Transcript

Management Conference

London
May 4, 1977

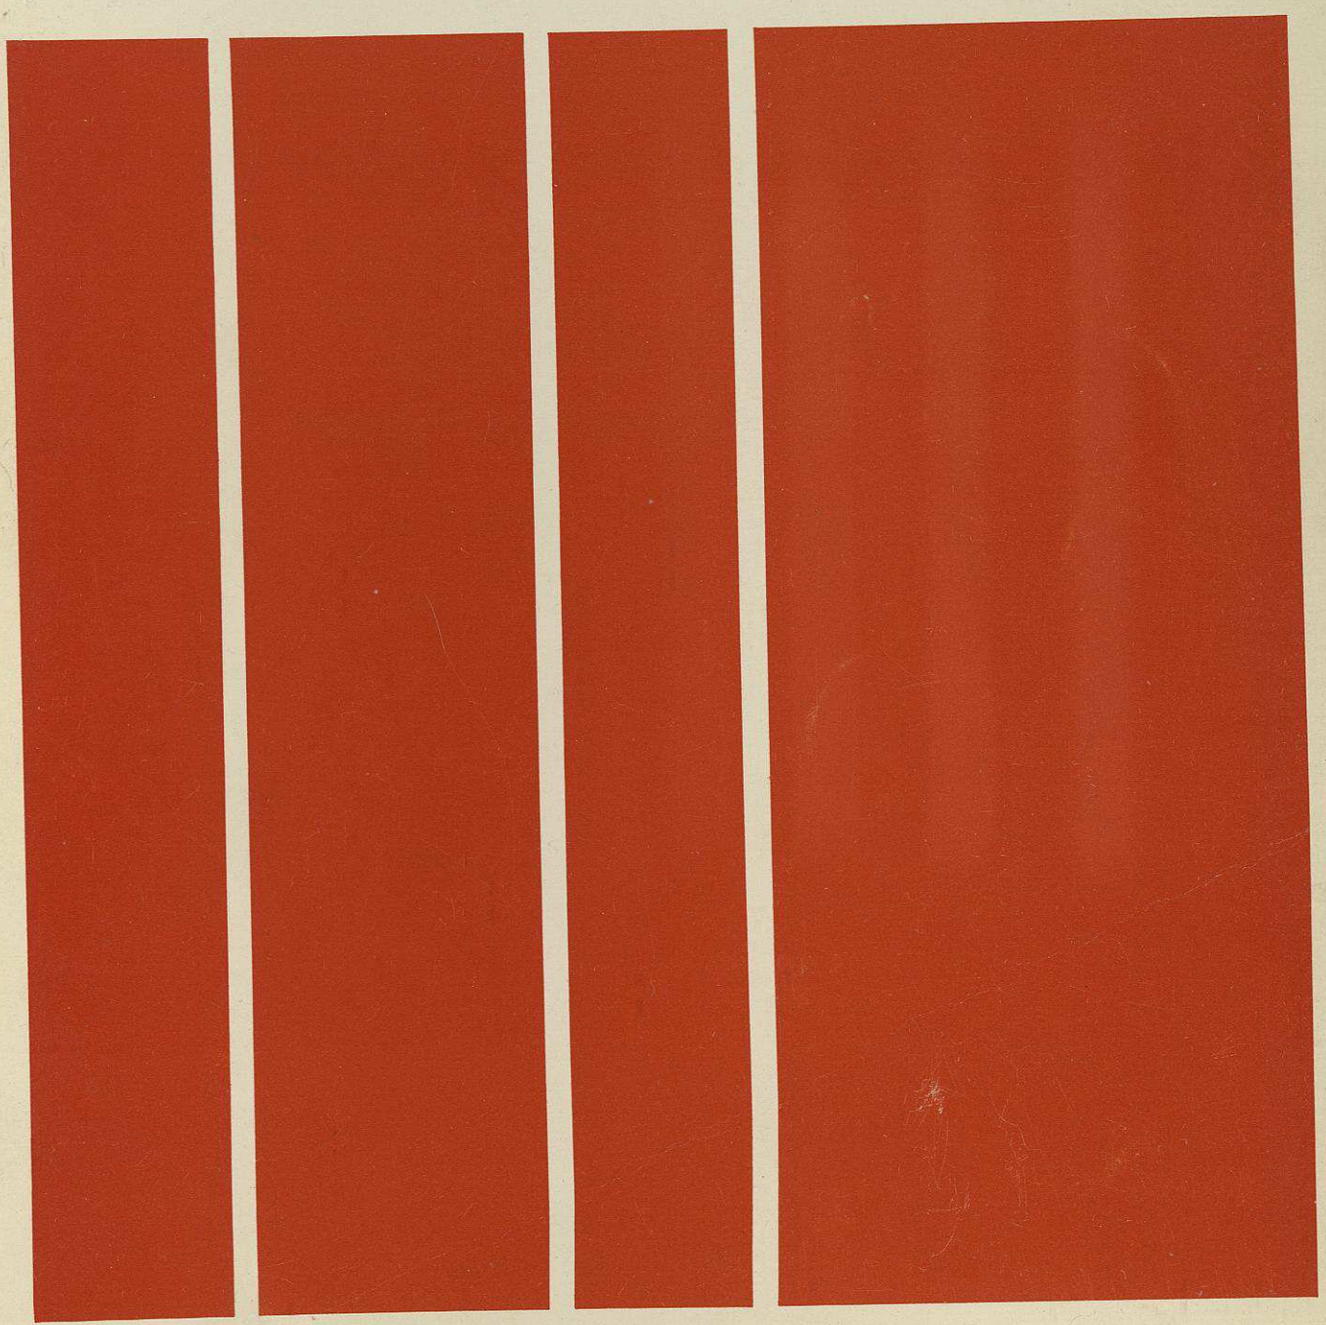

The Butler Cox Foundation

Butler Cox Foundation

MANAGEMENT CONFERENCE TRANSCRIPT

London, May 4, 1977

Butler Cox & Partners Limited

The Press Centre
11th Floor
76 Shoe Lane
LONDON
EC4A 3JB

CONTENTS

* Butler Cox and Partners Limited

DELEGATES LIST

Booker McConnell Ltd.	J.D. Golding
British Rail	O.R. Benz
British Airways	J. Lewis
British Oxygen	R.H. Tester
British Steel Corporation	T.G.W. Barber
Burmah Oil Trading Ltd.	P.H. Crosby
Cable & Wireless Ltd.	C. Jacob
Calor Gas Ltd.	M.J. Elliott
Central Computer Agency	P. Hearson
Central Electricity Generating Board	B.P. Buss
Conoco Ltd.	D.D. Vinestock
Dickinson Robinson Group Ltd.	M.G. Perrin
Alexander Duckham & Co. Ltd.	B.F. Watson
Freemans (London SW9) Ltd.	K. Hurle
GEC Telecommunications Ltd.	J.C. Townsend
General Motors Ltd.	K.R. Thurtell
Glaxo Laboratories Ltd.	I.G.F. Lee
Granada Group Services Ltd.	G.J. Cockbill R.L. Tyrrell
Grandmet Information Processing Ltd.	J.A. Corner L.F. Green
C.E. Heath & Co. Ltd.	M.J. Burton
House of Fraser Ltd.	T.W. Terrett
Hoover Ltd.	M. George
IBM United Kingdom Ltd.	J. Crickmer

Imperial Chemicals Industries Ltd.	K. Edwards
International Computers Ltd.	H.I.J. Goldberg
International Pinchin Johnson	A. Sinclair
Lamson Industries Ltd.	L.M.L. Coughtrey
Louis Newmark Ltd.	D.L.F. Wilkie
Marks & Spencer Ltd.	G.R. Hattam
Molins Ltd.	H.N. Wilkinson
National Computing Centre Ltd.	D. Firnberg
National Enterprise Board	J.A. Pearce
National Physical Laboratory	D.W. Davies
The Nestle Co. Ltd.	J. Wickstead
The Open University	J.J. Sparkes
Ozalid (U.K.) Ltd.	W.S. Paxton
P & O Steam Navigation Co.	F.E. Heenan
Plessey Telecommunication Ltd.	F.P. Whitlock
Post Office Telecommunications	T.D. Brougham
Prestige Group Ltd.	H. Washbrook
Prudential Assurance Co.Ltd.	L.G. Hickmott
Queen Mary College, Computer Systems Laboratory	G. Coulouris
Redland Ltd.	G. Chatterton
RHM Management Services Ltd.	J. Playfer
Spillers Ltd.	B.J. Gladwin
Standard Telephones & Cables Ltd.	P.A. Langston
Times Newspapers Ltd.	G. Consterdine
TKM	J.A. Wheatcroft

Tootal Ltd.	A.C. Smith
Unilever Ltd.	B.G. Maudsley
Unilever Ltd.	K.W. Sproul
United Biscuits Ltd.	D.R. Hoyle
Videomaster Ltd.	D.C. Martin
Weir Group Ltd.	R.G. Codd
Plandata, Sweden	J. Mattsson
Belgische Boerenbond, Belgium	G. Van Belle
Ahrend Groep, Holland	F. de Boorder
Rekencentrum, Holland	J. Maat
Akzo Pharma, Holland	Bollen
Akzo N.V.	A. Zylker
National-Nederlanden, Holland	N.J. Westpalm Van Hoorn Van Burgh
Computer Weekly	C. Dearn
Butler Cox & Partners	D. Butler (speaker) G. Cox (") A. Gunton (") K. Kozarsky (") R. Woolfe (")
	A. Brewer J. Kinnear L. Kehoe B. Lacey (Ms) J. Haffenden (Ms)
Speakers	H. Donaldson D. Holloway R. Malik

A

THE EMERGING PATTERN OF INFORMATION SYSTEMS

David Butler

BUTLER: My name is David Butler, Chairman of Butler Cox & Partners. I welcome you to this, the first management conference organised by the Butler Cox Foundation; and a special welcome to our visitors who have come from overseas.

First, may I briefly say a few words about our company. We were founded as a management consultancy company specialising in computers, communication and office automation, at the beginning of this year. We are independent in the sense that the company is wholly owned by its working partners without outside ownership. We were set up to tackle the three areas which I have mentioned as a consultancy service: computers, telecommunications and office automation. Six of us who were among the original founders of the company had worked together before for an American consultancy company; but as soon as we set up our company we decided to start handpicking new skills from different backgrounds to bring into the company, and we have now advanced fairly quickly in that direction.

The purpose of this part of our operations, the Foundation, is to analyse, evaluate and report to members on developments in the three areas which we are set up to cover. The Foundation is international in scope; hence the presence of a number of overseas invited guests here today. It is run jointly by ourselves and the members who subscribe to it; so it is an exercise in commercial democracy as opposed to industrial democracy.

The purpose of today's conference is to be the first of a series which will continue, we hope and expect, for a good many years. We are confident that, several years from now, many of us will be looking back and saying, "Do you remember that chilly, depressing day in May when we all got together at the Painters' and Stainers' Hall?" and some of the questions that were raised on that occasion will no doubt still continue to be defying total solution even then.

We intend during the course of the day to provide an initial report from our side on the main research projects which we have undertaken and on which we are now working; and to give an interim account of the approach which we have adopted to

these particular research projects. Most important of all, during the day we plan to lay on the table some of the main issues which we think are of vital importance to management services directors right now. I should like to invite you all to feel free, at the appropriate points in the conference, to comment, to question, to interrogate the speakers and to express your own points of view. This is intended to be a conference, and at a conference people confer; they talk to each other and do not just listen.

During the course of the day, the emphasis will be very much on the policy level decisions that management services directors have to take. We are not concerned today to go into very great technical detail; we are concerned to highlight at least some of the critical policy issues that management services directors have got to get right, if they are to succeed in tackling the problems and exploiting the opportunities that exist right now.

I should advise you all that after the conference a transcript of the conference will be prepared, so that those of you who ask questions or make comments from the floor are hereby deemed - even if you happen to be a late arrival and do not hear these words - to have given us permission to include your questions or your comments in the transcript of the conference.*

My task now is to try to set briefly the scene in which the other speakers, during the course of the day, will operate. First, I believe that there are four major factors which dominate the changed environment within which systems are now beginning to be put together.

None of them is new or will surprise you, but perhaps we can go some distance towards putting these factors together and understanding better their impact on systems. First, and perhaps most obvious, is the declining cost of hardware. During the course of the day several speakers, including Karl Kozarsky, Roger Woolfe and Dennis Holloway, will refer to the declining cost of hardware; of processors, of storage, of some terminal devices, and will draw different implications from these facts. Clearly, one has only to look in the windows of shops selling pocket calculators to see dramatic evidence of that particular change.

*Editor's Note: The questions were not recorded verbatim, but by hand in abbreviated form. However, all the answers were recorded in full, except K. Kozarsky's answers, which again were recorded in abbreviated form.

The second major factor is the higher cost of people - and by people I mean not only people working in offices, clerical workers in administrative offices, but also software development staff. During the course of the day, again reference will be made to this. Hamish Donaldson in particular will be drawing some lessons from that particular trend.

The third major factor in the situation - perhaps the most intriguing from a technical point of view - is what has been referred to as the convergence of technologies; which is to say, at its simplest and most brutal, that many of the skills which used to dominate, for example, the telecommunications industry, such as skills in electromechanical engineering and copper wire engineering, have become less important, while other skills such as the ability to design, produce and market electronic switching equipment which has a good deal in common with the design, fabrication and marketing of advanced computers, has become more serious. I take it that most of us in this room are probably familiar with that trend; again, it is a question of trying to learn some lessons from this.

The fourth major factor is concerned with the dispersion of systems power throughout organisations - what has been referred to as distributed processing; the ability to exploit some of the trends which I have already mentioned to put processing power, storage and retrieval facilities at the elbow of the user.

I think that these are four of the major strategic facts of life which management services directors will have to take into account over the next few years. The ways in which they may do this will form the subjects of a number of the talks which will follow later.

I should like to mention briefly some of the results of these four major factors which we see around us. One of the most stunning examples of the effect of these four factors at work is the evolution of the world's largest computer manufacturer - IBM. I believe that it is true to say that, a few years ago, when IBM announced that it was going into the business of voice telephony, most of the world's telecommunications manufacturers were not too alarmed by that announcement. They felt that they understood very well the business of designing, marketing and supporting voice switching equipment, and it seemed improbable to them that IBM could successfully establish itself in that particular market. We all know now that in Britain, and to a lesser extent in other countries in Europe, IBM has successfully established itself in the business of voice switching.

We also see another major factor in the evolution of the company: its diversification a few years ago into the business of launching and operating communications satellites through the medium of its subsidiary, Satellite Business Systems. We also see, more recently, enormous emphasis in the development of IBM's product range on systems designed to put beside the end user - word processing systems, the ink jet printers and such. The strategy of IBM is clear. IBM clearly sees itself as the first, the largest, and perhaps - who knows? - the only company which is able to offer a total range of service to the user of information. This is the strategic role which the company is attempting to define for itself.

Less clear - but in my view equally important in terms of the longer-term evolution of the market in which all of us have to sink or swim - is that other companies are responding in different ways to the same kind of pressure. It is no coincidence that the world's largest communications company - and indeed the world's largest company of any sort - AT&T - has changed its own policy during the period when IBM's strategy has become clearer.

For 50 years, AT&T did not pursue communications contracts outside mainland America. Shortly after IBM's announcement that it was going into the satellite business, AT&T announced that it was going to pitch for communications contracts worldwide. I don't think that was a coincidence, and neither does IBM. Secondly, AT&T is in the process of promulgating in the United States laws which would substantially alter its position as a monopoly supplier. I think that all of you are aware that, over the years, the American Government has looked at AT&T and said, "Well, the Bell system is a wonderful thing. We don't want to dismantle it or inhibit it, but we believe in free competition," and it has intervened, sometimes decisively, to encourage competition to AT&T in the sense of the development of specialised common carrier services and the like. The legislation which Bell is busily lobbying senators and congressmen in the United States to support would effectively put the clock back in communications in the United States, to a point before any of those competition-favouring moves were made, so that Bell would be re-established as the only effective innovatory force in communications in the United States.

Now I do not wish to judge the issue of whether that is a good or a bad thing. If you want to ask questions about that, I suggest that you address them to Karl Kozarsky, who knows much more about it than I do. But what I do see quite clearly is that this is a clear response of Bell to the challenge of IBM in this new world market which is emerging; and that if it is successful it will have a direct effect on everbody in this

room, because some of the competitive activities which have sprung up in the United States to challenge the AT&T monopoly have been seed beds for some of the very interesting examples which you will hear during the course of today. If those seed beds are dried up and no longer exist, it will have an effect - indirect but serious - on all of us in Europe as well as those in the United States.

Other manufacturers of both computers and communications equipment are responding in different ways. During the course of the day we will hear some examples of that.

Now I would forgive anyone for saying, "Well, in the last five minutes old Butler's been talking about some of these changes in companies like IBM, AT&T and so forth. It's all very interesting, but how the hell does it affect me and what the hell can I do about it?" I think that the first thing we have to do is to try to understand the strategic moves in the market place. In a moment or two, I will talk about how some of those strategic moves affect the user and some of the sensible things that he can do to try to equip himself to deal with them.

But first, I should like to give you two brief quotations which I think illustrate quite clearly both the necessity of understanding these strategic moves and the difficulty of doing so. The first quotation is from a legal action between the Federal Communications Commission in the United States and AT&T. The FCC judge asks an expert witness from AT&T whether the cost and profit mechanisms of business communication were more favourable to them than those of residential communication. "Do you make more money out of telephones in offices, or do you make more money out of telephones in homes?" The AT&T expert witness replies, "I can't, from my own knowledge, say that the business traffic in the day is more expensive to handle - again, I'm thinking of it on a unit basis - than residential traffic. I just don't know that. Such a study would require taking a look at what are the characteristics of business traffic, how much they affect the peaks of the various elements that go into making up the network, as against the residence traffic and how it affected the network. So, no, I can't answer that; and I think it would take a study, which I don't think I know having been made, to make such an answer." In spite of the inelegance of the language, his ignorance is clear!

The FCC judge, incredulity oozing from between the lines, says, "You can't say that the business traffic in total is more lucrative than the residential traffic?" The AT&T witness: "No, I can't say that, sir. No, because you're talking about the relationship of what the traffic costs you to what you get from it." He's right about that. "The business traffic to

the extent it occurs during the day gives us more revenues, but the relationship of the revenues you get to what it costs you to put it in, because you have to build the facilities for it, I can't answer. I mean, I've never looked at that."

Coming from the world's largest communications company, I find that to be an astounding admission, that in 1974 they did not know whether residential traffic was more profitable to them than business traffic. To my mind, this makes any discussion of tariffs, and their equity or otherwise, totally absurd in the United States. I should be interested to know how many PTTs in Europe could put themselves in a totally different position as far as costs and revenue are concerned.

The second quotation that I should like to make - again from AT&T - shows that it is not always easy to understand even what is happening in this world of systems. The first excerpt is from AT&T's annual report in the year 1971, published early in 1972. It says:

> "1971 saw a quickening of the pace at which the Bell Systems switching facilities, hitherto electro-mechanical, are being replaced by electronic switching systems, ESS, that are analogous in many of their operations to the most advanced computers. As the year drew to a close, new, multi-million dollar electronic central offices were being installed at the rate of about two every week."

That is from the annual report. An internal memorandum, dated a fortnight later, which would never have seen the light of day if it had not been subpoenaed by the Federal Communications Commission, bleakly states:

> "The reduced demand for ESS frames has been of great concern to both Western Electric and Bell Telephone Laboratories."

I make these points not to suggest that Bell is uniquely incompetent; nor to suggest that Bell is deliberately and systematically trying to deceive the world. It is, of course, but so is everybody else. I am making these points to underline that it is extremely difficult to follow the strategic moves in the market place; that these have a direct effect on the end user; and that it behoves us to try our best to understand them. Bell perhaps is in a slightly invidious position in that some of its most intimate secrets have been subpoenaed by the FCC, and therefore it is easier to see the moves in Bell than it is in some other companies.

So what are the results of this for the user, and for the management services director, who sees these strategic changes

going on around him and would like to adapt his policies accordingly? First, it seems to me that many of our precepts of good systems management were really evolved during an era when the computer, the main central processor, was the most expensive piece of equipment that we had. Many of our methods of managing systems depend on the assumption that it is a major objective to try to get high utilisation from the central machine. As soon as that ceases to be true, as increasingly it is, then I suspect that a new set of rules begins to evolve. It is understanding what some of those rules are which is one of the main purposes of this day, and is the subject of Hamish Donaldson's talk.

Secondly, caveat emptor - let the buyer beware - has always been a good maxim for the purchaser of systems. I think that it is even more valid today because some of these strategic moves teach us to understand that companies which seem to promise one thing are really delivering another; and companies which perhaps will be delivering some of the most interesting things in ten years time may not even exist today.

Thirdly, I believe that the users have to understand much better the total economic aspects of the systems about which we are talking. In the past, it has been traditional to talk about computer budgets which averaged some 2% of total company turnover; and frequently, those sums of money seemed relatively unimportant to top management. But once we start talking about the new kind of systems, embracing communications, computers and office automation, we are talking about having an impact on the total handling of information in an organisation, and we are very often talking much closer to 20% of the total company costs.

The final result for the user which I should like to mention concerns the human aspects of all this technology. Of course, it is possible for us to apply this technology in ways which are exciting and important. But the question is: are we going to simply assume that the people who have to operate those systems will somehow be able to fit in with them? Or are we going to take as a major criterion in the development of those systems the ergonomics and the hygiene of the jobs which people are trying to do? I think we already see today some examples of situations in which that hygiene is neglected.

Let me give a very brief rundown on each of the speakers and what he is attempting to cover during the course of the day. Dennis Holloway is a senior manager with responsibility in the area of strategic planning for Plessey Telecommunications. He is going to talk to us about some of the main points which should affect communications planning in companies today; some application areas; some bits of equipment; some policy points which should not be forgotton when you are drawing up your telecommunications plan for the next few years.

Karl Kozarsky, who is our associate in the United States, and whom, with all due deference, I can only describe as an expert in the field of private communications, particularly the area of computerised branch exchanges, will then develop one aspect of communications policy, namely the role of the computerised branch exchange. This is a fundamental piece of equipment which most people in this room are probably in the process of acquiring or have recently acquired. Some of the points which Karl is going to make will be illustrated by descriptions of the facilities of two new announcements in the United States. Apart from anything else, I think that may have a good corrective effect because IBM has very cleverly managed to create the impression in Europe that voice switching equals the 3750, and some of the points that Karl will be making will put that impression very much in context.

During the day, there will be numerous mentions of cost changes, increases in this cost and decreases in that cost. Anyone could be forgiven at the end of the day for saying, "So what? I see these cost changes but what effect do they have on me as a manager?" Hamish Donaldson, at Hill Samuel, has tried to adapt his overall strategy to the development of systems to try to swim with the tide of cost changes as opposed to trying to swim against it. He will tell you how he applied that philosophy in the bank, and what results it has had both on the quality and cost of the systems that he and his team have developed.

After lunch, another guest speaker, Rex Malik, who is a journalist, broadcaster and author, and whom I am sure is known to all of you, will present what will probably be the most forward-looking and futuristic session of the day. He will be talking about new horizons on this occasion in the area of office automation.

Dennis Holloway, earlier in the day, will have put emphasis on telecommunications at the policy level. But a crucial component in that policy is the network and the network operating system. Tony Gunton, from Butler Cox and Partners, will try to explain to you why he believes that networks are just about the most crucial component in communications policy over the next few years, and why he has very profound doubts about whether the way in which networking is moving right now is a productive area and the right direction in which to be going.

Finally, Roger Woolfe will talk about text processing applications. For most people text processing is the start point of this whole area of office automation. It seems a logical entrée into the area. Roger will be talking about some current applications; current technology; what is the experience of the first users of these types of devices.

At the end of the day my colleague, the Managing Director of Butler Cox & Partners, George Cox, will summarise and highlight the main policy points based on the speeches which have been given and also on the question and answer sessions.

Last night, I found myself short of a closing quote for this talk; and I got one from an unlikely source - a television play. Apparently Le Corbusier, the French architect, once said, "That which gives our dreams daring is that they can be realised."

Dennis Holloway.

B

SOME MAJOR FACTORS IN COMPANY COMMUNICATIONS POLICY

Dennis Holloway, Plessey

HOLLOWAY: Ladies and gentlemen, my intention today, as David has told you, is first to review some of the general case for believing that electronic office automation has a bright future; then to consider some of the particular items that will participate in that growth. I will be quoting a few numerical examples but I have no intention of arguing those in detail; I have put them in rather more for the principle of the thing than for the exact detail. Finally I should like to see what effect that may have; some of the problems that that will raise at the policy level, and perhaps some of the less obvious ones that will be raised. The examples that I have chosen are not intended to be in any way complete. I am assuming that my colleagues will be giving you a great deal more detail of some of those, so I have tended to choose those elements that perhaps get a little less publicity than others.

First, to summarise the case for believing that office automation has a future. Nearly a fifth of our GNP is spent on clerical and administrative functions (see Exhibit B1) and nearly half the working population is said to be employed in these tasks. What is more, these proportions are growing.

Of these expenditures, perhaps 7½% at the most is spent on equipment to aid these functions; equipment from computers, to filing cabinets, to pencils, to pocket calculators, whilst the largest part of this total, some 85%, goes on salaries and wages. A further 7½% at most goes on services, such as Post Office telephone and postal charges, computer bureaux and things of that sort. That is an average; I have no doubt that your individual companies and organisations may well be quite unlike that average. The capital expenditure per head in the office sector is only about a third of that found on average in the manufacturing sector, and productivity in the office sector has certainly not kept pace with productivity on the factory floor.

These factors lead to the conclusion that there is considerable scope for further investment in automation to assist the office workers' tasks. It has been estimated that perhaps 5% of office tasks have so far been impacted by data processing techniques. You will see the investment in computers as part of data processing and part of the total expenditure in the

7½% that we are talking about. This use of data processing has mainly impacted the more routine clerical and administrative functions, and seems to be holding down the growth rate of clerical staff.

Not all attempts to improve the office activity have been wholly successful, and one must remember that some activities such as the modern photocopier undoubtedly did a lot of good, but certainly resulted in a proliferation of paper, added to the postal bill, aggravated the filing space problem and made the updating of records an almost impossible task. Despite all the things that have been done, so far very little has been done to improve the efficiency of the professional and administrative staff who now represent the sector with the highest relative growth (see Exhibit B2). Their work tends automatically to be highly interactive, both among people and between people and a variety of information sources, both human and machine. Their demands fluctuate rapidly and, increasingly in our modern organisations, they tend to be needed in several places at the same time. I believe that this is the area where the concept of an Integrated Electronic Office or Integrated Automated Office promises the greatest future benefits.

Many new "stand alone" items of office equipment have already reached a high level of sophistication and new ones are coming along. Some of these can be justified in their own right, but we believe that the shortcomings of the traditional office over a broad front can only be improved by considerable integration between the various parts into total systems.

The various functions with which we are concerned are those of generating information, modifying information, communicating information, searching for information, sharing information and so on. The paperless office is a long way off, but perhaps not so the one where paper is regarded as a dangerous though necessary adjunct which is best thrown away since the information on it is almost certainly out of date.

An integrated system bringing together a lot of different functions will in reality be constructed of many individual pieces of equipment and subsystems, changing as new technology appears, as different problems have to be solved, and as organisational changes take place. Integration will need to be effective at two distinct levels. First, the user, the professional at his desk, will see one system, providing him with all the services that he needs to do his job reliably and simply. He does not need to understand the technology used to provide this service, nor to be aware of the hardware and software changes taking place. Hopefully he is aware only of a gradually improving service which is a pleasure to use and helps him to do his job.

At another level, the information system manager must be able to control the total system, to measure its performance, to plan, configure and re-configure it to meet changing circumstances and, probably most difficult of all, to assist his senior management in justifying its cost effectiveness.

To provide the cohesion necessary to develop and operate such a system will need a high level of control throughout the whole system. Fortunately, computing hardware is now cheap enough to distribute throughout the system and can provide the necessary flexibility at the three levels; the terminal level, the site level and the corporate level. The necessary national and international common carrier services are also becoming available to complete the picture. Tony Gunton will be addressing this important subject of network architecture later today.

The greatest problems lie, firstly, for the user to identify his true operational requirements; and secondly, to produce operating systems that can translate these requirements into the effective services that are wanted. This is a big task, but fortunately most of the individual elements which are to be integrated are already using programmable elements, such as the computer-controlled PABXs about which Karl Kozarsky will be talking. Communicating word processors nowadays often have microprocessors in them. Communication controllers are coming along. All the elements are there which have the degree of programmability that makes it possible to conceive of joining them together into systems that can develop with time.

Some of the elements which we have to consider linking together into a total system I have put on this chart (see Exhibit B3). I start at the top with the telephone, as it is the most generally used electronic aid; and so as not to forget it I have finished with the mail service, which is obviously still an important part of the communications system of most organisations. These are all candidates for being thought of as part of a total information system, and it is by no means comprehensive. These are mostly ones which are based on electronic communication and processing of information. All of these services can be made available and some have applications already where they can be made justifiable economically in their own right.

The possibilities, however, and in some cases the need for their integration are becoming apparent. What I hope to do this morning is to go over some of the individual items on this list, noticing their present situation and their present status, and at the same time noticing the opportunities and sometimes the need to consider them as part of a total system.

First, telephony is obviously extremely important. After all, the telephone, along with lighting, heating and toilet facilities, is now regarded as an essential requirement of any office structure or any office organisation. It is very seldom that anyone has to argue the case for having telephones. They may argue about how many, but they assume that a telephone system is necessary. It is so common that it is very easy to forget that the International Telephone Network can be described as the biggest and most distributed single machine in the world. The pervasiveness of the telephone and the intermittent use made of much of its plant due to the peaky traffic pattern makes telephone networks, both public and private, a natural target for further exploitation; that is, the use of trunk plant for carrying "off peak" mail and the use of PABX switching and cabling to carry additional non-voice signals. The modern PABX, even down to the smallest size, is now becoming available as a programmable device, so the opportunities to use this as the basis for integration are rapidly developing.

Computer multi-access and data storage and retrieval is an important candidate for integration. Multi-access systems that are in existence have already exploited the ready availability of telephone lines with the aid of modems, and there is a steady growth in the remote terminal population, in the knowledge to apply it and the software to carry it out. However, one of the most significant developments in this area is the development by the British Post Office of the VIEWDATA concept.

VIEWDATA is an idea which is being energetically pursued by the UK Post Office. It offers - or I should say, promises to offer - a simple, low cost and potentially pervasive system for communicating, storing and retrieving information. The Post Office, in co-operation with information purveyors, is emphasising the domestic market for VIEWDATA which, on the back of Teletext (in this country CEEFAX from the BBC and ORACLE from the Independent Authority) should provide the mass market to bring down prices and provide traffic on the telephone system outside peak hours, which is clearly one of the objectives of the Post Office in developing such a system. As the standards being used are compatible with the Teletext services one can expect the cost of the equipment of the chips, coders and so on to come down very rapidly.

As far as I know - and there may be Post Office people here today - the Post Office is certainly not discouraging the use of VIEWDATA for business purposes. It is necessary now to give very serious consideration to the possible use of VIEWDATA in developing business communication systems. One of the major limitations, due to the base being the standard commercial

television set, is the 40-column display format. The disadvantages of such a limited format may very well be offset by the low cost of terminals and their expected wide distribution. One of the decisions that information managers will soon have to address is the extent to which they can or wish to use this limited format - only limited in one way, of course, it has a great variety in other directions - and how to best integrate this type of system into their traditional data processing formats.

The present status of VIEWDATA (see Exhibit B4) is that an experimental system is now operating. If you have not seen a demonstration I can recommend that you do so. A market trial, or more properly a pilot operational service with some 1000 terminals is being planned by the Post Office to commence in about 18 months time, for which purpose adapted television sets will be procured at the terminals. These will be procured from television set manufacturers. A public service is likely to emerge some time around 1980. There is nothing to prevent the parallel development of corporate systems using similar equipment and standards. I understand that the Post Office would wish to encourage this to happen. The sort of areas that are clearly important to think of in this respect are those where one can foresee the need only only to communicate within one's own organisation, but where one needs to communicate with people outside, perhaps customers. One thinks of price lists, catalogue information, dealing with one's customers and suppliers and so on, where one can expect the use of VIEWDATA to be growing for other purposes.

The next on my list is facsimile. I think that facsimile is a term that got itself a bad name and I really should talk about document scanning. The currently available Group 1 analogue facsimile terminals, transmitters and receivers, have a very limited range of applications. They operate at about six minutes per page over a telephone line, so they are too slow for direct use over such lines for many purposes. The analogue signal is subject to degraduation due to noisy lines and is unsuitable for storage and processing.

The introduction of digital facsimile standards, which are now known as Group III facsimile standards and are rapidly obtaining international agreement, are for systems in which the scanned information is digitised and coded by a form of run length coding to remove redundancy. Not only do they operate faster over a telephone line - less than a minute - but because it is digital the signal is capable of being stored, processed, displayed and integrated into a more comprehensive system.

This (Exhibit B5) is one of the sets of figures, again not intended to be in any way a comprehensive justification for facsimile. This comparison was done about a year ago, comparing a guaranteed overnight delivery mail service using mail bags with a facsimile system, both analogue and digital. These figures assume that the transmission would take place overnight using Post Office "midnight line" tariffs. Clearly, if you have a network of your own which is not in use at night, then the cost would be even lower.

The interesting thing is that the figures here are intended to be comparative figures, comparative costs per unit of pages per night. You will see that as the number of pages of documents per night increases, so the cost of an analogue system will grow very rapidly, but the cost of a digital system becomes very competitive with the alternative of physical mail. Those comparisons will vary greatly depending on the particular circumstances. I think it unlikely that one would operate extensively a system of this sort in its own right, and I think I will show later that nevertheless document scanning will still be an important part of the total system.

The increasing cost, the lengthening delays and the reducing reliability of mail services have all conspired to encourage people to think of the direct transmission of mail by electronic means. Communicating word processors are generally billed as the simplest and cheapest form of electronic mail. Indeed, many forecasters suggest that the development of communicating word processors will in fact make facsimile or document scanning unnecessary. However, there are, as so often, problems of integration into the total system. A document that is suitable for typing and originates on the site is suitable for transmission by a communicating word processor mail service if you have the necessary machines installed. Documents originating elsewhere, such as catalogues, copies of incoming pages and so on, are not suitable for retransmission in this manner. A document scanning mail service, on the other hand, can handle both categories, as well as locally originated handwritten documents, graphs and so on.

We carried out a survey recently at one of the Plessey Company sites. This diagram (Exhibit B6) shows what we found about the mail going out from that site. Out of a total mail, we found that about half was suitable for electronic transmission. All of it therefore could conceivably be put on a facsimile type transmission system. Of that 47% of the total mail, again about half, 24% of the total, was generated in-house, and therefore might have been generated on word processors if we had them available on the site. That is the situation which exists at the moment. Clearly if word processing becomes more prevalent both in one's own organisation and in the organisa-

tions with which one communicates, the smaller of those bubbles (the word processing mail area - see Exhibit B6 again) will increase. But it will be very many years before the intermediate part, the one that requires document scanning, will disappear altogether.

I believe that an effective electronic mail system must include not only word processing or keyboard transmission, i.e. documents which are generated directly in electronic form, but document scanning must be included as well (Exhibit B7). This does not mean that you have a facsimile machine down the end of the corridor; these must be integrated together. One will need document scanners to deal with the items that have to be dealt with in that way. One must see that electronically one can pin a diagram to the text and keep it together throughout the system. At the receiving end one needs printers, but there is no need to separate the printer or the display that you use for graphic and written material, alpha numeric material. A number of printers are becoming available and facsimile receivers are available, which can in fact print equally well textual information, alpha numeric information or graphic information.

Another area which received a boost when the oil crisis came along was the use of remote conference facilities to save travel costs and time. The Post Office Long Range Intelligence Division and the Communication Studies Group of University College carried out some valuable work which showed that audio conference facilities can readily be justified in savings of travel time and cost and that in a high proportion of cases a remote conference is a satisfactory alternative to face to face meetings.

Suitably engineered systems, such as our own "Remote Conference Table", are now available, and savings can be demonstrated. Again, I have put some numbers down here (see Exhibit B8). I can tell you something of the basis on which they are made, but I certainly would not want to argue them in detail. However, they show that if you allow the cost of time and the cost of travel, then as distance increases and the number of people involved in meeting increases, the justification for using such conference facilities rapidly improves.

I believe that this particular study was based on the assumption that the conference terminals would be used for 50 meetings a year. It assumes that private wires are used to connect the conference tables together. The high cost of television transmission capacity has made it very much harder, though certainly not impossible, to justify television conferencing. Alex Reed's unit at the Post Office has demonstrated that meetings which are very short and which

require long distance travel even now could justify the cost of television; but these would obviously be for the most senior people in an organisation.

However we must remember that technology is changing as we go along and although the cost of television transmission is high at the moment optical fibre transmission is reaching a point where we have to consider seriously what impact it will have in the future. I have brought along a small sample of optical fibre. This is an optical fibre which will carry a very broadband signal with a very low attenuation, perhaps about 2 decibels per kilometre, which is very much lower than a co-axial cable. Systems using this type of cable are developing much more rapidly than many of the technologists even expected them to. We used to worry tremendously about how you would join these things together. We find that it is remarkably easy.

Transmission systems using this sort of cable will begin to be installed on the public network in some capacity round about 1980, and will be developing from then on. So we cannot ignore what that will do to the development of communications.

The last item on my list was Post. We must not forget that the Post Room is at present the store and forward node of the major message communication network, linking a local distribution system to a public switched network often augmented by private service. No electronic system can completely replace the physical mail system, but most of the electronic methods that we have been discussing will impact on the size and purpose of the postal network.

There are various ways in which this may happen. One solution to electronic mail may merely replace the intersite part of that network and still rely on the local messenger service. Another solution may be to go straight to the office desk, with terminals where people communicate directly. It is quite clear that many savings are available, but the task of transferring the load of information from the postal service from all or part of the present physical handling service presents problems which are certainly not all technical.

It is easy to see that there are many issues such as the development of standards and the understanding of operational needs which must be resolved. But as we review these things and look at their individual economics there is a great danger of concentrating on immediate cash justification, worrying about the immediate problem of getting the right number of telephones, getting the Post Office to install the right number of lines and modems, and losing sight of the real potential benefits of future systems.

The major impact of the next generation systems will be to help that growing army of professional and administrative workers in doing their work which is vital to those concerns in which they operate. We can measure the cost of equipment, but can we measure the quality of a professional decision? Can we calculate the value of improving that decision or the cost of getting it wrong? There will necessarily be judgments made, and these will be judgments which must be made by senior managers who understand what really matters to their enterprises.

Let me take an example of a manufacturing industry, the sort of which I have some personal experience, which perhaps like quite a lot of British industry finds itself short of exportable products. In order to put the matter right, this is the cycle of events (see Exhibit B9) which has to precede the launching of a successful product. One has to do market and technical research; look into the business options; prepare business and product plans; develop both the product and the market; go in for pilot production and marketing; and finally, put the thing into production, sell it and service it, and hopefully eventually make a profit out of it.

Each of those steps is quite a long one, and the total time involved is a very long cycle. Each step requires very high calibre staff and involves many hundreds of mini-decisions based on information from many sources and interaction between different functions. If we can improve the information aids available to the key staff in these functions, we could conceivably bring forward the time at which a new product is available by perhaps a year. How many of us can estimate the value of such a service to our enterprise?

Exhibit B10 shows the benefits of information automation in what I suggest to be an acceptable order. Probably the most generally available factor, the factor that is most common to all of our enterprises whether we are in manufacturing, banking or shopkeeping, is saving time. Saving time somewhere in the system will save - I was going to say money; it may well save the enterprise let alone the money. I think that is probably the most important thing, and it may be that one of the only ways we will have of justifying the systems is to find some way of totting up the total saving in time and then putting a value on that time at the end.

Other factors will be important; the ability to change geographical and organisational arrangements are clearly changed by the type of system about which we are talking. The quality of output and the motivation of staff is also important; and well down the list are the direct cash savings that you will make. I believe that if we can find a way of measuring the true benefits, so much the better; but just

because we may not be able to measure those benefits does not mean that they are not there and that they are not real.

It is also clear that managements may need to give attention to their organisation since the impact of the possibilities that we have been describing can cross many organisational boundaries. The communication manager in the company in the past perhaps has been concerned mainly with finding the cheapest way of providing 500 or 1000 telephones; arguing with the Post Office about the provision of plant and services; but he has had to do very little justification of the use to which those services are put. Office management, personnel management, telecommunications management and data processing management are all involved in implementing future systems.

Finally, the message is that only the very top management of a company can really identify those skills and resources which are vital to the operation of the organisation, and only top management can issue the instructions to see that those skills are supported.

BUTLER: We have time for one or two questions before the break.

THURTELL (General Motors): I am interested in your view that fibre optics are important for communications. Surely it's out of date - analogous to the cavalry in the 1914-18 War? Our transmission surely must not rely on pieces of wire linking our data centres.

HOLLOWAY: Thank you. I presume that you are suggesting that radio is the method that we must use to interconnect our services. I have not mentioned it but obviously the development of satellite communication radio is very important. But within a site for instance, perhaps a large site, one would certainly not expect to transmit information by radio. There will always be the problem of relatively short distance communication. I am sure you know that in the development of telecommunications networks only a very small part of the cost even of a telephone network is in the trunk transmission capacity; by far the greatest part of the cost is in the local distribution area. Although it is possible to foresee in the fairly near future the use of some main links on satellites, the use of links between perhaps company headquarters or main sites, there is very unlikely to be a replacement of all communications to small sites, distributed sites, without using something physically in the ground.

They both have a part to play and I am certainly not suggesting that optical fibres will come quickly, but I think it will make the possibility of an extensive broadbank network possible in the 1980's.

COULOURIS (Computer Systems Laboratory Queen Mary College): I am a researcher in the area of computers and office automation, I have two points. The first is this: I think we may be in danger of letting the user off the hook too easily. We tend to say that the user should not be involved with technical factors. Well I agree that he should be isolated from arbitrary technical considerations, but some technical concepts are most important. An example might be organisation - how best to organise for maximum efficiency. We need to educate users over issues like this. And in general I would say that information needs to be thought of as a process now, not something which is static.

My second point is to do with VIEWDATA. I was interested to hear your comments and to learn of your evident enthusiasm for the potential of VIEWDATA. However, my view is that this will be extremely limited, and that it will tend to constrain applications. I think that it won't have a significant future, compared with paper.

Perhaps you could comment on these two points.

HOLLOWAY: Yes, thank you. On the subject of information concepts I certainly did not intend to imply that the users, these highly skilled professionals in particular that I am talking about, will not be concerned with understanding the concept of information. After all, my job is mainly one of trying to understand that and organise my thoughts and the way that I collect and use information. What I do not wish to do with a communications system is to know, when I speak into a machine, whether it goes into a computer or a piece of hard-wired something; whether it goes over an optical fibre or a satellite link. I need to know as little as possible about the technology that manipulates that information. It may be going through one piece of equipment one day and quite a different piece of equipment another day, but conceptually the handling of the information is something that will be part of any professional's job. I quite agree with you about that.

On the subject of VIEWDATA, that's a very good question. All that I can do is to say that I believe that VIEWDATA or something very much on those principles will become widespread, simply because of the development; television sets are available and the costs will be low. It is a very good question as

to how far one will be able to use that limited format. If one looks at some of the things that are being put on VIEWDATA already, one is conscious of the fact that in this age when we tend to use shorthand a great deal, there are a great many things that can in fact very well be presented. People are very used to looking at football scores on their television screens. People are becoming conditioned; and we must remember that the A4 page is of historic interest to us.

COULOURIS: One technical detail. It is not true that we have 40 characters width on the TV screen!

HOLLOWAY: No, I don't want to argue the question of what a VIEWDATA will be. I don't know all the arguments. What I am worried about is that I saw in a recent New Scientist article that the Americans are talking about a different standard, with 36 characters. That worries me very much more than whether it should be 40 or something. Much more important is that there should be a standard which will be universally employed. I would welcome anybody's views on what you can and cannot do with a 40-column format.

BUTLER: May I, on your behalf, thank Dennis Holloway for a very illuminating presentation.

Exhibit B1

TOTAL ≙ 18% of G.N.P.

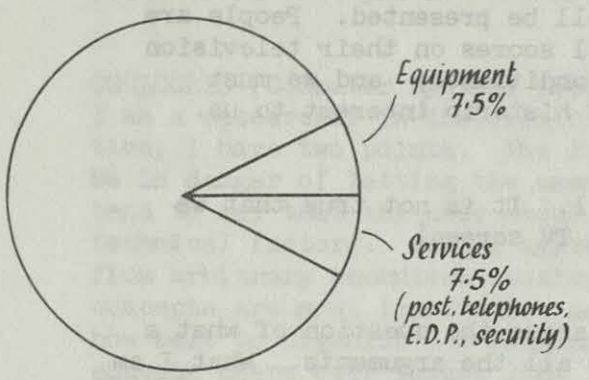

EXPENDITURE ON ADMINISTRATIVE & CLERICAL FUNCTIONS

Exhibit B2

BREAKDOWN OF U.K. OFFICE EQUIPMENT SPEND

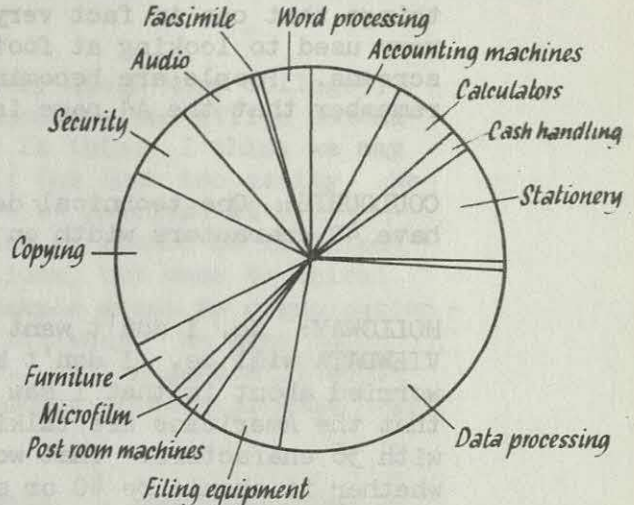

Exhibit B3

CANDIDATES FOR INTEGRATION

Telephone
Telex
Computer access
Data storage and retrieval
Document scanning
Document preparation
Word processing
Electronic mail
Data capture
Conferencing
Message transfer
Distributed processors
Mail

Exhibit B4

Viewdata

THE PERVASIVE PROBABILITY

Experimental system operating
Active participation by information vendors
1000 terminal Pilot Service by 1978
Compatible with teletext

How will it impact your business?

Exhibit B5

OVERNIGHT MAIL COST COMPARISONS

cost unit / pages / night

PAGES PER NIGHT PER STATION	COMPARATIVE COSTS		
	FACSIMILE		MAIL BAGS *
	ANALOGUE	DIGITAL	
30	7	10	8
60	12	11	8
120	24	13	8
240	50	17	15
480		27	30
960		54	60

* Guaranteed overnight delivery contract

Exhibit B6

EXTERNAL MAIL ANALYSIS

AT ONE PLESSEY SITE

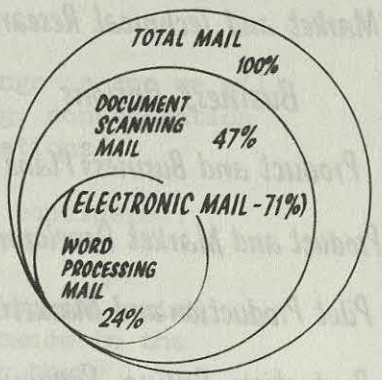

Staff - 1600
Telephones - 600
Total mail - 2,900 pages/day average

Exhibit B7

ELECTRONIC MAIL

needs

Integration of

Keyboard Word Processing

and

Document Scanning

Exhibit B8

REMOTE CONFERENCE TABLES ANNUAL SAVINGS

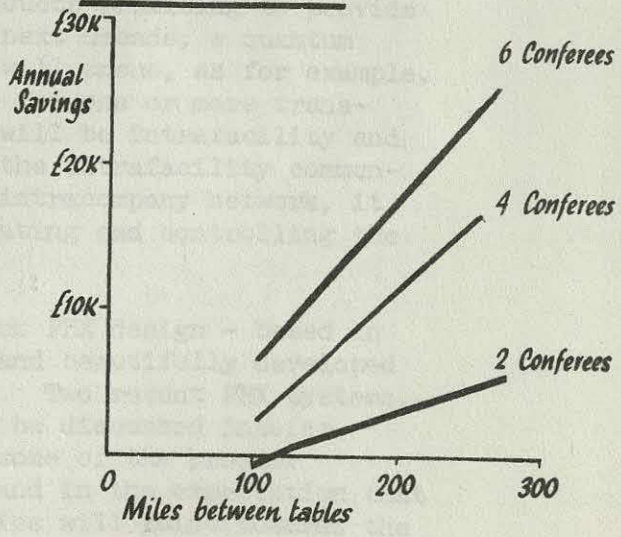

Exhibit B9

KEY CYCLE FOR A MANUFACTURER

Market and Technical Research

Business Options

Product and Business Plans

Product and Market Development

Pilot Production and Marketing

Production, Selling, Servicing

TIME

Exhibit B10

WHAT CAN INFORMATION AUTOMATION DO?

GAIN TIME

ALLOW GEOGRAPHICAL AND ORGANISATIONAL FLEXIBILITY

IMPROVE QUALITY OF OUTPUT AND MOTIVATION OF KEY STAFF

make direct cash savings

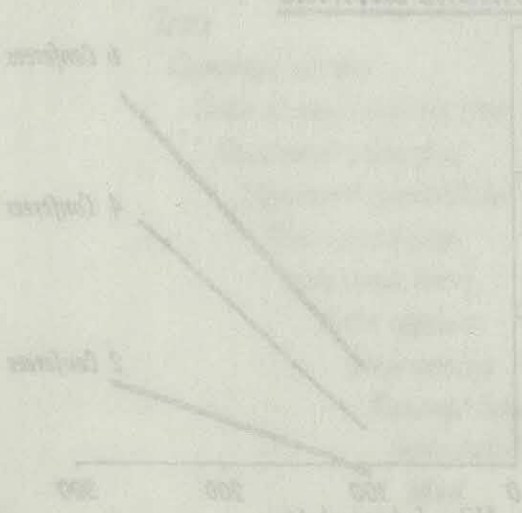

C

THE CENTRAL ROLE OF THE COMPUTERISED EXCHANGE

Karl Kozarsky

KOZARSKY: The role of the centralised exchange is one we will assume will be the central point of switching, concentration and control of voice, data and text communications.

The philosophy of this viewpoint will not be considered: that is another paper, or another conference.

The key to this central role is the stored program, which enabled the first enhancement of any significance in the ability of organisations to gain control over their communications.

Although computers added to conventional PBXs have been around for sometime, experience with integrated designs is only 2 or 3 years old in the United States, considerably lagging European experience primarily due to IBM's explorations of the voice switching field. However, the diversity of available products and larger user base in the United States tend to compensate for the relatively brief United States history. This morning I will draw upon that experience (see Exhibit C1) to describe those characteristics of computerised exchanges which users value - that they consider worth having paid for. The voice services are the core of these values - with only tentative efforts to date to embrace data traffic.

However, with office automation products promising to provide vitality to the marketplace in the next decade, a quantum increase in communications traffic will ensue, as for example, every keystroke becomes a candidate for one or more transmissions. Initially, such traffic will be intrafacility and intracompany; as the PBX provides the intrafacility communications and is often a node on an intracompany network, it is a leading contender for accommodating and controlling the added traffic.

This prospect must have its impact on PBX design - based in the past on the splendid isolation and beautifully developed statistics of voice telephone usage. Two recent PBX systems, the WESCOM 580 and Danray ADX, will be discussed insofar as they reflect efforts to address some of the broader problems of integrated switching - and in the expectation that a description of their characteristics will point towards the next stage of private exchange development.

There is general agreement on the virtues of the first 2 categories listed here (see Exhibit C2): usage control, effected by limiting the use of a telephone only to those services deemed appropriate for that user's conduct of business, with fine degrees of discrimination among classes of restriction made possible by stored program control.

A more positive aspect of usage control than abuse prevention is automatic route selection - which functions by displacing the choice of circuit group from the user to the stored program. It avoids the inefficient and vain procedures of training users to use tie-line groups, to be aware of time-of-day variation in circuit choice, etc.; nor of the codes needed to select a trunk group. Substantial savings have been reported from computerised route selection, up to about 30% reductions in cost per call.

A useful elaboration of this feature is trunk queuing, which at some inconvenience to the user, by delaying his call until a circuit in the trunk group of choice becomes free, squeezes more traffic over expensive, fixed cost trunk groups.

Capturing, then sorting and organising calling traffic and trunk usage data enables a number of useful results, including cost allocation by usage and prompt sizing of trunk groups according to needs. This is in sharp contrast with prior methods used by telephone administrations, which with great effort too often produced obsolete inaccurate data. Where account codes may be keyed by a user, the call and its charge is easily associated - taking a step into source data collection with information not already derivable within the PBX.

The third area of value, and this is not yet a universal opinion, is the electronic key telephone. A definition may not be presumptuous since one of the charms of this unique country is the virtual absence of the keyset. A conventional key telephone is an instrument used to allow one telephone to pick up a number of lines and also to allow one line to be answered, and possibly conferenced with a number of telephones. Given that definition, it is hard to see how one could do without it, and in many businesses virtually all the sets are keysets. But it has drawbacks - it is an awkward and monstrous piece of hardware.

It has been electromechanical, with locking pushbuttons and status lamps for each line, with separate electromechanical control boxes from which emanate cables to each keyset - with a minimum of 25 pairs of wire in each cable. Thus an overview of the wiring in a facility shows wire pairs spreading outwards from the PBX plus complex local wiring clusters around each key control unit.

An electronic key telephone is a generalisation of this concept. It has arrays of non-locking buttons and associated lights; each button is programmed to cause selection of a line - or, at the user's option - initiation of a feature, with the associated illumination providing status information. The signalling between the electronic keyset and its control is digital - coded instructions flowing from the telephone, and lamp control information back from the control. Only 2, 3, or 4 pairs of wires are needed due to the implicit multiplexing. The electronic unit, although a more expensive instrument than its electromechanical predecessor, compensates with its lower cost of cable and labour during installation, and then forges ahead in lower costs of maintenance, moves and changes.

An important point relating to PBXs is that when digital signalling is used between the PBX and its instruments, the only difference between a basic telephone and a keyset is the greater variety of codes passing along the signalling wires. Thus in principle, wiring to all sets can be uniform; once wired, standard and key telephones may be substituted or moved without struggling with copper.

Another reason for electronic keysets is indicated here (see Exhibit C3), and is due to the many user-initiated features now available: to select a feature with a conventional telephone system requires a switch flash to get dial tone, then depressing two or more buttons - if one remembers or notes the feature code number.

With an electronic keyphone a button can be dedicated to a feature and pressing that button suffices; furthermore, there is the reassurance of the associated light which confirms acceptance of the command. Some of the generally available features are shown on the Exhibit - the A column lists some features which once initiated are concluded more or less immediately - picking up a call, transferring a call, using an abbreviated dialling code, and so on. Under the B column are features in which a service is not immediately concluded after initiation. They may last for a few minutes, say for the Hold operation, or somewhat longer for Call Back, but may endure all day for Call Forwarding or Do Not Disturb. The lights on the electronic key telephone serve usefully to provide the user with a reminder of an existing longer duration operation - for example, to remember a line on hold or to cancel Call Forward or Do Not Disturb command. As features proliferate and require control to operate without adding confusion to the system, the more pressing is the need for an instrument rendering clear assistance to the user.

Of course the classic use of keys on a telephone is the ability to manipulate multiple lines, more than two being

beyond the capability of the simple instrument. Finally, the flexibility of the electronic keyphone in meeting the varied needs of users cannot even be approached in the standard telephone. As an example a single direct inward dialling number may have appearances on say 20 telephones; that number will not be busy unless 21 calls are in progress to that number - or a principal and a secretary who each have an appearance of both lines on their sets can be programmed so that four calls to the principal's number can appear at and be handled concurrently by the two telephones.

Other features of consequence include these (see Exhibit C4) which have been available for some time from Danray and IBM but have not been generally adopted into competitor's later model developments. A directory data base linked to an operator CRT displaying sorted personnel and classified listings is a particularly valuable service to organisations receiving many public inquiries. This is sometimes provided by a computer system independently of the PBX - however, Danray's CBX uses the standby minicomputer in its private exchange to provide this service. In addition to this on-line software, Danray also supplies a batch program which processes collected call data to produce trunk utilisation and cost allocation reports. There are good arguments for providing these capabilities in systems with stand-by minicomputers, but they are less persuasive in systems controlled by distributed micro-processors.

There are also the 3750 class of features with specialised data collection equipment and limited data validation software plus a communications link or media link to an EDP system. If this type of application endures, and it probably will, it is unlikely to do so where the data aspects are secondary to the voice orientation, so that data terminals require additional tone detectors and use MF tone signalling rates. Moving the voice response hardware to the switch has been initiated by AT&T in its Transaction Network Service. This is a switched service between transaction terminals and one or more ADP centres.

Facility monitoring and control appear to be capabilities which are most economically incorporated in computerised exchanges (see Exhibit C5). These are features which, promising as they were at first appearance, must receive a verdict of not proven today. But with several industries beginning their onslaught on the vast frontiers of the office automation market, an early target is the physical transportation of information that is usually carried to recipients on paper - substituting for it electronic communication with an optional transformation to paper at the receiving end. Since the appearance of items such as acoustic couplers or picturephones, the telephone has been moving from its stage as a

voice transceiver, and increasingly suppliers are viewing it as an element in an instrument complex providing a growing variety of communications options for the needs of organisation man.

The replacement and enhancement of office functions with electronic communications is a clearly discernible trend. One consequence of these changes is that significantly increasing amounts of traffic must be handled. Beyond the quantitative load there is the need for various controls including, but not limited to the same functions that have proved of value in computerised switches - usage control and traffic data collection and analysis. But there are major differences in the needs of text transmissions in terms of accuracy, priorities, time delays; that is for control needs more closely related to message switching functions than to voice communications. How those control needs are to be met is a high priority project at every office systems supplier.

Additional factors to be considered for communicating office terminals include the following (see Exhibit C6): wiring - does another in-plant wiring job need to be done in addition to the voice wiring, as is often done with data system wiring? Are moves and changes for these units to benefit from the data table entry changes in computerised switches? Will the additional traffic, if switched through PBXs, clog it to the detriment of voice traffic? Is the bandwidth capability adequate?...

At this point let us assume that voice, data and text communications will be integrated, in some sense, within computerised switches. Then we shall attempt to glimpse the outlines of PBXs designed to do so through two recent computerised exchange offerings, one a new design and the other a new configuration of an existing product. The first of these is the Wescom 580 which addresses in particular the traffic handling capabilities of a PBX. The other is the Danray Auxiliary Data Exchange, which integrates data systems with the voice switch, exploits its voice system to solve wiring problems, and takes a different attack on absorbing the data traffic. These products from non-traditional switching suppliers may or may not make their way to The European Market, but it is likely that any successful innovations they contain will.

The Wescom 580 is a time division switch converting the analogue signals to pulse code modulation in the North American T-carrier format (see Exhibit C7). At a maximum size of 2400 lines, 576 trunks, and a few more allocations for items such as tones, some 3072 time slots are provided; that is, one for everybody. First, each channel, consisting of a 64 KHz digitally encoded voice data stream, is multi-

plexed to 32 96-channel buses, and then shuffled rapidly among high speed solid state memory banks to perform the switching function by moving data about quickly enough to meet their time-slot appointments. This results in a non-blocking "all-time" switching capability, i.e., if everyone on the switch were to improbably pair off, all parties could be communicating at the same time.

More to the point, if a number of users put up data communications connections that last all day, there is no degradation of service to other parties using the switch.

This high traffic capability is of particular interest because of the price with which it is achieved; except in small switch sizes there appears to be only a small price premium over comparable switches of lower capacity, that is, with blocking networks. In traditional switches, the cost of non-blocking networks is prohibitive.

Part of the reason for Wescom's solution is that blocking networks require more housekeeping to control - as in tracking which paths or time slots are occupied and which are available. Externally, traffic balancing is an art that can be dispensed with. Further, using fast memory as the heart of the switch is hitching a ride on all the development money pouring into a market whose limits of elasticity have yet to be sighted. At any rate, the point of emphasis for our purposes is that a well featured switch in medium to large sizes offers a non-blocking capability at a reasonable price.

To illustrate the capability of their switch, Wescom shows the capacity in terms of CCS per Station in this table (see Exhibit C8). 36 CCS equals 1 Erlang, a more familiar unit, I believe, east of Boston.

The modest decline below 7% line-to-line - or internal traffic - reflects only the trunking capacity limit to 576 trunks, and occupancy per line declines with decreasing proportions of internal traffic. This is irrelevant to non-blocking and basically reflects a revocable packaging limitation on the number of trunks. But please recall, when looking at this chart, that most PBXs offer capacities somewhere in the range of 4 to 8 CCS per station. As for the status of the 580: preproduction units have just been completed and trial installations are about to begin.

My second example is in the Danray product line and is a direct approach to integrating voice and data traffic in an exchange (see Exhibit C9). This system is configured with an option on the Danray computerised switch called the Auxiliary Data Exchange. With this option the existing space division voice switching network is not impacted. A separate data switching

matrix, a time division unit selected as most appropriate for data, is added to handle such traffic. However, common areas remain, namely the computer control and common wiring.

A closer look at the wiring (see Exhibit C10) shows that the Danray switch employs three wire pairs per station: a receive pair which also shares digital transmission of the push button depressions; a transmit pair which shares digital lamp control information to the single LED on each telephone; and a third pair which carries power to the station. With the auxiliary data exchange, Terminal Interface Adaptors (TIA) can be interposed between the telephone and the wall telephone jack. This also permits a data terminal to be connected to the system with data transmitted over the power pair of wires. Any speed between 75 bps and 9600 bps is switch-selectable on the TIA, which is a low cost box for local connections up to about 4000 foot wire lengths. A major advantage is that existing wiring for any telephone is usable for including data terminals in the system.

The system diagram (Exhibit C11) shows TIAs joining telephones and data terminals along the standard wiring with the data signals diverted through the data switch module. Terminals may connect to one another or to an in-house computer shown here with a number of ports organised in a hunting group. To access remote computers, a shared modem pool may be used to considerably reduce the total quantity of modems that would otherwise be needed. Note the off-premises connection sequence is TIA to TIA to shared modem to remote modem. The data communications procedure is to use the telephone to establish the route for the data terminal connection, then to initiate the data transmission; the telephone may now be used concurrently for voice transmission.

In summary, the system's benefits include the ability to use a data terminal wherever there is a telephone without additional wiring, adding data without restricting voice usage or degrading traffic, and modem sharing for remote communications. It is a clever systems adaptation of an analogue switch.

The first user of this Auxiliary Data Exchange system (see Exhibit C12) is a Tektronix CRT terminal manufacturing plant in Wilsonville, Oregon - which is categorised as an experimental installation. There are 630 lines in the system with about 150 TIA's; these are, however, used with some 225 telephones. In many offices where user procedures permit, TIA's can be shared by moving the small unit from one wall jack to another.

A local large scale time sharing system has 30 ports on the data switch and is frequently used to down load 16K memory blocks into terminal microprocessors at 9600 bps. Another

minicomputer time sharing system with nine ports is also accessed through the switch.

Two groups of shared modems for remote computer connections have been configured, eight 300 bps units and four at 1200 bps - or a dozen for 150 terminals to indicate the potential savings by using the switch to concentrate modems. The environment at this facility is a factor that made this voice plus data switch appealing; 80% of the facility's population moves each year with a proportionate shifting in the need for locating data terminals. In the past the wire hauling for a separate data net has meant extensive delays and expense. With this Danray system any location with a telephone can be used for a data terminal connection.

The present status of the system is that the system is installed, the voice switch is working well and the Auxiliary Data Exchange is phasing into full operation. As of last week, some 90 users with TIAs were comfortably accessing the two time sharing systems through the data exchange. Modem sharing was functioning from remote points into the system but it is at the point where useful experience is beginning to be accumulated.

In summary these two systems - Danray's and Wescom's - show aspects of integrated switching beyond their contemporaries. They promise to be seminal designs for that portion of the future office market that will proceed with the switching hub and not the EDP center as the framework of information flow.

This final chart (Exhibit C13) is a composite switch proceeding from this discussion, representing various more or less proven elements. Both digital and analogue terminals are necessarily accommodated with emphasis on the trend towards digitising the world. Key telephones along with other types are shown with an eventual blurring of distinctive categories as the generalised approach allows more functions to be integrated in the user instrument. Signalling therefore is shown as digital - the usual tone detectors, dial pulse and tone generators are reserved for use with the less controllable outside world. Wiring is shared so that multiple terminals are available at each jack. Switching is non-blocking PCM time division - which is in the path of the technological main stream and is capable of interfacing increasing proportions of common carrier transmission systems without extensive conversion hardware. This part of the system is fairly firm; no one offers this yet, but it is likely to appear in the near future.

Finally, the optional processing and data base computer control is only indicated here as the point of executing many of the extensive control functions required of text and data communications.

The long list of problems associated with an integrated communications exchange are gradually being attacked as these Wescom and Danray examples show. Should their performance in the real world be reasonable, we may look forward to having the major suppliers of switches, computer and office systems offering integrated products for your offices in the near future.

Your responsibilities for information flow in your organisations will broaden along with the systems options available to you.

DAVIES (NPL): I have two questions. Firstly, although you have described systems which are physically integrated, nevertheless the services which are provided are separate. What applications are there for genuinely integrated services? And secondly, how will separate data and other services be incorporated in the future?

KOZARSKY: In answer to your first question, in the Danray design physical separation is an accident of history - it is not something which is fundamental. Secondly, information packets need to be controlled. This entails either expansion of the CBX, or the creation of an auxiliary function on an in-house computer.

DAVIES (NPL): Well, returning to my first question again, I still require some clarification. At the NPL, we have a separate data network. What I would like to know is the reason why I might want to make a functional change - engineering reasons excluded of course. Where do the two services of data and voice come together?

KOZARSKY: The answer is primarily in engineering - in engineering the equipment - and not in the functions which the equipment is designed to cater for.

Exhibit C1

- Evaluation of experience with Computerised Exchanges
- Implications of increasing telecommunications experience
- Two product examples addressing new requirements

Exhibit C2

VALUED CHARACTERISTICS

USAGE CONTROL

Restriction
Route Selection

TELEPHONE CALL DATA COLLECTION

Accounting Information
Traffic Analysis

ELECTRONIC KEY TELEPHONES

Exhibit C3

ELECTRONIC KEY TELEPHONES

USER FEATURES

A	B
Call pick-up	Call forwarding
Call transfer	Hold
Abbreviated dialling	Call back
Conferencing	Do not disturb
	Message waiting

Multiple Lines
Flexibility

Exhibit C4

OTHER FEATURES

- DIRECTORY
  - Operator terminal look-up
  - Printing

US

- TRAFFIC ANALYSIS

- DATA COLLECTION
- VOICE RESPONSE

Integration with EDP System

EUR

- FACILITY MONITORING

Exhibit C5

Paper creation

physical transportation and access } Communicating Terminals

TELEPHONES ⟶ SOPHISTICATION

Increasing intra- and inter-facility traffic

Exhibit C6

FACTORS

Wiring

Moves and changes

Traffic blocking

Bandwidth

Control

Exhibit C7

WESCOM 580

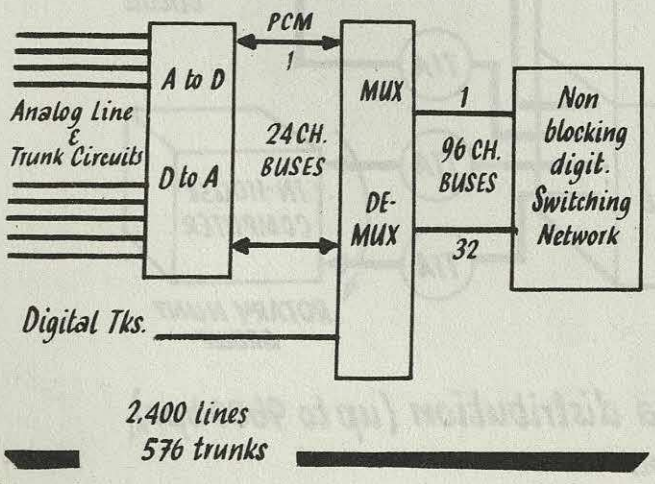

Exhibit C8

WESCOM 580

Number of stations	CCS PER STATION % line to line traffic 78%	70%	50%	30%
2400	36	26·5	15·9	11·4
2000	36	31·8	19.1	13·7
1600	36	36	23·9	17·1

·01 trunk grade of service

Exhibit C9

DANRAY

Separate switching matrix for data

Common computer control

Voice - Space Division
Data - Time Division

Common wiring

Exhibit C10

DANRAY 3 - PAIR WIRING

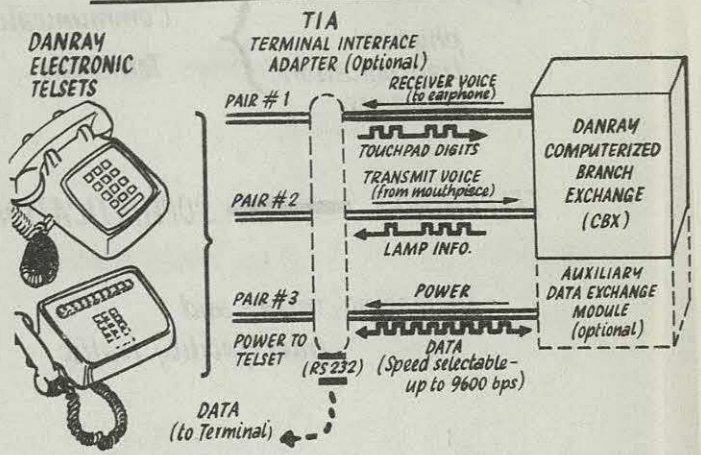

Exhibit C11

AUXILIARY DATA EXCHANGE (ADX)

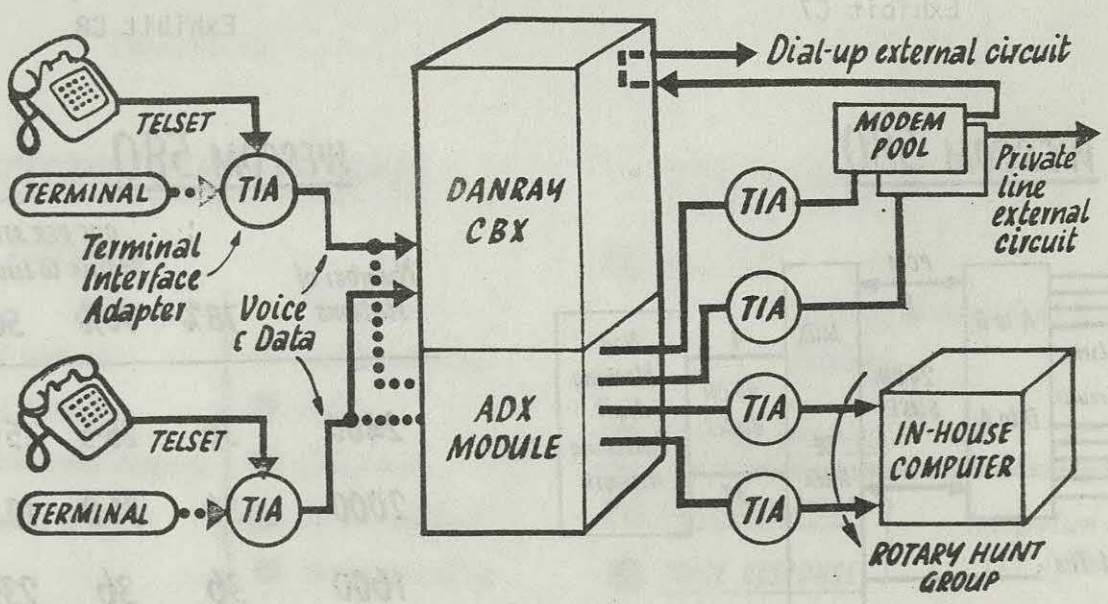

1. *Internal switched data distribution (up to 9600 bps)*
   - *without modems*
   - *without additional wiring*
   - *without disabling Telsets (Simultaneous Voice)*
2. *Terminals access a 'pool' of modems for external data communications.*

Exhibit C12

Exhibit C13

INITIAL INSTALLATION

630 lines

150 data ports –
225 with shared TIAs

Large local T/S system with 30 TIAs
down loads microprocessors at 9600 bps

Remote computers
Four 1200 bps
Eight 300 bps
} shared modems

80% moves/yr.

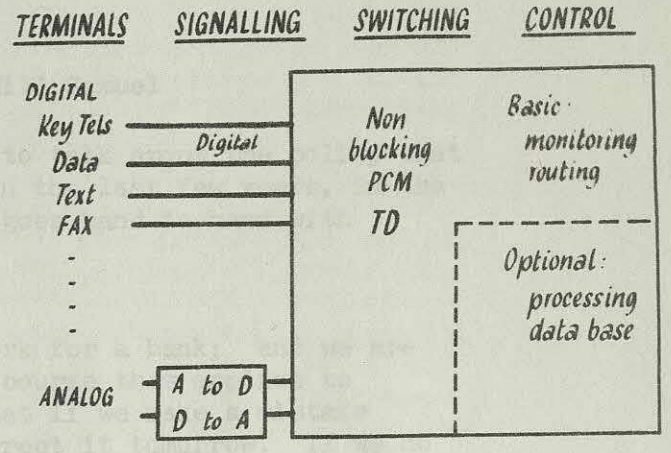

D

A SYSTEMS STRATEGY WHICH EXPLOITS COST TRENDS

Hamish Donaldson, Hill Samuel

BUTLER: Hamish Donaldson is going to talk about the policy that he has been following at the bank in the last few years, in the search for a system strategy which goes hand in hand with current cost trends.

DONALDSON: As David told you, I work for a bank; and we are supposed not to make mistakes. Of course this applies to everyone else; but it does mean that if we make a mistake today, it is not good enough to correct it tomorrow. If we do correct it tomorrow, it must be put right as from today; because if you don't, the interest calculations are wrong and you pay out wrong sums of money. And on £1 million, a day's interest does matter.

So the emphasis on data processing in our organisation has, from the early days, been on getting it right first time. You can imagine the shambles there was when we put the first batch systems in - unthinkable. All sorts of improvements were tried very quickly: data preparation with the users and so on. In fact we got a fairly good system working based on TC500s, which produced the initial contract notes and punched paper tapes to go into the mainframe. These TC500s were allocated to the user departments. So for quite a long time this idea of pushing out the data entry side at least of computing to the user departments has been a thing we have had to do. I don't say there is anything unique in this; although I think that possibly in the bank it is slightly more obvious than it might be in another business situation, but I think we will all see the need sooner or later. So we have been living with this thing, which is now called distributed processing, for some time.

I suppose that you get a gut feel about things and you only rationalise later. You go along the way you think seems to be a sensible route, and after a bit you've got to step back and say, "Why are we going down this route? What are the reasons for it?" We seem to have a mixture of mainframes and minis in our operation all with on-line capability as well as batch. We have a policy of pushing out the operation of computers to the users and getting them more involved.

But six months ago, I had the opportunity to do some research to try to establish the ground rules of what we are doing and try to see if there was a theoretical basis for what it is all about. I should like to go through the process with you. You will forgive me if I am telling you things which are either obvious or which you know already, but I do need to build up the case by developing it progressively. That will be the first part of my talk. The second part, when we have talked about cost components and cost trends, will be about the management implications: what were the management implications that we found and what were we likely to find in any future developments? So two parts to the talk: the cost components and cost trends; and, secondly, the management implications.

We start by defining distributed processing. The definition that we came up with (see Exhibit D1) says nothing about minicomputers and nothing about mainframe computers; and this discussion is not trying to sell you the virtues of one or the other - I think it is irrelevant to the discussion. What we see as being the important thing is pushing the computer power to where the people with business problems are. This is the theme of what we have called distributed processing.

To test the way that we might achieve this, we set up some seven computer alternatives (see Exhibit D2 and D3). First, a central batch processor - the traditional way: send your documents in, punch them up in some way and process them centrally. Second, remote job entry to the central computer. The third is multiple access to a central computer, maybe a time sharing computer, or an in-house mainframe of some sort. Next, the personal computer. This is the one where we have maybe a VDU, a couple of million bytes of storage, 64K bytes of memory, and a bit of printing ability - a powerful personal computer, a powerful mainframe of a year or two ago.

Number five is the local stand-alone computer, typified by a minicomputer PDP-11 with a few on-line terminals. Next is a local minicomputer with dial-up to the centre: the idea here is that quite a lot of the job may be done locally, but from time to time you need to consolidate a certain amount of information on a daily basis say in the evening, or maybe on a weekly or monthly basis. The final one is networks of computers.

We are unlikely to see a blanket solution to any organisation, at least that is our hypothesis; different situations will require different approaches. But what are the cost trends that are likely to influence a decision in favour of one or other of these approaches? The first thing to do was to work out the main cost components for our mainframe and our minicomputer.

The mainframe happens to be a Honeywell 66. We run it for three shifts and here (see Exhibit D4) are the main cost

components: hardware depreciation first, then operations staff. Twenty three may seem a large number, but we know what the numbers are. Our mainframe is installed in its own building, with its own security arrangements, with its own security guards at night, with no data preparation staff at all, with no analysts and programmers at all because they are in other locations, so the whole cost of the thing is very easy to determine precisely; hardware maintenance, consumables, data transmission and so on.

There are similar sorts of costs for our mini. It comes as no surprise to discover that the hardware costs a relatively larger proportion with the mini and that it requires fewer operations. The way we find that it works is that a chap has to come in early in the morning to switch it on and someone else has to stay late at night to switch off. This means that two people take it in turns to come in a bit earlier and leave a bit later. We try not to work our minicomputers on shifts. We go to quite a lot of trouble to say that they are run by our user departments (who do not understand how to run shifts), so if there is any night work to be done we will do that through management services.

If these are the main cost components, let us have a look at how they have moved over the last ten years and how they are likely to move over the next five. There are two more costs that I want to add. One is the cost of data input, and the other is the cost of analysts and programmers. With distributed computing, is data input cost likely to change? And if so, what sort of magnitude? Are analysts and programmers likely to change? If so, what sort of magnitude?

So let us have a look at the cost components. I will skip through them quickly because I am not trying to be exact or precise. I am trying to indicate trends, and I wouldn't like you to tell me that you think that the line of the graph ought be a shade higher or a shade lower. Either way, I agree!

The first thing was to establish what the cost of the staff was over the last few years. I took from Computer Economics Survey, to which we subscribe, three different jobs: the DP manager, the systems analyst, and the computer operator. I took average national basic salaries in 1968, 1972, and 1976, averaged them all out, took 1976 as 100, '72 became 58, and '68 became 39, (see Exhibit D5). They are plotted against the cost of living index over the same period (see Exhibit D6) from published statistics. I plotted the two on log paper. You can see that the cost of living index displays a disturbing upward trend on the logarithmic scale, until last year when it actually managed to flatten out and just be exponential. Staff salaries appear to have followed the trend in a similar sort of way, and we end up optimistically forecasting that it will not continue quite as

badly as the recent trend has been and that it will carry on at about this sort of rate. In five years time, we can expect staff salaries to have about doubled; and that seems to be consistent with our recent experience.

The next thing to do is to look at the hardware trends (see Exhibit D7). There are as many ways to calculate hardware trends as articles you read. I tried to do it in a pragmatic way by saying that a computer, at least in the immediate future, will consist of three main cost components. I have grouped them as shown because the rate of change of the costs is different for the three.

We got some figures from IBM, Honeywell and DEC, which were all pretty consistent. The cost of CPU memory has come down by a factor of 10 times in the last 10 years, but software has grown to fill the gap, so what we assumed was that it had gone down by a factor of about five times, which is perhaps more realistic and practical.

Disks. I thought the figures would be more dramatic than this, but disks are very price competitive at present and have been for a long time, on both mainframes and minis. It is the area where minis do not score. Disks seem to have come down by a factor of about four times over the last 10 years.

The remainder seems to have stayed about stable: VDUs cost the same as they used to; they do a bit more but they cost about the same. Printers cost about the same.

At current cost levels CPU memory seems to be about a third of the cost, disks about a quarter, and the remainder about 42%.

If the assumptions are right - and don't argue with me too much about them - then the costs of the three main cost components ten years ago were respectively five times than shown, four times that shown, and one times that shown. So overall the cost of computing has come down, in hardware terms, by about three times in the last ten years. When I got to that final figure, it seemed about right.

I then plotted the hardware line on the graph (see Exhibit D6 again). The forecast says something like halving the cost of hardware over the next five years, which again feels about right. At any given point in time, of course, the particular thing that you've got will be going up 10% in money terms every year, but if you buy in a couple of years time it would cost a bit less than now in real terms with a bit more capability. I am thinking of the PDP-11/34 which did not exist three years ago when we ordered our 11/45.

There is a deduction that you can draw from this immediately. For us to run our computer for four shifts five years ago was clearly pretty good use of expensive hardware and very cheap labour. To be using it on more than about one and a half shifts in five years time looks like pretty bad use of cheap hardware and expensive labour. That requires some thinking about. I had not appreciated that message quite so clearly before.

I don't know whether five years time is exactly right of course, but at some stage, and it can't be too far ahead, we shall be going to quite a lot of trouble not to run our computers at night. The point was put to me very eloquently by somebody recently, who said that he has three lavatories in his house and he doesn't use them all continuously, but they're very useful and cost effective. Actually, we don't even run our bank in three shifts, and the premises must cost more than the computer. We only run a bank on one shift, 9.30 to 5.30. Just to complete the cost picture, we reckoned that for accommodation and consumables, the picture was confused because the cost of accommodation was peaking out a few years ago in central London, but broadly we see the thing going up at the same sort of rate as the cost of living. Hardware maintenance we saw as staying about satisfactory, because the hardware costs are coming down but the engineer's costs are going up.

The cost of the telephone you may regard as being more argumentative (see Exhibit D8), particularly as everybody else here today appears to be forecasting dramatic drops in unit costs. Well, all I can say is that the graphs over the last five years have been horrific, and this is a very generous line to draw to the Post Office. I do appreciate that it could be that the accounting in the Post Office was a problem; that they were not depreciating in anticipation of replacing exchanges, they were just saying "They're all written off so we don't need to depreciate them any longer." They have made a profit this year, so possibly the recent rise has been unnecessarily dramatic.

The only way that I can justify this telephone cost trend is that it is the cost you cannot control. For all the other ones there is competition about; there is no competition for this. We saw the way that the Post Offices in Europe cunningly increased the cost for SWIFT, which they agreed not to do at the beginning. There is nothing that you can do about it, they are monopoly carriers. So I think that telephone costs will rise. The bit rate, in our experience, has very little to do with it, because we never want to send at very high speed; we want to send at some sort of speed which is compatible with our customers. We can't get modems to do that at decent speeds, although we want to do it at present. So we are limited by the current technology, not by the line speeds at all.

Those then seem to be some of the trends. How do these relate to our computer alternatives that we saw at the beginning? The cost of data input, when you begin to analyse it, is quite horrific under a batch system. I know that I am saying something that is familiar to you, but let me just labour the point (see Exhibit D9). From the work station, where the person who has the knowledge of the job sits, you go to some sort of data control function which buffers the computer against the real world. Often there is a coding function; then you have some sort of add-listing function and batch proofing; and all of these have got to be accurate. Then you go into some sort of punching and verifying; then into the computer where you do editing runs and so on, and various bits are recycled, and eventually your output comes bank to the work station. It adds a considerable amount of cost and delay.

Let me compare with you the way that a credit controller processing cash receipts might work with a VDU compared with working in a batch system. With a VDU, when the cash payment comes in, you want to call up the account. If it is W.H. Smith, you don't want to code it, you want WHS or something like that. If it is the Northampton branch, you put in NN1 and use the postal code to bring it up. If there happens to be two in central Northampton, then two come up, and you can say, "Well, that's the one I want." You can do that because the person at the work station knows what he's doing, as opposed to the person down the line who doesn't. So having got the account, with all the different transactions, we enter the cash amount and the discount amount, and then begin the allocation, "That one, that one, and that one." The computer, if it has a remainder, says, "Remainder", and you realise you've made a mess of the discount, so you go back over it again, or say, "Scrub that, let's just do it again."

Contrast that type of data entry with batch systems, where you had to fill out some sort of form, or hopefully you made it slick by using a remittance advice of some sort. You then passed it to punching and verifying and they did certain things to it; you stuck it in the computer, and when that was all done you computed the allocation. Then you discovered that the discount did not match. It could have been days before this cycle was completed.

We have analysed the source of our errors, and we have found that a third of all our errors are caused by correcting a previous error. It is extraordinary how error-prone is error correction, and huge amounts of effort have been expended on this to absolutely no purpose. With batch, if we have twenty people running our computer, then we are pretty confident that there are forty people outside, in data preparation of some sort, who are only there because of the system.

Operations staff here (Exhibit D4 again) were about the second largest cost of the whole operation. If we are talking about doubling that for the people outside just to run this batch computing operation, you will be talking about a cost which is far bigger than all the cost components we have analysed so far - half as much again. This is not an argument in favour of minicomputers particularly, it is an argument in favour of interactive data entry. It is not an argument against batch processing, it is an argument against batch data entry. I am postulating that the concept of using batch data entry was the blind alley of computing; it is the thing that has given computing its bad name, to my mind; and the thing which we will get right away from, I believe now that software exists for us all to be able to use it reliably without worrying all the time about the technology.

On the last main cost area, systems analysts and programmers, our survey suggests not much change. This is contrary to the belief of one or two people who visit manufacturers' glossy emporiums and are told that programmers will become like clerks in the future, using extraordinarily high level languages. As always, the higher the level of language, the cleverer the programmer has got to be to use it intelligently. Data entry systems are not places for analysts to be playing. This is where the system, if it is not designed right, grinds to a very visible halt. So I still think that we will be having our systems analysts and programmers, (which is quite good news for my chaps).

Returning then to our original trend summary (see Exhibits D2 and D3 again): central batch computer trends. Data input costs we see as going up, and they are expensive already. Operations staff costs we see as going up, and they are expensive already. Accommodation costs we see as going up, and it is fairly expensive already.

Hardware costs in this framework are coming down. So we see the trend as being against the central batch computer.

RJE. The trend is even worse as line costs are going up as well. I suppose in this analysis is an implication that there is no such thing as economies of scale these days. That's another argument in which I don't want to get involved.

Multiple access to the central computer. The data input cost is down because these chaps who have actually got the problems are able to put the thing in directly themselves. The operations staff costs are held, maybe even reduced a bit, because these people with their transaction processing systems drive the computer for you; you do not have to drive it. More things are on disk. Central accommodation is still going up. Lines only

go up if you've got a lot of external data transmission; if it is internal and it is under your own control, you do not have to worry so much. But already we are seeing the situation where the next one on the list, the personal computer, is already competing very favourably with external time sharing systems, if they are used at all intensively. Line costs come out to several thousand pounds per year which buys you your local computer. So we see this as being a very popular solution for quite a long time, but the trend in the long term is probably against it.

The personal computer. This came as a considerable surprise. I had not really been in favour of it, intuitively. But when you begin to analyse it, the data input cost is down because the man with the problem just sticks it in. The operations staff costs are down, because the man with the problem operates it. Accommodation costs are down: it's his desk. Line costs are down, because there aren't any. It is a very persuasive approach.

The thing that blocks it, of course, is if you need to share files. If you must have common access to data files, stock files or something like that, this is against it. But if you are talking about the scientist and perhaps the traditional time sharing user of computing, then it will be a very powerful alternative - if not today, then tomorrow, it will be extremely price competitive.

The local stand-alone computer. This is the one where you have a local small computer, say a minicomputer with on-line terminals, interactive data entry, accessing common files. We see this as being pretty attractive because the data input is done by the people who understand the problem. Operations, if you get it right, is done as a by-product of running the business. Accommodation costs are down because you are talking about marginal accommodation usually; about the corner of a room, not a three-storey building. Line costs are down, because you haven't got any.

Local computer with dial up to the centre. The idea here is that you could partition your application area so that you could do a lot of the work locally, with local files, and at the end of the day you could pump the stuff up the line intensively, using a switched network preferably, or alternatively a leased line in the evening or late at night when it's not being used for voice traffic. We see this as being a very satisfactory way of coping with the costs of lines and having very low line costs, together with local autonomy, but being able to get data into the centre. We include in this category producing magnetic tapes at the end of the day on our minicomputers, and taking them by taxi across to the mainframes. Files are brought back

next morning with the output for the local minis. So we have local autonomy with some sort of intensive link to the centre.

Finally, networks of computers. The trend is a little uncertain. Costs are down; operations staff costs are held; accommodation is up if you have to have a central computer; but the real thing is the risk associated today with doing it. This view may not be shared by any other speaker here, but personally I run miles from networks. I am sure that it is the thing of tomorrow, but I'm looking forward to a lot of other people experimenting first. I am frankly dubious about the business need for networks, but if the business need does exist then hold off for a couple of years while somebody else does it.

Given those cost trends, can we summarise the logical arguments in favour of either distributing our computing or centralising our computing? I will tell you our conclusions first. We ended up by thinking that almost all the arguments favoured pushing our computer power to the line departments and giving them their own machines; but equally the arguments favoured retaining central control in a management services sense.

Because there are some very large organisations present, I must be careful here about what I mean by centralisation. I suppose it is a sort of critical mass in terms of the DP function, and if it gets too big, like hundreds of people in the central management services, maybe that is too many. It may be that at that size you could afford to break it down. I think that in analysts and programmers a critical mass is somewhere round the twenty to thirty mark. Below that, you cannot build up flexibility of expertise; above it, you risk fragmentation and inability to control the thing. So my experience tells me that if you are a large organisation which can support twenty to thirty management services people in different divisions, then give each division its own staff, and make the centre a smaller division. If you are the sort of size that we are, which is twenty to thirty management services, analysts and programmers, then this seems to me to be quite a manageable chunk.

Our policy, on which you may think I am biased, is that we do not destroy the central management services; we keep that and we do the batch work, the night work and so on in central machines, but all the daytime work we are pushing out to our users, and they run minicomputers themselves.

If you analyse the literature, you will tend to find these sort of arguments (see Exhibit D10) in favour of distributed computing. Some of them are business directed and some are computer directed. These are familiar arguments.

Arguments in favour of centralisation (see Exhibit D11) seem to fall into the area of better systems development, better quality

control and less duplication; but there are other things which may determine it: for example, if you have got to share your data files, like an airline seat reservation system.

If we leave that backcloth and go now to the management implications (see Exhibit D12), we have tried to categorise management implications in these four areas: the way computers are used and the way that they are run; the way systems are developed and the way systems are implemented.

Starting with the first pair, data entry we have forecast to be interactive, not batch (see Exhibit D13). We see the key-to-disk systems as being a sort of red herring in the path; we see the demise of traditional punching and verifying and so on. We see everyone moving rapidly in the next few years into interactive data entry. I am not saying that everything has to be run in real time in the processing sense; batch processing will be with us for a long time but not batch data entry.

We see a move away from central mainframes. The cost performance comparisons that we have done (Exhibit D4 again); if you look at the totals our mainframe is about £465,000 a year and our minis are each about £42,000 a year. In terms of power, we reckon that four of those minis equals one mainframe. But we are running the mainframe three shifts and the minis about one and a half, so it is not an easy comparison to make. You normally see a very substantial cost reduction in using the relatively inexpensive hardware on minis.

I think that the guts of the point is that you can say now that business needs matter more than technical needs. You can say, "How did I do this system before computers were invented? What was the natural way of doing the job?" You can then, pretty certainly, provide a cost effective computer solution to that-today. In that sense we are moving away from central mainframes.

We see local operations developing with central support. I have no doubt that our local computers are better because we have a central management services to support them. I should like to make one more point here which is rather subtle. A couple of years ago, we put up a plan for computerising one of our retail banks: a clearing bank operation. The management services proposal was to use an on-line system for data entry, using the existing minicomputer of another department. Now the bankers turned this proposal down flat, It took me some time to work out what the reasons were. The reasons were entirely political. They were quite prepared to have management services run a shared computer; they were quite unprepared to have their work done by another line department, because of priority conflicts which they felt could be resolved with an impartial chap doing it, and certainly would not be resolved if he did not.

At the time, I thought their view was a bit cautious. I have come right away from that view now. I have seen all too clearly the parochialism of the people running their own local minicomputer, and how resistant they are to any change or any flexibility, or any priority scheduling. It is there to support their own needs, and that's what it is going to do.

So what looked like a cost effective use of hardware, using the same piece of hardware for two different things, I think was strategically wrong. It was interesting that our users were clever enough to see this, though I am not sure that they knew why. Trends in the way that systems are developed and implemented (see Exhibit D14). There is no data control function to buffer you from your mistakes. You have now got to put the terminal on the work station to actually do the job of work. Unlike batch systems, our experience with on-line systems is that you cannot get them wrong first time, and then improve them for the next two years. You have got to do a lot more testing and they have got to be right first time. The number of residual errors on our on-line systems is vanishingly small in comparison with our batch systems. You just can't get live with residual errors at all, so you don't have them. When you put it in, it's really got to work. This means that the business analysts have got to be very competent technically and very competent in the business sense. This animal was rare enough a few years ago with batch systems, but we see this as being the key problem area of developing on-line systems.

The interaction with the operating system is crucial because if you just design the thing in the abstract and decide to have a fantastic data base and so on, it just will not run. We found no substitute for duplicating certain elements of the data in one system. The first operation of the day was to regenerate and duplicate certain things in different sequences so that we could answer patterns of questions very rapidly as well doing our daily work very rapidly. We have not found a better way of doing it than that.

Systems reliability and availability. When we put in our first on-line system it was so much better than the previous system that when it broke down for a day or two, that was really regarded as a non-event. The previous system had been so dreadful that to be within two days of being right was a miraculous improvement. Now that we have had it in for three years, being down for a day is a major inconvenience that causes the business to stop dead. Users are no longer tolerant of hardware failures.

There is a paradox here in fact, because our users are also poor at dealing with manufacturers to ensure good hardware maintenance. When the engineer comes and explains the problems he has the user weeping on his shoulder and commiserating with

him, saying, "Of course it's reasonable that the machine's down all the time," and so on. You've got to adopt a very-headed approach to these suppliers, as you know, which users unfortunately are totally ill-equiped to deal with. They're tolerant, and they are taken advantage of.

Systems reliability and availability assumes new dimensions, and these minicomputers are less reliable than the mainframes that I am used to.

So you've got to adopt a new attitude to hardware availability and reliability. What we are now doing on the crucial systems is duplicating processors, and we have so arranged it that the second processor now is available for back-up anywhere else in the organisation. So there is a certain amount of redundancy.

Open-minded hardware policy. I don't mean by that laissez faire, you understand. It means that we have selected the best VDU that we can find. It is a very, very good VDU and very cost competitive; and it is not from any of the manufacturers from which we have hardware.

We want the best VDUs; it matters to us to be able to tune things. On the keying of foreign exchange deals, the £ sign, the $ sign, the millions and the thousands are the four most common keys we hit. So to be able to change the keyboard to accommodate these is of crucial importance. We reckon that it has cut four fours' keying a day out of the systems by halving those four keys in the numeric pad. Similarly we have function keys which we can tailor to our own specs giving large, visible layouts which are easy to use. These are things which the new generation of VDU users is getting. But I don't think that you can just accept that your standard hardware supplier has the best gear. We do have a policy for our mainframes, for our minis, for our micros. We have a policy for our overseas subsidiaries, a policy for our languages and so on; a policy for our communications - not the fastest but the simplest. So you must have a policy, but don't take for granted what your manufacturer tells you.

The last point is in some ways the most crucial and it is the point that I will finish with. It is the impact on the existing clerical functions when you change your systems. This sort of slide (see Exhibit D15) will be familiar to you. In a clerical operation you have your tasks to be done; you have people doing it; you have the organisation's structure within which you are working; and you have the technology that you are using to achieve it, in some sort of equilibrium which has been arrived at over a period of years.

If you change your technology dramatically, by putting on-line VDUs in, do not think, as your users may, that these other

things can remain the same. I should like to illustrate this point with an example based on our experience with our first minicomputer (see Exhibit D16). It was in our investment management company. We run portfolios, possibly for some of your own companies. We manage pension funds and the portfolios of private clients and invest on their behalf, hopefully very well. The dealers are the people who buy and sell on the stock market. They pass their instructions, which are in a very abbreviated form, through to contract clerks. The contract clerks encode the data on documents together with additional information.

Previously, when they had filled these documents out, they sent them to the TC500 operators, who prepared contract notes and also data for the central computer's accounting systems. Contract clerks then sent out the contracts to the clients.

Now when we decided to replace the TC500s with VDUs, the organisational question arose as to where to put the VDUs. I will postulate four alternatives. The first one is to say, "I've got four girls already. We'll train them each as a VDU operator. "We'll replace our existing batch TC operation with a batch VDU operation."

The second possibility says, "Why don't we divide our contract clerks into groups of two or three and give each little group a girl with a VDU." This is pretty efficient in VDU time; it would probably cut out certainly one or two of these operators and probably one or two of the contract clerks. A much more difficult management decision.

The third solution says, "Why don't we give every one of the contract clerks a VDU, and make the person who understands the problem put in the entry himself?"

The final solution is, "Why don't we scrap the contract clerks as well, and give VDUs to the dealers?" Then put voice response in and let them talk to the computer in some way.

These were the four alternatives that we could see, and we quickly ruled out the first and the fourth. The first one, which was just to retain the TC500 girls in a room and give them the VDUs we saw as giving us no benefits at all, it was just perpetuating the old system. It was converting a batch computing system to a batch on line system.

The fourth one we ruled out because it turns out that there are business reasons why these two functions are separate. They are dealing in different time cycles, they are doing a different job, and there is a business reason why they are split. If there is a business reason, there is probably a computer reason.

So we decided that we would not go for either the first or the last solution. The question was: should we go for the second or the third?

The second was to put a girl with each group of contract clerks; and the third was to give contract clerks VDUs. In fact the answer is that it depends. In this particular example, we gave each of the contract clerks a VDU, and our reason for it was that there was not a very large number of them, so that it would not cost us very much. A little inefficiency in the use of the VDUs just did not matter very much, and we had double the minimum number: double a small number still isn't very many. We have in fact eliminated the complete TC500 section. The contract clerks run the PDP-11 in their spare time and they do more deals a day than they used to. The number of errors now is virtually none. They cannot remember when there last was a mistake.

This is not the approach however that we adopted on our foreign exchange side. On our foreign exchange side we adopted the second approach, which is to group the mail transfers, commercial clients and so on, into little groups and put one VDU in each. The reasons were that there were more people involved in collecting the data; it is a more complex gathering-of-information job with less VDU time; and the size of the operation is bigger.

I am not trying to say that there is a given solution, I am saying that the solution depends on the situation; but most solutions involve bringing the VDUs very close to the work station.

This concept is not easy to put across even when you believe it, so here are two examples. The first is in our personnel department. We found that staff loans are handled by our accounts department which does the payroll. There is constant complaint historically that staff loans are often in a mess because the personnel department do not tell accounts that a loan should have been terminated by now and so on. We have been looking at an on-line system to replace the existing batch systems in the accounts department. We did not immediately see that we should put the terminal in the personnel department, not in the accounts department. The fact that the accounts department existed tempted us to say, "Let's improve the existing system and give them a VDU," but that was wrong. The answer was to put the VDU in the personnel department.

The second example was that we find ourselves entering journal entries from journal slips, and we found the analyst was designing the VDU form to correspond to the existing journal slip. As I have said, that is obviously wrong. You've got to cut out the journal slip and get back to the original cheque

or whatever. So it is not enough for me to say that this is our policy, you actually have to think the policy as well; and you have to question all the time why we are doing the existing system, because almost always we have created, because of our own failings as management services people, a whole bunch of middlemen who are not there to support the business at all. I think that the impact of this is substantial and far greater than I would have supposed.

This (see Exhibit D17) has been the impact on the management services budget of putting in the minicomputers, getting the new hardware, and going towards distributed processing. It shows the cost of living over the last nine years, and our budget over the same period. It has been stationary for the last four as we have been using hardware better and using people better. That is one side of it; what you don't see is the cost to the users. Over the next year, we are expecting to see staff costs halving in most of these clerical support areas because of the great reduction in inefficiencies associated with batch data entry.

In conclusion (see Exhibit D18), we see a move towards interactive systems, stand-alone minicomputers and central management services. We see big rewards; the DP department down by a third; clerical costs cut by a third to a half in these data entry areas. We see a better quality job being done much faster.

I was grateful for the opportunity of having to do this research, because it helped me to focus my own mind on these trends, whether you agree with the conclusions or not. We are where we are, and I am not suggesting that overnight we throw everything out and start again. We need to evolve towards these goals because we cannot guarantee that what is right for one organisation is right for another. I foresee a lot of trouble with our business analysts, getting the right people actually to do this job.

Exhibit D1

DEFINITION

Distributed Processing means putting computer power where the people and business problems are.

Exhibit D2

COMPUTER ALTERNATIVES

	COST TRENDS			
	DATA INPUT	OPS STAFF	ACCOMM-ODATION	LINES
A: CENTRAL BATCH COMPUTER	Up	Up	Up	-
TREND AGAINST				
B: RJE TO CENTRAL COMPUTER	Up	Up	Up	Up
TREND AGAINST				
C: MULTIACCESS TO CENTRAL COMPUTER	Down	Held	Up	(Up)
TREND SLOWLY AGAINST				

Exhibit D3

COMPUTER ALTERNATIVES

	COST TRENDS			
	DATA INPUT	OPS STAFF	ACCOMM-ODATION	LINES
D: PERSONAL COMPUTER	Down	Down	Down	Down
TREND STRONGLY FOR				
E: LOCAL STAND-ALONE COMPUTER	Down	Down	Down	Down
TREND STRONGLY FOR				
F: LOCAL COMPUTER WITH DIAL-UP TO CENTRE	Down	Down	Up	Low
TREND FOR				
G: NETWORK OF COMPUTERS	Down	Held	Up	Up + risk
TREND UNCERTAIN				

Exhibit D4

COMPUTER OPERATIONS COSTS

in £'000 p.a.

	MAINFRAME (3 shift)		%	MINI		%
Hardware depreciation	130		28	18		43
Operations staff (no.)	115	(23)*	24	5	(1)	12
Accommodation (sq.ft.)	110	(6000)	24	6	(200)	14
Hardware maintenance	50		11	7		17
Consumables	50		11	3		7
Transmission	10		2	3		7
	465			42		

* Mainframe operations staff : 23
Mini operations staff : 2 half time

Exhibit D5

COST OF STAFF

		NATIONAL AVERAGE BASIC SALARIES 1968	1972	1976
Job 01	DP Manager	3225	5128	8209
	(normalised)	39	62	100
Job 04	Systems Analyst	1662	2490	4286
	(normalised)	39	58	100
Job 13	Computer Operator	957	1406	2537
	(normalised)	38	55	100
Normalised average		39	58	100

Exhibit D6

STAFF COSTS (oS)

COST OF LIVING INDEX (•) plotted against time

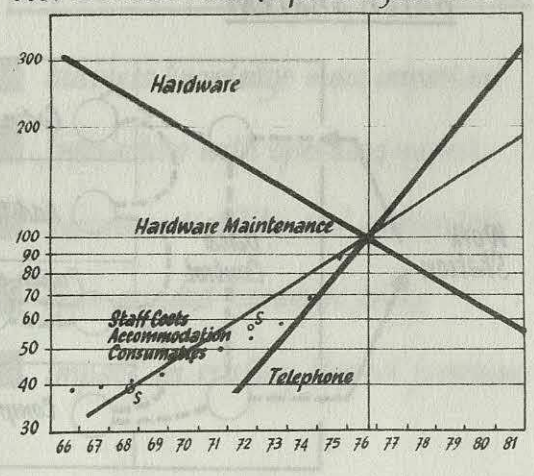

Exhibit D7

HARDWARE COST TRENDS

	CURRENT COST	COST 10 YEARS AGO
CPU and memory	33	165 (x5)
Discs	25	100 (x4)
Remainder	42	42 (x1)
	100	307

Exhibit D8

DATA TRANSMISSION COSTS

	1971 (Feb.) time for 1p	norm.	1974 (Sept.) time for 1·5p	norm.	1976 (Nov.) time for 3p	norm.
Local calls	6m.	17	3m.	50	3m.	100
Up to 35 miles	30s.	50	48s.	53	45s.	100
Up to 50 miles	15s.	33	15s.	50	15s.	100
Over 50 miles	10s.	50	15s.	50	15s.	100
Average		37·5		51		100

In April 1975 the Tariff T (speech-type private lines) was increased by about 40% (from 71 in 1975 to 100 today).

Exhibit D9

COST OF DATA INPUT-BATCH SYSTEM

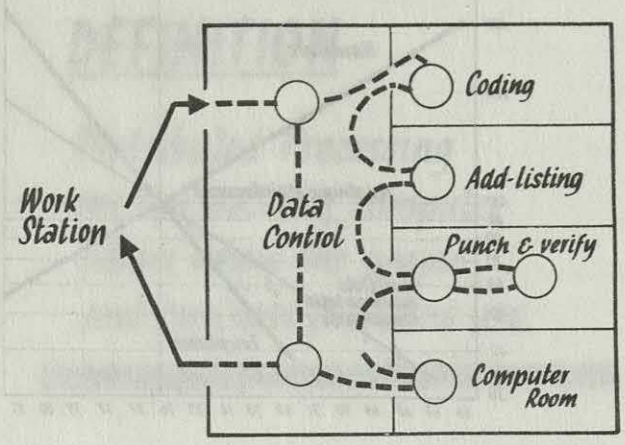

Exhibit D10

ARGUMENTS FOR DISTRIBUTED COMPUTING

BUSINESS DIRECTED	COMPUTER DIRECTED
Matching needs	■ More responsive
Use of people	■ Less rigid
Power near problem	■ Development risk
Local commitment	■ Operational risk
Flexibility for change	■ Failure recovery
Fraud / Privacy risk	
Incremental growth	■ Parallelism
Less complex	■ Hardware tailored
P. & L. accountability	■ Development time & cost
Cut central overheads	■ Data communications cost
Local priority scheduling	■ Simpler operating systems

Exhibit D11

ARGUMENTS FOR CENTRALISATION

BETTER SYSTEMS DEVELOPMENT	OTHER FACTORS
■ Avoid duplication	■ Necessity for shared data files
■ Reporting uniformity	■ Less data duplication / synchronisation
■ Equipment standardisation	■ Management and control of operations
■ Design quality control	■ Access to larger CPUs
	■ Current computer investment
■ + Economies of scale?	

Exhibit D12

MANAGEMENT IMPLICATIONS

The way computers are used

The way computers are run

The way systems are developed

The way systems are implemented

Exhibit D13

TRENDS IN THE WAY COMPUTERS ARE USED AND RUN

Data entry interactive, not batch.

Move away from central mainframes.

Local operations with central support.

Exhibit D14

TRENDS IN THE WAY SYSTEMS ARE DEVELOPED AND IMPLEMENTED

- business knowledge more important
- Interaction with operating system
- systems reliability and availability
- open-minded hardware policy
- impact on existing clerical functions

Exhibit D15

SYSTEM COMPONENTS IN EQUILIBRIUM

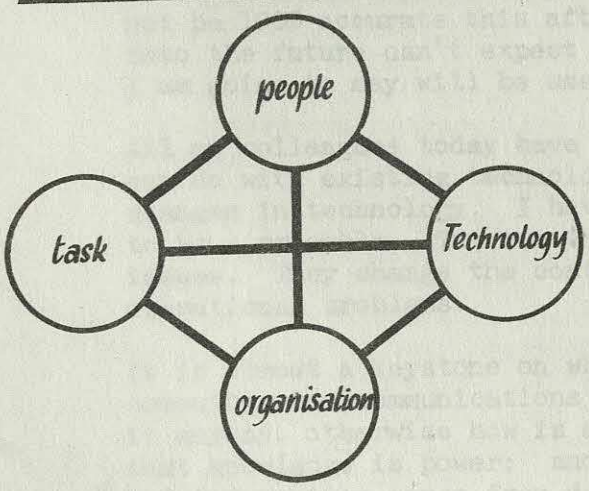

Exhibit D16

HOW SHOULD THE VDU's BE ORGANISED?

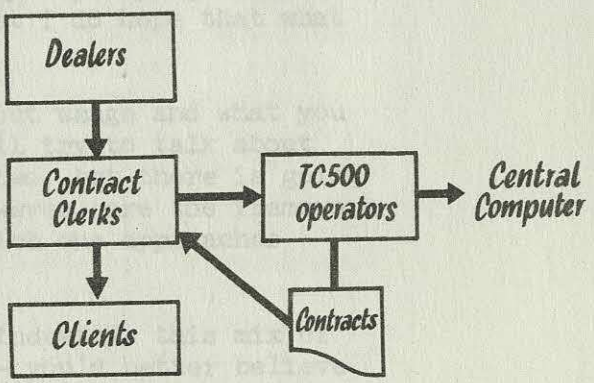

Exhibit D17

THE EFFECT OF DISTRIBUTED PROCESSING ON MANAGEMENT SERVICES BUDGET AT HILL SAMUEL

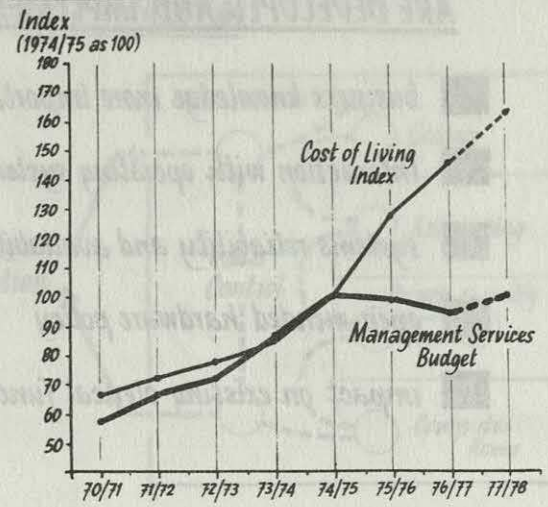

Exhibit D18

DISTRIBUTED PROCESSING - Summary

MOVE TOWARDS	Interactive Systems Stand alone minicomputers Central management services
BIG REWARDS	DP budgets down 1/3 Clerical costs down 1/3 to 1/2 Turnround and quality
NEEDS	Strategic awareness Planning now Business analysts

E

OFFICE AUTOMATION — NEW HORIZONS

Rex Malik

MALIK: I am not a consultant; I am not a management services controller; I am an observer; and so my views may not be structured in the way that you would expect from the speeches you have heard this morning. It is traditional for the first speaker after lunch or dinner to make a joke. Peter Goldman, when he was Director of the Conservative Political Centre, used always to start off his speeches by saying, "I have two jokes, and this afternoon you're going to hear the other one". This went flat a little because I heard him about four times; each time he used this; and each time he told the same joke.

I in fact have two stories, and this time you are going to hear the other one. It is one I found recently in the Financial Times, which may have some bearing on this afternoon's proceedings. It is the tale of the balloonist who came from France to England, landed in a field, did not know where he was, so he waited. Eventually, a man came along and the balloonist looked at him and said, very politely, "Can you tell me where I am?" "Yes," said the man, "you're standing in a wicker basket in the middle of a field." So the balloonist looked at him and said, "You must be an accountant, mustn't you?" The man said, "How did you know?" He said, "Well, it's quite simple. Like all accountants, your information is accurate but bloody useless!" I suspect that I will not be 100% accurate this afternoon, anybody trying to peer into the future can't expect to be; but I do hope that what I am going to say will be useful.

All my colleagues today have talked about usage and what you can do with existing technology; I will try to talk about changes in technology. I have picked two, but there is going to be a preamble. The two changes, however, are the framework issues. They change the context in which one approaches operational problems.

It is almost a keystone on which this industry, this mix of computing and communications is built - you'd better believe it anyway, otherwise how is all this kit going to be shifted? - that knowledge is power; and knowledge comes from information; and information comes from data; and where do you think data comes from? They don't tell you that there is a price to be paid for this.

I want to ask a question right at the start: is knowledge power in fact? Let us look at what has actually happened. This is an interesting chart (see Exhibit E1) which comes from a very remarkable study completed last year in the States, which talks about the growth of the information economy. What that chart shows is that for every dollar spent on information in its broadest sense within the US economy between '28 and '74 we have the productivity shown on the graph. You notice that the graph has not yet got horizontal, in fact it has started going down again. You will note one thing here, that what people call the "information rich" society is itself the rich society. The great big peak happens to be the years of the depression, and showed that you could still keep the machine reasonably operational in global terms, and have a very high added value for every dollar that you spent on information.

What this shows (see Exhibit E2) is that over 50% of the US labour force now, measured in terms of the wages and salaries paid (not in terms of GNP which has other factors in it) is in fact spent on information. The inclusive/exclusive ranges that you can see there - all this is based on the national income accounts; and what they had to do when they looked at the national income accounts was to separate those whose jobs are primarily concerned with information from those who are not, take the middle group and apportion the information content of the work and the non-information content of the work. The result is that you get something like 51%, as of the middle seventies, now spent in the US economy on information.

It is axiomatic that where America leads, we usually follow. I see no signs that the situation in Europe is any different. Indeed, nobody has ever done any studies in Europe. If anybody were to do the studies, given our long and differing histories, I would not be surprised to find that we are considerably over 50% already.

This (see Exhibit E3) is what we traditionally conceive of people as doing. You will notice that the information curve has flattened out. I refer you back to two slides ago, when I pointed out what happened in the Depression. We have had a recession. I don't regard that as any serious, long-term alteration to the trend.

I now want to look for a couple of minutes at one of the reasons that we have what you can call the "information society". I know that anecdotes are not research, but they can be usefully illuminating. I want to tell you a couple of things that have been going on which are reported out of the hearings of a Federal Paperwork Commission in the United States, in the land of the free and the brave. The State of Maryland refused to accept a $60,000 grant recently, because somebody did a quick analysis and found that the cost of completing the forms was

$45,000. In another example reported to the Federal Paperwork Commission, it took 800 lbs in weight of paper to inform an Indian tribal official who had to be informed of the change in one law.

If you look at the needs and requirements of government - whether you accept that the requirements are necessary is not at issue here - from the business community, take the examples of a couple of companies in the States. An oil company reported to the Federal Paperwork Commission that it spent $17 million and used 475 full-time workers to file government returns. This was in 1975. I suspect that with President Carter's new energy plan, that may well increase. I wish to point out here that taxes and associated matters were totally separate, and the reporting requirements were kept totally separate and extra.

Dow Chemical has reported that, again in 1975, it cost $147 million to fill out all the governmental forms and produce reports that the Government required. The chairman reported that $50 million to $60 million was totally unnecessary.

The chairman of Lillie has indicated that the company, putting more flesh on it, fill out 27,000 forms or reports a year, at a cost of $5 million. The research to back up that form-filling cost another $10 million a year. He gave some examples. An information index for one product submitted to a regulatory agency ran to 153 pages: and that is the index. Some of those entries referred to documents which were 3,000 pages long. I think that the best one is of another product where it was reported that the single application to sell one chemical ran to 120,000 pages.

Those are anecdotes, but there is a reason that they are there. It is worth pointing out, as if you didn't know already, that the tendency for governments to ask for more information is increasing. If you go on employing people in ever-increasing numbers to intervene in the name of society in other people's affairs, you must not be surprised if, to earn their salaries, they go along and intervene. There are no signs that this intervention will get less; indeed, President Carter is talking of more. I think that he is also talking of cutting back on some of the nonsensical paper work with which we all have to cope. I think that his chances of intervening in new areas is a lot easier. He will find it much easier to intervene in new areas than to cut down on the intervention that is already happening that he happens to dislike.

I would point out that this intervention will throw much greater stress on the business machine on both sides, both on the governmental side and on the corporate side. It will throw stress in many ways; not only volumes, but also the

data that is sought by the community is getting to be more and more what has been even in recent years considered sensitive. Even if it is not sensitive, a large part of it will be much more difficult to provide.

Let me put those remarks and anecdotes in a broader context. Business Week recently reported on how the tide of federal regulation is rising. With regard to spending by the major economic regulatory agencies in the United States, it said that in 1970 there were eight agencies and they spent $166 million; in 1975 there were ten agencies and they spent $428 million. These are the agencies which primarily concern business. But the social regulatory agencies, what with forms of national insurance, Medicare and so on, are also being increased. In 1970, 12 agencies in the United States spent $1.4 billion; in 1975, 17 agencies spent $4.3 billion. I may say that at the same time, in those five years, the number of pages in the Code of Federal Regulations which is kept on a lot of corporate shelves, rose from 54,000 to 72,000.

I use the United States as an example of the trend which we are following. In fact, it would be true to say that we precede them in many ways, but unlike the United States, in Europe we are not so statistically minded; our forms of organisation are different; and therefore our statistics in some respects are much more difficult to get at.

But let me try to give some more things which show the increase in government regulation, again using US figures. Between '58 and '70, which is just 12 years and which, incidentally, coincide with the boom years of traditional computing, other sources besides the study in Business Week that I have just used show that Federal Government information (economic planning information gathering) increased eightfold; industry regulation tenfold; diplomatic information gathering only sevenfold; and free information services, if you can call them free, such things as providing air traffic control, increased elevenfold.

There is an example of what this can mean in the field of industrial regulation. The Inter-State Commerce Commission, which regulates inter-state traffic, now has some 4 trillion inter-state trucking tariffs on file. What happens when you get too much regulation is this: you find that the rate of return goes up. The average rate of return for US big companies is 14%; the average rate of return for the 13 largest trucking companies in the United States is 22%.

All right, let's put some more cloth on this. Federal form-filling and reporting cost and US economy $40 billion a year. There are 12,000 laws which require the citizen, in one role or another, to complete 10,000 forms. You may well ask what

all this paper costs. It costs the Federal Government about $1 billion a year to print the forms for people to fill out; it costs another $1 billion a year to print the instructions which go with those forms to tell them how to fill them out. It also costs $1.7 billion a year to store those forms. Is it any wonder that although the Federal Government is trying to cut back on all this, the federal work force required still continues to grow, and so does the business work force?

I can put all the facts that I have shown in a very simple figure. Between 1960 and 1980, the last few years projected by the US Department of Labour, the information work force in private and government sectors has been growing at a compound rate of 3.87% a year. I have put this here because that really sets the context of what I want to talk about.

It does seem to me that much of the information that governments - whether American, the European Commission, or the various national governments of Europe - require is dead information. It is tax, accountancy, legal purpose information. It is basically the type of information which has been the backbone of commercial data processing. It is coincidentally almost a perfect example, government reporting to government, of the sort of problems which are perfectly suited to handling by well-organised word processing system. One can say this because, while perhaps the Government is not going to keep its forms exactly the same from year to year, they will not change with any great regularity; and also, there is always a common core to any form of government requirements. I may say that we can discount the likelihood of any government requiring less information, or even changing its mind with any great frequency.

People find this difficult to understand, but the problem is really the same as one faced in journalism. A French newspaper looked at its editorial columns and did an analysis; and it found that 60% of the news was predictable. It did not matter how far ahead you cared to look, 60% of the news was predictable. Of course, it is. The football season comes round every year, at roughly the same time of year. Christmas comes up with great regularity, and so do holidays. I have an old friend with a connection with that paper, who says that on two days of the year, the two famous weekends when the French public disappears off on holiday and comes back from holiday, he can predict the front-page headlines. He says that the only things that he cannot predict are the number of vehicles involved in the pile-up; the actual spots where the pile-ups occurred; and the number of dead. But he can be pretty certain that that is what he will face every year.

Let me change gear, and caution you for a moment on what I call the technology trap. There have been a lot of false promises

made, largely by members of my clan, who spotted a bright idea which makes good copy. Everybody thought, "It's coming," but unfortunately the 10 or 12 years of development work have not been done. We can all name them: robotics; machine translation; cryogenics. As we all know, the real test of cryogenics is: will it stand up to the boots of an Irish navvy? Colourgraphics, which is the early sixties. This problem is not restricted to computing. The solution to London's transport problems is obviously a monorail system. This is a solution which was strongly recommended and put forward by a Royal Commission. That Royal Commission reported in 1907, when the problem on the roads was not the problem that it is today, because London's transport was mostly horsedrawn. So why should we expect that some of the technological change that is going on today should be any different?

Secondly, any system which is going to reduce drudgery, one that is cheaper than what went before, and can, in the case of VIEWDATA to which I will refer in a minute, be run out of normal admin expenses and does not involve you in capital outflows; and given the likelihood of the Government changing capital allowances every five minutes, that sort of thing is liable to be a winner.

Last, we come to the technology. I am sorry, but I can't go along with all these people who go on about distributed computing and putting decisions back to where they are made. I think that one of the reasons that distributed computing is beginning to take off is not because people consciously wish to put the decisions back to the point where they tell us the decisions are made, it is because the systems that we now have are really much too complex, technologically too complex, to manage and to enable us to do what we wish with them.

I believe that we need a considerable amount of technological change with existing, non-distributed systems before we can get into a situation where we get a fair choice between the two. The key to that is very good, cheap, fast access memory technology. I think that it follows that we need to see an end to the thing that we accept as normal - the mixed memory within a system. I could have said this five years ago, and everybody would have said, "All right, Rex is discussing the 1990's and the year 2000," but now we happen to be on the brink of a technology which makes it possible. I am now talking about the system which does not look at all like the systems that we think of today; I am talking about the system based on the single level store. I think that the single-level store system is no more than about four to five years away. I may say that my forecasts have a habit of being wrong; I am usually too conservative. I may say also, having had a look this morning at the composition of the audience here, that, as ever, it will be the big corporation which will make the

changes first. So within this room are probably some of the people who will first feel the effects of this.

I talk about a single-level store. You don't need to take it from me. Maurice Wilkes, at the Future Systems conference recently, was talking about mainframe computers of current types. He said:

> "Many of the problems that arise with these systems do so at present because of disparity in speed between high-speed memory and fixed head disks or drums. If these disks or drums can be replaced by something faster, for example, bubble memories or charge-coupled memories, or, better still, be done away with altogether, and very large single-level semiconductor memories used instead, then these troubles will largely disappear."

I know that whether it is Wilkes or me saying that we are bound for one-level storage, I am going to upset quite a lot of people. DP managers are just beginning to feel comfortable. They have this mix of memories, disks and tapes as back up. People are just beginning to explore what you can really do with floppy disks. There are now alternative suppliers, with equipment at competitive prices. The DP manager has the cost structure well structured in his mind. And I'm now saying that we're due for change again. We have all been behaving as if the organisation of data processing was fixed. Most of what I have heard in the last couple of years where I have had communications intermingled with computing - extensions of what you can do with current systems, current architectures, current economics - it has seemed to most people to be a matter of rearranging the technology that we know.

We are witnessing something very unusual. We are in a situation where, whatever installation you have, there is a lot of competitive memory about. This will probably be my only remark about IBM: that memory is undercutting IBM in a very serious way, and IBM has not been doing much about it. Now any student of IBM knows that if a situation goes on for a couple of years and IBM is not doing anything about it, that is most unusual indeed. The buzzes that one hears on the market, both here and in the States, indicate that IBM is going to do something about it, probably in the bubble memory area. That alone would indicate that we are bound for change.

Let's go back and think about the one-level store. It should make it possible to create very large on-line systems. It will also bring some problems, for example, the problem of cooling. But what is it largely caused by? It is caused in large part by mixing technologies. It is caused by high-power requirements. Let me give you what people think are the two critical figures of the next few years, of the penalties for mixing

technologies. With current IC technology, the penalty for leaving a chip and coming back to another somewhere else is that the power requirement goes up 100 times, 100 times the power that would have been needed had we been able to stay on the chip. With the technologies that we are moving to, for instance, the electron beam addressed memory, the power requirement to move off the chip and move back again goes up quite dramatically: everybody seems to have a figure of about 1000 to 1.

It follows, whether you like it or not, that the computer manufacturing industry will move in the direction of systems which require you leave the chip as little as possible. There is, of course, an added advantage; not only can the industry talk to you about systems which cause you fewer problems, as far as the industry is concerned it is a very good reason to make you change your systems. Indeed, the industry has got to make you change your systems. There is no way that computer manufacturing industry can manage to keep us this rate of growth and these sort of profits unless it can make the major corporations go through another set of system changes; whether it does so by the method of introducing a Series 360, or a 1900 system, or whatever, or does so by a process that it calls evolution, is immaterial.

I could look at a large number of technologies. I want to concentrate on bubble for a few minutes, largely because there have been a number of announcements about bubble in the recent past, and I should like to try to show you what is going on. Currently, bubble memory bought on almost an experimental pre-production batch basis is on offer at about 10 times the cost per bit of current disk prices. If you take the best of the disk prices, the very large disk system, we are talking about 100 bits to the cent or 10K bits to the dollar. Bubble is currently on offer at about 25 bits to the cent or 2.5K bits to the dollar. Next year we are talking about 5K to the dollar, 50 bits to the cent. If you then apply the normal drops that we can expect now, you would say that disks are halving in costs per bit per annum, and that semiconductor as a general ballpark is about quartering in cost per bit per annum. If so, the progression towards 1981 looks something like this: 80K bits to the dollar for disks by 1980; but when you look at bubble you find exactly the same thing applying.

In bubble it looks as though it might well go faster than this initially. Texas Instruments will be offering this year a megabit for $530, that is 195 to 200 bits to the cent; and that is including interfaces. That would mean that bubble would be competitive in smallish systems by next year. When you get to that, I want to point out that the ball game has changed. I should also point out that when you talk about disks, there is a limit to how far we can increase packing

density. The best figure that I have been able to get my hands on indicates that we can go roughly 10 times what we are doing now.

Conversely, bubble technology is at its starting point. The price drops that I have just foreseen, are based on a current type of bubble technology; they do not really envisage much technological increase.

I said that I was not going to talk about beam memory, but I would say that beam memory will probably be competitive by about 1979-80. Certainly I know of a trillion bit memory, a B-MOS type memory, for delivery in the States by 1979-80. When we are talking about a trillion bit memory, for the price scales we are talking about it is a megabit per dollar. Anybody who tries to tell you that the curve is slowing down is not living in the technological, real world. The pace is just now beginning to pick up. I would point out that on that trillion bit memory you could store the Bible for less than $50, and that is certainly competitive with anything that I can think of.

Single-level storage brought to us by something based on the IC semiconductor technology does not just mean storage which is different from the way that we have approached storage in the past, it does lead to the possibility of totally new architectures. There are different ways of doing the things that we have been doing, and some of them are almost alien to our way of looking at systems. I just want to mention three possibilities here. One is the notion of the array technology. There is, after all, no basic reason why storage and processing power should be separated. This is a process which has happened because of technological problems.

Caxton Foster came over here recently and was discussing a system which you can think of as a mesh of computer systems, which had storage at each processing point; which was used as a filing system - a fast access filing system. He projected that a system like this, which would cost no more than $2 million and could be built from existing technology - and the $2 million includes as many contingencies as you like to put in there - would have a capability of holding 4 million directory enquiries. He said simply that the access time was such that, if you set it up against a normal directory enquiry system, you should be able to handle between 70 and 80 calls a second, with an average response rate of 20 seconds - which is much better than Directory Enquiries can do now. He also said that if anybody in a directory enquiry situation had 1500 staff to spare, his calculations were that it could save 1500 staff if you had them.

That is a conventional way of looking at it. There is another system that has just been proposed by Bob Barton, of Burroughs. He calls it "The Ultimate Stat Machine". He is going to write a book about it. He has said that Burroughs will build it within three years. The Ultimate Stat Machine is far removed from anything that we can conceive of as a computer. It might seem unbelievable, but it doesn't start with electronics and it doesn't start with software; it starts deep in the area of cognitive research. It is a machine organised according to the best ideas that we have of roughly how we handle material in the brain. Very simple material, but considering the things that we are doing mostly with computers, we do not have to worry too much about very complicated material.

I put those two forward because the point that I am trying to make is that technological change is bringing with it new opportunities to do things which are not constrained by the technological limitations of the past 20 to 25 years.

Now let me move on to VIEWDATA. We have heard a certain number of comments. For those of you who do not know it, it would be worth describing what VIEWDATA is about. This (see Exhibit E4) is the Post Office's drawing of VIEWDATA, no doubt a little idealistic. But basically VIEWDATA consists of the following things: a low-cost, rugged - it has to be for the sort of public service it is - reliable terminal, which is cheap and is based on the domestic TV receiver.

It is a set of interconnected computers, ideally situated within local telephone call reach of the majority of users; an information structure which is quite simple to understand and easy to use; computer protocol naturally adapted for use by human beings, which must require no training whatsoever on the part of the users. Obviously, the software must minimise the cost of information retrieval by minimising the amount of computation for each transaction; and, of course, you must have the dial-up telephone network.

Now within that system, as the Post Office put it forward, privacy is possible. This is what the Post Office tell me. I have some doubts that in fact the sort of privacy that an individual company requires is likely on the type of system lime VIEWDATA within the near future. What I have no doubts about is that one can get the privacy of what one can call a "special subscription club", a number of users paying a preferential rate with restricted access to the system.

I said just now that there was a way in which you do not need to worry about VIEWDATA and capital expenditure. It is noticeable that the Post Office are talking not just of a television receiver terminal, an adaptation of the colour set that you have at home, they are talking also of a special-purpose desk

terminal at a rental of £8 per month. If you must go for the colour receiver, I would suspect that VIEWDATA would not fly in the way one would like it to fly were the receivers to cost much more than £50 to £75 above the existing colour TV set rental.

The one thing that nobody considers about VIEWDATA, which I regard as initially probably the most important use of the VIEWDATA type system within business, at least economically most important, is the use of the system as a sort of instant office telex, without having the bother of putting something down to the telex girl, who will then find that the line is clogged up; the ability to do it on your own, almost as if you were using the telephone.

This paper is called "Office Automation - New Horizons", so what is really new? I am saying that storage costs in relation to what we wish to put into a system will now become trivial. I did not, you notice, say that data prep was going to be trivial, and Hamish Donaldson this morning talked about that. But we are now really beginning to turn a corner. It will now be getting likely to become cheaper over the next few years to keep things up on a system - I almost said, "At last".

What about VIEWDATA? Why should I put VIEWDATA into a session like this? I could run through with you the list of the possible uses that the Post Office have put forward; they are many. But they have one characteristic: they deal in large part with the rag-bag of interests that we have never really been able to afford to put up on the system; that we have never really wanted up till now to put up on a system. I can give you all sorts of almost idiot examples. A VIEWDATA type system used internally in a large manufacturing organisation as a method of non-economically but totally vital information dissemination; about the company club meetings; the things that you like to put up on notice boards which sometimes get taken down when somebody thinks, "Well, I'm applying for that one to transfer to that one before somebody else does." This whole rag-bag of bits and pieces of information which people with current conventional ways of distributing information round large organisations use and where the likelihood of everybody who wants to see it actually seeing it is pretty remote.

But what will make VIEWDATA fly - and I think this is the critical thing - is the fact that, at long last, we have an idiot, cheap, maybe limited, but widely available terminal.

May I leave you with one famous, well-used quote about those who try to forecast the future, which is very difficult indeed? It comes from Carl Sandburg's "The People, Yes".

"The white man drew a circle in the sand and told the red man, "This is what the Indian knows." And drawing a big circle round the small one, said, "This is what the white man knows." The Indian took the stick and swept an immense ring around both circles and said, "This is where the white man and the red man know nothing."

Exhibit E1

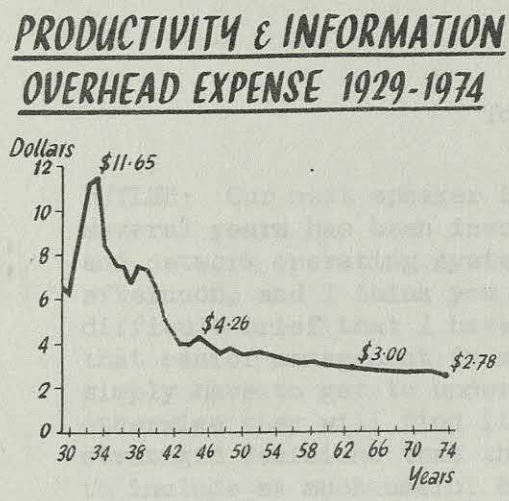

Exhibit E2

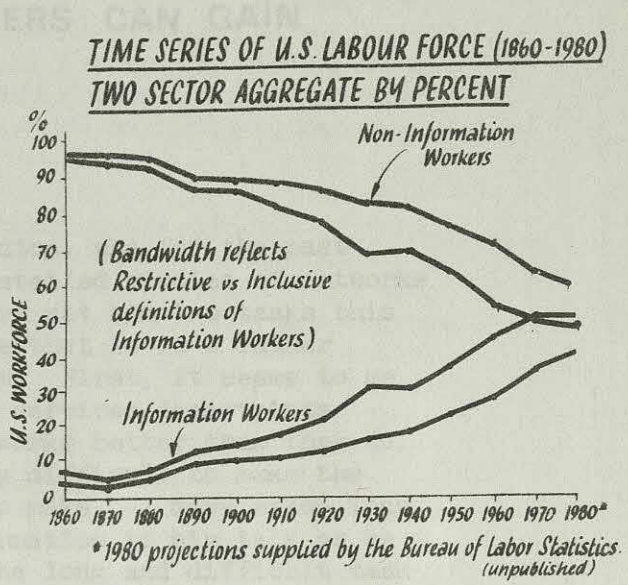

Exhibit E3

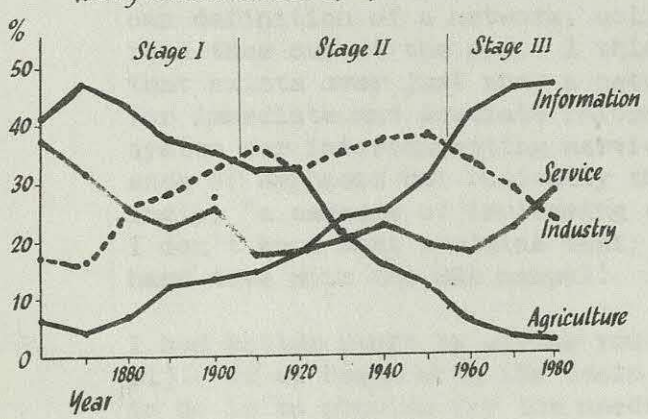

Exhibit E4

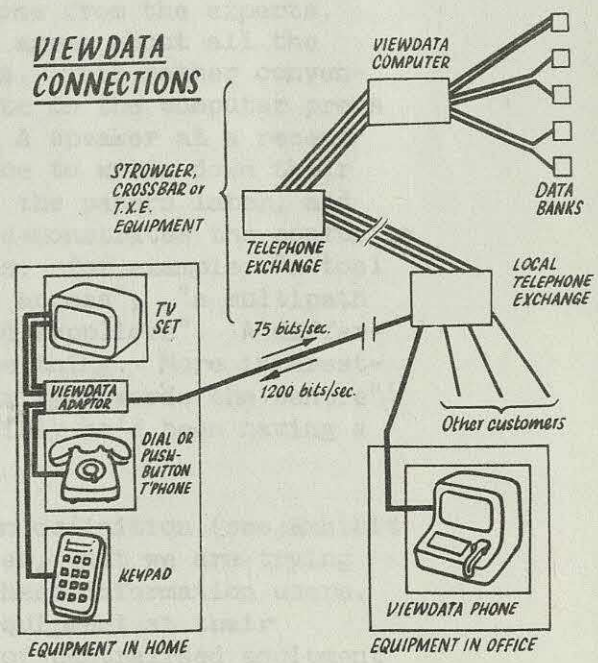

F

NETWORKS – WHAT USERS CAN GAIN

Tony Gunton

BUTLER: Our next speaker is Tony Gunton, who for the past several years has been involved in detailed studies of networks and network operating systems. I have set him two tasks this afternoon, and I think you will agree that it is a rather difficult brief that I have given him. First, it seems to me that senior management in management services departments simply have to get to understand networks better than they do, otherwise they will find it extremely difficult to make the strategic decisions that they have to make. I have asked Tony to include as much useful basic information in his talk as he can. Secondly, we are starting on the long and difficult task of sorting out exactly what is important about network operating systems for users. Neither of those subjects would be easy to tackle on their own, and I have asked Tony to begin at least a preliminary analysis of both.

GUNTON: As David has implied, networks tend to be a rather daunting topic, for a number of reasons. I think that there is a certain amount of technical mystique surrounding them still because of the stage of development. They also seem to suffer, like economics, from conflicting opinions from the experts. Two experts typically will tend not to agree about all the issues that are concerned with networks. This rather conveniently was demonstrated by a little note in the computer press last week, which you might have seen. A speaker at a recent conference apparently asked the audience to write down their own definition of a network, collected the papers later, and read them out at the end. I think it demonstrates the confusion that exists over just what a network is: for example, "a tool for immediate and accurate information access"; "a multipath system for interconnecting services and suppliers". A difference of emphasis but basically the same thing. More interestingly, "a cascade of increasing complexity towards the centre"! I don't know what explains that; possibly he'd been having a hard time with the SNA manual!

I had better start by giving you my own definition (see Exhibit F1). If we begin with the basic problem, what we are trying to do is to provide for the needs of these information users. They are likely to have a variety of equipment at their immediate disposal; a certain amount of centralised equipment probably as well, available for everybody; probably a manage-

ment services organisation with some responsibility for meeting their needs. One inadvertent omission from the equipment down here is a telephone, which of course is the most common device to meet these needs. It perhaps demonstrates my bias in the matter as a computer man rather than as a telecomms man. It is probably best that you should know about that bias.

To my mind, what the network has to do is to provide for the distribution and the management of information in this context. Distribution perhaps is a problem that has been given most attention; management perhaps not enough. So I feel that it is probably in the area of management information that the new networks that are being developed have most to offer. The PTT services that I have shown there are obviously an essential component of networks. I am not necessarily implying that all networks providing for this basic problem of handling of information consist of copper wires, optical fibres and so on, although it seems to me likely that the electronic transmission of information will become increasingly the way of handling the problem. I hope that I can demonstrate what electronic handling of information can add by way of controlling the situation.

Before I talk about what is happening in the field of networks, I should like to spend some time going over the background. It is rather commonplace to talk of the rise of data communications as a feature of computer systems and of the use of computing facilities; but I feel that the potential is perhaps such as to surprise many people.

Consider for a moment two parallel developments. On the one hand I think the need of users for more powerful and more sophisticated terminal-to-computer communications is already substantial and likely to become increasingly important. There are a number of reasons, many of which have already been reviewed today; for example, the cost of the physical transport of information and of people which can be replaced by electronic means, using networks; the inadequacy of batch methods - or perhaps I should say batch data entry methods, as Hamish Donaldson illustrated; and, paradoxically, distributed computing, because if it does not increase the volume of information that you have to transmit about the organisation, it will tend to increase the sophistication of the means needed to distribute and control the information.

So on the one hand, the need for electronic means of communication and the diversity of the forms of communication is likely to increase substantially. On the other hand, suppliers are beginning to develop standard products and services which are likely to go at least some way towards meeting these needs. I want to go into these later on.

I talk of a potential sharp increase in the use of data communications for transmission of information of all kinds. There are a number of important conditions. The first, and obvious, one is the involvement of the PTTs in the situation, their regulatory and tariff policy. Of course, that is a vast subject on its own which I do not want to go into now. I am sure you are all aware of its significance.

Secondly, it is conditional on whether network architecture, which is becoming the new orthodoxy of the computer manufacturers, actually delivers on its promises. At the moment it is very early days and there is not a great deal more than promises. I think that most of the time is still spent coping with the past rather than taking us on into the future.

Thirdly, and almost as important, it depends on to what extent usable standards can be established, because as the potential for communication increases, I feel also so does the potential for confusion and for misunderstanding. There will be far more subtle ways of getting things wrong with the sort of techniques that are beginning to be introduced. To clarify this last aspect, I want to look at some of the underlying technical issues in networking. The tool that I am discussing I am sure will be familiar. I hope that I can indicate that they are not just buzz words for those of you who have perhaps dismissed them as that, quite understandably. I hope that those who are familiar with them will bear with me.

First, protocols. It seems to me that the best way to illustrate protocols is like this (see Exhibit F2). The protocols that are shown over here, organised in a hierarchy, reflect the transformations through which the information has to go in its progress through a network. The essential purpose is for information exchange to take place. To be able to do so, it is necessary, for engineering reasons, to take it through one or more of these transformations. To do so, a sub-language is established which is appropriate for the particular medium that is being dealt with at each level.

If you feel that this organisation is rather too elaborate, that it is just another clever idea that is being imposed on us by people who are more keen to sell pieces of equipment than to solve the basic problem, I should like to give you an analogy, and that is of language communication itself, normal language communication between people. If you consider the processes that language has to go through before it becomes an agreed concept in two minds, it starts off as a sequence of noises; at various stages it will be composed of words, of phrases and so on, the construction of which is directed by syntax and a whole host of conventions in language that people really are not aware of in normal use. The point is that, without this structure, and even with it, the potential for communication

and its accuracy is very much reduced. It is clear to me that very much the same is true of communication using electronic techniques, that without this elaborate structure the opportunities for confusion are substantially magnified.

I use the analogy of language to illustrate two points. First, even if you know all the words in a language, unless you can put them together in the right way, and even above that, unless you happen to know the idioms that the people to whom you are talking understand; unless you know all these elements then, even if you know all the words, your listener will get it wrong. This is just as true of this process. Unless there is agreement on the rules at every level, all the way up to the top, the chances of accurate communication taking place are not very great.

If we look at the standards that are being developed internationally, typically they will reach this far, up to message level. When we get up here, there is not exactly chaos but certainly it is not very ordered. The point that I should like to make is that, without agreement right to the top, the chances for accurate communication are substantially reduced.

Secondly, language and the way it develops illustrates how many rules tend to be made by usage. Whether or not you have a French Academy which says, "You'll do things this way," in the end it is the way that people talk to each other and their particular needs that determine how the language develops. Special interest groups will develop their own jargon that they need for their special purposes. Again, I think that the same is true here and at the level at which this language is used, which is at the terminal or at the work station, the rules in the end have to be made by the user, or at least the basic rules have to be modified to meet the immediate needs of the users. Although broadly based standards like the CCITT standards and the sort of standards that the computer manufacturers promulgate are invaluable as a foundation, without this additional level of understanding as to the rules I feel that the chances of networks meeting their potential are that much reduced.

If you want an example of the way it can work, the airlines SITA network reflects the needs of people in a common interest group who have evolved their own standards to meet their own particular needs. It is a network that has been operating very successfully for some considerable time.

This protocol structure also leads to another valuable quality of modern networks, and that is so-called functional layering. This is not new. You can choose your term: functional modularity or whatever you like to describe it; but I think that it is beginning to have a formal recognition that is long overdue

in modern network structures. Essentially, each of these layers which corresponds to the levels of protocol already has its own sub-language to operate with. They are designed in such a way that they are functionally independent, one from another; each layer then supports and provides a service for the ones above it. What this should mean is that, as technology changes, the way these layers are mapped on to the hardware can be changed without affecting the protocol. This means that the user is left with a constant environment in which to develop his applications; that the changes are absorbed in this structure and are not apparent at the user level at all.

Another concept with a very similar objective is the so-called virtual terminal. It could be shown something like this (see Exhibit F3). The theory is that all applications here see one or a number of virtual terminals. These are terminals with an arbitrarily defined set of characteristics that would meet certain specific purposes. The applications will only ever see these virtual terminals and out in the network, conversion will take place to meet the local conventions of the actual physical devices that are being used. The objective is to buffer the applications from the changes in technology at the terminal level so that these can be carried over. The constant in the equation is the virtual terminal protocol.

That is all very tidy and elegant and perhaps you would be forgiven for believing that it was not much more than an elegant concept. I am not suggesting that it has limitless power to resolve imcompatibility between terminals and the devices that they are addressing, but if it is only a partial solution to that problem which tends to bulk very large in the problems of many data communications users, it is valuable enough. The essential point that I want to make with all these concepts that I have outlined is that they are all concerned with providing the user with freedom from considerations which are not really relevant to the basic problem. The basic problem is meeting the needs of information users. If you spend 90% of your time wondering whether terminal brand Y can be installed without a 50% re-write, then obviously your intention is diverted from what is important in the situation.

If we move on from the theory to what has been achieved in practice, the first thing that strikes one when looking at the practical developments in networking is that there is not a great deal happening that is entirely new. The SITA network that I mentioned really has most of the features that are commonly associated with the packet switching networks like ARPANET which are mostly in the news; and also some of the American time sharing bureau networks similarly have all these kinds of features, alternate routing and various reliability features. What I think has happened is that what in these cases are really application specific systems have been made

into a general purpose scheme, a scheme that can be applied generally to the solution of communications problems rather than specifically to airlines or a time sharing bureau or whatever. The sort of techniques that made that possible are the ones that I have just been discussing.

The question is how much more will a general purpose solution of that type cost than the specific ones, because I think they will always cost more in the end. The new networks, particularly in the States, are starting to demonstrate the commercial viability of these techniques, but mainly in the public domain. Whether the private user can actually afford them as well is a very different question, an open question still. Probably the best local model that will be available very shortly is the British Steel network. It is probably fair to say that the particular combination of circumstances that enabled British Steel to adopt this solution will never occur again for anybody else. It should also be noted that they are pioneers and will have all the costs associated with being pioneers. Taking these two together, it is likely to be a very valuable demonstration of the basic viability of these techniques within a private organisation.

But what seems to me most interesting about the British Steel network is not so much whether or not their network will prove economically viable compared with the conventional data communication techniques that it replaces, but the new possibilities that it will open up. There are new control facilities available to the management services organisation running the network. It seems to me that the number of situations in which British Steel management services, when a user says, "Can I have that?", can say "Yes, we can quite easily do it," will be that much greater. It is an enormously powerful tool which, once the basic problems of communication of information are solved, will enable a very flexible view to be taken about what solutions are used for the computing problems with which everybody is familiar.

If I can generalise about the changes that are taking place in the structure of networks, there are probably three significant changes. First, a network has become an autonomous system in its own right, with its own dignity, if you like, and independent of the devices that it supports. If we return to the network diagram that we started with (Exhibit F1), that seems to me entirely natural. Just because these are central and these are distributed does not seem to me to be a good reason why the controls should all be exercised from here. The only place to exercise the control and management of information, from a philosophical and logical point of view, is here in the network. It seems overdue that that realisation should come upon the DP community.

The changes in the structure of the network are motivated by three needs (see Exhibit F4). First, reliability in all its aspects including availability, security and so on. The presence of an intelligent network means that these sort of problems can be handled in a much more flexible manner. Secondly, freedom of access. I think this is a crucial factor. What has been brought into computer systems is the same sort of flexibility that is taken for granted in the voice systems that everbody uses every day. It does not seem unreasonable to expect the same sort of flexibility in an information system, whatever medium it happens to be using. It means that the availability of the services to disperse divisions and so on is much less location dependent. It means that if you choose to have your data centre specialise, as British Steel has, that the mechanisms for doing so are already there.

Thirdly, a better use of resources. The new techniques enable a greater mix of traffic to be carried. They allow flexibility in timing. For example, you can use the intelligence of a free-standing network to spool your bulk data overnight. I have no doubt that everybody has the same problem, that their systems development people want to put in all their program amendments at four o'clock in the afternoon, and then expect all the results there at 9.30 in the morning. If everybody packs up and goes home overnight, that is a very difficult requirement to meet; whereas an intelligent network, with the capability to do these things on its own without other than minimal supervision, removes the difficulty.

Finally, it is no longer a transmission system alone, it has taken on a whole range of new functions. It has become a means of providing a service. If you want a rustic analogy, it has stopped being a railway and it has now become a travel service. I can quote to you from an IBM source about systems network architecture:

> "It's a unified network structure to develop and support user requests for services."

That seems to me to be a key definition of what the network should be.

Before I go on to look at one or two of the computer manufacturers' products who have really been the traditional suppliers of a complete travel service for data, I think it is useful to go over some of the background. As far as the computer manufacturers are concerned, the increased effectiveness of the PTTs in setting standards and their obvious ambition to extend the scope of services that they provide for data, and also the ability of these new networks to resolve incompatibilities, seems to me to pose a serious threat to the total systems ability of computer manufacturers. Traditionally, they have

advanced the argument that if they were to invest in providing a total systems capability to their users, then it was legitimate for them to try to protect what they were installing against the attentions of predatory small suppliers who were very keen to pinch the lucrative bits and leave them supporting the rather dull and not so profitable sections. Whatever you feel about that argument and its merits, it seems to me important to understand how the computer manufacturers must view developments which potentially threaten to break down the defences that they have erected around their systems.

The point is that the network products that they are bringing out are rather ambivalent, because on the one hand they are a continuation of this defence against the outsiders who are breaking into their systems, and on the other hand they are partly innovatory and I think they have to be understood as such. Systems network architecture illustrates this very well because it promises substantial new freedoms, but at present it delivers very little over and above what has always been available. It is very much in the tradition of mainframe dominated communication systems (see Exhibit F5).

Of course, IBM is in a position where it has to provide some continuity for its users, and the present release of SNA is not the end of the story; we will have to wait on IBM for further instalments. But if I could return to the point that I made earlier, a lot of what is going on in networks is about constructing new freedoms, which in one sense is against the interests of the computer manufacturers to supply within their own systems, and I think that this conflict cannot be avoided.

It is interesting to compare the approach in SNA with the approach of the giant in the minicomputer field, DEC. If you want to characterise SNA, the control is still centralised; some of the tasks are distributed to the extremities of the network, but very much under close supervision of the mainframe system. Essentially it is about terminal networks. These terminals are intelligent but their intelligence is limited.

DECNET, on the other hand, distributes the control to the components of the system - the computers that form the network. Essentially, it provides for computer networks where the units that make up the network have stand-alone capability. They are autonomous; they do not depend, as these terminals do, on the mainframe for various essential services. DECNET has also begun to approach the problem of high level protocols that I touched on earlier. The protocols are designed to enable any component in the network to share the data, to share the files of the other components connected into the network. It is only a beginning, but it is an extension of the language that is available to the user, just one step higher; so that he does

not need to write a program to access data on another system, he has a command which does it for him.

I have gone over these because I feel that they are really only beginnings, but it is important to understand what they offer, because these products are necessary to judge the direction that the computer manufacturers are likely to take in the new situation. They are the important first indicators of what the computer manufacturers are likely to do.

Following that whirlwind tour of networking, I should like to try to summarise the conclusions that I think can be drawn from what is going on (see Exhibit F6). First, I have no doubt that networks of the type that I have been discussing, which primarily have been data communications networks, will also carry other forms of information - text information, electronic mail and so on, the applications in the office automation field that have been discussed. They have the capability and the flexibility. They have all the basic features necessary. So I feel that, first, networks can be seen as a step towards introducing coherence into the situation. I think that we have to face the interchangeability of all forms of information; and a lot of effort is going into trying to find some compromise between the digital forms of data and analogue forms used for speech, as Karl Kozarsky was describing. But it seems to me that the problems go far beyond that, and there is no logical reason why a distinction should be made between any of the forms of information that we have been talking about this afternoon, as far as distribution and management is concerned. That is at a conceptual level. Granted there are substantial engineering problems, but at the strategic level it seems to me that if information systems are not constructed to provide some degree of coherence, then the potential that they offer will be largely wasted.

Secondly, it enables a much better service to be provided to information users. It combines the flexibility of voice telecommunications with the higher speeds and the other characteristics of digital communications. It seems to me that information users have every right to expect just the same sort of flexibility that they get from their voice switching systems and from data systems.

Thirdly, it provides greater flexibility in the deployment of resources. There are limits to this, of course, but an intelligent network essentially is able to buffer incompatibility to the extent that it is possible to make it invisible to the user.

Finally - and the importance of this point will vary depending on the particular organisation - while I agree to a very large extent with what Hamish Donaldson was saying earlier, it seems to me that one point that was overlooked is that perhaps not

all management services organisations are as fortunate as his, in that I have no doubt that when he gives the users a minicomputer then they do with it what he wants them to do with it, and not a great deal more. If that is not necessarily the case, then the introduction of minicomputers all over an organisation is a potential recipe for chaos - perhaps not tomorrow, perhaps not next year, but in 10 years when, for some reason, you want to integrate systems that currently it is fashionable to distribute, then you will begin to understand the problems. Networks are a medium for control and standardisation in that sort of situation.

I do not say this in any authoritarian sense necessarily that they are going to dictate the rules, just that you have a medium here; it is a basic essential discipline; there are some rules that you have to obey to get access to the benefits. That seems to me, in some circumstances, a very important weapon to have at your disposal. I think that at the present time only the bold or the very fortunate will venture into networks - fortunate in that circumstances conspire to make it convenient for them. But it seems that in a very short period of time standard products will be available which will remove a lot of the risk, and that the reducing cost of hardware will take a lot of the cost out of it as well. So it seems to me that these sort of benefits, which I am convinced can easily be made available if the right attitudes are adopted, are likely to be available in the very short term. I should have thought that they would be of great enough significance for most management services organisations to demand attention.

Perhaps I can finish by reminding you of the important conditions that I feel exist before the potential of networks can be realised. These are conditions which apply to users, to management services organisations. First, it is a very diffuse market in which many different suppliers are competing and nibbling away at different parts of it. I feel also, as I have tried to explain, that there are historical limits on the traditional suppliers of data communications services; which means that I think the overall solutions to the problems involved are unlikely to be supplier led. I feel that unless users are able to bring forward their requirements, the overall solutions that are needed are unlikely to evolve. For that, the user must understand the concepts involved; he must understand the possible applications; and he must understand the important issues involved in networking.

Secondly, the communication, using networks or whatever, will only be as precise as the rules governing the language that is used for that purpose. I feel that the user community in general has an essential part to play in the formulation of those rules. For many people the present situation must not be a satisfactory one, where resolution of incompatibility will

tend to bulk as large in the solving of data communications problems as the fundamental user problem with which you are trying to deal.

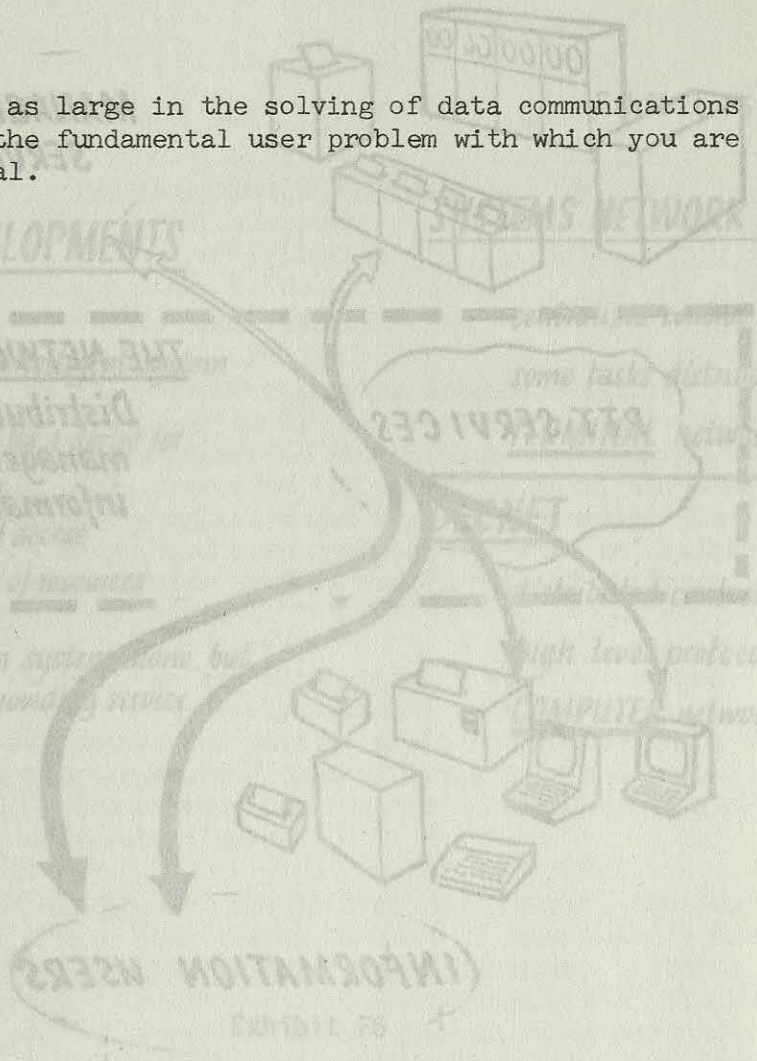

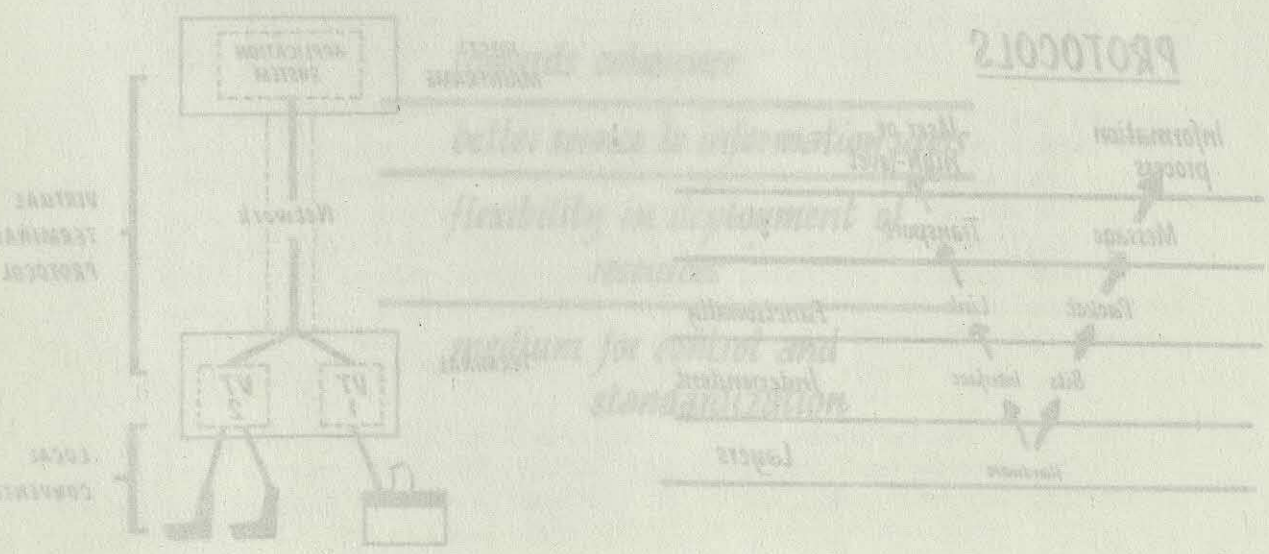

Exhibit F1

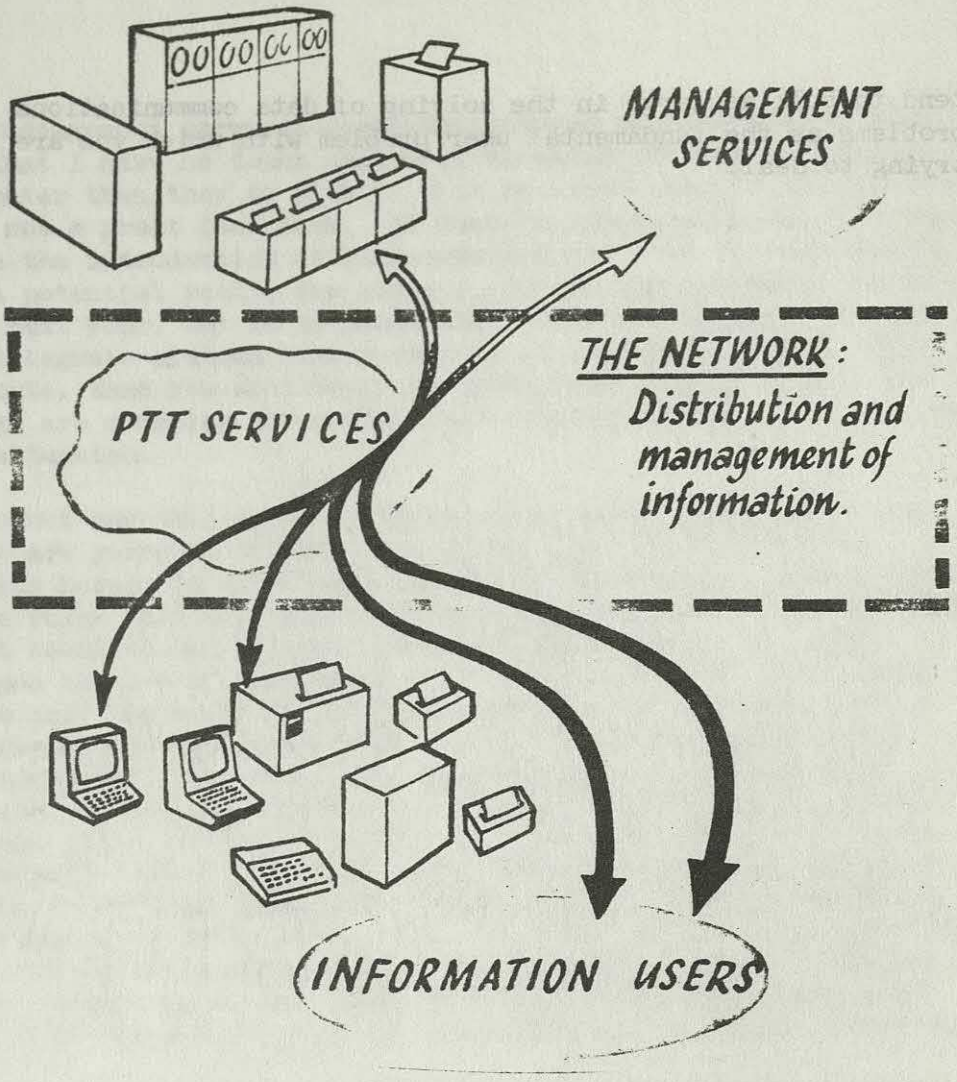

Exhibit F2

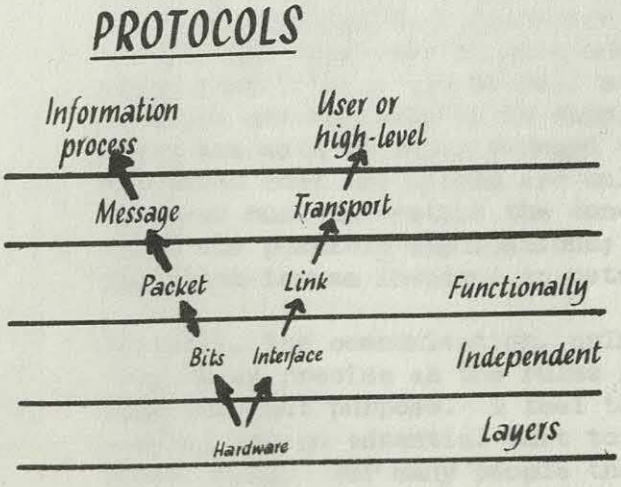

Exhibit F3

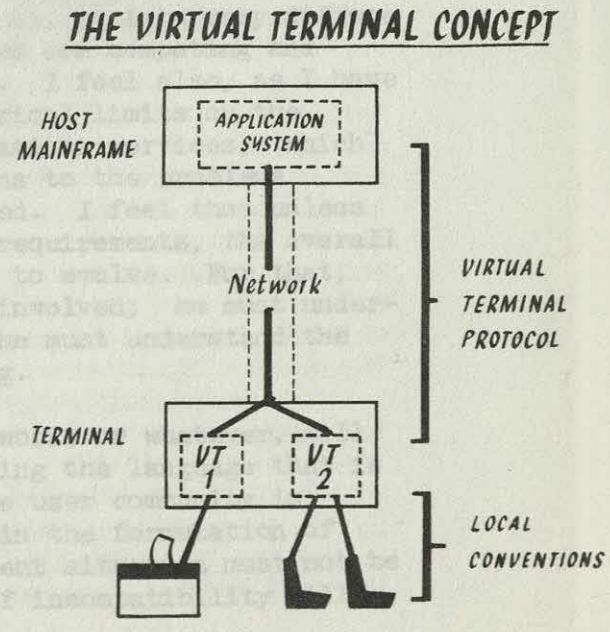

Exhibit F4

NETWORK DEVELOPMENTS

1. An autonomous, intelligent system
2. Change motivated by a desire for
   - reliability
   - freedom of access
   - better use of resources
3. Not a transmission system alone, but a means of providing service.

Exhibit F5

SYSTEMS NETWORK ARCHITECTURE

centralized control
some tasks distributed
TERMINAL networks

DECNET

distributed control
high level protocols
COMPUTER networks

Exhibit F6

THE POTENTIAL BENEFITS

towards coherence

---

better service to information users

---

flexibility in deployment of resources

---

medium for control and standardization

G

CURRENT APPLICATIONS OF WORD PROCESSING

Roger Woolfe

WOOLFE: The structure of my talk this afternoon will be in four sections. First, why we are interested in word processing. Then I will look at the benefits of automatic typewriters, setting the scene for the third section of my talk, which will be on more advanced, computer-based systems. Finally, I will attempt to draw some conclusions and perhaps some guidelines for potential implementers of such computer-based word processing systems.

Let us start with why we are interested in word processing. I think that we know, intuitively perhaps, that the office is a place of low efficiency and waste; maybe not yours, but I am sure you know that many are. Capital investment and productivity are low, and equipment tends to be a conglomeration of non-standard, stand-alone units.

Just by way of background (see Exhibit G1): first, the cost and volume of office text communication. I have heard a number of figures quoted for the cost of text production: 1,000 characters cost $2; a one-page letter typically $5. Figures vary from company to company, but in general those will not be too far out. As for volume, in the US I have seen some figures that indicate that during the period 1960 to 1970 text grew some 3% per annum. I believe that the figures were much the same in Europe. Coming up to date, I suspect that in the UK at least there is a temporary hold off in the growth of office text communication as a natural consequence of the economic climate. But equally, I suspect that if the UK emerges from this position, volume growth will take off again.

Secondly, labour content, labour cost and productivity. Labour content is perhaps some 70% of office administration costs, as Dennis Holloway indicated this morning. We all know that labour costs are growing. Next, productivity: in the US between 1960 and 1970, a recent SRI report indicated that production worker productivity grew some 83%; during the same period, office productivity some 4%. I saw some equivalent figures for West Germany a few weeks ago. They were not quite so dramatic but they were in that kind of field.

Then, low per capita investment. US figures again: for the farm worker, $35,000 capital invested per farm worker; for the factory worker in the US, about $24,000; for the typical typist in the office, some $2,000.

So managers are searching to find a solution; office managers, occasionally ADP managers, specialists in this field. They know that the overall extent of administration costs in their businesses may be as high as 15% or 20%, as David Butler was saying this morning. They know that, with increased competition, they have to have faster, more accurate response; so it is not unnatural that vendors at the moment are applying pressure to managers to adopt a solution.

The main focus of word processing attention in the last few years has been automatic typewritters. At the moment in Europe I think there are about 60,000 installed. Automatic typewriters typically consist of a heavy-duty typewriter with a recording medium such as magnetic cards, paper tape, magnetic tape or floppy disks on the more advanced units. They have been used in a variety of different installations.

What benefits have they brought? The main one that is quoted by suppliers of this equipment is an improvement in typist productivity (see Exhibit G2). This is a justification quite contrary to what Dennis was talking about this morning. In the past, vendors have been after real, measurable benefits; and what they say, typically, is that the actual average typing rate in a typical office today is something like 4,700 words per day, and that is true of Europe as a whole. But they quote a fairly arbitrary, achievable standard of something double that, like 8,500 per day. They have been able to quote that to get from actual to standard requires more efficient, more effective loading of regular, ordinary typewriters. But to get from a standard level of some 8,500 words up to a theoretical level of 20,000 which corresponds with a typist's free typing rate of something like 300 characters a minute, requires a change in technology.

Automatic typewriters have been effective in many installations in getting the standard output rate moved a little bit up towards 20,000. Success has varied, but there are only a few areas where 20,000 has actually been reached. I want to come back to that when we look at computer-based word processing systems.

That is just the first of the benefit opportunities. The second one that vendors tend to quote a good deal is improving the originator's productivity; by originator, I mean the executive, the producer of the original text. By that I mean reduced need for proof checking; faster turn-around; and faster text origination in the first place, because of the stimulus of having high speed support in a word processing area.

The third one is improving the final text product, in two main ways: quality and delivery. First, quality. The appearance of the output document can be improved: top copies only;

uniformity; fewer errors. These are all widely quoted by suppliers of this kind of equipment. As for content, there is some evidence to support the view that text originators feel less inhibited about making changes to text that is produced on automatic typewriters, and consequently they will make changes which otherwise they would not have done and the effect can be a better content in the final piece of text. Delivery; that naturally arises from the higher throughput rate of automatic typewriters. Finally, improving job satisfaction. This is a contentious area and one to which I will come back later on.

So these benefit opportunities have led to justifications in many instances, the prime justification being cost displacement (see Exhibit G3). Typically, vendors will quote a cost per automatic typewriter of some $3,000 to $4,000 per annum, but benefits in the range from $3,000 to as high as $20,000 per annum; that is with increased productivity from girls whose cost can be anything from $4,000 to $9,000 per annum each, depending on which part of Europe you come from.

But such justification is not quite as simple as it may seem at first sight, because there are problems. The first one is the big investment jump, typically from a few hundred dollars per typewriter to the low thousands for automatic typewriters. That takes a bit of wearing. Secondly, the risk involved with using equipment which is new and untried, which may involve vendors who are suspect for one reason or another, may cause people problems amongst the users, and problems amongst executives who have to be educated to make use of it and to understand what its capabilities are.

Thirdly, training. Operator training in the past has been a big problem. I have heard of automatic typewriter installations where the operators have not come up to full output until after a training period - an on-the-job training period admittedly - which has lasted three, six, nine, even 12 months. Training manuals can be very long and very difficult to understand; and the training support from certain vendors has been pretty poor in the past, although that is being corrected today.

But the main problem, in my view, is the need to apply these machines very intensively in order to get a pay-back. This need for intensive use has resulted in the establishment of typing pools, where loading can be controlled; but this has led to difficulties of effective management and also the problems of specialisation. In many cases, it has led to the identification of two types of typing staff: correspondence secretaries and administrative secretaries. Correspondence secretaries are those girls who work on the automatic typewriters in pool situations, and very often they find that they have dissatisfactions of drudgery, impersonality, non-identity with the work; they find themselves anonymous and in a remote

area. However, I have seen installations where correspondence secretaries have found satisfactions; satisfactions of the complexities of the equipment mastered; a showcase working environment; being members of a team; producing work that is higher in quality and quantity than before.

In the case of administrative secretaries, the problem that I have come up against frequently is that of genuine tasks for them to do. It seems that the days when administrative secretaries will grow into executives in their own right are not so close as the pundits would have us think. I suspect that the "liberation" that we might read about in some of the literature is not so much a result of the right word processing equipment, but more a result of the right management.

More recently, there has been a move away from correspondence secretaries in a pool, to the establishment of smaller clusters or puddles of secretaries engaged in this kind of fairly intensive work, in order to get a pay-back from the equipment.

I will briefly look at some case studies (see Exhibit G4) of automatic typewriters, point by point, in order to lead into computer-based word processing equipment. In the number of case studies with which I have been involved in the US and in Europe, the first point that I would bring out is that the motivation for using word processing has been increased throughput and reduced costs. That might seem obvious to you, but it is a little different from what Dennis Holloway was identifying this morning.

One US bank, for example, quoted to me that the only way it could increase its throughput was by making use of this word processing equipment. Secondly, fairly conventional functions; for example, law firms tend to produce deeds, making use of paragraph selection and revision capability on their equipment. Utilities widely use standard letters. Accounting firms have made use of the full editing capabilities of this kind of equipment on long documents. Hospitals turn out medical reports which require very fast turn-around. Commercial companies use this kind of equipment for a variety of purposes.

In general, the benefits that they have achieved are hard to evaluate. I have been surprised by the relatively small number of companies which are able to give me firm information on the real cost/benefits that they have achieved. However, that is not the case all the time, and some companies are able to quote that productivity has been improved by 25% or 50%. A company quoted to me that their turn-around had dropped from 48 hours to 4½ hours; this was in a big manufacturing company. They claim reduced errors. One or two companies even claimed reduced labour turnover, which is interesting; and better executive support - very hard to substantiate, but that has

been claimed in a number of cases. Interestingly, in general, companies which have made use of automatic typewriters seem to me to be developing with caution; by developing, I mean that they are looking ahead to further investment in this kind of equipment. But they are being more careful about equipment tests; about staff selection; about training; and finally, about planning and work flow.

Having painted the background to the needs for word processing and a brief look at automatic typewriters, where they are used, and why, I plan now to turn my attention to computer-based word processing equipment. I have four categories here (see Exhibit G5). First, stand-alone equipment, equivalent to the personal computer that we touched on this morning. Typically, a piece of stand-alone equipment will have a display unit; a control unit; perhaps a floppy disk store; some communications capability; a keyboard; and a printer. Shared logic systems are aimed at larger companies where the price per work station can be reduced simply by sharing a mini processor, and generally cartridge disk storage which tends to be the more expensive part of this kind of equipment. Stand-alone devices cost typically around $7,000 to $8,000 per annum at the moment; shared logic equipment, depending on the number of work stations that are hooked up to the central processor, can fall down in the range of $5,000, $6,000 or $7,000 per annum per work station.

The third category is the time sharing kind of configuration where we have work stations attached to a remote time sharing computer. There are one or two services of that sort available in London, a few on the Continent, and quite a number in the US at the moment. One bank to which I was talking the other day, which had been making use of this kind of time sharing service, claimed that it was a very good way of introducing computer-based word processing equipment into the office area, minimising the original investment and the risk.

Finally, the mainframe approach used by some big companies in this country, and certainly in the US, where a big, central, in-house mainframe is supplied with the necessary software, terminals and so on, and quality printers, in order to look after text processing.

All of these four types of equipment have in common the following features: some kind of logical control facility; nearly always a display device, either large or small, which I will come on to later; storage, which can vary from magnetic cards to floppy disks, up to cartridge disks and large disks on time sharing and mainframe computers. Obviously they have printing capability, and they must have the ability to produce high quality output print via the typewriter. Finally, nearly all of them have some kind of communication capability.

What is the justification for this more expensive computer equipment? The justification that is being pushed hardest at the moment is further improved typist productivity. Now we have talked briefly about improving typist productivity from an actual rate of about 4,700 words per day to a standard rate. That can be achieved either with automatic typewriters or even without them in a pool environment, with careful attention to loading. What we want to try to do is to get from the standard output rate of some 8,500 words per day up to a theoretical rate of more like 20,000.

Here, (see Exhibit G6) I show a breakdown of a typical company's work mix in a typing pool. You will see that in the column headed "Standard" something like 30% of the time is idle; 15% to 20% of the time retyping; error correction takes 20% to 25%; paper handling 5% to 10%; new words, only some 20% to 25%. But targets can be set at this sort of level as shown alongside, so what we are after is a new word figure of something like 55%, an improvement of some 2 to 3 to 1 on the standard output rate, which can be achieved on regular typewriters or fairly simple automatics.

By new words, what I mean is words that have not been keyed in ever before; words that are not part of a standard letter; perhaps words in a new report, something that is completely new. Take an example of a new report; you may find in your office that a complex report has to be retyped three or four times, because the author makes corrections to it. It could be that when the author reviews the draft the first time, he wants to make changes to only 10% of the words, though if he were using a regular typewriter, every single word might have to be retyped because of the major changes to the format of the text. Only 10% of the words are genuine new words; the other ones have already been keyed on the previous occasion.

Computer-based word processing equipment can show further productivity gains, but it is expensive; it is more expensive than automatic typewriters, therefore there is a continued need for heavy loading in order to justify the equipment. That gets us back to the problems that we already have with automatic typewriters. So the vendors are tending to try to do two things: first, cut the prices; and secondly, they are trying to extend into new areas.

Here (see Exhibit G7) is a quick look at some of the new areas that are being looked at at the moment, and which we have discussed briefly this morning. First, archiving; new types of storage devices allow the opportunity to store information much more cheaply than in the past. Retrieval systems are becoming available, both indexing retrieval systems and also inverted file type retrieval systems, which make the information, once stored, much easier to get hold of. There remains the

problem of incoming mail, to which Dennis Holloway referred briefly this morning. In most cases it is not worthwhile to re-key incoming mail because it is too expensive. There are various ways around this. One is to use OCR, but that is limited in terms of its font handling capability at the moment. It is also quite expensive. The second way is to keyboard in an abstract of the information, and to put it under an identity number which could be retrieved at a later date. But the method that I put my money on is digital scanning which is becoming more and more efficient. In my opinion, digital scanning with some kind of an identity number attached to each document will be an area of some importance in the future, particularly as that will be linked in with facsimile technology.

Secondly, printing. It is interesting the proportion of typed work that finds its way into printing in businesses today. In fact, not so long ago a figure was quoted to me that is really quite surprising, and that is that the amount of printing that is done in businesses, for business purposes, in the form of sales brochures, catalogues and other forms of literature, exceeds the amount of printed material available to the public in the form of newspapers, magazines and books, by several times. It is clear that it is inefficient to re-key information that needs to be printed. Typically, a business at the moment will put its printing work either into an in-house printing shop, using a photosetting machine, or else it will sub-contract the work to a photosetting sub-contractor. But by making the link between the word processor and the photosetter, re-keying can be cut out.

This link has been made in a number of cases in the US. I understand that something like 8% of information that is keyed into word processing machines in the US is passed across, either by magnetic tape or by paper tape to photosetters which are geared to take the information. There are benefits of an improved point size range on printed information, a better type style range, the ability to use proportional spacing, graphic art quality and so on, but there are difficulties in embedding the necessary control characters into the transfer tape. It is not a thing that has been done very much in the UK. What little research I have done on it indicates that only 3% or 4% of the information that is keyboarded into word processing equipment goes automatically across into a photo-typesetter. But I think it will be an area of growth in the future.

The next point is calendar and follow up. Clearly, given the kind of equipment that a computer-based word processing device contains, it will not be too long before calendar information showing name, date and times will be incorporated on the system; follow up files perhaps with the ability to sort information by name, subject and due date could also be installed on the system, and displayed through the screen.

Electronic mail we looked at briefly this morning; and all of this leads to the multifunction work station, which is more than a gleam in the eye of several vendors at the moment, who are pushing the idea that "Data is a key company resource". On that topic, a figure which many of you may know indicates that something like 30% of the information in an average business is numeric, but 70% is textual.

What is it that will allow these extensions and what is actually happening on the scene of computer-based word processing equipment today? Coming back to the key components that we looked at earlier, that is logic/control; display; the keyboard itself; the printer; and storage, let us look at the situation as it stands and what the trends in the next year or two might be.

Here (see Exhibit G8) is a list of some of the capabilities of the equipment. On the editing side, there is the ability to insert, delete, alter, at the character, word or sentence level; to transfer, duplicate, and merge information (for example, to take a file of names and addresses and to automatically type out standard letters). Then the ability to search for words or sentences throughout a fairly lengthy text. Then, the ability to look after text format; by that I mean tabs, margins, heading centering, column alignment for tabular information or numeric information with decimal points. Then hyphenation, which is semi-automatic on some equipment, and even automatic with dictionary look-ups on some more advanced equipment. And justification, both left and right-hand justification, right-hand justification with variable inter-word gap lengths. Finally the ability to look after paragraph starts, pagination, footnotes, positions of names and addresses, and so on.

Additional functions that are becoming available include accounting; simple invoicing; and the ability to do simple calculations on some of this kind of equipment. The penultimate heading on this chart is "maintenance". I was going to draw the distinction between RAM and ROM systems and the benefits and disadvantages of these two approaches on software maintenance, but I think that time will prevent me from going into that in any detail. Finally, on the cost front, many small computer-based word processing pieces of equipment at the moment are using microprocessors, which are available in the $20 or $30 range. This indicates that there is not a great deal of further meaningful cost reduction to be made in that area. I think that the importance of this will be reflected as we go towards 1980, with a move away from shared logic systems more to stand-alone systems.

The second feature that I want to look at is display. Here (see Exhibit G9) I have listed a number of factors. The first is window size. Some computer-based word processing equipments

use full page size displays; the VYDEC for example, which has a 64 x 96 display; the REDACTOR 11 which has about 60 x 84. More typical, though, is a half-page size display, like WORD-PLEX with 24 lines x 80 characters. IBM's new entry, the Office System 6, has a much smaller window size, about 6 x 85. I have little doubt that IBM will extend that size later. They have a very good sales pitch which tells you why it is not important to have a big display size at the moment, but I think they will change as the months go by. With a window size that is smaller than full page size you would need both horizontal and vertical scrolling on the screen.

Character formation. Here I mean the range of character sizes that are available. Character shape is very important. Some display screens are able to display the same shape and the same type style as the printer is handling at the time; others are not able to do that. Definition is important here, and character spacing too. Typically you can get 10 or 12 inch pitches equivalent to typewriters.

Screen prompt is most important. Usually, one or two screen lines are devoted entirely to prompt in order to help lead an operator through the task; in order to help her know where she is, what line number she is on, what paragraph number, perhaps what page number, and where she is going.

The response time is most important. This depends on the display technology in use, and also the speed of the line connector to the central processing unit, which may be divorced from the equipment.

I think that fatigue-reducing factors will come under increasing attention in the near future. There is a big difference between an operator staring into the word processing screen all day long, and an operator who is using a VDU simply to retrieve information from a computer that is displaying information perhaps about airline seat reservations. Not nearly enough research has been done in this area so far. We need to think in terms of screen brightness and position relative to the keyboard, character positioning, colours, flicker-free screens, the problems of ghosting with text movement and so on.

Finally, the cost of displays. Virtually all the displays available at the moment are CRT type displays, which typically cost something like a dollar per character. In my opinion, these prices will come down, primarily due to increased output volume; and then in the early '80s, because of competition from other technologies such as plasma discharge and liquid crystal displays. I also expect to see limited growth in multi-colour displays and displays with graphic capability.

Moving on to the keyboard area (see Exhibit G10), most new keyboards are electronic, and they are going to stay that way. It is interesting to speculate about what changes in keyboard layout might arise in the next few years. We are all standardised on the well-known QWERTY keyboard at the moment which has its disadvantages, but the one big advantage that everybody knows it. Now that we have electronic keyboards, I see no reason why more ergonomic layouts might not arise; keyboards do not have to be rectangular in shape. The PCD MALTRON keyboard, which you might have seen photographed in the Press recently, could be an indicator of the way that keyboard layouts will change to suit finger lengths, roll-over rates and so on in the next few years, while retaining the basic QWERTY format. Keyboard costs I think will reduce just slightly, perhaps 20% between now and 1980. That is in fixed currency terms.

As for the printer, we are obviously concerned with speed and quality. I don't want to go into the ramifications of daisy wheel printers, and whether they should run at 45 or 55 characters per second; but clearly it is most important to a potential buyer to ensure that the speed and quality of the print out is correct, and that the font and form range available on the printer is right. As for pitch and proportional spacing, I expect to see more equipment that is able to handle proportional spacing. After all, the daisy wheel is generally driven by a stepping motor which goes in one-one hundred and twentieth of an inch steps, and is therefore theoretically able to handle proportional spacing without much trouble.

In the area of impact versus non-impact printers, I see the impact printer surviving for some time, but there is no doubt that non-impact printers, particularly ink jet printers, are making great strides at the moment, and the resolutions of ink jet printers are improving almost daily. I suspect that, as prices drop, within three or four years there will be a range of non-impact printers available on the market, some of which will be programmable and able to change their type style ranges quite considerably. One impact of that will be perhaps a reduced need to use photosetting equipment in future. Printer costs will drop as volumes and competition increase.

As for storage (see Exhibit G11), we have heard already, and I do not want to labour this, how storage costs are dropping. This simple chart shows metal oxide semiconductor storage is now turning out at around 5 cents per thousand characters per month. Floppy disk storage is well below that, around the 3 cents level. Medium disks can store information at less that one-tenth of a cent per thousand characters per month. But the opportunity that is available from technologies in the future will make storage even cheaper. As Rex was saying, if we have a device that can store information at the cost of

something like 10^{-4} cents per bit, that means that you could store some thousand pages of typed A4 on line for about 4 cents a month, or the whole of the Bible for about 40 cents per month.

Another quick look at the way promising storage technologies hold opportunities for us in the future (see Exhibit G12). This chart shows the cost in cents per bit against access time. I think that we need to be looking more at the right-hand side of the chart rather than the left-hand side. Again, Rex was talking about bubble memories; this shows one forecast of the way, during the period 1975 through 1985, the cost in cents per bit of bubble storage is likely to drop. Compare that with disk technology which we have at the moment. Floppy disks come in at about the top of the range, medium disks at about here at the moment. With laser and holographic type memories, with a capability of holding very large quantities of information, perhaps 10^{12} or even 10^{15} bits of information, I expect access times to be less speedy, but costs down here somewhere in the 10^{-4} cents per bit range in the not too distant future.

That is a very quick run through the trends that I see in the major parts of computer-based word processing configurations. Now on to the directions and what this means to the potential user (see Exhibit G13). First, I think that today users are more cautious and sophisticated about their investment plans for word processing. I think that in the short term only centralised equipment is likely to remain justified, except for a few exceptional instances. So I see the shared logic kind of configuration being justified in companies which are large enough to make use of a number of terminals. I think that will continue to be the scene for a year or two, but then, because of the changes in cost at which we have looked already, and the increase in costs of labour, I think we will see a new era where decentralised, stand-alone units will be justified at relatively low loading rates.

Users will be more concerned about cost effectiveness measurement and simplicity, the need to have equipment that is very simple to use; prompt display is most important here. They will look for much better service and technical support from suppliers. They will be more careful about planning, education and training. It will be a vendor-led market growth area, with IBM and Xerox leading in the vanguard.

Finally, a chart (see Exhibit G14) on guidelines for establishing or developing a word processing investment. First, make sure that you establish the real aims; work out what you are trying to do before you invest the money and attempt to do it. Secondly, gain familiarity with the range of equipment that is available at the moment; and try to use your crystal ball to decide what will come up in the near future. It is an area that is changing

very rapidly. Thirdly, study your work mix. Everybody is different; I could quote you figures which indicate average work mixes in average companies, but that may not match your own case.

Fourthly, you would be surprised at what benefits you can obtain by just streamlining the existing methods before investing in new kinds of equipment. Fifthly, consider all the costs. Be warned that there are a number of hidden costs in the area of automatic word processing. I suggest you limit the pay-back period to something really short, such as two or three years, such is the high rate of change in this area, so if you are going to do a justification, do it over a short period. Manage the organisational consequences. There will be a number of them; they may surprise you. In many cases, it can be beneficial to limit the extent of the involvement and prescribe specific applications for word processing. Do not necessarily create an area which will handle all the text jobs in the business, but devote one or two limited functions to the word processing area. Finally, gain commitment not only from the operators but from the executives who will have to make use of this equipment.

Many suppliers see word processing as the entry point into the office of the future. They do not see it as an extension of ADP, which is often isolated from the main thrust of the business. So I suppose that my final point would be a warning to ADP managers who are here to be aware of that, recognise it, and take notice.

QUESTION (Post Office): Do you think, since telex is developing along a similar road to the one you have described, that some developments will arise as a result of sophisticated telex? Will some of the word processing market growth result from telex replacements?

WOOLFE: I feel that at least some of the market growth will arise as a result of regular telex replacement, but what I can't do is to quantify the importance of that. It is equally plausible that new types of photosetting machines will also encourage market growth. Take one example, the Linotronic machine, recently available on the UK market. It is more expensive than your Monotype 96 or IBM System 6, it is nearly three times the money. But prices are dropping fast and it has a lot more capability. I don't know the answer to the question: will this be a major part of the market growth? Perhaps 10% is the closest I could get to it.

CRICKMER (IBM): You said that data is a key company resource, but your talk concentrated on word processing and only hinted at the opportunities which can arise in future as a result of the realisation and exploitation of this view.

WOOLFE: Yes, quite right, John. I'm sorry if I gave you that incorrect impression. Do you remember the slide that I showed, extending into the future, including archiving and retrieval systems, electronic mail and so on? Undoubtedly suppliers such as IBM recognise this as a major area for revenue growth, and will be encouraging users to thrust forward into that area.

ZIJLKER (Akzo): Is any supplier currently able to offer integrated word processing and data processing in one system?

WOOLFE: As far as I know, the straight answer is, "No, not yet". I think that we are just beginning to see a little bit of cross-over, a tiny grey area in the middle; but we are only just beginning to see it at the moment.

Of course, if you're talking about work stations connected to a mainframe, that is a different story, though the proportion of mainframe systems that is up and running in the UK, is really very limited at the moment. But if the question is constrained to just stand-alone units or shared logic systems, then my answer stands.

Exhibit G1

BACKGROUND

cost and volume of office text communication

labour content, labour cost and productivity

low per capita investment

managers: search to find a solution

vendors: pressure to adopt a solution

Exhibit G2

BENEFIT OPPORTUNITIES

- improving typist productivity

  words per day: actual 4·7k → stnd 8·5k → theory 20k

- improving originator productivity
- improving the product
- improving job satisfaction

Exhibit G3

JUSTIFICATION

- cost displacement:

  cost $3-4k pa/station
  benefit $3-20k pa/station

- but

  investment jump
  risk
  training
  need for intensive use

Exhibit G4

CASE STUDIES

increased throughput, reduced costs
conventional functions
benefits hard to evaluate
developing, with caution

- equipment tests
- staff selection
- training
- planning and work flow

Exhibit G5

COMPUTER BASED W. P. EQUIPMENT

stand alone	logic/control
shared logic	display
timesharing	storage
mainframe	printing
	communication

Exhibit G6

IMPROVING TYPIST PRODUCTIVITY

■ actual → stnd

■ stnd → theoretical

	stnd	target
idle	30%	30%
retyping	15-20%	5%
error correction	20-25%	5%
paper handling	5-10%	5%
new words	20-25%	55%

Exhibit G7

EXTENSION INTO NEW AREAS

archiving
- storage devices
- retrieval systems
- incoming mail

printing

calendar, follow up

electronic mail

multi-function work station

'Data is a key company resource'

Exhibit G8

LOGIC/CONTROL

editing
- insert, delete, alter: character, word, sentence
- transfer, duplicate, merge
- search, select

format
- margin
- centering
- column alignment
- hyphenation, justification
- paragraph starts, pagination

additional functions

maintenance

cost

Exhibit G9

DISPLAY

- window size
- vertical/horizontal scroll
- character formation
- prompt
- response
- fatigue-reducing factors
- cost

Exhibit G10

KEYBOARD

- layout
- cost

PRINT

- speed, quality
- font, form range
- pitch, proportional spacing
- impact, non-impact
- composition
- cost

Exhibit G11

STORAGE: COMPARATIVE COSTS

	capacity million chars.	cents/ bit	cents/ 1000 char./month
MOS	10	5×10^{-2}	5
floppy disc	1	3×10^{-2}	3
medium disc	50	8×10^{-4}	0·08
laser	200	1×10^{-4}	0·01

Exhibit G12

PROMISING STORAGE TECHNOLOGIES

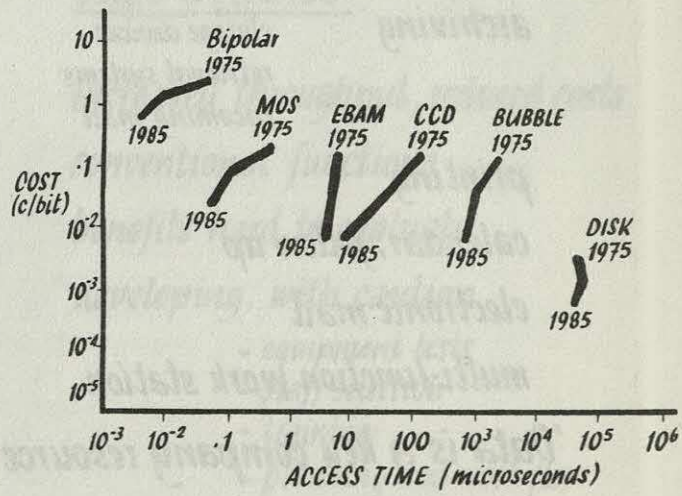

Exhibit G13

DIRECTIONS

user caution and sophistication
short-term: centralised
cost effectiveness measurement
simplicity
service, technical support
planning, education, training
vendor led market growth

Exhibit G14

GUIDELINES

establish real aims
gain familiarity
study work mix
streamline existing methods
consider all the costs
limit the payback period
manage the organisational consequences
limit extent, and prescribe applications
gain commitment
monitor performance

H

SUMMARY AND CONCLUSIONS

George Cox

COX: Gentlemen, there are many things that I would like to do that I have yet to achieve, but one thing I have never sought to do was to end up giving the summary at the end of a conference or seminar. It's a "no win" situation as far as I can see. Everyone I have seen who has ever been called upon to do this has fallen into one of two traps. He either cheats and writes a talk in advance, hoping that he can pick out one or two quotations which fit in during the day. But that's not a summary, and cheating is far from my heart. Or he attempts to give a potted summary of each talk, usually succeeding in picking out the trivia and distorting a lot of what has been said as he goes along.

I can assure our speakers today that I'm not going to do either of these.

Let me tell you what I've got out of the day, and see if I can give you the implications as I see them. What has been illustrated is that the general information systems area is being bombarded with change.

Let us first summarise some of the changes that are taking place. First, said again and it is still an important factor, the tumbling cost of processing power; not the total cost of computing but the cost of processing power. I think just to settle any arguments on that so that we don't get involved in it, when electronic calculators cost £100 a time my expenditure on them was zero; it was a negligible cost in my family's budget. Now they are down to under £10, I see that we have four of them; as far as I can see, my wife doesn't realise that the batteries are replaceable! The cheaper they get, the more we'll have.

Secondly, mass storage is becoming available and memory is likely to become cheaper. Roger Woolfe illustrated a number of advanced memory technologies. Gentlemen, like many of you, I've seen these. You've seen them, sitting on the shelves in manufacturers' labs, waiting to go. I think that the most impressive is magnetic bubble, because when you shine ultra violet light on it you can see the bubbles, and you can crank it round by hand. There is something very reassuring in actually seeing your information jump from T-junction to T-junction. But these are here, and what they are really waiting

for is the manufacturers to decide that the moment is ripe to pick them up and push them and offer them.

Computer technology is migrating into the telecommunications area. People talk about these technologies converging, and I think that is possibly misleading. So far, if you see the evidence, it is like saying, when a mini was crossing a level crossing and it got hit by an express, that it "converged" with the 11.30 out of Kings Cross, because so far the convergence has all come one way, from one direction. What we have seen is the computer companies migrating, the large ones and the small ones, into a telecommunications area; and I think, without being unkind, rather taking the telecomms companies off guard.

We see, too, that the computer companies and the telecommunications companies, in many ways seeing their own markets in some areas saturated and threatened, are now beginning to look very actively at these other areas and how they come together. We will see changes in transmission techniques; a change from analogue to digital transmission; the change in transmission media that Dennis Holloway mentioned this morning. We have to ally these changes that are taking place to the increased people costs to which a number of our speakers have referred.

Rex Malik introduced the figures on the growing number of people in our society who are concerned with handling information - the target area for our systems. I think that these changes will not only bring about changes in the way we do things internally, but for many companies here they will actually take us into new areas of business, in ways that we cannot foresee. People could not have foreseen that credit cards would come about for banks; that a very staid industry like banking, because of an advance in technology, would be taken into a whole new area of business. If you think of the entertainment business, the publishing business, the information access firm, it will take organisations into quite new areas of business. So we are not just looking at doing things differently and reducing typing costs.

In a nutshell, one of the things that I have got out of a recent interest in this area is that systems in the future will take note of information which is in the form it is generated and free form, written, sketched, spoken, free form text, instead of being preoccupied, as we have been, with codifiable and quantifiable information. That is an important change.

So we are seeing major developments in technology; this is not crystal-ball gazing, these are certain developments in technology. How important they are we might debate, but they are certain to come. Many of them are coming now. We are seeing major and significant developments which are being driven by the commercial interests and the need to satisfy

growth of the major computer companies and the major telecommunications companies and the major office product companies. When you marry those two things together, it spells change for all of us - substantial change.

It is also interesting that because of the nature of that same technology which is being moved and shifted by the large companies, you are also getting small companies coming in on it, people we have not heard of before. Just as in the computer business, many of the forces now are companies which you would not have dreamed, eight years ago, would be commercial forces at all. Hearing what Roger Woolfe was saying tended to reinforce my view that what is required for word processing to take off - and it is now happening - is not for the products themselves to be better developed and the software to be better developed, it is actually requiring several hundred salesmen, employed by IBM or Xerox, out there, telling people about word processing. Roger's talk, even though it was a dispassionate, objective talk, was awfully persuasive. I found myself being sold on word processing early in the talk. When you have the sales forces of these large companies out talking to many companies about word processing, that is when it takes off.

What we are saying is that this change is coming, but I think that its speed and direction in some ways is unclear. There are some things that we have discussed which will be non-starters. We will be looking back in five or ten years' time saying, "Whatever happened to this?" or "Wasn't it funny to think that would come about?"

It is useful to reflect on what actually dictates the pace and direction of such changes. I think there are three components. Partly, it is dependent on feasibility, whether certain technical problems are solved and whether it is then proved that you can make it reliably at the right cost, as with fibre optics. Certain developments had to take place before it could come into widespread use. Partly, it is dependent upon market strategy. It is partly dependent upon what the big companies decide to push. That is realistic and it is not even a criticism of them. They are great agents of change. If IBM had not introduced the 3750, or the 2750 before it, I daresay that the whole telecomms industry would still not be thinking electronic; we would still be some way off it.

Partly, it will be dictated by what proves sensible and acceptable to people in practice. Let me illustrate this. We talk about electronic mail. We can easily get enthusiastic about it. You say, "Right, we can originate it here and send it there, and it never needs to be typed. It can be sent to any destination." That's fine. But it was comforting to hear what Dennis Holloway was saying about the amount of mail that does

not even originate inside a company. Before we start gleefully talking about getting rid of our internal post, we have got to think about the way in which people use letters and memos.

There is something proprietary about a document. When I get something in my mail tray that's really good, I very often have to run to somebody with it. Or I get a memo from someone on my staff, and as I go over it, "Rubbish!" People use it that way. I like to take it around with me. Similarly, you hear people talking about conferencing facilities, but it is just two people exchanging information. Well, I have a lot of sympathy with ideas on teleconferencing, but that is not the sole purpose of a conference. Part of it is the sheer getting out of the office and speaking to somebody else; things that you can say, quietly, choosing your moment, that you would not say over some network.

I look back in embarrassment at some of the things that I was saying 12 years ago, and I am sure that one or two other people here feel that as well. I was a great subscriber to the integrated management information system where all information in the firm would eventually be stored, and every manager would have a VDU on his desk. I used to draw VDU pictures up on the board and explain how it worked. It never came about. The reason that it did not come about was not that we did not develop VDUs capable of doing it, but it was a fundamental misunderstanding about the kind of information that a manager wanted on his desk, and the way he did his job. These are the kind of things which in many ways will dictate the pace at which this technology gets picked up.

You might think that if this technology is here now, why can't we predict a little better? I think that we are notoriously bad at predicting its use, partly because of the factors that I have just mentioned. Appalling. We kid ourselves on our ability to forecast, I mean humanity in general. If we take an example of this which always fascinates me: man on the moon, "Ah," you can say, "everyone saw that coming." Back in the thirties people were talking about rockets to the moon, and many of the stories you read indeed had three people in it. In point of fact a lot of the technologies forecast were not even used. Gravity shutters and atomic motors didn't have to come into it, all they did was fill the thing with chemicals. You say, "Ah, that's just what everyone was predicting even in the fifties." Rubbish! Where everyone got it wrong in predicting, even the more learned writers, was in the whole scale of the enterprise.

If you recall your boyhood thinking, the rocket, looking remarkably like the Saturn, was always built in the back yard of some crazy professor, or some little organisation; and it was a quest for knowledge that drove it. To have painted a picture

of a whole industry employing hundreds of thousands of people to get in half a dozen moon shots would have been the most ridiculous scenario you could have painted. What actually got people there was not the quest for knowledge or the quest for prestige by a particular rocket designer, or even a quest for military advantage; it was locked into a complex political situation between two major nations. What stopped us going on to Mars a year later, was not the lack of technology, it was the fact that it was unacceptable to spend the money doing it. So even when you have the technology, working out how people will react to it and whether they will support it is quite difficult.

If you look at the computer business and look back at the way you could have seen, in the late sixties, processing power was coming down and memory size was going up, perhaps you could have predicted what would happen. I submit that few people predicted what would really happen, certainly not the big computer manufacturers, otherwise people like DEC would never have got a look in. Of the people who were able to exploit the mini boom that became possible, many of them got into it by accident. They were pushing their computers into process control and scientific applications, and it was enterprising third parties who suddenly recognised that there was a slump there and caused the growth. So even knowing what was happening, interpreting how it would be used was wrong.

What we have been talking about today, the convergence of technologies, communications and computing coming together, you could say, "Ah, that's a very planned thing, a logical thing." It may appear that way now, but that is not the way it appeared to the people going into it. I think that it is an openly-known fact that the exchange that came from IBM was a bootleg project initially; it was not part of a grand strategy at all. I think that, even now, it is only beginning to form part of a grand strategy.

So what I am saying is that it is no good sitting here trying to paint a scenario of the future, because there are so many uncertainties that are unrelated to the technological facts that we already have; they are concerned with actions by major groups; actions by the Post Office; they are concerned with the way that people will react and use some of these technologies when they are put in front of them. Nevertheless, it remains clear that in our systems area we will be bombarded with change which we will either have to contend with or exploit.

What are the implications that have come out of today as far as I can see? First, I think there is a need within every organisation for someone to take hold of this as a total area. I am not making a plea for the management services directors

to go out and seize more power, but I think that clearly there are so many things which need to be integrated, not necessarily physically but within a policy, that it must come from somewhere; otherwise you will be back to rag-bag buying a lot of bits which just do not fit all the way round the organisation.

A very good question asked to Roger, which I wish we had more time to explore, is where does telex fit into this? I know that Karl Kozarsky, has very strong views on an increasing role for telex.

When people are looking at their new private networks, where does data fit in? What will the volumes be? How does it fit in with voice? How does a choice of this equipment here affect what we can do in the future? How does the choice of this manufacturer's equipment influence what we might be able to do in an entirely unrelated area, or apparently unrelated area?

I think that we will be into a situation where we are dealing with new standards. Tony was talking about networks, we were talking about interfaces: are we going to adopt the standard of our hardware supplier, or a national standard, or what? I think there are a number of major choices there, as we have had in the past with computer languages and software.

I think that we will have to contend with changed rules, the point that Hamish Donaldson made, using that very fine example of utilisation of hardware. Up to now, we have measured hardware by its throughput because of the capital involved. That will change; it will not be the main criterion at all. Perhaps efficiency of programming will change; with core and memory being less expensive perhaps efficiency goes out the window. I don't know, but the rules will change.

We will also be concerned with the evaluation of new products and new suppliers, partly because we are moving into areas in which we have not operated before, and partly because there are new suppliers. Karl Kozarsky this morning was describing a couple of devices; one of them was the DANRAY device. Who in this room had ever heard of DANRAY? I know that, charged last year with looking at an acquisition situation for a major company, we were looking for a particular product. We found one company which had just such a product, which we thought could be either licensed from them or even taken over. We went to the company and said, "You should look at them," and they said, "Who the hell are they? You're joking." But there it was, the most advanced product, from a completely new company. So we are talking about new products and new suppliers.

We are talking - and I think this is important - about the management of new skills. I am sure that it is of real significance for people in this room, because many of us will be

managing skills with which we have not grown up. I am sure that many people here have come to their present position through early computing. You may be removed from programming techniques now but you have, through decisions taken, kept well in touch. People are basically operating in an area where you know the supplier and you know the disciplines. Ah, but now we are talking about bringing in telecomms - a different game. We will be managing new types of software. Tony was talking earlier about network operating systems which I am sure are new to most people here, even the implications are new.

When we are talking about managing telecomms I think this gives us a real problem, because the way that one or two people are trying to do this already is to call someone in and say, "Right, young Cox, I'll make you telecomms manager," as if you can sweep it into a nice little bag. It isn't like that at all. Furthermore, many of the people at the moment who are concerned as telecommunications manager in many quite large organisations are people from a straightforward telephone background. The thought of the telephone system being used as anything else is quite new to them. So we are talking about managing new skills.

An important new skill, which I was very pleased to hear a gentleman raise this morning, is the question of better skills in information use; particularly now we are talking about free form data, speech, drawing, facsimile transmission and so on. There is a need for greater understanding of how people use this information. I do not think that we have thought these skills through and we certainly have not taught them. Most of one's training as a systems analyst is concerned with looking at data and information from a computer viewpoint, articulating requirements and the way they can be processed. We look at a piece of information and get a picture of it; field size, alpha numeric factors and so on. But it is a long way removed from thinking how people use information. What do people do when they get it? They walk around with it. How do they use it? How much is secondary? How much is passed on to somebody else? What is better in sketched form and so on?

I think it is a whole new skill area. Hamish Donaldson said that we want to school our people to be business analysts. That's fine, and with most computer systems analysts you hang a new badge on them saying, "You're now a business analyst." That is a good step. Where do you train these skills? So we are involved in managing new skills there. I think that we are also faced in that area with a reactionary attitude.

When I was putting my first systems in, I was convinced that one of the obstacles was the reactionary attitude of the users; could not or would not see the benefits in the system; could not or would not see why it was worth flogging through some of the problems. At the time I could not work out whether this

reactionary attitude was a question of age, as I then thought, or background. I now know that it was neither: it was a question of human nature. I think that you will find that demonstrated more strongly amongst computer people, funnily enough, than amongst anybody else.

There is a great tendency - I hope that one or two people in the audience can feel it now - to denigrate the importance of anything that you do not like or understand. "I don't think that will be important. I don't think it's really going to disturb this." Particularly as in many areas, although the skills are becoming quite sophisticated, we are moving away from some of the technical sophistications which involved people before.

I was giving a talk in an Eastern European country in January this year. The first half of the talk was on really quite well-proven ground; I was talking about management of systems development, good old stuff. I thought, "Rather than go through that I'll give them a glimpse of some of the things that are taking place in the United States and Europe in systems. That's bound to fascinate them," and it also interested me rather more anyway. I was speaking to one of the people afterwards and I said, "Did you enjoy the talk?" He said, "Oh, Mr. Cox, the first half was very good." I said, "Didn't you like the second half?" and he said, "Oh, all that stuff about telephones and typewriters, that's very boring. You must understand, Mr. Cox, there's some very senior computer people here today." It was just swept aside, quite emotionally. That is an extreme example, but I find it now in organisations everywhere at the moment, and I am sure you do. It is quite an obstacle to contend with.

Those are some of the implications of the change; the need to take control of this area rather than have it fragment; the need to make decisions on new standards; the awareness that certain of the rules will have to change. We need to be able to evaluate new products and new suppliers; evaluate not just the product but their ability to support it, or even their continued viability in the market; the question of managing new skills which are not the skills with which many of us have grown up. Software, telecomms; different attitudes to the use of information; and to overcome possibly some reactionary problems as we re-train our staff and add new ones.

The role that we see for the Foundation is certainly not to crystal-gaze; to sit here and paint pictures of the future is just not on, and it really is not what interests us. The role of the Foundation is threefold: first, it is to track these developments; what is being launched; what is happening; what looks exciting; what seems to be in trouble; what effect will this have. Secondly, identifying their implications, both

technically and from a management viewpoint. Thirdly, to monitor experience in their use; to pick up early difficulties that people are experiencing; what people are finding in practice.

What we have done today is to lay a number of the issues on the table with which we think the systems area should be concerned. It was not exhaustive; there are many important things that we have not touched on; and some of the things that we have touched on today may, in time, turn out to be non-starters or simply hares. We have laid a number of the issues on the table, to illustrate the changes that we think are taking place. One of our aims today has been quite unashamedly, to stimulate interest in the Foundation, and to stimulate your interest in participating in the future.

We will be sending everyone here a guest list so that everyone knows who else he has been in contact with. Most important, I hope that you have found the day of interest and have enjoyed it. I should like to close by thanking our guest speakers very much for their contribution, and to thank you for your participation.