a DBMS, but I am concerned about security and recovery. The rules for our conventional files are straightforward; every file has a control record which contains record counts, has totals of key fields and so on, and the whole thing balances. We have a requirement that every time we read a file sequentially we check its controls, and we will always read a file sequentially on an appropriate cycle. What do you do in your data base?" He thought for a bit and said, "Well, we could do the same."

It then became clear that they had given little attention to the problem of what happens if records get corrupted — nor to detecting corruption. Further discussion revealed that verifying the Data Base would be difficult and time consuming — much worse than a conventional file structure.

Access to data. Access to data is easy if you have thought about the problem in advance, and absolutely impossible if you have not; the idea that you can throw everything into the computer is wrong. Information has an inherent structure and the problem is accessing it in different ways. We get this problem, for example, in our foreign exchange procedure where there are two patterns of questions. One pattern is: what are the deals outstanding for a customer? so we need to go to a customer and look at his deals. Another pattern of questions is: what is our net position on a given day in a given currency? So we need to access the deals by currency. One of the methods that we looked at was chaining. Should we hold all the deals by currency within value date and chain customers together? Or should we do it the other way round and hold them by customer and chain the value dates together? When we actually worked out the overhead we ended up by doing neither. We decided to hold the file in two structures. The outstanding deals file is held by customer, and it is effectively an ISAM file. We hold our net position on any day in a direct access file — a set of pigeonholes 400 days across and 25 currencies down. The file is kept up to date during the day and recalculated every evening. Given a date, the computer calculates the offset and gets the position in one disc access. We only need about a megabyte of relatively inexpensive disc storage to hold these positions. Disc accesses tend to be expensive (particularly in an on-line environment) so this is a good solution. Chaining would give us an unacceptable number of disc accesses and poor response time.

A key problem with any file structure is inserting new records. ISAM files are abysmal at it. When many new records are inserted the accesses get slower and slower and slower. You try creating a new ISAM file and compare that with writing a sequential file and backing it down. Last time we tried we had to stop the ISAM run after eight hours; we started again and wrote a sequential file and backed it down, and the job was done in three minutes. That is the order of inefficiency when you are inserting records — even with quite a good ISAM.

How much more is it true of a DBMS — equivalent to several ISAM files rolled into one. So data bases may work for applications where there is no movement of the files — for example housing records in local government. But if there is no movement there probably is no real problem anyway. If there is file movement the DBMS is therefore doubly suspect.

Currency and consistency. The trouble with data is that the currency is different for every user. The accounts department want to know a position at the end of the month, and that is absolutely useless to the person doing stock control every day. If we do a valuation of a portfolio, the valuation is on a given day and must be done some time later, because all the details have not yet arrived. We need to get them all in before we can do the valuation, so we need to freeze history on a given day. I suggest that there is no such thing as currency of data; it depends on the context in which you are using it.

I am not sure even that an item of data is very important in itself — it is the context that is just as important. An overdue account is a fact — what you do about it depends on many other factors. Historical data is required for different purposes for different people. So I am very dubious about the concept that data can be held just once.

Implicit in a DBMS is a desire to centralise. The argument is that if data is centralised then it is available to all. But it seems to me that the reverse argument is more valid. If information is decentralised then it will suit the local need and other users (including the centre) can always dial in when they need to. In this way the local data set is always current in the context of the local user — and in this sense the centre is just another local user.

Efficient use of the computer. Disc space is only relevant when the data files are very large (as some of ours are). We have not found however that a DBMS saves space. We certainly know that the DBMS adds processing and access time.

Ease of use. No. The DBMS tends to demand armies of software specialists. It is not easy to maintain or tune because it is not simple to understand. At present, they are not portable and are very expensive.

I have said enough. Data Base Management Systems appeal to the intellectually lazy. In almost every practical situation we have shown them to be inflexible, inefficient and 'insecure'. I suppose that there must be exceptions, but I find it very difficult to think of one. I understand that the CCA has a requirement that all minicomputers that are tendered must have a data base management system available or they will not be considered. My contribution to the CCA, to improve the effective use of computers, would be to ban the use of a DBMS in any government installation.

You could argue with me that the really skilled specialist can play tunes on a DBMS. But there are not very many of those people and anyway, what we are trying to do is to run our business efficiently. I cannot see that adding complexity, and slowing the system down and pandering to lack of analysis, is solving any problem.

Why do you not use my checklist for evaluating your own DBMS situation?

However, we do have a design strategy which we use in Hill Samuel. I do not want you to think that we do not actually think about file structures, because we believe that file structures and design strategy are a key to success. Our approach can be thought of as a double funnel; an input funnel and an output funnel. In a batch environment we recognise only three types of input data. One type is monetary transactions, movements, orders and so on. The second type is static data, that is changes to names and addresses and new customers. The third type is generated entries — which I will come back to in a moment.

We insist that one program only updates each master file. Bear in mind that in a bank, when things go wrong, it is important to put them back right, as they were, and recalculate all the interest. Limiting the number of programs that fiddle with the main files limits the risk — true in every application, not just banking. Once the files have been updated they are then frozen for the accounting cycle.

The business procedures are run as a series of 'back end' suites. The example on the slide is taken from our Investment

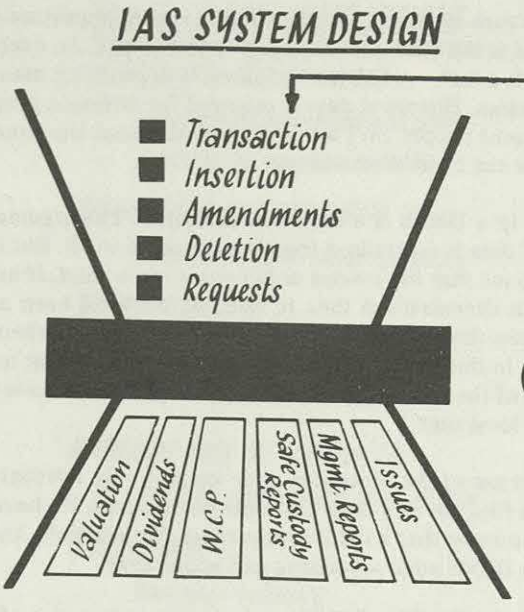

HILL SAMUEL MANAGEMENT SERVICES

Accounting System. Back end procedures include valuations, dividends, working copy portfolios, safe custody reports, management reports, issues. All of these subsystems can read any one of these master files and may write to none. That is the first rule.

The second rule is that they may not talk to each other at all; they are separate logical business procedures. The third rule is that they may not keep any permanent files.

But some procedures do need to change the master files. If we have an issue, say, a 1 for 4 BP issue, we solve the problem by generating input entries. For every four BP shares, in the portfolio, you are entitled to one more. So this is my third type of input; we call it "generated entries". The rule is that they must go round the system and are subject to all the audit controls of any form of input.

Broadly, all our systems are designed in this way; they have a very clear file structure. We go to a lot of trouble to freeze the file structure at an early date; and then we build our system round it. That is our strategy; and it works for on-line systems as well.

If you get the design strategy right, it does not really matter if some programs are written incorrectly — because you can re-write them. If you get your strategy wrong, however, it does not matter how clever your programmers are, you can never recover.

So we do have a data base philosophy; but we do not believe in a DBMS.

Word Processing. Now I want to move on to word processing, which is my third main topic. I was in a train the other day, when I heard a man talking to his colleague about word processing. He had seen the light. These were the phrases that he was using. "Office productivity has only risen 7% over the last 10 years". "Information is the most imporant asset in the business". Emotive statements, both of them.

He said, "Look, if you are typing a report using an ordinary typist, it's 3,000 words the first draft. Then it's corrected, and it's 3,000 words the second draft. Then it's corrected, and it's 3,000 words the third draft. 9,000 words typed in total. But if you've got word processing, it's 3,000 words the first draft; it's only 600 the second; and 36 the third." He was able to prove, you see, that the productivity improvement was 2½ times; actually he said 2.47!

Now I work for our management services department and I used to work for a consultancy. Reports are the only end product that we produce; yet we have never found it necessary to type everything three times. You will end up by Xeroxing the thing anyway if it is a report, and Snopake, scissors and paste are powerful typing aids. If the girl is sick or the machine has broken down, you can still work it all out and get the report out.

He went on to say, "It will enforce standards throughout the organisation." That seems to me to be the last defence of the manager who is scared of telling the girls how to organise their typing; 'it' will enforce standards instead of the manager.

By the way, he had not actually got word processing yet, he had only talked to the salesman.

I describe all this as the cleaning lady analogy. One of our neighbours interviewed a new cleaning lady recently; and, in prospect, she was magnificent; reasonable rates; industrious and so on. Our neighbour talked to all her friends about it and, within a very short space of time, her friends had snapped up all the cleaning lady's spare time. When she started work it turned out that she stole the teaspoons! It is this selling in prospect which bothers me; it is before you have got it that it is going to be good.

We do use word processing in our bank, and I am a believer in it in the right situation; but I do not think that it is a blanket solution to every problem. So what I have prepared is an analysis kit, to help us to look at the sort of jobs that we are likely to want to do in the office — and to see what the alternatives are. We might want to do an individual letter, for example.

WORD PROCESSING APPLICATION AREAS

		ELECTRIC TYPEWRITER alone	ELECTRIC TYPEWRITER +Xerox	ELECTRIC TYPEWRITER +offset	MTST	DISPLAY WP	WP in DP
Individual Letter	Q	H	—	—	H	H	H
	C	L			L	M	M
Circular Letter	Q	H	L	M	H	H	H
	C	H	M	M	M	M	L
Standard Paragraphs	Q	H	L	M	H	H	H
	C	H	L	L	L	M	M
Reports	Q	H	—	—	H	H	H
	C	L			L	M	M
Invoices/ P.O.s	Q	H			H	H	H
	C	H	—	—	H	H	L
Photo-composition	Q						H
	C	—	—	—	—	—	M
Flexibility in use		H	M	L	H	M	M

HILL SAMUEL

Across the top of the slide are the alternative ways we can do it: we could use an electric typewriter; or we could use MTST (which I have used to describe a magnetic card typewriter, the modern Flexowriter); a typewriter with magnetic cards so that you can backspace and correct. By display word processor, I am talking here about a device with a screen where you key in text, manipulate it, and then print it on a quality typewriter beside you. It is quite fast, about 40 characters per second, and the quality is just about as good as an electric typewriter.

By word processing and DP, I mean a computer, mini or mainframe, for doing data processing work with word processing as a by-product. Notice also that Xerox and offset litho are included as alternatives. Obviously, we would not use them for a single letter. For circular letters we might, for example when writing to all the sales customers so that only the names and addresses need to by typed every time.

Standard paragraphs is a third application area, say for a legal document where you are building the text from standard paragraphs. We will also consider a one-off report; a report on a conference that you attended recently, or a survey report or something like that.

Invoices and purchase orders are included; in most small businesses the individual letter and the invoices and purchase orders are perhaps most common in terms of typists' time. Photocomposition, the final area, can be a by-product of word processing also.

Let me analyse the chart rapidly. In each of the boxes we will assess quality and relative cost on a high, medium, low scale. On quality the electric typewriter is quite high; the MTST equally high; and word processing printers also. Some printers are not high quality but let us assume that if you are going to invest in word processing you will do it properly. So in general, all of these methods can give us a high quality end product.

What about cost? The cost of the *individual letter* is low when typed on an electric or MTST. It is medium on WP equipment in the sense that you have quite a lot of kit which will be under-utilised

With a *circular letter* the cost of doing it on an ordinary typewriter is high, because every letter must by typed individually and the names and addresses as well. If you Xerox all the standard text and just type in the names and addresses, then the quality will be lower but the cost will be more modest. It will not be as cheap, for example, as doing it by merging in a customer file on a data processing system. If you have your customers all on file and can merge the file with a standard letter the unit cost will be very low. Offset litho is similar to Xerox, but with higher quality.

Standard paragraphs will be high in cost on the typewriter because you have to type every paragraph again. Using the Xerox, it is low cost because you can use scissors and paste, but it is lower quality (which may or may not be important). The MTST can do standard paragraphs well but they are more messy than word processing machines.

For *one-off reports,* it is difficult to find a cheaper method than conventional typing. By report, I am assuming that you do not have to have 20 perfect tops; that you produce a document which is pretty good and then duplicate it in some way. The MTST is equally low cost and high quality. Word Processing machines I rate as medium cost for one-off reports— rather more expensive than using a typewriter intelligently.

Invoices and purchase orders are only low cost if you have computing power. What we are now seeing is that a word processing device and a baby computer, are moving together.

Photocomposition can be a by-product of a computer but it is expensive in comparison to typing. I rate it medium cost.

Flexibility in use. Typewriters are highly flexible because you can put them almost anywhere. A Xerox is only moderately flexible, but you can have one on each floor. Offset litho is not very flexible because you tend to have one in each division or group. MTSTs are pretty flexible, because they do not cost much more than conventional typewriters. Display word processors are only moderately flexible; either they are all put in one place (which means centralised typing and lost flexibility) or extensive cabling is required. The same sort of thing is true of DP.

So what do we conclude? What we are looking for is H for high quality with an L for low cost. That is the ideal. We see that the MTST is good for many of the things we do, such as individual letters, circulars, standard paragraphs and reports. However, the solution that scores consistently well is word processing as a by-product of data processing. But I am not sure it is right to ever generalise; I think that you should look at your individual problem and work out what suits it, using an analysis like this one.

WORD PROCESSING OBJECTIVES?

- Better use of secretarial time
- More productive typing
- Less typing
- Lower costs

HILL SAMUEL

Decide what are your objectives for word processing. Is it to make better use of secretarial time? How, then, does the secretary use her time?

TYPICAL SECRETARY'S DAY

Away from desk	30·6%
Typing	19·4%
Clerical work	19·2%
Communicating in person	7·3%
Communicating via telephone	4·3%
Waiting for work	8·9%
Taking shorthand	3·3%
Filing*	2·6%
Handling	2·6%
Self-authoring documents	1·8%

* an unmeasured percentage of filing is included in 'away from desk'.

I found the figures on the slide published in a WP article — and they look plausible. Notice that the secretary spends only 20% of her time typing, so doubling that productivity will not do us much good. Display word processors cost about ten times as much as a conventional typewriter, and they are no better at spelling. So I do not think that we will solve many of our secretarial problems by introducing word processing.

If the aim is more productive typing how will you measure it? I can get my girls to type everything twice — will that show up as more productive or less productive? Surely we should improve productivity by *removing work.* For example, we used to type our system specifications. They are terribly difficult to type and we found they were always full of mistakes (how can a typist tell whether BASCOM is a valid data name?). We used to have extensive dictation equipment also, but have thrown it out. In our business, getting the reports right is what is important. So we asked the staff not to dictate, but to write out their report, get it into shape and then have it typed once.

We improved productivity dramatically by removing the dictation equipment, two word processing machines, and two out of the four secretaries. The productivity of the principals has not been affected at all — they just think a bit more than before.

We can also improve productivity by *grouping the work.* Personal secretaries are poor in terms of typing productivity, and yet their support role can be invaluable. In our department we have solved the problem by having two secretaries to support the 30 analysts and programmers in the building — for all typing and administrative support. This ratio of 15 to 1 is good for a clerical department but it will be wrong to over generalise.

I find the idea of local working groups very attractive. After all, in computing we are trying to distribute the processing. Centralised dictation systems and centralised typing are the things we are trying to avoid. Perhaps what we ought to be doing is arranging that working groups have the right amount of service to make them largely autonomous; the right balance will improve the productivity and also be very acceptable.

Of course all these ideas can be wrecked if the boss is not concerned to get the balance right. Perhaps we should train the bosses before going much further.

TECHNICAL COMPLEXITY

Results from solving the wrong problem

Technical Complexity. I see that I have come to the end of my time, which allows me to put up my last slide which is on technical complexity. We are in an age where, because of fascinating technology, we get trapped into ever increasing complexity. I would like to suggest to you that whenever our approach gets complicated, we are running the risk of solving the wrong problem. If you need a computer to control your project management system, for example, you may well be approaching project management in the wrong way. If you need Critical Path Analysis to control activities, then you may well have the wrong breakdown of activities. If you need electronic mail and complex networks you may have the wrong business strategy. We should not be seeking better ways of solving the wrong problem.

Over centralised businesses do not work as well as those which are partitional into logical groups; groups which have responsibility with accountability and where systems are designed to support this logical group.

Data bases, electronic mail and networks are all there to solve a problem that was only created by central computers, and no longer actually exists. We can solve technical complexity by partitioning the problem — that is the trick. Lewis Carroll, of course, had the solution: what happened to the Jabberwock? I would like to tell you about the vorpal sword:

> "One, two! One, two! and through and through
> The vorpal blade went snicker-snack!
> He left it dead, and with its head
> He went galumphing back."

Here, then, is the solution. Break our problems up into manageable chunks; keep our solutions simple — that is the way that we can get results, minimising the risks.

COX: I thank Hamish for at least three good talks in an hour. Unfortunately, pressure of time prevents us from taking questions or abuse from the floor, but no doubt you will have an opportunity over lunch to put alternative views to Hamish, to quiz him further and to stab him with your forks. Hamish, thank you very much indeed for a most engaging presentation.

MYTHS AND REALITIES OF PROGRAMMABLE MICRO CIRCUITS

A. d'Agapeyeff

BUTLER: Alex d'Agapeyeff is going to speak about Myths And Realities of Programmable Microcircuits as the second in our series of talks on "Advanced User Experience"

Mr. d'Agapeyeff has asked us not to print the recorded conference transcript. The summary below uses the notes and slides he prepared for his talk.

d'AGAPEYEFF: Programmable microcircuits are the most potent, the most wide-ranging and most unbelievable advance ever made in computing and control systems. They will so change consumer and industrial products, manufacturing and administrative procedures that they form a threat to the on-going viability of many well-established companies, to levels of national employment and to the balance of payments of Western Europe. They are like a tidal wave sweeping toward a beach at an acute angle, drowning the succession of old products, practices, and other bulwarks of a former era. However, round the next headland everything is calm water, so no preparations are in hand for the coming deluge.

Now, you may feel that to be a fine piece of purple nonsense. Yet, if it were true it clearly would be a matter for managerial attention and understanding. Surprisingly, perhaps, for some industries it can be shown to be true. For example, from the production orders of microcircuits which have recently been placed, we know that both American and Japanese car manufacturers are planning to produce a new kind of car by 1979. There is little chance of an equally enhanced European car being produced in volume, prior to 1981. Similar threats exist for TV sets, radios, computers, industrial pumps and machine tools. If you can put a microprocessor successfully into a tiny, hand-held, micrometer — as PA have done — it would seem likely you can put it into almost anything.

Let us suppose, therefore, that the tidal wave suggestion may be true. I want to consider some of the questions which need to be asked.

PROGRAMMABLE MICROCIRCUITS

- **What a manager needs to know**
- **Judging their impact on a business**
- **How to get started**
- **Myths and Problems**

You will not, of course, expect full answers to these questions. In such a volatile field only a fool indulges in unqualified predictions and concrete solutions. Nevertheless we may arrive at some helpful hints and steps.

In order to grab your attention allow me to ask you a few questions which are admittedly intended to have a shock effect. Here in my pocket I have a wafer of silicon containing some 150 processors; how long ago would you have believed that to be possible? Most of you have a computer in your business costing, ignoring all peripherals and disc storage, between £50-250,000. Did you expect CPUs of equal power and working storage to be available in 1978 for about 1/10 - 1/20 of that cost (albeit with rather less software)? Did you know that my company demonstrated at Datafair a version of COBOL plus a small, file-processing, operating system on a 16K byte Z80 microcomputer whose CPU costs about £500?

Now is no time for amateurs. In Britain there is a convention that a manager should not appear to be too technical. A typical remark made prior to an important technical discussion might be "But I, old boy, know nothing about programming" — or whatever might be the subject in hand. Happily, this remark is not always true; it is part of a defensive posture. In any event this whole attitude must be changed if the right decisions are to be taken in the right time-frame, within this period of radical and rapid change.

It is absolutely necessary for the relevant decision-taking manager to have a thorough grasp of the fundamentals, in order to ensure he is being kept properly advised; to maintain an on-going debate, investigation or project continuously aimed at the main-line objective and adequately resourced for that purpose; and to enable *him* to explain to the rest of the management what the implications are of both internal and external developments.

THE FUNDAMENTALS

- **Nature of Integrated Electronics**
  (See 1977 September issue of 'Scientific American'.)
- **Progress in fabrication + products**
  (Skim 'Electronics')
- **Nature of programming**
  (?)
- **Progress in circuit exploitation**
  ('Fortune' and 'Business Week').
- **Progress in system development**
  (Local computing periodicals).

What are the fundamentals? This next slide provides a summary of this knowledge and where it might be obtained. I would draw your attention to the following points. First, excessive summarisation can be misleading — my aim here is to give you a taste for the topics. Next, this talk will provide you with something of the nature of integrated electronics, but that issue of *Scientific American* is vital reading*. Third, fabrication is the making of circuits through one of a number of competing and quickly evolving methodologies. In this field the American magazine *Electronics* is the best, and is well worth a regular glance.

The nature of programming is an important issue, but is rather obscure: we will look at this briefly later.

Fifth, the magazine *American Management* regularly reports new forms of exploitation and can be a helpful guide. Finally, 'systems development' here (on the slide) refers to computing applications *per se*.

NATURE OF INTEGRATED ELECTRONICS

- *Thousands of components (e.g. transistors) condensed into a single circuit module.*
- *Finger-nail size circuits mass produced from cheap materials (e.g. silicon).*
- *Such circuits are robust and readily inter-connected through standard interfaces.*
- *A single circuit can contain all the components of a limited power central processing unit.*

(i.e. a circuit can be programmed like a conventional computer).

I now want to talk briefly about the nature of integrated circuits. The history of electronics can be related to the improvement and domination of valves or their replacements. Valves were once the minority items in lists of components normally found in a circuit. Being fragile and expensive their numbers were minimised. The transistor put them into solid state form and made them more robust and reliable. Integrated circuits largely consist of thousands of transistors configured to form a primary component within a compound circuit.

The CPU of a large computer of the 1950s could consist of 20,000 valves, an odd mile of wiring and innumerable resistors and capacitors. The equivalent today might be just a handful of chips.

* *Editor's Note: to be available in book form in early 1978 according to the UK agents, Messrs WH Freeman & Co. Ltd. of 58 Kings Road, Reading (0734 - 583250) who will supply the details.*

SOME REMARKABLE PROPERTIES OF MICROCIRCUITS

ALL the parameters (e.g. speed, density) are improving rapidly.

No physical limitation is in sight.

Circuit cost depends on volume of production rather than its content.

Circuit reliability depends on quality of production rather than its content.

(e.g. 1978 16K bit memories could be CHEAPER + MORE RELIABLE than 1977 4K bit memories).

Put crudely, if unbelievably, the smaller a transistor is made, the faster it switches, the less it costs, the lower its power consumption and the greater its reliability. Electron beam lithography is expected to reduce the unit of detail in current circuit etching from 4 microns down to 1 micron by about 1979. This would increase the level component integration by a factor of 16 and should lead, for example, to memory modules of 16x16K bits or 32K bytes by 1980.

NATURE OF PROGRAMMING

Imagine blind, deaf mutes who have to navigate round London...

then a Braille set of movement + test instructions (e.g. curbs + bus stops).

Each step and action must be anticipated for every route.

Provision necessary for the unexpected (e.g. road works) + recovery from mistakes.

No way of assuring all route programs were always absolutely correct.

This is the best example I have been able to invent to communicate the nature of programming. Please try and imagine the task of planning each route in this level of detail; working out the dual language messages (ie braille and English) for handing, say, to bus-conductors and anticipating the possibility of error in their response (for example touch left arm for yes, right for no) — such as when the wrong person is identified as being the conductor. My object here is to give a manager some idea of what is difficult to program (for example because it is outside the programmers' previous experience) and what the needs of verification really entail.

EXAMPLES OF PROGRAMMABLE MICROCIRCUITS

A MICROPROCESSOR	- part or whole of CPU
A USART	- parameter driven I/O controller
A PROM	- programmable read only memory
A PLA	- input signals variable connected to output signals
A FLOPPY DISK CONTROLLER	- parameter driven
A VECTOR INTERRUPT CONTROLLER	- a variable list of interrupt priorities and their addresses

A MICROPROCESSOR

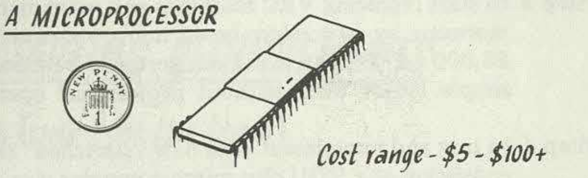

A MICROCOMPUTER BOARD

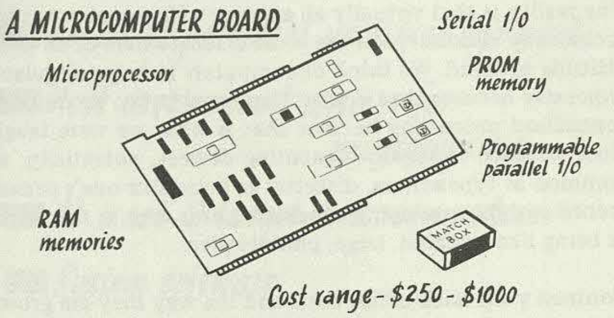

A MICROCOMPUTER CONFIGURATION (with Floppy Disks)

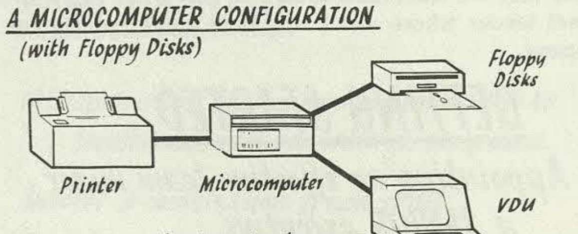

A MICROCOMPUTER CONFIGURATION (with Cassette Tapes)

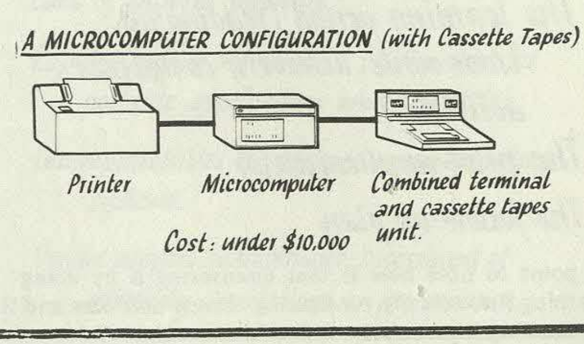

At this stage you might very well ask the question: is it all too much?

It might be objected that this implies too much reading for a manager, or is impossible anyway for some companies to do because none of the managers has the necessary education. Certainly, a large amount of reading is involved, but unfortunately it is unavoidable. If left to a subordinate one of two things will happen, both of which are damaging. Either the subordinate will effectively take decisions, without having proper authority to see them through, or the manager will vet his proposals in the light of his commonsense — which actually means a distillation of his own past and partially obsolete experience — so that the probability of gaining approval will depend on the *credibility* of the proposal rather than its *correctness.*

The lack of requisite technical education is a reasonable fear, but it is entirely unwarranted. There is nothing in the fundamentals that cannot be explained to an interested and patient manager, whatever might be his background. Finding the right person to do the explaining though is quite another matter.

I now want to move on to the topic of the impact which microprocessors can have on your business.

It is customary among experts to dodge this question on the valid but unhelpful grounds that it all depends on your business and all businesses are different. Thus it is likely that you, and only you, can judge their probable impact, and then only after you have properly understood the implications of the general trends in microcircuits.

A common method of making this judgment is to wait and see what one's competitors are doing — since every manager has been warned of the dangers of pioneering. Seemingly there is not the same warning given over the dangers of becoming uncompetitive or out-of-date. But that is a warning of microcircuits. Very often they do not simply change a product, process or procedure: instead they transform it into something quite new insofar as the customer or user is concerned. This is what makes new markets emerge, like calculators and digital watches, which are soon to be followed by pocket radios and 3 inch TVs plus music centres, for sale to our children.

Out of a legion of possible transformations we will only consider two: namely, control systems and office networks.

CONTROL SYSTEMS

On-board end products.

In or for machine tools, power, test and monitoring equipment.

In or for warehousing + transporting equipment.

In or for environmental control, security and similar equipment.

In or for display and process control equipment.

This slide addresses control systems in the broadest possible sense, both within a product and in the making of it. Consider a motor car, or a washing machine or radio (note that 10 million circuits are being delivered in 1979 by Motorola to General Motors. And already we have one-chip timers for washing machines and one chip AM radios that require only a loudspeaker).

Within any factory consider also the potential for use in power distribution, in displays on monitoring systems and in ways to prevent materials, part assemblies etc being simply lost on works floors.

In warehousing everyone has read about automatic stackers, but much simpler systems are also now available (eg to guide picking).

It is also worth looking again at plant heating and security systems where the equipment is often expensive and inflexible (eg sprinklers can do more damage than a fire).

One problem is getting the engineers to believe it. End products are naturally a key area of impact. In essence any product which functions through the measurement of time or through the feedback of sensors is a potential application of microcircuits. Nevertheless, the greatest difficulty may come in persuading your engineers that this is both feasible and desirable.

Common objections are that the circuits are too expensive or unreliable, unable to sense a low-level signal or to drive a significant force. The difficulty is that sometimes these objections are justly made, but often they only appear to be so out of ignorance. Since micros currently can sense a few milli-volts, control aero-engines, trip steam-hammers and guide surgical probes it may be they are worth another look after all!

A less competitive area arises from the discovery made from recent events that many industrial processes are very exposed to sudden power cuts (ie have no power-down capability). This often suggests that companies have been unaware of risks and probably unaware how they consume electricity. Remedying such a lack of information can be a small start to the use of new control systems.

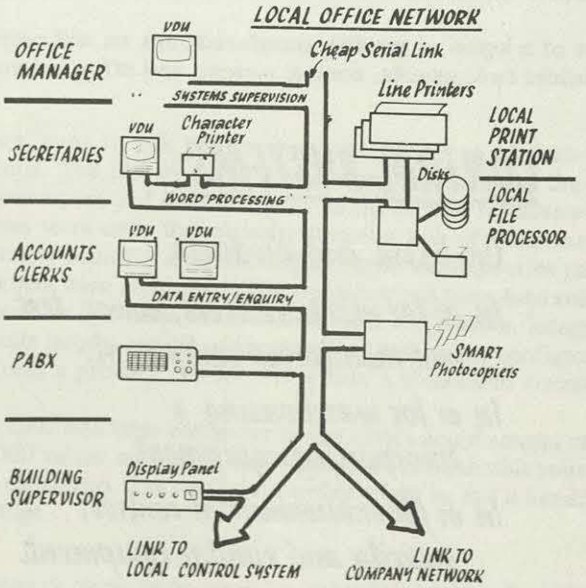

This visual looks like some futuristic office of a Dr. Strangelove but much of it is directly relevant now to quite small offices in most commercial companies. It is also applicable to those myriads of offices found in factories busy filling and entering black books whose existence may or may not have official recognition. (I have a personal suspicion that there are more clerks in British industry than there are in British commerce).

In fact this is the kind of system already working in executive offices of Citicorp — the giant American bank. It is not possible to reach such a system in one jump, of course, but rather in a succession of small steps:

Step 1 to get your DP Department involved in word processing, or so-called unstructured data systems, because despite appearances these are normally general purpose systems which can be linked into simple networks;

Step 2 to recognise that you do not need a complex communications set up to get started, since for initial purposes everything can be made to look like a teletype and be linked by a pair of wires;

Step 3 to start replacing VRC machines and other mechanical monsters by new micro-based minis which only cost £6,000 - £7,000 each and can certainly be made to do simple ledger work without professional operators;

Step 4 to buy and experiment with new "stretched" terminals consisting of a VDU plus micro-computer plus, say, 1 diskette available for £2,000 - £4,000 as an individual's data entry/data enquiry machine.

The reality is that virtually all existing office equipment is now technically obsolete, but the most obsolete item of all is our attitude of mind. We think of computers as being singular and expensive because that is what they used to be; we think of centralised processing because that is what we were taught. Now we have to regard computing devices, potentially, as common as typewriters, diskettes as being like one's personal record and the remaining expensive items such as fast printers as being like a shared, large, photocopier.

Contrast your back-office costs and the way they are growing with what is happening in hardware costs. I realise that I risk sounding like a computer salesman but don't you need more automation in your offices? I do in mine.

Notice that the team itself should be lead by an established project leader whose views will have weight within the company.

GETTING STARTED

Appointing an effective team under a SENIOR executive.

The learning period (reading up, trade visits, discovery competitor activities).

The pilot application(s)

The follow-up plan

The point to note here is that discovering is by doing something (i.e. actually configuring chosen hardware and then programming it preferably at first in assembly code).

PURPOSE OF PILOT APPLICATION(S)

- Training the team
- Discovering the scope of programmable microcircuits.
- Internal demonstrations of that scope.
- Laying down subsequent guidelines for subsequent usages.
- A useful application (but avoid specific promises and onerous deadlines).

STARTING REQUIREMENTS

- Time (for the team).
- Money (external expenditure £20,000 - £60,000).
- Access to potential user management.
- Host development system for programming.
- At least 2 target microcomputers.
- System software

SOME PROBLEMS

Microprocessors/microcomputers should not be treated as small conventional computers.

Involve a combination of new skills.

Lack of uniform software.

Rapid change + obsolescence (hence need for portable applications where possible).

Incompatibilities of files and communication systems.

Vendor support, maintenance, assessment of reliability, limitations on performance.

SOME MYTHS

- 'We are going to hardware-ise (sic) out the programming'
- 'You must be an expert in Integrated Electronics to use them properly'
- 'The fall in computer prices means hardware will become a small part of DP budgets.'
- 'Our next computer will be exactly compatible with...'
- 'There is no such thing as true program portability.'

BUTLER: Thank you, Alex, for a characteristically brilliant performance. I should like to start the discussion by asking Graham Hawker, from Grandmet, to talk about a variant in the approach of getting people into the area of microcircuits, because I think that what Grandmet are doing in that area is very interesting.

HAWKER (Grandmet): Our opinion is that the only way you can really appreciate the true nature of the micro revolution is to get hard hands-on experience. By nature Grandmet is not a revolutionary company but it has recognised the huge potential of micros. A good example which highlights this is the August issue of *Interface* magazine, which comes with a plastic disc, free, on the front cover. The disc holds a program which you can read on your home hi-fi and transfer to cassette, which can then be loaded onto your micro. It provides a complete sales ledger and accounting package for a hotel.

It was clear to us that if anyone from our Hotels Division got hold of a copy of the magazine there would be trouble! We bought a £200 micro and made it available to all our programming staff, with the idea that in their own time staff could gain skill and experience with these devices. We now plan to buy a £2,000 kit and a full scale implementation. I strongly recommend this approach to others.

BUTLER: Thank you very much, Graham. Questions and comments from the floor?

FAUVRE (Digital): We have some experience manufacturing these devices at Digital. One problem was that engineering kept designing their own. But the situation is better now that we have standardised on the PDP8-11 instruction set.

QUESTION: Could you explain the difference between the £200 kit, and requirements for a host processor?

d'AGAPEYEFF: The thing is that you can get a wide choice of device but they have no business software. None of them has for example index sequential access. We may be the only people in the world to have index sequential on three of the world's top processors: the Motorola 6800, the Z80, and the Intel 8080.

If you do not buy a host development system where somebody has built such software — and there will soon be a lot of choice because this sort of thing is mushrooming fast — you will have to do it on the micro, for which you haven't the tools. When you are doing the initial experiments, when you are doing what I suspect Grandmet did (ie trying to demonstrate that something was possible), you are not immediately making a tool. It is certainly arguable however that you must go through that process first.

I do encourage people to go further than that, because the other danger that no doubt they have anticipated and escaped from is that you start buying a few of these, and then you buy different ones, and they are all incompatible. Somebody else comes along and says, "Oh well, I grant you they're all cheap, but we don't seem to be doing anything practical with these things." I hate to give DEC a plug because their equipment actually infuriates me. It is the world's best instruction set for a mini, and yet I find it a little expensive. Nevertheless, we put all our micro development software on a DEC PDP/11. They are about the best host that you can have. You can run links from a tiny LSI 11, as the host into a variety of target packaged systems. As long as you stick to standard board microcomputers as the target a single host machine can be practical for all application development.

BUTLER: Thank you very much, Alex. In closing this session I should like to mention one particular merit in Alex d'Agapeyeff's talk, which I think is particularly valuable to all of us. Sometimes we have sessions at these conferences which cover the technical ground extremely well, but which seem to be a little bit blind to the management implications; and sometimes we have sessions which concentrate exclusively on the management problems while tending to gloss over some of the technical realities underlying them. I think that Alex's talk was a very good example of one which covered both the technical realities and the management problems of getting value for money out of those technical changes. Thank you very much, Alex.

THE EFFECTIVE USE OF SYSTEMS HOUSES

H.C. Zedlitz

BUTLER: Gentlemen, welcome back to the second session in the part of the conference which is concerned with the experience of users. Increasingly, large users are finding themselves inclined to make use of systems houses which can provide equipment, systems, or even a turnkey operation; and I think that this is a phenomenon which is likely to be with us for some time. In fact, a computer manufacturer said to me the other day, "We're all systems houses now," and I believe that there is some truth in that comment. But on the part of the user, of course, some management skill is required to ensure that the services and products of the systems house are used in ways which are effective.

We thought that it would be a good idea to identify some of the opportunities and problems which arise for the user in this field. To do so, I invited my friend, Chris Zedlitz, to come and talk to us about his experience in this field. Chris is the Director in charge of Systems and Automation for the pharmaceutical company of the AKZO conglomerate, and he is based in Holland; so we are very pleased that he has been able to find the time to come over and address us today.

ZEDLITZ: The title of my speech suggests implicitly four facts, namely that the speaker knows the difference between the attributes "effective" and "efficient". I didn't know myself until one year ago when I attended a senior executive international management course, and I was told by some bright Harvard people "to forget about efficiency, only focus on effectiveness". After a beer, I made my own opinion and I said, "We'll have to combine them both."

But we really are using systems houses so we are talking with some experience. We see a future for systems houses and we intend to continue the use of them. I should like to elaborate on the use of systems houses in the framework of our own DP policy, to analyse the areas of use of systems houses, the types of services in each area, and the experience that we have gained with them up till now. Then I should like to focus on the expected role of systems houses in the future, giving as a sort of reference how we see the future and the types of services provided: will they change or will they remain the same?

Let me first give you my view of the criteria of effectivity. Effectivity, in my opinion, is to do the right thing, not things right; and to do the right thing *economically* (there we have the cost benefit aspect), at the right *time* (the time aspect), with a *calculated risk* (the risk aspect), and aware of the *social impacts* (which we call the social aspect). All those criteria are interdependent.

Some words now on our DP organisation. I am responsible for a budget of 13 million Dutch guilders, and for 142 DP people. We are linked with other AKZO divisional DP activities on a cooperative basis; that means that my colleagues and I define the overall DP policy of AKZO, and the tools and functions that we pool together for service centres and technical expertise.

We plan our activities on a cyclic basis, and this planning activity is linked with business planning, so that each business plan of a local company (eg of a department in our division), and of the division itself, reflects DP activities in the main organisation. So we have got links between our own activities and the business activities, our goals and the business goals, our costs and the business costs, and also our expected results and the business results.

We are focusing on R&D as well as on business processes; and as you may know, the pharmaceutical business is highly research oriented on the one hand, and highly marketing oriented, and rather less production oriented, on the other.

We work as a cost centre; we have to charge out any activity. We are controlled by status reports. We are controlled by project and we have to provide top management with a cost benefit control calculation for any of the projects that we have implemented. It provides a very good feedback for our planning activities.

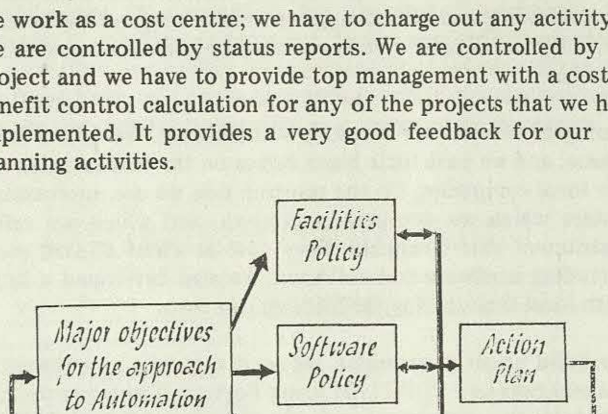

Now for our policy, because our policy might enable us to place *a priori* restrictions on the use of systems houses, or the reverse. We agree major objectives for our approach to automation with our top management. From these we derive our policies for facilities, software and security; and these form the basis for our actual plans.

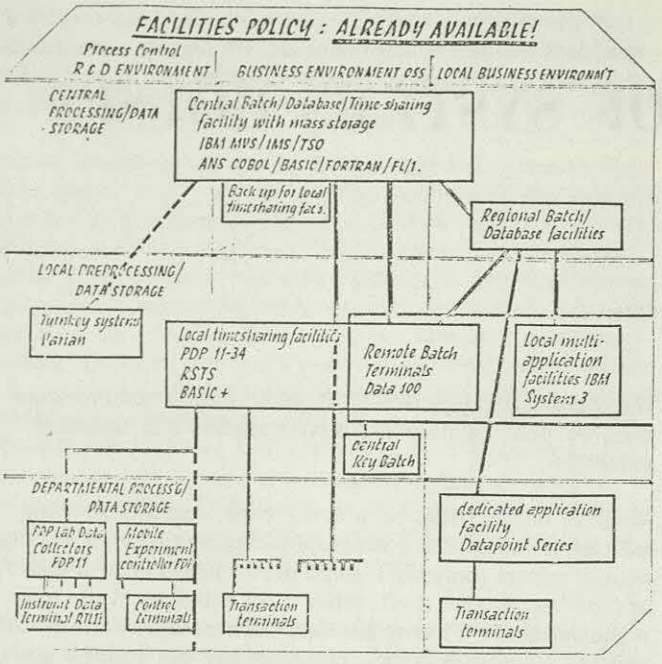

Our facilities policy recognises a distinction between our process control and R&D environments (we know the environment of our headquarters and the environment of our local companies). It also distinguishes between the departmental processing and data storage environments; between the local preprocessing and data storage environments; and between the central processing and preprocessing and data storage environments.

Our systems software is mainly standardised. We have a mainframe, and we have little black boxes on the business side of the local companies. On the research side we use microcomputers which we developed ourselves, and which we call instrument data terminals. They cost us about £3,000 each, including hardware and software. We also developed a link with local time sharing facilities on our own.

To build up an experiment, we send a mobile experiment control over to a PDP-11/04 using Fortran. The lab technician builds up the experiment with the experimental data. If the methodology has been prepared, we screen and decide whether this type of test should become the standard test or not. If not, we forget it. In the latter case we use quick and dirty programming, and throw it away. Alternatively if the test becomes the standard, we re-program on the £3,000 microcomputer, and it becomes a black box in the research environment.

We are now following the same approach in the business area, implementing a lot of such black boxes in our local companies in the UK. Implementation goes very quickly for a local company, with an elapsed time of about two months for general ledger, order entry and stock control systems.

Why am I telling you this? There are some areas on which I shall elaborate, where systems houses are used very well, and other areas where you have to forget them. The same holds true for our software policy. Confused by the slogans, "You should buy packages on the market", "You should try to focus on the development of common systems", "You should prototype", or "You should focus only on logical data bases", I asked myself "What should I do now? Shall I buy a package? Shall I do common systems development? Should I buy hardware and software?" So I tried to analyse and structure the situation.

SOFTWARE POLICY

Type of applicat'n / Degree of specification	Operations	Control	Planning	Change
Structured	Packages/ Common systems	Packages/ Common systems	Prototyping Packages/ Common systems	Modules
Semi-structured	Prototyping Database Packages	Modules Database Packages	Individual development Database Packages	Individual developments
Unstructured	Individual developments	Individual developments	Individual developments	Individual developments

BUTLER COX & PARTNERS

We analysed it like this. We recognise four types of application: *operational* or *transaction-oriented* applications; *control* applications; *planning* applications; and applications which we call *change.* These application types receive the degree of specification advocated by Hamish, increasing programmer productivity by improving analysis skills. But this is only half the problem; because we know a lot of problems, especially in research, where you also have to build up a methodology. The specs evolve and you cannot just ask your client, "Please specify." You have to give him a very good analyst.

So we have to distinguish a degree of specification, between structured, semi-structured and unstructured. We decided that, in the structured area, packages and common systems are somewhat in competition. Whereas in an area of change we should never use packages but just focus on modules.

SECURITY POLICY

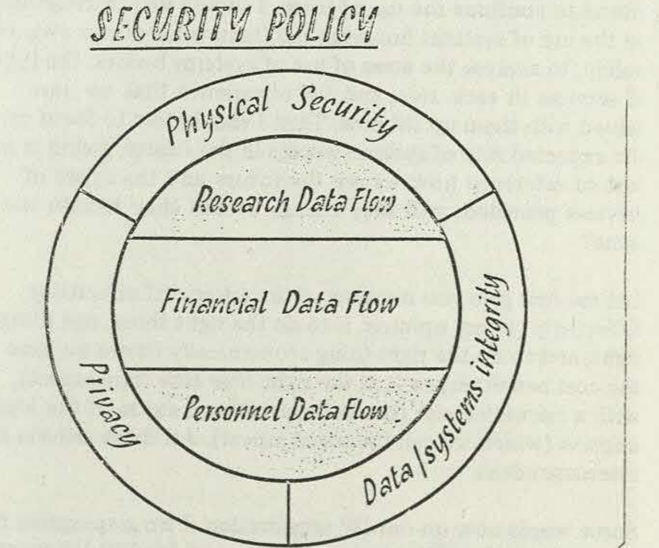

In the semi-structured environment, we should try to focus on prototyping and mainly use data bases. By putting everything on a data base I mean a very easy, quick to access, logical data file which you can throw away again. In the non-structured areas, we need very good people from the user side, from other disciplines, and from the DP side to tackle individual developments.

Another factor is our security policy. We have built up a security policy distinguishing between our research data flow, our financial data flow and our personal data flow. The research data flow is covered by research or health authorities. They prescribe very strictly — and they are improving in their restrictions — what we have to do. On our financial data flow our accountants are very active building up a series of rules on what we have to do and what not to do. And finally the different privacy laws in Europe require us to do something on personal data flow. Therefore we have to build up an overall security system which helps us to manage all these data flows, with as low overheads as possible

The security policy highlights how you should be aware of the way a systems house can damage you in an area. If the unions have you on their list of bad records, don't ask for a systems house to help you with personal data flow.

In what areas do we make use of systems houses? We use them or have used them for the development and provision of hardware and standard software; for systems development; for risk management and security management and the facility operations side; and also as DP management consultants.

USED SERVICES PER AREA

- **Development/provision of hardware & standard software**
  - development of a microcomputer system
  - provision of turnkey systems
- **Risk management/Security management**
  - concepts
  - individual consultancy
- **Facility operations**
  - fall back if own staff leaves
  - meeting peaks
  - service centres
- **DP Management Consultancy**
  - interpreter between Use Management and DP Mgmt.
  - temporary DP Management in critical periods
  - consultancy on organisational matters

How did we use them in these areas? Regarding the development and provision of hardware and standard software, we asked a systems house to develop a microcomputer system for us. We also bought a lot of turnkey systems. In the risk management and security management area we required first class design, and we had individual consultancy. In the area of facilities operations we used them as a fall-back when our own staff left (not every DP man in Western Europe is happy to be confronted with a strict DP policy, and sometimes they leave). We use them in order to meet peaks and, of course, as service centres. As DP management consultants we used them as interpreters between user management and DP management, and as temporary DP management in critical periods in some local companies, as well as consultants on organisational matters.

USED SERVICES in SYSTEMS DEVELOPMENT

Degree of specification \ Type of applicat'n	Operations	Control	Planning	Change
Structured	Package conversion Programming		Package	Package
Semi-structured	Project managem't; Systems realisation at a fixed price; Programming			
Unstructured	Programming			

In the system development area, we use them in the structured problem area for conversion, for packages — we bought packages — and for programming. In the semi-structured area, we used them for project management, for systems realisation at a fixed price, and for programming; and in the unstructured area we use systems houses just as a source of programmers.

EXPERIENCES MADE WITH SYSTEMS HOUSES

TYPE OF SERVICES \ EFFECTIVITY ASPECTS	COST-BENEFIT ASPECTS	TIME ASPECTS	RISK ASPECTS	SOCIAL ASPECTS	QUALIT. TOTAL
Development of microcomp. systems	++	-	0	-	+
Provision of turnkey systems	++	++	++	+	++
Security mgmt. concepts/indiv. cons.	+	++	+	+	++
Facility operations	–	++	–	0	+
Service centres	++	+	+	-	+
DP Management Consultancy					
Interpreter	++	+	+	–	++
Temporary DP Mgmt	–	++	+	–	+
Consultancy	+	+	0	–	+
Software Development:					
Packages	++	++	–	+	+
Conversion	–	++	0	+	+
Project management	–	+	–	–	–
Fixed price realisation	+	+	–	--	–
Programming	0	++	–	0	0
Qualitative Total	+	++	0	–	+

What has been our experience? I have prepared a value matrix and the overall result, as you can see, is positive. Let me now elaborate a little bit on each item.

The development of a microcomputer system saved us a lot of money. If we had done it on our own, it would have cost us three times the price. On the time aspect, we did not gain so much. The risk aspect was zero. The social aspect was negative because, of course, I have people on my payroll who are very good engineers and electronic technicians, and to build up a microcomputer system would have pleased them very much, so they were de-motivated.

The provision of turnkey systems saved us a lot of money. The time aspect was very positive; we got it very quickly. The risk aspect was very positive. You can see it work, you can try it, and then you buy it; so no risk. The social aspect is positive because a good DP man is not interested in reinventing the wheel which he can buy and can see working.

The security management concepts and individual consultancy in this area saved us some costs; digging into the different privacy laws of Europe can cost you a lot of time. The time aspect was very positive; we got the results very early. The risk aspect is positive because you know that you have got legal advice and DP management from the country which has the experience. If you transform it into your own organisation, it helps you with less risk. The social aspect is positive because DP people are not very interested in security problems.

Facility operations. From a cost benefit point of view this is negative. From the aspect of time it is very positive. The risk is negative; you run the risk. From the social aspect it is about zero because the very fact that you have facilities operations assumes that you have got troubles.

Service centres are very cost beneficial for us. We use them in the UK, in Germany and in Spain; and they cost us less than the minimum requirements we should have to operate on our own, with our standards. The time aspect is positive; things get done very quickly. The risk aspect is positive because we always in our contracts ask those people to inform us of any hardware and software system change, and we have the right of first refusal. Our risk managers have access at any time of the day to control these DP operations.

DP management consultancy. The interpreter function between the user management and DP management we see as very positive. Sometimes you get to a deadlock, and then you need an outside man to help you to get out of it. The time aspect: you come out of this deadlock at an earlier stage, so you win time. The risk aspect: you are the deadlock. What can you lose? The social aspect is somewhat negative. Our own DP staff cannot see the benefit ('why do these outside people get a response from top management and not our own boss?').

Temporary DP management costs a lot, but we have to do it sometimes. It is time against money. The risk aspect is very positive. This man is more controllable; you can clearly state what he is going to do and what he is not going to do. If he is not fulfilling, you just kick him out; with a new man on your own payroll, you can't do that. The social aspect is negative. Why don't your own DP people get the same cheques?

Consultancy we regard as very cost beneficial. We win time. We have the last say, so no risk. The social aspect is negative. ('Why do we need consultants? Why don't we build up a new work group?')

Software development packages, of course, are very cost beneficial. We win time, but there is a risk. I have not seen one package that we have bought which was really well debugged. You can still get problems, even after seeking assurance by checking out a list of thirty or more current accounts. The social aspect is positive. If you can get software from outside, your own people are not interested in developing it. I can't imagine a man on my staff who is interested in re-developing a payroll system.

Conversions; Well, you have to do them; with outside people they cost you more, but you save time. The risk aspect is about zero. The social aspect is positive. Your own people are not involved in this dirty conversion work; they can focus on new developments.

Project management costs you a lot. You win some time. The risk aspect is very high, because under your own flag the project manager from an outside company can do a lot of damage. Our conclusion is: no more project management from third parties.

With fixed price realisation we had some positive experiences but only in areas where we provided the systems house with a clear set of functional specs and our own standards — and then it worked very well. Our congratulations to you: we have had especially good experience with UK-based systems houses.

Now for programming. First, the cost benefit aspects. British software houses are fairly competitive with our own internal tariffs, but Dutch, French or German systems houses cost a lot more than our own people. The time aspect: you use programmers if you do not have the skills available yourself, so you save time. The risk aspect is negative. From one point of view you have restrictions: you cannot let these people work on all programs, so we have some research applications where we say, "No third party". We have some personal applications where the same holds true: also some financial application areas where the same holds true. Social aspects are zero.

If you keep third party programmers too long with your own staff, then you get questions — some silly and some good. "Why don't you take this chap on your payroll?", "What is the average time this chap is staying with you?" If it does not exceed one year, you get no problems with your own staff. That is our experience. The social aspects then are zero. We would say that we need them if we get peaks and we do not have our own skills, but the overall value is zero.

To sum up, from a cost benefit point of view, our experience with systems houses is positive. This is particularly true with regard to the time aspect. Regarding the risk aspect, there is a balance; we say zero. But the social aspects are negative and we have to do a lot to motivate our own staff. But the overall result until now has been positive.

How do we intend to use systems houses in the future? Again, you must define their role in your overall strategy for the future. I wrote a strategic paper for the board and we had a discussion on it, analysing our automation resources consisting of our specialist hardware/software methods and tools. How do we transform those resources in the best way to help the research developments, the logistics and the marketing processes of our companies? And not only in the light of technological developments, but also international standardisation, the labour market, education and so on. What is management going to be about in 10 years? Will we know more about automation or the same as now? How will the main markets change? (This is especially interesting in the pharmaceutical industry).

Emerging communication technologies such as Viewdata might

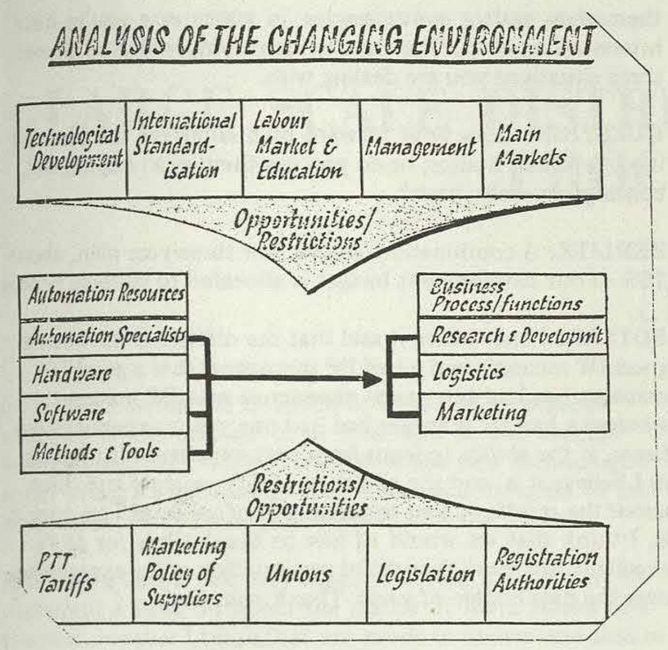

help us or might substitute for contracts which now cost us a lot of money. So main markets might change. The general practitioner might disappear in groups; or PTT tariffs might still be very restricted. The marketing policies of suppliers: will they become more aggressive? Will minis and micros proliferate without control? The unions, and legislation: will legislation go on playing a restrictive role? In the pharmaceutical business, we have a lot to do with the registration authorities.

We made a careful analysis of how these factors will change in the 10 years ahead. We came to some conclusions regarding the changed nature of our automation resources, especially about the priorities of future developments regarding our processors. What does this mean for our use of systems houses?

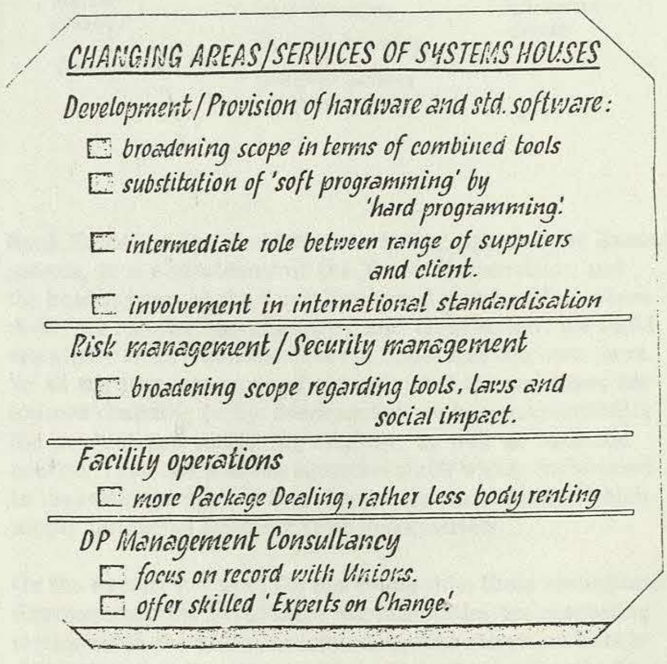

In the area of development and provision of hardware and standard software, we see a broadening scope for *combined tools*. They will not offer us computers and applications software, but maybe combined word processors, computers, image processors and the like, each focusing on one application area.

We see also the substitution of soft programming by hard programming, not so much in the area of application programs, but more in the area of standard programs. And as for the role of intermediary between the range of suppliers and the client; I should personally very much appreciate some systems houses which you can trust, playing an intermediate role on our behalf between the suppliers of micros, minis, word processors and so on, while we focus on applications. As the business grows bigger and their role becomes more intensive they might get some influence over international standardisation, which would be good because until now the suppliers have had some influence but in my opinion not a very positive one.

Regarding the risk and security management, it might be that they will broaden their scope regarding tools, rules and the issue of social impact. Regarding facilities operations, I presume that they will offer more package dealing and rather less body renting than they are doing now.

DP management consultancy; I think that we both will have to focus not only on the skills of the consultants but also on their record with the unions. What we really do need is skilled experts of change. If you keep those many-skilled experts of change on your own payroll, they will knock at your door after two years wanting to become the general manager of one of your local companies. But we are not growing so fast that we can afford that many general managers!

As for systems development, we foresee an emerging market for packages, turnkey systems and tools which will be provided by the systems houses, and more contract developments on a fixed price basis. Our clients are getting more and more conscious. They are also now trying to change me from a cost centre to a profit centre, though I have to offer projects at a fixed price.

In conclusion I think that DP departments on the payroll of a company will in the future focus more on data flow in a communication network, and less on algorithms. So systems houses will focus more on providing us with tools and algorithms.

In this sense, we expect systems houses to be used effectively now and in the future; but we have to be aware of the restrictions of our own risk policy; and they have to follow our own standards. In some areas they are forbidden, and we have to look at their record with the unions. Certainly, we must still be in the chair; no project management from third parties.

Finally, "Don't ask them for a fish, but ask them to teach you how to fish, or to work as a fisherman under your control."

BUTLER: We have time for a few questions or comments from the floor.

QUESTION: Several times you emphasised the influence of the unions — something that hadn't crossed my mind before. Could you elaborate on this?

ZEDLITZ: Yes, because if you are starting a new project the unions, of course, are focusing on it. I have seen this in Germany and in France.

QUESTION: Do you mean unions in your own organisation or in the software house?

ZEDLITZ: No, the unions in my organisation.

BUTLER: Are they unions in the data processing department or unions in the user department? Or both?

ZEDLITZ: The unions in the user department. I know one example where Volkswagen had to drop a big project because their unions did not agree with the systems house that they contracted.

QUESTION: I think that the union question is particularly important at the time of system implementation.

ZEDLITZ: Yes, but you have to inform your unions how you are going to solve the problem or to realise the project. Then you have to inform them of the project organisation. If you say, "I am hiring some staff, some expertise, from this system house," they might object.

BUTLER: In at least one country in Europe to my knowledge, which is Sweden, there is now an official legislated Code of Practice, which means that companies have to submit plans for reorganisation and systems projects to the representative unions in advance for their agreement. If the unions object to schemes on the grounds of reduction in the work force or job impoverishment, then the company may in some cases not actually be able to proceed with the scheme. Although it is perhaps irrational, I have little doubt that the presence of outside skills in the planned implementation would tend to prejudice the unions against rather than in favour.

ZEDLITZ: We always try, especially in the Netherlands, not reorganisation but reshuffling of the organisation. Reshuffling of the organisation does not need the approval of the workers' councils; reorganisation does. That is why, with any new, big projects, we try to convince the workers' council that we are only reshuffling!

QUESTION: Have you any advice on whether to go for fixed price contracts, or time and materials? On the one hand the danger is that quality can suffer; on the other you're writing a blank cheque.

BUTLER: If I could just interpose a word from our own experience, we quite often find ourselves in our consultancy practice standing between a client and a systems house, and sometimes negotiating on behalf of the client. I would say that, in our experience, there are three different sets of reasons why a software house will sometimes seek a fixed price contract. One is where they have, as a result of their past work, developed tool kits which will enable them to do the job very quickly and very effectively. In that case, you can be pretty certain that they will do the job within the price. They may be making a terrific profit on it because of the prior work that they have done, but if it is saving you money, why should you care what profit they make?

The second set of circumstances is where they desperately need the work. In that case, you run the risk that their survival problem may become your problem later on. The third case is where, as you said, they are loading the price to cover themselves against contingencies; in which case you're not interested anyway. The trick is identifying which of those three situations you are dealing with.

QUESTION: Does your forward plan allow for the planned use of systems houses, or do you use them as an unplanned, contingency-only, basis?

ZEDLITZ: A combination. In our new three-year plan, about 15% of our development budget is allocated to systems houses.

BUTLER: I once heard it said that the difference between a good DP manager and a bad DP manager is that a good DP manager has had ten years' experience as a DP manager, whereas a bad DP manager had had one year's experience ten times. If the ability to learn from past experience is of value, as I believe it is, and the ability to codify, analyse and think about the results of that experience is of value, as I am sure it is, I think that we would all like to thank Chris for an excellent and most thoughtful presentation of his experience over the past couple of years. Thank you, Chris.

LARGE SCALE DISTRIBUTED INFORMATION PROCESSING

L. Elstein

COX: I think that the question of how best to organise ourselves is one that all of us face, and it will become compounded when we go to some of the new technologies about which we are talking. To give us some thoughts on this area, we have asked along Les Elstein, of Rank Xerox.

ELSTEIN: Thank you very much. I should like to start by explaining a little bit about the nature of Rank Xerox as a business, because I think that one needs to understand that to realise why we have done the things we have done, why we have the problems, and why we have adopted our particular solution to them.

ORGANISATION OF RANK XEROX

XEROX CORPORATION (USA)

RX LTD. CHIEF EXECUTIVE

REGIONAL DIRECTORS (MARKETING)	CHIEF STAFF OFFICER	MANUFACTURING
General Manager-Country A	Product Strategy	Plants
Country B	Marketing Strategy	Distribution Centres
	Legal Accounting	
	Financial Accounting	
	Functional guidance to operating units	

Rank Xerox has the worldwide marketing rights of the Xerox patents. It is a subsidiary of the Xerox Corporation; and the headquarters of the Rank Xerox end is in London. From there, we run the manufacturing and distribution. We build essentially all our own machines in house and our own parts. So all the parts are sourced centrally and the machines are sourced centrally. In the headquarters we have responsibility for product and marketing strategy, as well as financial control. There are also headquarters staffs which correspond to the main operational functions of the business and which supply functional guidance from headquarters.

On the operational side, the marketing side, there are regional directors, and the geographic responsibilities for marketing service are in the individual countries. Not shown on here is the engineering function which in fact reports directly to Xerox; in other words, it is not even consolidated in the Rank Xerox accounts. The engineering, manufacturing and distribution centres are in England and Holland.

KEY FEATURES OF BUSINESS

Same products world-wide

Similar marketing approach world-wide
eg. consistent pricing philosophy
own service force

Strongly profit-conscious
analytic approach to decisions

Strong element of central control

The key features of the business: we are essentially marketing the same products across the world. Obviously, the products that we market in Rank Xerox may be marginally different from the Xerox ones — the power supply and that type of thing — and there may be specific marketing conditions which demand some slight modification; but they are close enough that you can say the products are the same. Not only is the product the same, but our marketing approach is the same. We operate consistent pricing philosophies; across the world we have our own service force, our approach to whether we are leasing or selling; and the way in which we are handling the consumables (the pieces that are needed to keep the machine going) is the same across the world.

It is a strongly profit-conscious, American style company; the approach to decisions essentially analytic and numeric. Needless to say, having said all that, there is a strong element of central control of the business decisions.

CHARACTERISTICS OF RANK XEROX DATA PROCESSING....

Computer utility in Europe based on large data centres (IDC's) for batch processing.

Minicomputers/intelligent terminals with constrained functions linked to IDC network for transaction processing.

External vendor with network linked to IDCs for timesharing.

As for data processing, over the last three years we have set up a computer utility in Europe. We have replaced individual, small computers which were widespread in our marketing and manufacturing centres by three large data centres for batch processing. At the same time, we are introducing minicomputers, intelligent terminals and so on; but they are treated as part of the utility in the planning sense. We also have very large usage of time sharing, and I will come back to that later.

. . . and RX TELECOMMUNICATIONS

Private data network for data centre traffic.

Private voice and message-switching networks linking HQ and major operating units.

Transatlantic leased lines carrying mixed data, telex and voice traffic between RX and XEROX locations.

On the communications side, a private data network for the data centre traffic; private voice and message switching networks; transatlantic leased lines.

I am not going to describe to you the detailed technical basis of all that. I think that what we are really concerned with is: why have we gone to such a complex and expensive aproach? I will start by talking about the data processing side. The first point that you might ask about international data centres is: do they actually save money against having 10, 12 or however many it is, computers scattered around the place?

A. DATA PROCESSING

1. INTERNATIONAL DATA CENTRES

large data centres can give economies of scale if environment tightly controlled.

---

main justification:

- *reduced application software costs by developing/acquiring once and using in all operating units.*

Certainly the arguments of economies of scale on hardware are not easy to sustain. The real point is that, if we are trying to run the business as one business, the last thing that we want is proliferation of software tools, application tools where each marketing or manufacturing unit believes that it needs the complete range of software tools. Essentially, the data centre is a concept of a utility for processing; but the real pay-off is in the economies on software development or software acquisition. You have to develop it, in theory, only once.

We have implemented an application within one week in every one of our marketing companies, that is 15 or 18 countries.

2. MINICOMPUTERS

- *more economic for simple transaction processing*
  - *reduced telecommunications costs*
  - *avoid excessive overhead on batch-oriented mainframes.*
- *still want data to be a company resource*
  - *avoid private data bases for individual functions*
  - *information consistent at various organisation levels.*

On the other hand, we are not blind to minicomputers — and for minis I could have read micros, but for the moment we will say minis. There is no doubt that for certain types of transaction processing, minicomputers are more economic. For a start, you can cut down the telecommunication costs; if you do not have to send each transaction all the way to the centre, simply to register it, edit it and send back the errors, for example, there is undoubtedly an economy to be gained there. The other thing is that the sort of processors that one typically has for large batch applications may not be in any way economic for transaction processing; the overheads involved can be very high.

However, we want data to be a company resource. While we are not talking about what Hamish would term "one massive data base", nonetheless we do not want private databases. We want consistency at the various levels of the organisation. So at the headquarters level we do not need to know exactly which transactions took place yesterday, but we certainly want the record available at the head office of last month's transactions to be absolutely capable of being tracked back to those transactions. So it is consistency, not total ability to get into the data base which is the chief thing.

3. TIMESHARING

recognise importance and special characteristics

- *role in planning/decision processes*
- *personal computing*
- *specialised offerings*

---

standardise on language and vendor

- *sharing of applications*
- *financial leverage*

---

integrate timesharing and batch environments

- *T/S access to batch files*
- *high-volume print for T/S jobs.*

Why do we regard time sharing as important? Let me start by saying that especially in a company with this heavy concentration on an analytic and numerical analytic process, which is typical of American style management, then the ability to analyse alternative strategies in a short time, answer the so-called "What if . . .?" questions and look at alternatives is seen as very important. There is a large amount of that going on.

We have, of course, the question of personal computing; the idea that you, sitting at your desk, can churn some numbers around quickly for a purpose. For argument, we can call that time sharing. Then there are specialised offerings which are available, things like PERT or whatever, which outside bureaux offer and which are seen as cost justifiable.

Clearly, there are problems in controlling time sharing. Typically, for example, you can go to a time sharing bureau and open an account with no financial commitment. The question then is: how can anyone control the opening of such accounts or the usage of them? What we are trying to do essentially is to permit the sharing of applications in the time sharing domain; we are not talking now about personal computing, but for example financial planning applications where, on the one hand, you want to do the analysis at the local country level, and very soon afterwards analyse the consolidated file at the head office. So the sharing of applications is important then.

By standardising on the language and vendor, one can achieve financial leverage and get a good level of discount and also a good level of service.

We also see a requirement to integrate time sharing and batch computing. On the one hand, you may want to drive some sort of analytic process from your batch files, from the history of the actual transactions of the company, rather than create a special file *ad hoc* for time sharing. You also may have applications where you initiate the job on a terminal at your desk, but it prints out pages and pages of data; and you clearly do not want those pages to come back via a Teletype terminal. So there are applications which require integration.

B. TELECOMMUNICATIONS

1. LEASED DATA NETWORK

- only currently available method for international transmission of high-volume data.

2. LEASED VOICE AND MESSAGE NETWORK

- cost saving compared to single call tariffs.
- permits cost/service level trade-off.

Coming now to the telecommunications side, you saw that we had a pretty widespread leased data network, 9600 bauds. It is the only way to transmit on an international scale; you cannot get 9600 baud public switched networks throughout Europe. As far as the voice and message network is concerned, essentially we are talking about, first, possible cost savings with the heavy telephone traffic that we have between specific centres and, perhaps an arguable one but important to us, the ability to control the level of spending against the level of service. I am sure that the PTTs would say, "Why do you want to suppress demand?" But there may be situations where it is not so much a question of suppressing demand but degrading deliberately control service levels so that the traffic is forced away from the peaks or something like that. Typically, on Telex traffic for example, we are prepared to accept a higher average delay, and keep down our capacity and therefore our costs.

Decisions and trade-offs in the technical environment

Since we have been talking technology in this conference, I should just like to say a few words about the sort of technical trade-offs and technical decisions that we have been facing or will be facing. One that exercises a lot of people's minds just now is: to what extent should one distribute the processing power in a computing network? One of the first questions that one has to ask is, "What is the real need for processing power at the most remote part, at the furthest part from the centre of the entire network?" because that finally determines your architecture.

1. DISTRIBUTED vs. MAINFRAME PROCESSING: QUESTIONS...

- What applications require on-line processing:
  - user preference not always cost-justified
  - on-line update vs. on-line data capture and edit.
- Should applications share the distributed network at level of
  - terminals
  - minicomputers
  - telecommunications

Note that there you get into some rather interesting arguments about what does the user really need or can he really justify. There is no doubt that users are switched on by technology, by gimmicks, by the latest thing — whatever it may be — or simply the feeling of having more and more power at their fingertips. The fact that that is not always cost justified may not be easy to track in your decision process; especially as the units of power become easier to acquire — microprocessors, minicomputers or whatever. It may be very difficult indeed to keep these trade-offs in view.

One of the key decisions that you have to take is: when a transaction occurs, do you really need to update a central file, or a main file of the business at the time that transaction occurs? If you are in the airline reservation business, there is no doubt that you do, because the validity of the next transaction depends absolutely on the previous one. But how many of us are actually in that type of business? How many of us could survive quite happily if you could capture and edit the data and get it more or less right, 95% of the errors out at the time of capture, and then, overnight, transmit it to some central point; do a batch run; get all the controls there; send back the erroneous data; and operate in that mode? I assure you that the costs are very, very different for the two types of application.

The other thing that you have to decide in talking about distributed networks is what are you going to share? To what extent are you creating utilities in a distributed network? Should a terminal be able to access several different types of application, or be available to several different departments? Should your minicomputers be specialised for application, or should they be mini-utilities like the mainframes at the big data centres? Should you be sharing your lines, your modems and so on?

... AND RX CURRENT THINKING

- very few parts of our operation merit on-line processing.
- applications with complex on-line processing (complex logic, file manipulations) should have dedicated minicomputers.
- economic case may exist for utility serving simple needs (data capture, edit, store/forward).
- sharing of telecommunications is practical and desirable.

Well, our current thinking is that we do not believe that the sort of operation that we are in, on the whole, is an on-line business. We think that if you find an application which needs complex on-line processing, with complex logic and file manipulations at the time that the transaction takes place, and you can really justify it, then you should dedicate a minicomputer. You should not attempt to share a minicomputer among several complex types of applications, because you are back in the game of the excessive overhead. Your minicomputer has now become a data centre; the cost goes sky high.

We do have one or two applications of this type. For example, in our distribution work we have hierarchical levels of stocking of parts. Bear in mind that we are in the rental business; if a machine breaks down, it is in our interest to get that machine repaired as fast as possible and get the spare part out. The parts are stocked in echelon stocking; and it is also possible to substitute a part. There are some parts where you demand part A and in fact you could use part B as an alternative. Therefore, you need a fair amount of logic and possibly access to files to answer two questions. If the part is not in the store where you first try, is it somewhere else? If it is not in either of those two places, is there some alternative part?

On the other hand, for very simple needs (for example the data capture and edit which I described as an alternative for transaction processing) you may well be able to share the minicomputers, or the terminals, everything—just by arranging that transactions of type A are entered and validated between 9 and 11 a.m., the next lot between 11 and 1, or whatever you want to do. So you can construct the idea of a utility in that sort of area. But as for telecommunications, absolutely share it if you can.

2. DESIGN OF DATA COMMUNICATIONS NETWORK .... QUESTIONS

should different computing components share the international network

timesharing, batch

should network control be from a central mainframe or through distributed processors.

Looking now at how we design data networks, the sort of questions that we have to ask are these. We have an international network: what should we run on that? Should we share just on the data side time sharing and batch? Bear in mind that you cannot just throw VDUs, Teletypes, or remote batch terminals straight into a network; you have to multiplex time sharing and batch. The other thing is, in designing a network, to what extent should the processing power to control the network be distributed, or be from the mainframe (the SNA argument)?

...and RX CURRENT THINKING

timesharing terminals hard-wired and multiplexed through existing data network could cut dial-up cost significantly:

dedication of fixed amount of line capacity gives problems on level of service.

network control through distributed processors economic approach in short and medium term.

We believe that significant cost advantages can be gained by multiplexing time sharing and batch. You may find the problem that, as you increase the amount of line capacity that you are giving say to time sharing, then the other side of the service will start to suffer. But there are quite a lot of advantages in that approach. We do not think that the approach is SNA; we think that we will have to distribute the control of the network.

3. SHARING OF DATA AND VOICE NETWORKS

can pay off if one facility lightly loaded or insensitive on level of service.

otherwise operational complexity makes the approach suspect

- *we may discontinue sharing on transatlantic lines as data volume builds up.*
- *we will not use our voice lines as fallback against breakdown of our data network.*

Another point which you may wish to consider is whether the data and voice network should be integrated. It can pay off, especially if one of those facilities is lightly loaded, which means that you can just run it more or less on the back of the other and it is invisible. But if they are both tightly loaded or the line is at its capacity, then you can get operational complexities.

I mentioned to you that we have leased lines across the Atlantic which share voice data and Telex; we think that as we build up the data volume we will have to disaggregate that and dedicate a line to data, leaving the voice and Telex to do whatever we will.

Again, if our data network breaks down, we do not think that we will use our voice network for fallback. There are problems. First, you have to handle that manually from an operational point of view, but you are also back to this game of the sensitivity of the voice users to a degradation when the data comes on top. We do not think that is the way to do it.

4. ADVENT OF NEW PTT OFFERINGS

- **should we use packet-switched networks rather than private data networks?**

- **will we have an option?**

Talking for a moment about what the PTTs may be offering, one could talk for an hour on this subject but I will mention only one example. We are currently a very large user of private data networks. Should we in fact go to packet switched networks? A more interesting question perhaps is: will we actually have an option? We know that in certain countries the PTTs are considering cutting out the availability of leased lines.

MANAGING A UTILITY

I have talked about utilities. What I mean there is the management of resources on an international scale. There are certain problems that this causes.

MANAGING A UTILITY THE PROBLEMS....

- *Organisationally diffuse*
  - *dependency on operation reporting to remote part of organisation*
  - *responsiveness to local demands*
- *Monopoly powers*
  - *how are prices set*
  - *is there a commitment on level of service.*
  - *do users have a choice*
  - *who arbitrates in disputes*

First, there are organisational problems. The fact is that some remote part of the organisation is having to depend on something which it does not see every day, which perhaps reports to a different part of the organisation. In that sort of situation, how responsive is this central utility going to be to the actual needs at the local end?

Again, what we are creating here is perhaps some sort of monopoly power within the company for those utilities. The sort of questions that one should ask a monopoly are: how are the prices set? How do they relate to the outside market? Is there an actual commitment on the level of service to be provided? Are the users able to choose between the internal "monopoly" and an external service? There may be reasons for it. It may be a specialised service which is not economic to provide on the inside. It may be that the external service appears, certainly to the user, to be far cheaper. In that sort of situation, who will arbitrate in a dispute?

Let me start by saying how we approach those problems on our data processing utilities. The data centres are operationally responsible to marketing companies. They are data centres specifically to serve marketing companies; they are located in marketing companies and organisationally they report to them. But their policies and their technical strategies are determined centrally, so planning is a central responsibility.

. . . AND RX SOLUTIONS

A. DATA PROCESSING

- Data centres operationally responsible to marketing companies.
- Operate within commercial policies and technical strategies co-ordinated by HQ.
- Service contracts with users.
- User committees at I.S. Manager level (meet 3-6 times/year).
- Steering committees at General Manager level (meet 1-2 times/year).

They have service contracts with the users. They actually state the level of service; and that contract is treated as a legal contract, for purposes of satisfying the external authorities that the money that flows between the users and the data centre is justified.

In terms of how we assure ourselves that the services provided are in line with what the users need and continuously are up to scratch, there are user committees at the information system manager level which meet quite frequently, and steering committees at the level of the general managers of those marketing companies who are the users. They meet once or twice a year.

B. SYSTEMS DEVELOPMENT

- Comprehensive standards govern product quality.
- Specification driven by HQ functions, concurrence and commitment to benefits by operating units.
- System creation by special teams ('centres of competence') some located in HQ, some in operating units.
- System implementation responsibility of local I.S. groups.
- High-level steering committee resolves local/shared development trade-offs.

Systems development again we are approaching on a shared basis; that is, the purpose of the processing utility is to avoid duplication of application development or acquisition costs. It is not so much productivity we are after, but assurance that we are dealing with the right problem. So we make sure the functional users have correctly defined the problem; that they agree to the systems development that is taking place; and, at the time of a decision, they are committing to achieve the benefits from the introduction of that system.

The only people, in my view, who can commit to any benefits at all are the users of the programs. If they are not committing, don't develop it; it does not matter whether technically it works or not, it will never benefit the business. Of course, there has to be a standard governing product quality. Documentation is perhaps one of them. Perhaps far more important are the standards around, say, base case testing. There is a concept that you have a test package which is a systems assurance package and which tells the user that it does meet his business needs and will deal with the transactions that he has to deal with every day.

The system creation is done by special teams. We call them "centres of competence". It is not so much that their skills are special as that they have special organisational relationships to make sure that they are correctly related to the headquarters users, to the local users, to the different information system groups and so on. But the implementation of the systems is the responsibility of the local information system groups which report to the general management of those companies. Again, it is an operational responsibility. The general managers have the profit responsibility, and they must have, as far as possible, operational responsibilities for themselves.

Clearly, that leads you to questions of how do you resolve priority conflicts? For example at the centre we might be trying to solve a companywide problem, and therefore give the highest priority to, say, system A; but in one individual operating company they have a local problem, and they think that some local development should have priority. How do we resolve that? There is a steering committee and, if I refer back to my first slide of the organisation, that steering committee operates at the level of the chief staff officer, the directors of these functions, and the regional directors.

C. TELECOMMUNICATIONS

- not currently treated as a utility.
- ad hoc agreements on payment for part shares of lines and equipment.
- no standards on level of service.

but...

we are working towards a utility approach

XEROX already have it.

we are introducing measurements on voice network service level.

On the telecommunications side we are not so far advanced in terms of treating the area as a utility, in the sense of a user acquiring his telecommunications services not from the external PTT but from an internal telecommunications utility. If we have an internal network of leased lines and the payment for that is based on an *ad hoc* agreement — people perhaps pay half at each end of the line, or somehow the line cost is distributed among the data processing cost for the applications which are being processed. It is essentially *ad hoc*. There is an absolute lack of standards, if you will, for the level of service to be achieved by the telecommunications side.

But we are working towards the utility approach. In other words the movement of data, messages, the carrying of voice traffic, the planning of the capacity for that and so on, will be carried out by a central group, and in many cases the user will be totally unaware of how that is being carried. Bear in mind that on the larger scene, I am sure that we will soon see a case where the companies themselves are totally unaware; of course, you are for your voice traffic, which is not private. You don't know how that traffic is carried from one location to another when you dial up a circuit. With an internal utility exactly the same argument applies; arrangements to shunt the units of data or of voice traffic are invisible to the end user. Our parent company already has gone far deeper in that way.

The other thing that we are doing is to put minicomputers on to our voice network to measure the level of services, for example, how many calls cannot be carried because the service is busy. It is very easy to measure how many calls are carried; a test of the level of service of a voice network is how many calls fail. That is more difficult to measure, but we are introducing that.

PROJECT MANAGEMENT

With the sort of concepts that I have mentioned, we are talking about large investments; in many cases, millions of dollars on networks, hardware or application programs. I would like briefly now to mention how we manage such projects with a reasonable level of success.

PROJECT MANAGEMENT

THE PROBLEMS....

Technical projects tend to over-run on cost and time

The claimed benefits are often not achieved.

Line management are not controlling technical investment decisions.

I do not think that I am saying anything new in putting up that slide. Technical projects usually overrun, whether it be cost or time. How many of us can put hand on heart and say that we have achieved the claimed benefits, whatever they were when the project started? Really, is the management of the company in control of these decisions? Quite frequently, not. Quite frequently, the real decisions are taken by Joe, the programmer down there; or some person working in a back room. The technical investments are not being controlled. Again, one could spend a long time talking about this.

... AND THE RX APPROACH

- *Clear definition of roles and responsibilities.*

  e.g. user responsibility for system objectives, base case, implementation planning, achievement of benefits.

  concurrence responsibilities at different stages of project.

  role of internal audit.

- *Formal criteria for investment decisions.*

I think that the key is the definition of who is responsible for what; and above all, I would say that the user responsibilities are the key. They are responsible for the objectives of a system, for planning the implementation and, finally, achieving the benefits. If they said, at the start of the project, that this will save 10% on the stock level, or 52 people or whatever it is, they are going to be measured as to whether it did; whether they finally got rid of 10% of the stock or 52 people.

But at different stages of the project different departments may get involved. As you get into these large organisations all sorts of people, different departments or functions need to have their responsibility defined. An interesting one is the role of the internal audit. Should they be involved during the development or when the thing is finished? We in fact have a mixed approach on that. We say that on certain major projects we involve internal audit during the development of the project; in other cases they will be involved after the implementation.

SYSTEMS INVESTMENT CRITERIA

Pre-tax ROI 40%

based on tangible benefits only

Payback 30 months

Development cycle 18 months

plus

Hierarchical investment sign-offs

We have formal criteria for investment decisions, especially on large projects. I should like to finish by telling you what those criteria are. We look for a return on investment of 40% based on tangible benefits only. In other words, if someone comes to you and says, "I have a marketing information system which will increase our penetration of the market by 1%," that is not a tangible benefit. There is no way that he can prove that the information actually will yield a 1% improvement in market penetration.

Tangible benefits mean: can you save staff? Can you save physical assets? Can you improve the direct contributions to the profit and loss account of a business? We want a pay back in 30 months. And for large projects a development cycle of 18 months, that is from the end of the feasibility study through to the actual live implementation. That's a maximum, not an average, for large projects. For small projects we want pay back in 18 months and implementation within one year. So that is pretty tough. I don't know whether it is necessary to say it, but obviously there are hierarchical investment sign-offs, that is, an investment of a million dollars has to go to a certain level anyway in the company, irrespective of whether it is a systems investment or any other sort of investment.

I hope that has been helpful in explaining what we do at Rank Xerox and, more important, why we are doing it. I will be pleased to have your questions. Thank you very much for listening.

COX: Les, perhaps I could exercise my chairman's prerogative to ask the first question. I think that when the last slide went up, a number of eyebrows were raised; and people felt that if those standards were rigidly applied, they could stop all development now because they would never get through that gate. When you talk about return on investment and the life cycle of the project, this depends at that stage on your estimates. You are right at the start of a project, and you can have a very strong case for going ahead then, which in practice does not work out because your estimating was way adrift. You say yourself that in the past it has been the norm for costs to escalate and so on. So you can have a project which appears to give that return, to pay back within 18 months, to be implemented within 12 months, and you go ahead on that basis but do not meet those criteria. In practice, how good are your people at estimating?

ELSTEIN: The question is absolutely valid. One of the key aspects of controlling the systems programs is that we continuously monitor actual against estimated. In other words, one of the operating parameters of our systems groups is their performance to plan. Assuming that all the projects that were started met those criteria that I have just shown, then the operating performance asks how good were you at fulfilling your performance against plan on all of the parameters. Well, not so much return on investment, but certainly on performance to cost, performance to time scale and performance to the resources needed. So we then measure them against that. Clearly, as we notice that certain units are just deviating wildly, then you would apply a bias against their estimates in the first place. It is difficult to give a more precise answer, it is a thing which is different perhaps in almost every case. But at the stage of what we call "business proposals", that is the investment decision after feasibility study, we apply sensitivity parameters which say that if the cost escalated by 30% or 40% or whatever, or if the time scales got longer, if the benefits were reduced by a certain amount, what would the return on investment then be? So we are judging not only the so-called stated return, but also the likelihood or the sensitivity of that return against things that could happen.

Clearly, if someone has a project which will yield a 100% return on investment, but if it is delayed by two months the return on investment will fall to zero or 12%, then you have got to judge that against your alternative project which has say a 42% return but is almost unshakable. In other words, when you get that project in it will give you 42% and there is almost no risk involved in doing that. So the risk and the return are assessed at the same time, and that is important.

QUESTION: These are obviously very stringent criteria. What happens if in practice the actual ROI falls well below the expected ROI — well below the criterion for ROI — what do you do?

ELSTEIN: Cancel it. We are talking now about things where the business could survive without the thing. We are talking essentially about an investment which has a pay-off but not, for example, the need to satisfy a legal requirement. Clearly, if the Government comes in and says that VAT is going up, there is no point in evaluating the return on investment if you cannot meet the government legislation and you are out of business or something, so you have no choice. But we are talking here about situations where there is a choice, where business would survive without the new system. The objective of a new system is to save money, but its costs are something else. Clearly, if it does not meet what I would regard as reasonable business criteria for an investment, you should not make it.

QUESTION: But suppose all the investment in system development has been made before the low ROI is revealed?

ELSTEIN: I think that the first thing is to note that the estimating is erroneous and, as I said in response to a previous question, keep that in mind for next time the person comes to you for an investment. But you still have one course, which is to say, "OK, bear in mind sunk money is lost. If you've sunk $200,000 in that project you'll never get it back." I assume that you cannot sell the software. So the only thing that is left to you now is to judge whether, based on the forward going costs only (the operational costs plus maintenance of the system) you are still in business. Because if you are, then what you do now is say, "I will still continue because my return on investment on the costs that I have not yet incurred is still 40% or whatever we are calling for." If it is not, then you say, "Stop". But you have to disregard at that point the sunk cost, you cannot get it back.

We have stopped systems that had come to implementation, where looking forward, even though we had completely developed the software product, we would not get back the forward cost; as we operated it, it would not give an adequate benefit. We said, "Tough cookies. People have laboured night and day and finally developed this incredible system, and we're not going to use it." Tough.

QUESTION: How much is your approach influenced by possible future Xerox developments?

ELSTEIN: As far as the use of internal Xerox products are concerned, there are clearly slightly special rules as to the price. The way the rules are set effectively is the transfer price for the product so that if, for example, we were talking about using some piece of hardware which had been developed by Xerox, the price that we use in evaluating that might be less than the price published on the market. There is some transfer price in there. But certainly we would always also make the evaluation at the market price. It is just the same game as when you say, "Should I use the internal utility or not?" The end user, whatever price he is being asked to pay, that is the price on which he makes his evaluation. If there is another level of the company that says, "But there is an advantage to the company to use some special product," or some service which is already available and therefore its overhead is being recovered some other way, then you have another organisational tier in the decision. So you can get into a slightly more complex situation with internal products. I did not dwell on that because, as you can imagine, that is another long topic.

QUESTION: Returning to your criteria, I'd like to know first how good you are at achieving them; and second, if a project looks like overrunning the 18 month criterion do you abandon it, or throw in more resources?

ELSTEIN: Let me take the first question first: how successful are we? We do a post-implementation review as a standard part of our project control procedure; in other words, the phase that follows implementation, six months later, is post-implementation review. One of the things that is reviewed is the benefits. Clearly, that is part of this feedback process of saying, "How good is the whole estimating and decision process?" If you do not do that, you are at great risk because you are only using forward projections, never feedback on actual performance. You have got to do that, at least on your major projects; check that you are getting those benefits.

The second question asks, "Do we reapply the criteria in the same form as the project proceeds?" I really gave you the answer in terms of the return on investment. You discount the sunk cost. But you have asked it again now in terms of the duration of the project. The reason that duration of project is important is that the business objectives change. The technology changes; and frankly, as the project extends, its life cycle may not change. It may still be obsolete at the same point in time as when it was first conceived. Therefore, what is really happening as you push back the implementation date, is that you are shortening the period at which you can count the benefits, and that, in turn, will kill your return on investment. Therefore, you may well find that on a lot of projects, for those reasons, you still want to try to make the implementation if you can. It may mean putting more resource in; but you will find, if you analyse the return on investment, taking into account what I have just said, it may still be worth your while. You have got to look at the two alternatives.

QUESTION: And returning to the first question, when you perform your reviews do you find that the majority of projects meet the criteria?

ELSTEIN: I would say that the majority do not, or have not. I should say that the criteria that we have put up here have been tightened recently. We were not operating to such strict criteria in the past. What was happening, as we checked back against performance, was that we were effectively getting a comparatively poor return.

QUESTION: Did the move to tighten these criteria come from within your area, or from general management and the user?

ELSTEIN: It was a perception of the senior general management that the return on investment on the total systems portfolio was not good enough, which led to a sharpening of these criteria.

QUESTION: We have evaluated the results of 20 projects from our 3 year plan, 1976 — 1978. Deviations from plan were: 30% on development costs, 15% on production costs, and zero on benefits. More interesting perhaps, all the deviations were on unstructured projects; structured projects gave no deviation at all.

ELSTEIN: That's very good. Let me just say that we think that we should be able to get to that point, certainly on the development costs. We are historically around the 30% that you have said, between 25% and 30%; and the target that we have set for our development people is to get closer to the 10%. We think that is achievable. Above all, it requires very much a tightening of the systems specification stage, to make sure that you know before you start what you are going to build. I think that the biggest single cause of overrunning is when at the start of the project you do not know what you are going to end up with; the product that you finish with is not the product that you started with.

QUESTION: Have you any experience of word processing, and if so has it entailed any reorganisation and has it been socially acceptable?

ELSTEIN: I should say that I am not personally responsible for that aspect, which is why it was not included in my talk, but I can give you a few comments. Bear in mind that we are in the business.

The experiment that was tried — I should say "imposed" — in my own shop was effectively to remove the idea of the personal secretary and set up an administrative service centre which comprised specialised typists operating word processing equipment, together with administrative assistants serving a number of managers or senior people and handling things like the mail, telephone, and the setting up of conferences.

I think that the key to this type of approach is that you must pool the resources to be able to specialise. If you do not specialise it is a disaster; if you just put one girl on the word processing unit and she is away, that's the end of that. You must have a pool of people to provide cover. Then there is this idea of specialising them within the pool so that some are doing the typing and others the administrative support, and maybe they switch over from time to time for more job interest. I think that is very necessary, because there are some tasks which are not word processing but which are, from the point of view of the people being served, just as important as typing.

Word processing caused problems in that the people who were there before we put it in have not stayed. They have in fact moved to other jobs in the organisation. We started by saying, "Can we do it with the department?" and clearly not. People have joined you as a secretary and are now told, "You're now a word processor." That is impossible, and therefore had to be changed. We had different sorts of people to do these jobs. It required a rotation of staff or recruiting of new people, allowing other people to leave. So there were implementation problems from a people point of view. There were other problems, mentioned by a previous speaker, in terms of would the managers accept it? From the point of view of the staff who did not previously have personal secretaries it was a big improvement, because the turnaround on their typed documents and so on improved remarkably. The managers who had previously had some sort of personal support, which was now taken away and only supplied on a pool basis, felt deprived. Then you have got to balance who gained and who lost.

QUESTION: Returning to the criteria we were discussing earlier, and given that you don't take account of intangible benefits but only quantifiable savings, two questions: when did you last do a project for the marketing department, and when do you expect to do another?

ELSTEIN: Well, it is clear that the marketing projects cannot satisfy the criteria of savings, at least not in normal terms. What I didn't put in there was: what do you do with projects? Are there any exceptions to those rules? Clearly, any investment in what you might call marketing information, how effective will the sales force be? If we have it, how effective will the sales force be? Those types of investment therefore do not satisfy those criteria, and they are then subject to senior management review. The only thing that you can then say to management is, "All right, here is the portfolio of investment for 1978. Here is my list of projects, and my resource that I want to apply to this is, say, $10 million." I will then rank the projects in terms of their return on investment and all the other criteria. There will be some projects which are in there, like marketing information projects, which do not satisfy the criteria, and yet the business says, "Gosh, if we don't have those, our decision making will be a disaster." So at least management knows as it decides to do a marketing information project rather than a spare parts control project that they have dropped tangible benefits for the sake of intangibles. Provided that they are conscious of that, I think that you have to give them the prerogative to do it. So we have not stopped doing marketing projects, but what we have said is "At least let's know very well that those have no tangible benefits and what tangible benefits we're giving up by doing them."

QUESTION: What happens if a user refuses to renew a service contract?

ELSTEIN: I asked, "Who arbitrates for disputes? Do the users have a choice?" As a policy statement, we say that for all practical purposes computing is an in-house facility. Therefore, disputes on service level are handled usually by the two levels of user committee which I showed on the chart. First at the level of the information systems managers, who have the operational responsibility; and second at the level of the general managers who are paying the bill and who are responsible for the policies under which these data centres are run; they are steering that. We have not had a case where that finally exploded and they could not reach agreement. I would say that it is always possible that you could get that, because to some extent the payments between departments of the same company are regarded as "funny" money. In that situation, the whole concept of service contracts, price differentials and so on may become quite warped and become a political as opposed to a factual discussion.

So I can only say that we recognise the risk and that we attempt to handle it through the steering committees; that the issue is open and does not just become a unilateral abrogation.

That may sound just like pious talk. The one principle which is very important there is to make the whole thing an open management process. The policies under which those data centres are run are openly published; they are openly discussed with the users. The meetings are run on a democratic basis; anyone can introduce an agenda item in advance. The reference to senior management and so on is quite open. That is the way that we run our company. If we were running it in hermetically sealed units, where the thing just festered away and finally appeared at the senior management level, saying, "Decision tomorrow. We want to abrogate," I think that could be extremely difficult to handle; and the best way to handle that is not to get there in the first place. If you did get there, I think that you could be in real trouble.

COX: Thank you very much. At that point we must cut off questions. Whenever a speaker comes along and puts forward a very clear-cut, hard-line approach, as Les has, people look a little bit suspicious and think, "Does it really operate like that in practice?" I am sure, from Les's conviction, that it really does. It would be very interesting to have him back in a year or so's time, reviewing the progress and whether people have really met those criteria. Thank you very much for a very frank discussion.

THE LESSONS OF THE LAST FEW DAYS

G.E. Cox

COX: To go through the speakers that we have heard and pick odd points from their talks I think might impress us with the number of topics covered, but would hardly contribute to our overall understanding. If you are like me, you will have sat through the last couple of days and experienced a number of different feelings. Excitement: when one hears Viewdata being discussed, one is excited about it. One starts thinking of opportunities. It does not really threaten many businesses here, and it is wholly positive in its aspects. Reassurance: when you hear people predicting that in the future we shall have unlimited capability for transmitting data and that many problems that we have now will disappear, it is reassuring to know that, in time, things are certain to get better and restrictions are certain to be removed. Certainly when Derek Roberts spoke so convincingly about memory in the future, you know that there is more to come and it will be cheaper, again it is reassuring. Technology will be there to assist us.

In other areas, I think that we experienced apprehension, perhaps even fear. I found Alex's talk earlier today quite frightening, because he spoke of a discontinuity in development which I am sure is correct. He spoke of developments which overtake things that we have at present and obsolete them, which means that perhaps some of the lines that we are going up might even be cul de sacs. That is something that you cannot ignore. I thought that his analogy that the impact of microprocessors was like a wave hitting one bay and swamping it, while the next remains calm and you don't even suspect what is coming, was a very good one.

I suggest, too, that whilst one is listening to a conference such as this, one's attitude oscillates. There are times when one thinks back and says, "How do I get hold of all this?" and other times one retreats into thinking, "Perhaps it's being exaggerated. Perhaps the impact isn't quite as strong as that or as quick as that." It is very difficult, therefore, to sum up the overall situation. If one had come in for only one of the talks, one would have gone away with a different feeling, I believe.

So what is the situation? Perhaps I can contribute to the overall proceedings by standing back and trying to give you a perspective as I see it. Think of the environment in which all of us here operate. Firstly, there is the enormous growth in people who work solely in information. Figures that we have quoted this year indicate that round about half the population already is employed as information workers, or professional people. Clerical workers, civil servants and so on; in other words, people who push nothing harder than a button; people who lift nothing heavier than a piece of paper or a telephone.

This is partly brought about by automation in other areas; automation in manufacturing, distribution, agriculture and so on; and it is partly brought about by society's apparently endless ability to generate information and data — in education, business, government, our private lives — and our insatiable appetite for it.

There is a rise in administrative costs. Whilst we have a migration of half the population into non-physical work, we also have a rise in the demands that are made upon them in terms of the flexibility and response required of systems, and their general complexity. We have growth in remuneration; pressures which have been off for the last year or so, which will come back on very hard. And with the shift of the population into information working, I think that you also see a shift in industrial power into that area.

Couple this picture of the working population with the capital investment picture. I mentioned a moment ago the much higher investment that has been made in automation in these other industries. I have seen many figures on the subject. Ones that I tend to believe are somewhere near the truth come from a survey done for the States, which shows the average capital investment for a worker in agriculture was \$35,000; in manufacturing \$25,000; in information handling \$2,500.

If that surprises you, you had a very good example in the newspapers this week. You must have read that Barclays are cutting back on branches. One of the major clearing banks is being forced to cut back on the number of branches that it operates. I believe that it is still correct to say that Barclays are the largest single commercial data processing user in Europe, the biggest private customer for IBM in Europe, a very advanced computer user. I regard them as being highly automated. The banks themselves have been very advanced computer users. Yet, as they quoted in their Press statement, 74% of their costs were labour. That was the figure quoted in the papers this week. Those labour costs are forcing them to cut back on branches and their services. I think that is an illustration of the statistical trend that one picks up from these surveys.

So against this demand where more people are working with information and the cost of those people is getting higher, we have on the other hand an explosion or revolution — and I don't think either term is too strong — in available technology for handling information. We have been reminded here of the increase in processing power and decrease in processing costs that are available, figures that are quite bewildering. They reach the point where they just strike you as being fantastic and no longer have a physical meaning. The possibility of storage with ever larger capacity and ever smaller physical size. Costs which come down to the ridiculously small compared with our past experience. We see developments in transmission and switching, which again are remarkable.

Consider for a moment word processing. We may debate whether it really gives us an improvement in typing performance, we may debate how it will be used, we may debate whether it is cost effective yet. I tend to go along with certain people who believe that it will become widespread in the future as it gets cheaper and as we overcome certain of the social problems.

We have the question of integrating these facilities: voice; data; telex; and the many services that are available. We have the question of something like Viewdata, surprising and unlooked-for on our part; not just as a new way of communicating with the public and leaving messages for your wife, but as a cheap, idiot-proof terminal for accessing the company data base: it is enormously exciting.

So on the one hand you have this enormous market for supporting people handling information, and on the other you have this revolution in what is being made available. I think this means that the effect will be dramatic, inevitable and fairly fast.

The question that we have to face then is: can we meet the onrush of these technologies? I have heard persuasive arguments put forward for carrying out reviews of information in companies. You can say, "If we are the people concerned with company information systems, surely we should have at our fingertips questions about how much information is entirely internal and never crosses the company's boundary," the way that one of our speakers has obviously analysed it. "Surely we should have a feel already for how much mail could be transmitted electronically," and most of us do not. People say that what is really required, if you like, is analysis on a grand scale; looking at the various areas of information transmission in the company and communication, and breaking them down.

Again, you hear persuasive arguments for putting in someone to head up all information flows in the business, a sort of information supremo.

I have heard both these possibilities put forward. I really think that neither is attractive nor practical for most of us here. Even if one fancied the job of being an information supremo, to say, "It's required and I'm here to fill it," is not going to endear one to the rest of the organisation. I think that to talk about carrying out studies on a grand scale also does not move us forward.

Alex d'Agapeyeff made a highly practical suggestion earlier in the day when he said, "Here's what's happening in microprocessors: here's what you need to do about them. Get yourself a microprocessor and a bit of money, put a jolly good man on it with some good people under him, and try it." I think that is good advice, and I think that almost no one here will accept it. As soon as he said, "What you've got to do is take one of your better project managers and put him on it," the thing lost reality, I'm afraid, in terms of the way that people react. Our company makes a good business, and so does Alex, out of supplying project managers because people cannot even meet their current demands. To put people off, to experiment, would be a bold move which would be carried out by a very few organisations.

At the same time, if one cannot go for a grand policy to accommodate what is happening, neither can one sit back and nibble; nor adopt the attitude which is very tempting, which is to say, "Well, I don't really have to go away and worry too much about any of this, the more important bits will hit me anyway; and the rest of it, why speculate on it?"

I think that what is really required is to think of action in four areas. One is to start pushing for at least an extension of one's corporate policies to accommodate some of the issues that you hit fairly quickly. I had a friend who was at our May 4 conference. He went away and he said, "It's most important that my organisation do something about this, at least in one or two areas. One is we've been planning our networks and I think we've taken far too limited a viewpoint. We must widen it out and look at the other things we'll be doing with our networks in the future, rather than just what we're planning them for now, and at least make people aware that we'll be using this resource in a different way." The other thing that he tried to do was to get some action on word processing. He had a very sensible approach. He said, "We have a technical team looking into the technical merits, and we'll set up a team to look at the personnel problems, a team constructed of people like the chairman's secretary and the like." That is a very sensible approach, but it got squashed. The network one went ahead, because they were actually planning a network and were about to sign a cheque. It was really felt that word processing was not in his bag. The management said, "Yes, we ought to formulate a policy on this, but computing and networks are yours and someone else ought to think about word processing." I think that already in organisations policies are laid down for buying computers and minicomputers, but office equipment does not fall within it. I think that one will have to push to at least get recognition of where a policy on this is required by the company or group.

The second thing that needs to be done is to start modifying our future systems plans now — the kind of people we employ and the kind of skills that we employ. When we meet something like microprocessors, are we then really going to devote a resource to it next year? Are we going to look for a project for it? Where does this come in our plans? I think that it is fair to say that most plans in the data processing area turn out in practice largely to be extrapolations of past plans. We have been able to follow a fairly well-defined route in recent years where we build on this plan. I wonder now whether we don't have to question our rolling five-year plan for some of the things that it probably excludes altogether at present.

This leads me on to two other points. There is clearly a need to educate users over what is happening. One of the pleasant things that I do at present is that about once a month I go and talk to one of our clients, an interesting and a very fast growing group. They have very little in the way of technology already, and a lot of money, which makes them overpoweringly attractive to consultants. They are putting all their executives through an executive training programme. I give them a talk on computers and also on some of these other areas, and I end up discussing things like Viewdata. It is most enjoyable because I can feel them warming to it as I go through, and getting positively excited. I really bring them to the boil in the end. If I time it right, I can be like Billy Graham: "All of you who believe in this new technology and what it can do for your group, come forward."

They all go away and think of the impact that these new information technologies could have on the business, and they come back and list some of the things out. What is tremendous is to witness the impact on people who are very commercially minded — it is a very commercial group — who were previously totally unaware of the kind of information tools that are becoming available. I really do think that there is a lot for the user to enthuse about and appreciate in this area, and pitfalls for him to be aware of that should come from informal advisers, not from salesmen. The capabilities of word processing should be something that is understood internally, not when it is offered as a machine. We are in danger of losing our grip in this area. Most people would naturally talk to us about computers, but there are many other related subjects as well which do not quite appear to fit into our bag.

That leads me to another problem. We are in no position to educate the user until we have educated better our own analysts and designers. A week or so ago someone rang me up and said, "We have a problem of communicating price changes every day, and it strikes me that Viewdata would be an ideal way of keeping our supermarkets up to date with price changes. George, what worries me is that my team didn't even think of using Viewdata for this. Can we just chat it over?" We chatted it over, and really the best answer to the problem was not Viewdata in this particular case; nor was it the computer terminal system that they were considering implementing; it was more likely to be facsimile transmission or Telex, or even dictating to a girl over the telephone. One or two solutions were mundane, one or two they hadn't thought about. The fact is that it hadn't even entered their thinking.

I think that emphasises the fact that the majority of today's systems analysts are really used as, and largely trained as, computer systems designers. Like most of us, I went through analyst training, many years back. What it really taught me to do was to articulate requirements so that we could put them on a computer.

I was involved in the application of computers to machine shop control. I looked at the requirements of each section, and the only requirements in which I was really interested — like shop routings — were those that could be put on to the computer. We tried to produce some shop documentation on the machine. We did our scheduling, which was legitimate. But in terms of documentation of the information that the foreman required we were way adrift. For example, drawings were of no consequence to me: "We can put them in the plastic bag with the cards." Thinking that one could present the information all in telegraphese and capital letters was a joke. But now we have many more tools genuinely to tackle information requirements and systems analysis. We are moving towards the capability of genuine, real systems analysis at last. The thing that strikes me is that the big gap in knowledge, even awareness, exists within our own departments now. I think that is probably the priority in tackling the question of how we meet and cope with the technologies that we are hearing about.

Gentlemen, you have listened very patiently. I should now like to ask David Butler to conclude the conference.

BUTLER: Gentlemen, a few brief announcements before you depart. First, you will have noticed that we have chosen not to distribute session evaluation sheets. We are confident that you, the delegates, will provide us with the kind of feedback we need to ensure that the session subjects and formats are tailored to extract the most from those conferences in the future.

The transcript of the conference, including the visuals, we are hoping to publish before the end of this year; but if there is a delay with Christmas mail you should get it very early in January.

I should like, on your behalf, to thank all the speakers who have come to talk to us during the past two days. It has been self-evident that they have put a lot of thought, care and attention into the preparation of their speeches, and I am sure that we appreciate that enormously.

I should also like to thank the hotel for the effort that they have made to make us feel at home. I should also like to thank you, the audience. A great deal of the success or failure of a conference of this type depends on the readiness of the audience to become involved in the sessions and, from my point of view, I must say that you've been a smashing audience in that respect.

With that, may I wish you a good and safe journey back to your various destinations and say that we look forward to seeing as many of you as possible at the next conference in April. Thank you very much.

The Butler Cox Foundation

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

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

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

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

Printed in England