Transcript

Management Conference No. 11

Davos
24-27 May 1982

The Butler Cox Foundation

ELEVENTH MANAGEMENT CONFERENCE

DAVOS
24-27 MAY 1982

Introduction

This document contains the transcripts of the presentations made at the eleventh Foundation Management Conference held in Davos, Switzerland between 24 and 27 May 1982. The theme of the conference was "Future Directions for Information Technology", and the programme was designed to address a broad range of issues that the senior manager needs to be aware of in order to plan successfully into the next decade.

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

THE BUTLER COX FOUNDATION

Butler Cox & Partners

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

Objectives of the Foundation

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

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

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

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

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

Membership of the Foundation

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

The Foundation Research Programme

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

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

The Report Series

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

ELEVENTH MANAGEMENT CONFERENCE

CONTENTS

ELEVENTH MANAGEMENT CONFERENCE

LIST OF DELEGATES

AIR FRANCE	Mr G Lavergne Mr Muller
AKZO	Mr A W Zijlker
AKZO SYSTEMS	Dr C Zedlitz
AKZO ZOUT CHEMIE NEDERLAND	Mr J C R Bus
ALGEMENE BANK NEDERLAND	Mr J M Mudde
BANCA NAZIONALE DELL'AGRICOLTURA	Mr O Cavallini Mr G Gorini
BL SYSTEMS	Mr M Colin
BROOKE BOND OXO	Mr D Lee
BURMAH OIL TRADING	Mr P L Thompson
CADBURY SCHWEPPES	Mr R W Pipe
CENTRAL COMPUTER AND TELECOMMUNICATIONS AGENCY	Mr D Butler Ms E Conn Mr K Murphy Mr G Watson
DEBENHAMS	Mr P Howard Mr L Jeffery
DEPARTMENT OF HEALTH AND SOCIAL SECURITY	Mr B Gibb Mr K C Keen Mr J M Wray
DOMANVERKET	Mr T Sundin
EASTERN GAS	Mr G W Bennett
ELECTRICITE DE FRANCE	Mr Doucet Mr de Latourne
THE ELECTRICITY COUNCIL	Mr C C M Parish Mr C A Thompson
ENAIP	Ms Martuti Mr Tomai Mr Mazotti
ENIDATA	Ms Zoppe
GEC COMPUTERS	Mr L J Bental
GRAND METROPOLITAN GROUP	Mr L F Green
HEINEKEN REKEN CENTRUM	Mr C Rietdijk Mr C G M Scholten Mr H P van Gent
IMPERIAL BREWING & LEISURE	Dr C Montagnon
IMPERIAL CHEMICAL INDUSTRIES	Mr K Edwards

JEUMONT SCHNEIDER	Mr Renaudie
MOBIL OIL FRANCAISE	Mr Mizrahi
NATIONALE NEDERLANDEN	Mr C T C Heyning
PHILIPS BUSINESS SYSTEMS	Mr A C Pick Mr K Tweedale
REGIE AUTONOME DES TRANSPORTS PARISIENS	Mr des Rochettes Mr Gedin Mr Soulat
RECKITT & COLMAN	Mr J T Fields Mr G Gregory
RHM MANAGEMENT SERVICES	Mr J Playfer
ROWNTREE MACKINTOSH	Mr D J Fearnley
SAE	Mr Barbaro
SCHLUMBERGER	Mr C Baudoin
SCHWEIZERISCHE BANKGESELLSCHAFT	Mr B R Bachmann
SENATO DELLA REPUBBLICA	Mr Guelfi
SHELL UK OIL	Mr R R Knowles Mr T A Tobin
SIP	Mr Pomella
SOCIETA PNEUMATICI PIRELLI	Mr Ariatta
SWEDISH NATIONAL ROAD ADMINISTRATION	Mr R Bydler
SWEDISH TELECOM ADMINISTRATION	Mr A Carlsson
THOMSON CSF	Mr Bastie Mr Gaugenot

SPEAKERS

EDWARD DE BONO	Independent consultant
EARL JOSEPH	Anticipatory Sciences Incorporated
PHILIP KRAFT	University of New York State
MARC PORAT	North Star Communications
LOUIS POUZIN	INRIA
RAJ REDDY	Carnegie-Mellon University
BENGT ROSENBERG	A/B SKF
GORDON THOMPSON	Bell Northern Research

STAFF

Butler Cox & Partners Limited	Mr D Butler Mr G E Cox Mr J M Kinnear Mr D S Seabrook Mr E Towbin
Butler Cox, France	Mr A Loize
SA Butler Cox NV	Mr C Jackson
SISDOCONSULT	Mr R Bellini Mr G Franchina
STATSKONSULT	Mr H Johansson

THE IMPACT OF INFORMATION TECHNOLOGY

David Butler, Butler Cox & Partners Limited

David Butler is Chairman and co-founder of Butler Cox & Partners Limited and of its research group the Butler Cox Foundation.

After attending Mill Hill School Mr Butler won an open scholarship to Keble College, Oxford, where he read Greats.

He entered data processing in 1962 as a trainee programmer. After working as a systems analyst and programmer he joined the Urwick Group as a consultant/researcher. From 1970 to 1976 he filled a number of senior posts with a well-known American consulting firm. Butler Cox was set up in early 1977.

He is a Vice-President of the British Computer Society and has won two national prizes for essays on computing. He has published over a hundred articles in magazines and newspapers and is an occasional radio and TV broadcaster. He has lectured widely in Britain and abroad and led the UK team which presented viewdata at the White House, Washington D.C. He is the author of "Britain and the Information Society".

In the early days of the year of our Lord 1960, a young man left college. He did not know what he wanted to do for a career. He only knew what he did not want to do. He did not want to join the commercial rat race. He thought that a job in some kind of public service would be morally superior.

He joined a unit of the local administration in the country where he lived. He became a trainee in the accounts department. He studied in the evenings in order to become a qualified accountant. In the daytime he examined invoices and wrote an accounting code on them. He examined invoices from the gas company, from the electricity company and from the water company. Many of the accounts were wrong and, when they were wrong they always charged too much, never too little. By correcting them he saved many thousands of dollars a year for his employer. His salary was $1,200 a year. "Why don't we have holidays, books, a car?" asked his wife. He had no answer.

In the early days of the year of our Lord 1962, the young man was close to despair. Invoices, children and the salary of $1,200 a year had combined to undo him. "What shall I do?" he asked himself. But he had no answer.

One day the chief financial officer of the administration wrote a memorandum to all the many and intelligent people who worked for him. It had been decided to buy a computer. Applications were invited from those who wished to become computer programmers. The young man did not know what a computer was, much less a computer programmer. But he thought that to be a computer programmer sounded futuristic and laid back — this was still the 1960s — and so he applied.

The competition was fierce, both from other young graduates tired of writing account codes on invoices, and from older employees who had discovered that what the future held for them was no different from writing account codes on invoices, although their salaries were more than $1,200 a year.

The other candidates to become computer programmers studied books on the principles of the electronic digital computer. He studied the reports that his boss had written on why to buy a computer. He gave the selection committee smug answers based on everything they had already agreed. He got one of the jobs.

A month later, everything had changed. He was much closer to despair than he had ever been checking invoices. The language of the computer was incomprehensible. How was it possible to read one character at a time from paper tape? A paper tape might contain a million characters. At this rate his program would take a hundred years to write, let alone test.

One night he had a dream. He dreamed that he was called to the computer manager's office. "We're firing you from the computer team," said the manager. "Why?" he asked. "Because you don't think like a man."

Next day he understood about subroutines. He wrote, once again from scratch, his program, compiled it clean and tested it.

In the year of our Lord 1980, another young man took his car from his garage and told his young and pretty wife that he was going for a drink with his friends, a respite he had well earned from his heavy weekly work. But he did not join his friends and he did not take a drink. Instead, he placed in his car the tools of his hobby, which was murdering women. He drove into the red light district of a town near where he lived and quickly attracted the attention of a prostitute. In a recession prostitutes are easily attracted. The prostitute knew how many women had been murdered in the district. In order to avoid alarming her it was necessary for him to pretend to go through the routine of such commercial negotiations — his requirements, her price.

As it happened, his car was illegally parked. Two young police officers on routine traffic duty noticed the car and radioed its registration number to their headquarters. A computer file had been prepared of the registration numbers of all cars seen in the vicinity of earlier murders. The computer found a match, the car was searched and the hobby instruments found. The lives of an unknown number of women were saved by a computer.

In the year of our Lord 1982, the elected government of one European country decided to join the growing number of nations which have legislation to prevent the abuse of data fed into computers. They decided to have a registrar who would tell people what they might and might not do with data in computers, and who would punish savagely those who did what they were not allowed to do. In this way, they hoped to satisfy other countries which already had laws about computers.

The young man who had been a programmer in the year of our Lord 1962 (but not — I repeat not — the same one who had been a mass murderer in 1980) had by now become the chairman of a consultancy company. He was even vice-president of a national computer society. And his salary was no longer $1,200 a year. But he had joined the commercial rat race. Now he was old enough to be asked his opinion concerning the new law governing what information about people should and should not be put into computers. But many people in other organisations were also asked.

The new law proposed by the government did not cover much of what the government itself would put into computers about people. "It's a sham," said many people, "I don't care what companies put into computers about me, I only care what people put into government computers about me. They are the enemy. I need to know what they think they know about me."

But the once young programmer thought of the murderer in his car, with his hammer and his knife and his shears, and reflected that if the murderer had known what the government knew about his car then he would have changed his car. And the vice-president of the computer society was very, very confused.

In the year of our Lord 1970, a young man sat at the controls of an unusual machine. It was a flying machine, without wings or rotors, because it was designed to fly without air. It was intended to land upon a landing strip, prepared by no hand but God's, that was 280,000 miles from its base. The young man was relaxed because he knew that the computers would do their job. Only they did not. Twenty feet from the realisation of the dream of Plato, Jules Verne, and Herbert George Wells, the computers failed.

The young man had enough power to blast himself back into orbit. He knew that he would never have another chance to be the first man on the moon and so he gritted his perfect teeth, strained his 20-20 vision through the dust clouds, and switched to manual control. In the end, 280,000 miles down range, the hand and eye co-ordination that enabled our ancestors to stuff a spear in a sabre-toothed tiger's eye put Neil Armstrong in the history books.

Do you think you know the first words spoken by a man standing on the surface of another heavenly body? If you think you do, you are probably wrong. Examine the videotape carefully. "It's one small step for man but a giant stride for mankind," that sickeningly predictable concoction of NASA's public relations department, was uttered while Armstrong's spaceboot was in mid-air — sorry, mid-vacuum. His first words on terra firma — sorry, lunar firma — were: "It's sorft and kinda crumbly".

In the south Atlantic, in the year of our Lord 1982, as we sit comfortably here, sipping our mineral water and wondering whether the next speaker will be more interesting than the present one, young men crouch, frozen and apprehensive, above their weapons. Few will show their fear, although all must feel it. Both nationalities are convinced of the rightness of their cause. Both are aware, as the missiles and the smart bombs come winding through the winter air, that their lives depend not just on courage, rectitude and a willingness to die in a just cause, but on the quality of the integrated, interactive, heat-seeking, decoy-dodging software systems that are deployed on either side. It is called the competitive edge of information technology.

What have we done? We can use computers to create for ourselves more interesting and rewarding jobs. We can use them to catch mass murderers. We can use them to blow other people's sons, brothers and friends to pieces. We can use them to explore the universe.

To find out what we want to do and how we can do it is why we have come to this conference.

THE INFORMATION SOCIETY: POSITIVE ACTION

Edward de Bono, Independent consultant

Born in Malta, Edward de Bono later proceeded as a Rhodes Scholar to Oxford and has held faculty appointments at the universities of Oxford, London, Cambridge and Harvard.

He is author of 22 books which are in the general area of thinking and especially the thinking concerned with change and innovation. Dr. de Bono is the originator of the term 'lateral thinking' which is now officially part of the English language with an entry in the Oxford English Dictionary. The books have been best-sellers in Japan, Germany and the USA.

Dr. de Bono has lectured extensively throughout the world. His instruction in thinking has been sought by such corporations as IBM, Shell, Unilever, Bank of America, Westinghouse, Northern Telecom, ICI, Prudential, Marsh McLennan, Ciba-Geigy, Monsanto, General Dynamics and many others.

He has made two TV series: 'The Greatest Thinkers' in 13 parts and a 10 part series for the BBC, 'de Bono's thinking course'.

He is the founder and director of the Cognitive Research Trust at Cambridge which runs what is the largest programme in the world for the direct teaching of thinking as a subject in schools. Several thousand schools are involved world-wide and 100,000 teachers have been trained in Venezuela which has put the material into every school.

I shall be talking about different aspects of thinking. It might be as well right at the beginning to say why I coined the term 'lateral thinking''. The reason is that ''creativity'' as a word is very inadequate. Many so-called creative people are not creative at all.

If we imagine that most of us look at the world in one particular way, there may be a highly creative person who has a different way of looking at the world. If he is effective in communicating with the rest of us through his writing, music, painting, or whatever, he will cause some of us to look at the world through his eyes. As a result, he is of great value to society — he enlarges our perception, our vision, our experience. So we call him creative.

But that person may be very rigid. He may be able to look at the world only in his own particular way. He may be quite unable to look at it the way other people look at it and quite unable to change his perceptual structuring of the world. I do a lot of work with artists, designers and architects, and there is no doubt that many of them are very rigid people. The same certainly applies to many creative scientists. In other words, it is possible for a person to have an idea that is different and valuable but to lack completely the ability to change his ideas.

This can often be the case with young children. If you give a youngster a problem, because he does not have an established approach to that problem, he is likely to come up with a highly original approach. But if you then say to that youngster, ''That is very interesting. What about another approach?'' then very likely he will say, ''No, no, that is the only possible one.'' So again we have an original, creative person, but one that is also rigid.

This is one reason why it was necessary to invent the term ''lateral thinking''. The term describes the ability to change our perceptions and to keep changing them, with the emphasis on that movement rather than just on the difference between perceptions.

Another reason why it was necessary to invent the term is that the word ''creativity'' is only a value judgment on a result. No one ever uses the word ''creative'' to describe a new idea that he personally does not like.

If we look only at the result, it tells us nothing about the process. We know that Darwin spent 20 years of study and visits to the Galapagos Islands to come up with his idea of evolution through survival of the best adapted ('survival of the best adapted' is of course a tautology). However, another gentleman, Alfred Russel Wallace, came up with exactly the same idea after three days of high fever when he had malaria out in Borneo. If we look only at the result, we must conclude that malaria is really rather more effective than 20 years of study and field research.

In other words, attempts to describe creativity by working back from the result are virtually useless, because you can achieve a result in a number of different ways. Discussion of 'creative thinking' can therefore be very confusing.

That was why it was necessary to devise something called lateral thinking, which is a neutral process that takes place in some particular information universe. You may use lateral thinking and come up with no result, which is usually the case. You may use it and come up with a result that is good but no better than one you already have, or the cost of transition may be too high. Or you may come up with a result that is better than the one you already have.

In each of these cases one is equally using the neutral process of lateral thinking. The process is neutral in the sense that it is not defined in terms of the result we may get. The advantage is that we can learn and use the process. Later in my talk I will define the information universe in which the process takes place.

Let us now look at three broad aspects of how we handle information (see figure 1). The first is where there is an information source, passive or active, which sends a signal to some thinking being, who then interrelates the information with his other experience and takes some decision or action.

Figure 1
Three types of information system

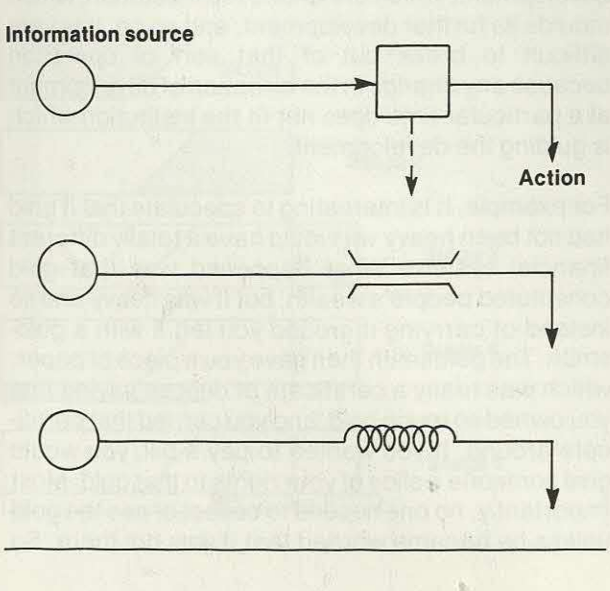

The second is where the information flows from the source straight through some sensor device or other linkage (which has been originally designed by our thinking being) to action.

The aspect in which I am particularly interested is the third, where the information creates its own reality which flows on into action. That is the particular type of information system that I shall refer to from time to time in this talk — what we might call active information systems rather than passive systems, which the first two are.

Let us now look at the three basic ways we have for getting something done. I shall use the model of a ball rolling down a slope.

The first approach is to cut a channel or a groove in the surface of the slope. That corresponds to using procedures, formulae, or algorithms. It is very efficient and very effective. Most of our personal or corporate life is spent in trying to set up just such channels so that when we enter them we end up where we want to go.

This approach has many advantages, but it also has some disadvantages. You cannot start from somewhere new unless you cut a linking channel. You cannot shift your objective unless you cut a connecting channel.

The second method is also one we use a great deal. In my model, this consists of placing a light source (for example, a light bulb) at the target and equipping the ball with some sensing device so that it finds its way towards the target. This is goal-directed behaviour, which in biological terms is known as trophic behaviour, and in management terms is known as management by objectives.

Clearly this is less efficient than the first method because it wastes a lot of energy. Because it involves taking decisions, it requires sensitivity and a higher calibre of person. But it does have the advantage that you can start at any point and that it is relatively easy to change your target if you want to. So the second method has some advantages and some disadvantages relative to the first method.

The third method is also one we use a great deal, although we do not often acknowledge it. This method consists of letting the ball go and then, when it has arrived somewhere, we decide that that was where we wanted it to go. This method applies to the way we run businesses, and it applies to the way we develop ideas. In particular, it applies to the way our minds form concepts. We have virtually no control over the way we form concepts — the nature of the landscape or the terrain, the pressures, the opportunities, the technology available at any moment, shape where we are, and then we justify that as being the way of looking at things.

Now let us consider the two basic models of progress. The first is the technological model, which is shown diagrammatically at the top of figure 2. At some stage of progress there is a technological input that accelerates progress. Then later there is another technological change that further accelerates progress. Aviation provides an example of this model of progress. There are many people alive today who were born before the first aeroplane ever flew. Some time ago, crossing the Atlantic by air, I was reflecting that the spoonful of mashed potatoes that I was about to put in my mouth was travelling faster than a rifle bullet. So, within a lifetime, we have progressed from nothing to huge achievements in aviation. In data processing electronics we have seen the same sort of technological, geometrically increasing type of progress.

Figure 2
Two basic models of progress

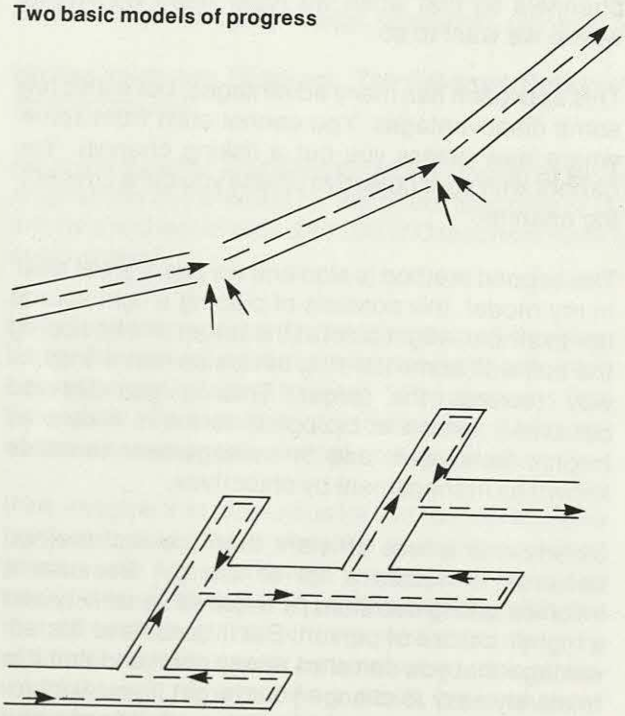

In contrast, the lower portion of figure 2 shows another type of progress. Here, our initial experience in the field sets up a certain pattern, concept, arrangement, or organisation of information, which then carries us along with its own momentum. We may then reach a point from which the only way to progress is to undo that particular concept and go back to find a point from which we can progress with a different concept. This is an extremely slow process, which is why in human affairs in general our progress is so extremely slow. This sort of progress occurs in technical fields as well, but primarily it is to be found in human affairs.

The undoing of concepts is extremely difficult. We have no natural graveyard for concepts, and no satisfactory way of changing concepts. Later, I describe different ways of changing concepts, but before that let us look at the different forms of continuity.

There are two sorts of continuity. The first is continuity by neglect. To give an example, for 40 years a little solenoid-operated arm which flipped up was used on motor vehicles to indicate the driver's intention to turn left or right. It was perfectly useless — it was always breaking off and it could not be seen from many angles. Yet for 40 years it was regarded as satisfactory. Then someone thought of using flashing lights as indicators. The technology for such lights had been available all along, but artificial arm indicators were used because they were a continuation of the way things had been done before. When you turned right in a carriage you put your whip hand out. In the first motor cars, when you turned right you put your hand out. When cars became enclosed the natural thing to do was to use an artificial, mechanical arm. That is what I mean by continuity by neglect.

In data processing, and many other areas, much of our thinking is to use new technology to do things in the same way as they have been done previously.

The second sort of continuity I call continuity by ricochet. Something comes into being, and that creates an institution, a procedure, or a set of mechanics, which once it exists determines the way things will continue to develop. Each successive development is then influenced by the way the previous development was carried out. This produces a ricochet effect where each development creates a setting which then moulds its further development, which enhances the institution, which moulds its further development, and so on. It is very difficult to break out of that sort of operation because any change in the direction of development at a particular time does not fit the institution which is guiding the development.

For example, it is interesting to speculate that if gold had not been heavy we would have a totally different financial system. What happened was that gold constituted people's wealth, but it was heavy and so instead of carrying it around you left it with a goldsmith. The goldsmith then gave you a piece of paper, which was really a certificate of deposit saying that you owned so much gold, and you carried that certificate around. If you wanted to pay a bill, you would give someone a slice of your rights to that gold. Most importantly, no one needed to collect or see the gold unless he became worried that it was not there. So the whole system of money and credit expansion, and our whole concept of how we handle economics, developed in this way.

Had gold not been heavy, or had we had some notational way of indicating a person's wealth, we might have had a very different economic system. In the future I think that we will get back to rather different economic systems based on different concepts from those of the current system. For example, with our data processing capacity, we may have several different, almost closed-loop economies working within the same economic system without requiring a multipotential flow in different directions. But that is something in the future.

To summarise, there are two types of continuity. One occurs through neglect and the other through a ricochet effect where the happening creates the institution which preserves the happening, and so on.

To continue this brief look at continuity, I shall use a simple analogy. We live over time and time passes in a particular direction. As time passes, we get different inputs, different experiences. Whether they are technical, marketing, political or whatever, we try to make the best use of what we have at the moment. We cannot say that we would like to stop existing at a particular point in time and take up existing when the picture has become clearer (much as British Leyland would like to do from time to time!). You need to put the inputs, the experiences, together to maximise what you have at the moment.

Figure 3 provides an illustration of this. Imagine that you are given shapes, which represent your input, one at a time and are asked to try to make the best use of what you have. Best use in these terms is a

Figure 3
Examples of maximising the use of successive inputs

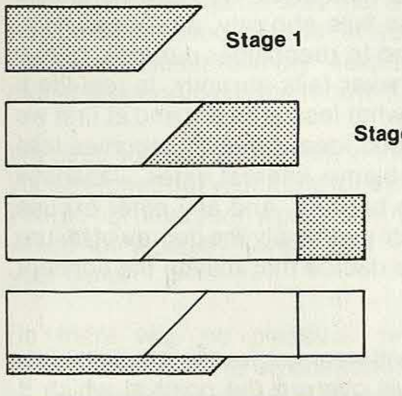

shape that is easy to describe to someone who cannot see what you are doing. At each of the five stages shown in figure 3, the new piece, the new input, is shown shaded.

Most people put the first two pieces together to form a rectangle as shown. For each of these first two pieces, the long side is equal to two widths, and the completed rectangle is three times as long as it is wide.

Time passes and we are given a further input (stage 3). We can build on our existing system fairly easily as shown, and we get a rectangle four times as long as it is wide.

Time passes and we are given a further input (stage 4). Quite sensibly, we build on what we have. But then, when we receive the final piece (stage 5), we find that it does not fit. More than likely we then decide that it is a good enough fit and leave it at that.

There are two points to be made here. First, whether or not we like it, if we are in a sequential system where we try to maximise our use of input at each stage, our ideas at any moment will inevitably be less than if we were to make maximum use of our experience or of the technology available to us. That is inevitable. There is no way of escaping that, because if you are dependent on a sequence then you cannot maximise in the same way as if you were not dependent on that sequence.

The second, important point is that being right at each stage is not enough. Up to and including stage 3 we clearly had a correct, economical arrangement. In order to proceed from there, we may have to go back and change what in its day was a perfectly valid, and probably the best, arrangement. If we do this then we could achieve the arrangement (using the same inputs) shown in figure 4. With this arrangement,

Figure 4
Different use of the input in figure 3

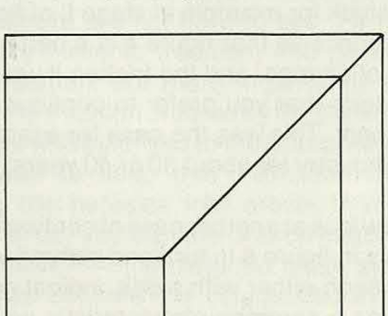

adding the last two pieces is extremely easy, and we do not have to worry about the ratio of the sides.

In any audience there are some people who feel that, with their superior intelligence, they would at the second stage have put the pieces together as in figure 4 and so have avoided running into trouble. However, had they done that then I, in my role as Fate, God, Providence, or whatever, would have given them as the fourth input the shaded piece shown in figure 5, in which case they would have been better off with the first, more straightforward arrangement.

Figure 5
Alternative input at stage 4 of figure 3

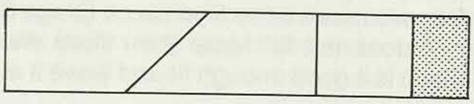

In other words, unless we have complete knowledge of the future, which is somewhat unlikely, there is no way that we can so arrange our conceptions or configurations to make the best use of whatever may come along. We can plan for contingency, we can have back up positions, we can have fluid positions, but whatever we do we are committed to some particular direction, whether we like it or not. Even if you keep all your money in the bank at a particular moment, you are still not uncommitted — you are committed to the continuance of Paul Volcker as head of the Federal Reserve in the United States. So whatever you do you are always committed in some direction. And whatever concept you have, inevitably there will come a time when usefully you ought to go back to change that concept in order to move forward.

This does not mean that because you ought to change it will necessarily be feasible to change. If you are stuck for example at stage 5 of figure 3 and you can conceive that figure 4 is a better concept, the cost of change, and the friction it would cause, may be such that you prefer to continue with your first concept. This was the case for example in the printing industry for about 30 or 40 years.

Let us now look at another case of continuity. Look at the letters in figure 6 in turn and commit yourself to marking each either with a tick, indicating that you think it has a common characteristic with the first letter (already marked by a tick), or with a cross, indicating that it does not have that characteristic (as in the case of the second letter which is already marked with a cross).

Figure 6
A further example of continuity

P A C H G J F R
√ x

If I put these letters up one at a time on a projector, and mark them myself with ticks and crosses, asking the audience to call out 'Yes' if they think I am right or 'No' if they think I am wrong, then the audience is always in agreement when I mark as follows:

C √
H ×
G √
J √

When I reach F, however, and mark it with a tick, either people think that I have made a mistake and say "No" more loudly, or they think that I have cheated and changed the rules, or they just get confused. The same happens when I then mark R with a cross. In fact, I have not made a mistake and I have not changed the rules — there is a consistent theme running through the crosses and ticks with which I have marked the letters.

What happens is that up to J people are usually led to think in terms of straight lines and curves, and up to that point, that is a perfectly correct hypothesis. I will tell you in a moment what hypothesis I am using in marking the letters with ticks and crosses.

The concept of this exercise is trivial but the principle behind it is extremely important. The principle is that if we have an idea that works, and works, and works, and works, then we have no choice but to consider that as a valid idea. We have no choice whatever. If it then fails abruptly, as it does here, then we are caused to reconsider our idea. In real life, of course, it never fails abruptly. In real life it starts being somewhat less efficient and at first we still hold it as a good idea. When it becomes less efficient we now blame interest rates, Japanese competition, trade barriers, and any other excuse we can think of, until eventually the degree of failure is so great that we decide that maybe the concept needs changing.

In other words, whether or not we like it, inevitably a successful idea will overrun the point at which it could usefully have been changed, reconceptualised and rethought.

The particular concept used in the exercise in figure

6 is trivial. All the letters that I tick would be unstable if they were solid objects. There are many other concepts that would fit my pattern of ticks and crosses. The important point is the principle that repeated success of an idea causes the idea to become so established that, outside of mathematics which, being a closed system, can be used to validate ideas, it will inevitably go on being used beyond the point at which it could usefully have been changed.

What mechanisms do we have for changing concepts? We might say that one of the mechanisms we have for changing concepts is evidence. If we pile up the evidence and this shows that the existing idea is wrong, then we will change that idea. This happens sometimes, but the way that it is handled in science is extremely inefficient.

I will give you an example. Some years ago, in Canada, there was a famous experiment carried out on rats. The experimenters wired up a rat so that when it pressed a bar it closed a contact and that stimulated part of the rat's brain. They found that when they put the electrode in at a certain point, the rat went on pressing the bar time and time again. Even if they put some food, or a rat of the opposite sex, alongside the wired-up rat, it went on pressing the bar. As a result, the experimenters called that part of the rat's brain 'the pleasure centre'.

For years, all around the world, people in laboratories tried to find the pleasure centre in humans, hoping at the same time that they would never find it, because, if they did, we would simply put in an electrode, wire ourselves up, and spend the rest of our lives in corners, pressing contacts. It was only 22 years later that it was discovered that the pleasure centre hypothesis is probably not the explanation for the rat's behaviour. For 22 years all the data had been looked at through that hypothesis and all the data had confirmed that hypothesis. What was discovered 22 years later was that a certain chemical was being released into the rat's brain, and that this chemical release simply keeps going whatever circuit is operating at the moment.

We need such a mechanism in the brain because if you stretched out your hand for a glass of whisky and you did not have such a persevering mechanism you would immediately forget why you stretched out your hand.

So there was no pleasure centre at all. The behaviour was more an obsession or a compulsion — one particular circuit in the brain was kept going by the release of a chemical.

That is a classic example of how our normal scientific method, consisting of choosing the right hypothesis, looking at the evidence, and trying to support the hypothesis, is a very limited way of proceeding, and why science could move much faster if we got into the habit of using alternative hypothesis windows through which to look at information.

So that is the first mechanism for changing concepts, and it is not all that efficient.

The second mechanism is what we might call the dialectic method. This means that one person has an idea, someone else opposes it, they argue it out, and in the end one of them either changes his idea or simply loses the argument. This is very much the Western tradition. Again it is a very inefficient method because you cannot change an idea unless you can prove it deficient or inadequate. Also, when you attack an idea, someone defends it, and the two of you then spend a long time attacking and defending. Moreover, you cannot seek to change something which really is adequate, because if it is adequate someone else will say, "It is good. You cannot prove it bad. Why should you change it?" Also, in this tradition there is less willingness to explore, because you cannot explore unless you have proved the need to explore.

Contrast that tradition with the non-Western traditions. The Japanese, for example, have not had the dialectic tradition. They in general have very rigid societies and the paradox is that rigidity gives you great freedom to explore. The reason that rigidity gives you freedom to explore is that you can explore knowing that, if you do not find anything, you can come back on stream. In the dialectic tradition, you cannot explore unless you have destroyed the base, so you are reluctant to do so.

In non-Western traditions you explore and, if you hit a channel that makes sense, you immediately switch to that. So you get exploration and switching rather than dialectic clashes between people adopting fixed ideas.

The reasons for there being a dialectic tradition in Western culture are not hard to see. Partly it is a result of the Hellenic tradition, and partly of the Socratic tradition, but more importantly it is the residue of the tradition created by the thinkers of the Middle Ages who belonged to the Church and whose purpose was to keep that institution intact by destroying the heresies that arose. If you could destroy a heresy you kept the Church intact. So that type of thinking — the dialectic clash system as such — has become the established thinking of Western society ever since the Church ran the universities, schools, and so on. For its original purpose it was a valid way of thinking. For the purpose of being able to change concepts constructively it is very inefficient.

Let us look at other methods of change, such as mistake or accident. In medicine, for example, virtually all advances have come about through mistake or accident. The more recent trend in medicine of using heavy number crunching for measurement analysis has produced very little — in terms of the investment in time and money that has been made, it has produced horrifyingly little. The reason is very simply that analysis will always be inadequate as a tool in a complex system, because one can never know all the parameters.

Let us look at one classic example of a mistake. This happened when Pasteur was working with chicken cholera. He went away one weekend and before he left he told his assistant to put a beaker of cholera germs into the drinking water of the chickens. He came back after the weekend, expecting to find most of the chickens dead. He was then going to examine those that had survived to see why they were resistant to the cholera germs. To his surprise, all the chickens were still running around. He asked his assistant whether his instructions had been carried out. His assistant replied that they had. However, when Pasteur saw the bucket in which his assistant had diluted the cholera germs, he was angry, saying that the bucket was much too big. The following weekend he went away again and told his assistant to be sure to use the correct, smaller bucket. He came back after the weekend expecting to find most of the chickens dead. Again, every one of them was still running around.

From that mistake of using the wrong bucket the first time came the whole concept of immunisation, whereby by giving a low dose of an infecting agent you protect against a full dose. In hindsight we can see that it makes sense, but it would have been very difficult to see in foresight by means of analysis. I will discuss the reasons why analysis is inadequate in complex systems in a moment.

Let us move on to the type of information universe that I am talking about. Once, in order to illustrate a quite simple point, I invented a little jigsaw puzzle. It has 16 little squares, each having two faces, which you put together to form a big square of 4 x 4. If you try to put in one piece every second, day and night, without stopping to eat, drink or sleep, it could take you more than a thousand million years to complete the jigsaw. That is a rather long time, particularly when you consider that man has been on earth for only ten million years, and even the termites for only a hundred million years. It is the sort of thing you give to your enemies for Christmas!

The reason why it takes so long to solve this puzzle is that the design does not emerge unless you insert every piece in the correct sequence. There is no sorting strategy or hierarchical strategy to guide you in putting one piece against another, so you have to solve the complete puzzle by trial and error. With 16 pieces, each having two faces, the mathematical combinations are so huge that you cannot possibly hope to try them all.

In other words, if we ever believe that in any thinking or information system we can operate by putting the elements together in different configurations and choosing the combination that suits, then we are fooling ourselves — it is utterly impossible simply because of the mathematics of combination.

If we had to cross the road by genuinely analysing all the relevant data available to us, it would take us about a month just to cross the road. We do not take a month to cross the road, because the brain is designed to be brilliantly uncreative. That is the main purpose of the brain. If it was anything else, it would be useless.

By that I mean that the brain is designed as an environment in which incoming experience can organise itself into patterns. Once we have such a pattern, a certain shape, for example, moving at a certain angle to the eye, will trigger that pattern and we will then read out all the characteristics associated with it. In other words, we create the scene, we do not just receive the information. What is more, anything at all similar to that pattern is treated by the brain, automatically, by means of one of its simplest operations, exactly as that pattern. The brain finds pattern recognition the easiest possible function to perform. Essentially it cannot do anything other than pattern recognition. In contrast, in normal data processing, pattern recognition is one of the most difficult things to achieve.

How does the brain do this? How does the brain create an environment in which incoming information organises itself into a pattern? The answer is that the brain uses active information systems as opposed to passive systems.

Let us look at some models of active information systems. First, imagine that we have a tray of sand and that we drop a steel ball on to the surface of the sand. The ball stays where we drop it. If we drop it somewhere else, it stays there. That is a normal, accurate, passive, recording system. If we put a grid over the surface of the sand and assign co-ordinates (letters along one axis, numbers along the other) to the input then when we drop the ball in position A2 it stays in A2, and when we drop it in C4 it stays in C4. That is the normal way we like to treat information when we want an accurate record for communication, storage, and reacting purposes.

Let us contrast that with another system. Imagine a tray that has a moulded plastic surface such that when we drop the ball it always ends up in exactly the same place, the lowest point of the curved surface. Such a surface is no longer an accurate record, it is in essence curving the information. It is an active system, not a passive system — it is creating its own reality. From where the ball ends up, there is no way of knowing the position of the point onto which it was dropped.

Let us move on to look at the next system, which consists of a tray containing a heavy viscous fluid — silicon, putty, heavy oil or whatever — with some sort of membrane on top. When we drop the first ball on the surface, it sinks in and changes the contours of the surface. If we now drop a subsequent ball on the surface, because the fluid is viscous, this second ball will roll down and cluster near the first one.

So here we have an environment in which information is self-organising into some sort of grouping or clustering. The contours of the surface are not preformed, they are formed by the arriving information.

Let us take this a stage further and look at another model. Let us imagine a towel laid on the surface, with a bowl of ink alongside. When we put a spoonful of ink on the surface, it leaves a stain. This happens with every spoonful of ink we put on the surface, and so we have an accurate recording system. If we want to use it we need an outside processor to measure and relate the stored data. This is the traditional view of the human mind, the view concerned with memory thinking. For perception, however, it is probably the wrong system.

The next system is a dish of gelatine. We now heat up the bowl of ink on a little fire. When we put a spoonful of the hot ink on the surface, it dissolves the gelatine. When we pour away the cooled ink and dissolved gelatine, we are left with a depression in the surface. If we place a second spoonful of hot ink near the first, it will flow into the initial depression, creating a shallow depression of its own and deepening the depression created by the first spoonful of ink. If we continue with spoonfuls of hot ink, each placed near to the previous one we end up with a channel eroded into the surface of the gelatine.

Here we have in essence the formation of a pattern, the definition of a pattern being that if you move from one state to another state with a probability greater than pure chance you are part of a pattern. If a ball is placed on any point in the depression, it will automatically get moved to the deepest part of the depression. In this system we do not need an outside processor — by having different layers of patterns we can get any type of information processing we like.

Let us move on now to what happens in nerve networks. The impulse travels along a nerve in digital form. This releases a chemical at the synapse and if the concentration of the chemical is sufficient (note that this is now analogue processing) then another nerve is triggered. This whole operation is set against a background of varying chemical fields, producing a field effect. So the network involves digital, analogue and field effect processing.

If we imagine a sheet of interconnected units then, when we stimulate one of them and it becomes active, this spreads to all the neighbours which also become active. If this were all there is to the system, you would get an epileptic fit whenever you looked at anything, which would not be much use. You can get the same effect by taking strychnine. So the system as described so far is not much use.

With a slight modification, however, we can turn it into an extremely effective system. If each time we activate a unit we release a certain chemical that builds up a background of chemical effect that in turn inhibits each other unit then we can create a balance of excitation and inhibition. As the number of active units increases, so the excitation pressure on any unit rises, because it has more neighbours that are active. The inhibition pressure, however, will rise much more quickly, because it is determined by the total number of units that are excited, not just the neighbours.

So a point is reached at which inhibition exceeds excitation, which makes contained, coherent representation possible. This provides the elements of a pattern-forming, pattern-using system.

This was all written up years ago in a book of mine called "The Mechanism of Mind" which has now suddenly become of great interest to people involved in artificial intelligence and the like. The mechanism has also been simulated on a computer and does behave as predicted. It learns, discriminates, shows humour, and has insight.

What does all this amount to? It amounts to the fact that in the mind is a system that provides an environment for incoming information to organise itself into patterns.

Such a patterning system has many characteristics that make it a completely different information universe from the discrete systems we normally use. There are many predictions that can be made from the behaviour of information in a patterning system which are counter intuitive and quite different from behaviour in discrete systems. For instance, you can show that it is much easier to learn something backwards than to learn it forwards. There are many experiments that show quite startling differences.

Let us now look at two particular features of a patterning system. First, when one particular pattern is established, there may exist many other tracks and patterns that at this moment in time are invisible. When an input creates an area of excitation, that area will expand by spreading the excitation to the area immediately surrounding it, but smaller areas of excitation elsewhere will collapse for the moment as a result of the general inhibition effect created by the larger area of excitation.

This is why patterning systems are viable. This is why we can cross the road, why we can read, why we can shave, why we can recognise people, and so on. It is also why many of the old philosophical problems of free will determination are non-problems in a patterning universe. They simply do not exist as problems.

Figure 7
Asymmetry of patterning systems

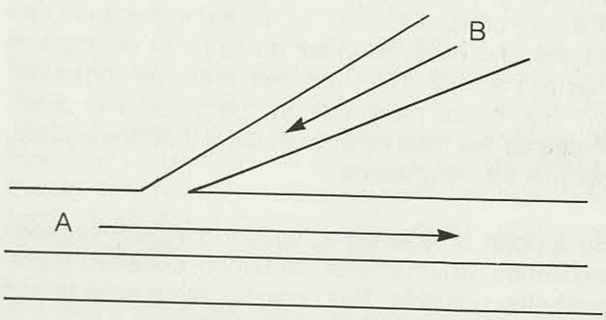

The second feature is that a patterning system necessarily contains asymmetry. For example, in figure 7, the route from A to B may be very roundabout, but the route from B to A may be very direct. It is simply that the distribution of probabilities is asymmetric. This has immense uses. On the other hand, it makes the changing of ideas and concepts very difficult, because it can be very difficult to move from A to B.

An interesting point about this asymmetry is that it is the basis of humour. It has always amazed me how little attention philosophers, psychologists, and information theorists have paid to humour, because it is by far the most significant characteristic of the human mind — far more significant than reason. Reason in information terms is a fairly cheap commodity, it is simply a matter of running a sorting system backwards, and can be achieved by means of pebbles, cog wheels, semiconductors, or whatever you like.

Humour can occur only in a self-educating patterning system that has some sort of interface, or language, with which to communicate with the outside world. What happens in humour is one of two things. Either you are going along the main track and someone now heavily emphasises a point down a side track, which provides the double meaning, pun type humour. The other is where we are taken away from the main track and in hindsight we can see it makes sense. I shall give an example of each.

A classic example of the first kind of humour, the pun, was provided by Bob Hope when he complained that he had a very poor Christmas because he had been given only three clubs and, what was worse, only two of these had swimming pools.

The other type of humour is illustrated by the story of a fellow sitting in a train compartment. The ticket inspector comes in and asks for the tickets. The man starts searching frantically for his ticket in his pockets, his case, his coat, and so on. After a while, the ticket inspector has mercy on him and takes the ticket out of the man's mouth where it has been the whole time, punches it and gives it back. When the inspector has left the compartment, the man's companions turn to him and say, "Didn't you feel very foolish looking for your ticket when it was in your mouth the whole time?" He says, "No, that wasn't foolish at all. I was chewing the date off the ticket."

Basically, the humour model is the same as the hindsight model and the same as the insight model. What we are trying to do in lateral thinking is to move from A to B in figure 7. But to do so we may have to employ a special process.

To illustrate this I will give you one of my favourite examples of lateral thinking, which concerns a Wimbledon tennis tournament. It is a singles tournament and 131 people have entered. Two people play, the winner plays another winner and so on, and in the end two people play a final match to decide the tournament winner. It rains the first week and the organiser has to put all the matches into the second week, so he wants to work out how to arrange to have the minimum number of matches and what that minimum number will be. There are a number of ways of organising the matches, including having some byes in the first and second rounds. There is however one very simple way of working it out if we make a little creative move, and having got to the answer we can see in hindsight that it is obvious. The creative move we make is this. Normally we focus our attention on getting a winner of the tournament. There is one winner. If we shift our attention to the losers and we regard the tournament as a way of producing losers, not winners, we can see that there will be 130 losers. How are losers produced? Each loser is produced by a loser-producing match. So, if there are 130 losers, there must be 130 loser-producing matches. We then add in the other sorts of match there would be in the tournament — in fact there are none, because there are no draws. So the answer is that the minimum and maximum number of matches is 130. So just by making a creative step at the beginning, a shift of attention, we get the insight we need to solve the problem.

I am making several points here. One is that in a patterning system we need ways of crossing patterns. Traditional experiences such as induction are all pattern-using methods. Judgment is a pattern-using method in terms, first, of locating the pattern and, second, of stopping us from wandering off it. What we need is a whole range of different thinking strategies where we are concerned with moving across patterns rather than up and down established patterns.

I do not have time to go into this in detail but I will give a few examples of the sort of things we can do. One is to create the idiom of movement. In the judgment idiom, when we come to an idea where there is a mismatch with what we expect, we back away from the idea or try to alter it. In the movement idiom, we use an idea for its movement value irrespective of its judgment value.

To indicate the movement idiom, I created the word "Po" as distinct from the word "No" which belongs to a judgment system. Po is an indicator that we are using the movement system.

I will give a couple of examples. At one time I was in the United States talking to people involved in pollution legislation. When a factory puts out pollution, people downstream suffer. As a provocation we made the statement, "Po, the factory should be downstream of itself." That sounds pretty illogical, because how can a factory be downstream of itself? But from that provocation comes the following idea. Normally a factory takes in water and puts out pollution. The idea is to legislate that when a factory is built, its input must always be downstream of its own output, so that the factory is the first to get a sample of what it is putting out. That has since become standard legislation in Russia and all east European countries, but is not so well-known in the West.

That is an example of using a provocation as a way of crossing patterns. If we want to cross patterns, we have to use a temporary, intermediate, unstable state, just as in chemistry or nuclear physics it may be necessary to use an intermediate, unstable state in order to move on to a new state. This is quite different from following patterns in a sequential system.

I will give you another example of a provocation. One of the ways of getting a provocation is just to reverse something. I was doing an exercise with a group once, and the statement was made, "Po, aeroplanes should land upside down." Obviously, this was not meant to be a serious suggestion, but from it quickly came two concepts.

First, someone said that if planes landed upside down the pilots would get a much better view. So from that arose consideration of the positioning of the pilot. Is he where he is because that is the best place or simply because when planes were very small it was obvious to sit the pilot on top, and then as planes got bigger and bigger the pilot was put in the same relative position?

Another point made was that if planes landed upside down they would land positively — that is, they would be lifted into the ground. From that came the concept of having a landing surface that is in some way retractable, or maybe of variable geometry, so that when the aeroplane comes in to land the downward lift could be balanced against the normal upward lift. This would provide a much finer landing control than just cutting the engine power.

So from what at first sight is a deliberate provocation, we were able to move off into two interesting directions.

To give another example, the statement was made, "Po, cars should have square wheels." This produced about ten different ideas, one of which is now being worked on by Firestone Tyres. I will tell you about just two of them.

First, someone said that square wheels would give the advantage of a better adhesion surface. The disadvantage of course is that the wheels cannot roll easily. The normal method would be to back away from such a silly idea. Using the movement idiom, we ask how we could achieve the advantage and get rid of the disadvantage. From that came the concept of having a hub with an inner tyre at a pressure of say 28 psi, and then an outer tyre at a pressure of say only 7 psi. So the wheel rolls on the inner tyre but also gets tremendous adhesion from the outer tyre.

A second idea that arose from the square wheels provocation came from thinking that the car's suspension could be adjusted so that the wheels are lifted as the corners of the square are rolled over and so a smooth ride could still be achieved. From that came an interesting idea of a vehicle that does not bump over rough ground. Imagine that the vehicle has three pairs of wheels plus a jockey wheel at the front. In fact the jockey wheel could be replaced by some other sensing device such as a sonar beam. When the vehicle approaches a bump, the jockey wheel senses the profile of the bump, which is then computed, taking into account the vehicle's speed. When each pair of wheels approaches the bump, the suspension simply lifts it over the bump, with the vehicle continuing to run on the other two pairs of wheels, and puts it down on the other side. This produces a vehicle that flows over the ground instead of bumping over the ground, and the energy use is only one-twentieth that of a bumping vehicle, because you do not have to raise the whole vehicle over bumps, you only have to raise the wheels. So here we have developed a quite different concept of transport, arising again from a simple provocation.

If we look at the way the brain handles information, inevitably our experience causes us to use the established patterns. There is no occasion on which we could analyse the field of potential concepts — the mathematics of combination are so huge that we just could not begin to analyse all the concepts that are potentially available. Traditionally we have never developed ways of cutting across patterns. We have always moved up and down them, moving in one direction with induction and in the other with deduction. Once we appreciate that we are working in a patterning universe, very different approaches start to emerge. Here, I have only scratched the surface of the field. There are whole areas of other ideas associated with the realisation that we operate in a patterning universe.

One such aspect is concerned with perception in general. One of the biggest problems we have in thinking is that of point-to-point thinking, which means that we hit upon a particular pattern and then at each point we follow the widest pattern. This means for example that if we are looking for something on a street map, we would follow what at each moment is the widest and best-established pattern, and we would not scan the whole area.

During an exercise with 24 groups of schoolchildren in London, the suggestion was made that bread, fish and milk be free. Some of the children came from very poor backgrounds and said that they only had milk on the few occasions their parents could afford it. Yet 23 out of the 24 groups decided that the suggestion was a bad idea. They were using classical point-to-point thinking. For instance, they would say that if such food were free then everyone would want it, therefore the shops would be crowded, therefore the buses going to the shops would be crowded, therefore the bus drivers would want more money, the bus drivers would not get more money and so they would go on strike, then other people would go on strike, and so it is a bad idea. This is classical point-to-point thinking of the kind you see happening all the time.

In order to get round this problem we need to create some metacognition or scanning strategies that allow us to scan the field more completely. In a huge programme of teaching thinking in schools in Venezuela, England, Australia, Canada, the United States, and other countries (a programme we call PMI) the first lesson we teach is simply to get the children to scan in the plus direction, in the minus direction, and in the interesting direction. They scan in all these directions, and of course once they have thought something then they cannot unthink it — each thought has a lasting effect. When they have scanned in all directions, they make their decision.

Once in Sydney I was talking to a group of 30 children. I asked them if it would be a good idea if they were each given five dollars a week for going to school. Every one of them said that it was a great idea — they could use the money to buy sweets, chewing gum and comics. Then, during the course of the lesson, in separate groups of five, they went through this scanning procedure of PMI. I did not say another word about the idea of being paid to attend school. At the end of the lesson they reported back to me and they said things such as, "The bigger boys would beat us up and take the money off us," "The school would raise its charge for meals," "Parents wouldn't give us so many presents," "There would be less money for teachers and school equipment," and 29 out of 30 had totally changed their minds about it being a good idea.

This sort of experiment is highly repeatable, with sophisticated adults as well as with children.

An important point about scanning is that if you record the characteristics in a situation and then only afterwards decide which goes into the plus, minus, and interesting categories then you are not using the same process — you are using a judgment process instead. The scanning process is quite different from a judgment process. There is growing evidence that when you are looking in say the plus direction you are actually using, as it were, a different brain — the chemical setting in your brain is actually different from what it is when you are looking in the minus direction.

Several aspects are now emerging where what we know about neurochemistry, neurophysiology, and self-organising systems are coming together to give us practical guidance to help us in our thinking.

My next point relates to something known as the intelligence trap. One of the biggest problems in education is the notion that if you are intelligent you will be a good thinker. That is probably one of the most disastrous fallacies in education. The reason why it is disastrous is that people in education feel that if you are highly intelligent then nothing needs to be done about your thinking, and if you are moderately intelligent then nothing can be done about your thinking. The result is that nothing is done about developing anybody's thinking skills.

Not only are intelligent people often not especially good thinkers but they may actually be worse thinkers. The reasons for this are complex — some are physiological and some are sociological. For example, intelligent people can take up a position on something and use their thinking to back up that position. The more coherent, rational and sound that position is, the less they ever see the need to explore the subject. Since you can construct evidence arrangement through a hypothesis, you may never get to explore the subject — instead you stay locked into one particular view.

The speed of processing of an intelligent mind is such that from a few signals it can interpret a situation. A slower mind, which has to take in more signals, may get a better view. An intelligent person bases his self esteem on being right and so is much more unwilling to speculate and be creative. The best form of achievement for an intelligent mind is to prove someone else wrong. The reason that is the best form of achievement is that it is immediate and complete. If you have a constructive idea, that has no value unless your listener thinks it is any good. It is rather like telling a joke — if your listener does not laugh, you have not told a joke.

There are about 14 different reasons why highly intelligent people, outside a limited sphere, tend to be rather ineffective thinkers. Some of the world's leading schools for gifted children are now deliberately teaching thinking as an heuristic strategy rather than relying on intelligence as such.

So, within our world of information, there is the traditional way in which we handle information, and there is also the perceptual area of thinking. Most of the traditional ways in which we handle information (for example induction and deduction) will be taken over by data processing systems, information nets and so on. The perceptual area we will still have to perform ourselves. We have only just begun to realise the parameters of perceptual processing in a self-organising patterning system. We are only just beginning to scratch the surface. In the future, development of our understanding of perceptual processing will make a huge difference to our lives.

Let us now look at some of the aspects of data processing. One aspect of interest is that, although we can make information available, sort it, store it, and transfer it, in the end there is still some sort of human interface. The automated systems do not run the information right through into a decision making process. We actually make the decisions ourselves.

Another aspect is that the sheer volume of information available will be a limiting factor. We may have all the data we need available, but how are we going to interact with it? In this context, development of the word processor may prove to have been a very bad thing. If we are travelling down the main track in figure 7 but really should be travelling up the side track towards B, the word processor has allowed us to move further in the wrong direction. In other words, there will come a time, where because of the sheer volume of information with which we need to deal, we will have to create a newer, higher order language for dealing with situations, but the word processor has delayed our doing this. Dealing with situations in ordinary language is much too cumbersome. We need a higher order of language, probably about two orders of magnitude more concise than ordinary language. This is something I am working on — a special, learnable language for human interaction with events — not with machines but with events — so that our thinking and expression of them can be much more effective. In general though this has been delayed by the development of better ways of dealing with cumbersome, ordinary languages.

Another example is that our ability to store records has kept us on the record storing track. For example, an insurance company keeps a record of each person who has an insurance policy with them. With the current, improved filing systems, insurance companies can keep a lot of records. If we did not have such improved storage methods we may have changed to a different concept of insurance, for example where there are no longer one-to-one matched records. Instead some sort of actuarial principle could be applied under which insurance companies pay out the claim without having to keep matched records, which is a whole different way of handling information than on the one-to-one basis.

That is an example of where the availability of support for doing things the way we have done them traditionally has delayed the point at which we could conceptually restructure our approach.

Another interesting example is that of privacy. About 18 months ago in America there was a big fuss because one of the data networks in Canada was broken into and the records of some of the larger companies, including a cement company, were obliterated. Fortunately, duplicate records had been kept. It was traced back to a couple of schoolboys in New York, who had used an ordinary classroom terminal and had broken through all the security codes and accessed the information. I think it is pretty well accepted that there is no such thing as a security code that cannot be broken. So to rely on security codes is rather out of date.

In the future, we may well arrange privacy — in terms of the public interfacing with systems — by giving each individual four or five identities. For example, one person could be treated as Mr. Smith (health), Mr. Jones (financial), and Mr. Brown (legal). All these would be completely different people in terms of their records, with no cross-linking at all. All this person's health records would be assembled under Smith, all his financial records under Jones, and all his legal records under Brown, with no cross-linking between these three sets of records at all. That way privacy can be defined while retaining the benefits of data collection.

In some areas the idea of free data carrying systems is already eroding. There is a growing national protectionism. For example, I am under the impression that Lufthansa and Varig in Brazil have backed out of the SWIFT air reservation system. They seem to have decided that they no longer want to be part of an international network — they would rather concentrate on their own national airline and, if this means that it takes a couple of days to book on for example Pan-Am, that is no bad thing because it will increase the business for their national airlines. So there is a trend towards national protection in the area of data transfer and data sources.

I shall now look at another aspect of people situations. If you are dealing with someone and he does not agree with you, you have two options. You can either think that he is stupid, ignorant and prejudiced. Or you can assume that he is highly intelligent, but working within what we could call a different logic bubble. A logic bubble is the set of perceptions of circumstance and structure within which a person acts. If you take the first option then you are not going to achieve very much. If you take the second then you will be able to achieve something, even if it is only a different construct.

I will give an example. A motor company was once experiencing a lot of wildcat strikes. People would just walk off the assembly line, go and sit on the grass for an hour or two, and then come back. The chief executive there knew something of my work and became interested in the logic bubble idea. As a result, he instituted a very small payment — I think it was about $10 a week, which was a fraction of a normal week's pay — and you qualified for this payment if you completed the week without walking off the job. This payment was in no sense a bribe — he could have raised the people's pay and it would not have helped with the problem, because giving people more money is not in itself a solution to the problem. The difference the $10 payment made was that when someone suggested that they should walk off the job, people were motivated to ask why. This so slowed down the previously explosive reaction that the frequency of wildcat strikes dropped to one-sixth of what it had been.

Let us look at it another way. Let us ask why people in the civil service or the public service do not innovate more. Suppose such a person does innovate and makes a mistake. That mistake will be round his neck for the rest of his career — he will be known as "that guy who made a mistake". In the public service you cannot recover from making a mistake. As an entrepreneur you can have successes and failures. In Europe you are still visible as a failure, but in America you are invisible as a failure because American society is so success-oriented. Indeed, if in America you fail and then come back, venture capitalists are more willing to back you because they feel it is advantageous for you to have gained some experience of failure. So there is no difficulty in moving from failure to success in an entrepreneurial environment. In the public service, however, this is impossible.

Suppose someone in the public service does have a good idea. Several things can happen. One is that people say, "That's a great idea, why didn't you have it five or ten years ago? Think of all the money we have wasted because we have been doing it the wrong way!"

Another possibility is that he has the idea, it works, it is timely, and it could not have been done before because the technology was not available. When the selection group comes to appointing the head of department, it says, "Yes, Joe did have an idea. He is an ideas man, but we have no guarantee that his other ideas will be any good, so let us appoint a sound man to be head of the department, someone who doesn't have ideas." In the logic bubble of the people saying that, it is considered that it is not intelligent behaviour to innovate.

To counter that, I once made a suggestion that is perfectly logical, but also totally unacceptable. My suggestion was that, if in a service organisation you could genuinely abolish your own job, you should be given full pay to retirement age and pension thereafter. The reason behind this is very straightforward. If you did not abolish your job, you would have received such income through being in the job. If you abolish your job, however, then you abolish all the support and ancillary costs that your being in the job would generate, and so there would be a saving.

Having abolished that job, you could then move into another job and abolish that as well, and have the two salaries for the rest of your life. You could sweep through organisations, munching up jobs. Economically, this is perfectly logical. For moral and other reasons however we would never accept it.

Another point concerning this whole area of structure is what I call Catch 24. As you probably know, Catch 22 comes from Joseph Heller's book of that name which is about a fighter squadron where there is a high mortality rate. One of the pilots does not want to be killed, so he tells the medical officer that he cannot fly, because he is going mad. The medical officer replies that not to want to fly is such a sane thing that the pilot cannot possibly be going mad, and therefore has no excuse not to fly. That is Catch 22 — in order to achieve one thing you have to do another thing that makes it impossible to achieve the first thing. It is rather like when I talk to education audiences, who want me to make what I say so complicated they will be impressed but unable to use it.

Catch 24 states that something may be a good idea, except that it is not a good idea at any particular point in time. That is a perfectly logical thing to say — an idea can be a good idea, but not at any particular point in time.

I will give you some examples. In the academic world everyone says that we need generalists to cut across all the lines of speciality. That is a good idea, and everyone agrees that it is necessary. But at any particular moment in time, on a particular date, at a particular place, or for a particular position, the specialist is superior, and so he gets the job. So there is an example where in general the idea is a good one, but at any particular point in time it is not logical to use the idea.

Another example is provided in the search for opportunities. At the top of a business cycle, where you have all the market you need, your problems are concerned with production capacity and you are not really looking for more ideas and opportunities. At the bottom of the cycle, survival is the name of the game, and you do not have money to invest in speculative efforts. On the way up and on the way down the uncertainty is such that you think you should wait until the business plateaus and stabilises. So, although in general people in business regard searching for opportunities as essential, there is no one point in time at which it is logical behaviour. I am not being sarcastic. I am simply saying that the structure of systems is such that something may be regarded as a good idea, but never at any particular point in time.

My final point I call Catch 23. This says that in order to reach the senior position in an organisation you have to keep hidden, or be without, exactly the qualities you will need when you get there. In other words, on the way up you have got to be a problem solver, a fire fighter, the sort of person who plugs in standard solutions. If you are not then you do not survive, the company may not survive, and you do not get promoted. When you are at the top, you find yourself in a conceptual and strategic area, which contains a quite different set of idioms from those found at lower levels.

If we look at the normal, logical structure of organisations many shortfalls and inadequacies arise. One of our unfortunate hang-ups is that we regard such shortfalls or inadequacies as being due to ignorance or ill will. Such shortfalls and inadequacies in fact very rarely arise from stupidity, rigidity, malice, or such like, and we would usually be much better off to consider the logic bubble in which the people involved are operating.

THE INFORMATION SOCIETY: A STRUCTURAL DEFINITION

Marc Porat, North Star Communications

Marc Porat is an associate director of Program on Communications and Society at Aspen Institute for Humanistic Studies and a principal of his own company. He has been principal investigator on major studies in communications and information policies for many organisations. He is the United States' representative and consultant to the Computer, Communications and Information Policy Group of the OECD.

He has written on "The Information Economy" and has developed a TV documentary on "The Information Society".

It is now a widely held, and valid, view that the United States has evolved from a manufacturing to an information-based economy.

In neo-classical economics, value was created by certain technical combinations of labour, capital, and natural resources. In fact that was the way in which all economic organisation was understood. As a result, all of management science dealt with using rational combinations of labour, capital, and natural resources to produce value, which in turn became income, profits, and exports. We got very good at that and started building ever more sophisticated and complex models of how to take labour, capital, and natural resources and produce income, profits, and exports. We gave such modelling work names like operations research. This work dealt with orders of efficiency and precision that were quite staggering. The work of obtaining definitive answers from large-scale number crunching in operations research also proved to be a lot of fun.

In the midst of all this, there is the insight that information is a resource in itself, a resource that can be divided into information labour and information capital. These are not inputs that are generally measured — certainly not in neo-classical economics. There are two reasons for this. First, they have never been conceptualised as such. Second, they are very difficult to measure as inputs, although they are much easier to measure as outputs.

So the basic idea is that information is a vital input. It gets mixed in with the non-information resources of labour and capital. Then the engine of production creates from this mixture employment, profits, and so on, which presumably are good things.

Let us take a closer look at what we mean by information labour and see whether we can quantify it. I did some work a few years ago which was replicated in many countries, for example, Canada, England, France, Germany, Australia, Japan, and Spain. These countries asked themselves a very simple question. They asked themselves how many people in their workforce are engaged in manipulating symbols and information rather than working with material things. It is impossible to have a definition of an information worker that everyone is happy with. So I will not even try to develop one. Our researchers however did develop a critical test in trying to figure out how much of the workforce is involved in information. They decided that if a person is mostly sitting around all day long, taking in information as an input, doing something with it, and his or her product is information also, that person would be categorised as an information worker. Information workers are people who are outside the sphere of manipulating matter or energy.

With that simple test, we then went through all of the national statistics and data, having defined hundreds and hundreds of different categories of work. Some categories were not particularly concerned with intellectual or knowledge-based tasks; many were processing-oriented jobs such as clerical tasks of one sort or another; some were very transaction oriented (where, for example, a manager would have to say to a person, "Do this on the basis of that."). Some people however were involved in intellectual property — they were involved in what, in a non-value sense, is called the higher uses of information, such as knowledge, judgment, decision making, wisdom, poetry, and so on.

This gave us a very broad brush with which to paint the picture called the information workforce. I might add, however, that the brush had within it a multitude of very fine hairs — there was a multitude of definitions based on trying to decide whether or not a secretary or a manager is an information worker. In making these decisions, we always took the cautious, conservative approach — for example, we threw out all physicians even though they are clearly people who deal with knowledge and use it in their work.

This gave us a taxonomy that was comfortable not only within the United States, but internationally.

The data for the United States produced results as shown in Figure 1. Many of you will be familiar with the concept that half the workforce in the United States is now involved in information. In 1900, agricultural workers swamped other types of workers in numbers. We could characterise that period, and everything that happened up to that period in the United States, as being an agricultural society.

Figure 1

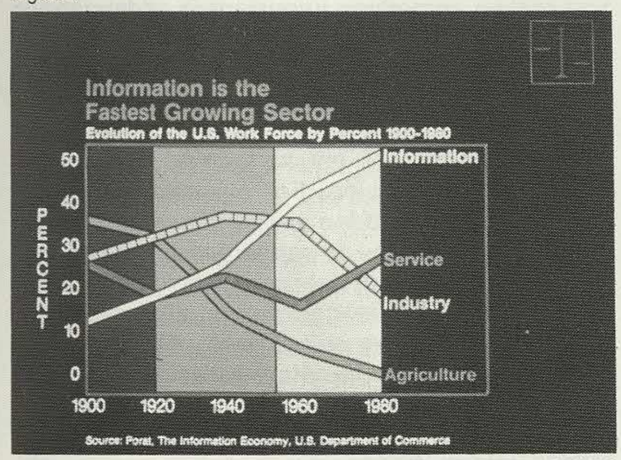

In the next band (1940), the United States (and many other countries) was at the height of what was then called the industrial society. The preponderance of workers were based in factories spending all day manipulating matter or energy, or both. The number of service workers declined somewhat and the number of information workers started to grow.

In 1980, there were very few agricultural workers. Industrial workers as a group were shrinking. Personal service workers were increasing, mostly because of medical industry workers such as nurses and technicians. The information workforce, as defined before, is exploding.

Over 50 per cent of the workforce today are people who fall into the category of information workers. Most of them are bureaucrats in both private and public bureaucracies. Most of them are educated. Most of them work with machines. Most of them have an enormous amount of support staff, equipment and personnel behind them.

You can see that the industrial workforce has collapsed, although in the industrial workforce there are many jobs now that are much more information intensive. I think that is something of great concern to us all and I will talk about it further in a moment.

In services, we differentiate between personal service on the one hand and information service on the other. That is important, because whereas there is no way of improving the productivity of a personal service worker, such as someone who works in a restaurant or in the transportation sector, information technology has a great deal to do with improving or changing the productivity of people in the information sector. So information technology addresses the needs of information workers, but only tangentially addresses the needs of personal service workers and industrial workers.

Agriculture now accounts for about 2 per cent to 2½ per cent of the workforce.

This data has been replicated for the various OECD nations that I mentioned. You can draw a family of curves that are almost parallel to those for the United States. All these countries are changing equally rapidly, although their stage of development lags somewhat behind that of the United States.

As a result of this approach, which says that one can characterise where a nation is at any particular time by what people do, we have reached the conclusion that the United States has an information-based workforce and that the requirements — human skills, organisational skills, communication skills, and numerical skills — of the majority of the workforce now have to do with their ability to manipulate information. The extent to which our people are good at manipulating information will determine the future for America.

The emphasis that we as a nation have put on the service industries, manufacturing industries, and agricultural industries has diminished. And this is reflected in the way that the laws, the legal structures, the regulatory structures, and so on, have developed (I shall talk about these later).

The issue of what people do to earn a living is very important because 80 per cent or so of the national income arises from wages. We can see now that most of America in earning its livelihood is dealing with information.

There is another way of characterising what we mean by an information economy, aside from just looking at what people do for a living. We can look at what the different sectors of industry are doing. There is a prejudice that most of the people in the information workforce are government bureaucrats and that the growth is in the government area. This is a fiction. It turns out that government administration, education, and R&D are not growing very quickly. In fact they are diminishing as a percentage of the information workforce. Most of the growth in the information workforce is taking place in large corporate organisations and, to some extent, in small businesses that find the need to have a very small bureaucracy within the business, for example, an accounting office. These trends are shown in figure 2.

Figure 2

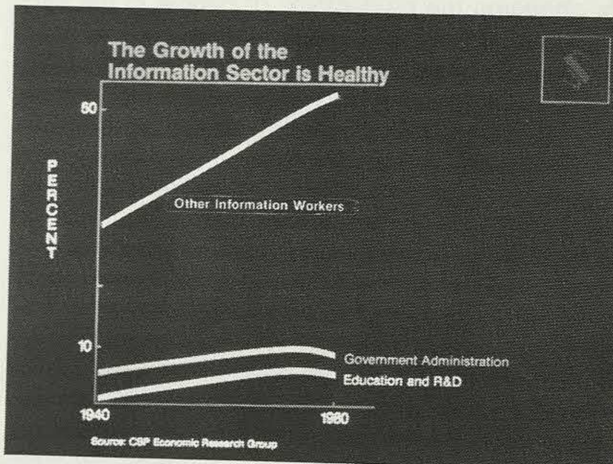

As another way of attacking the question of what constitutes an information economy, we defined categories of industrial activity that are generally involved in information. The first category is intuitively clear. It is those industries that are per se engaged in the production, processing and distribution of an information product or an information service.

This is obviously the computer equipment and communications equipment suppliers, whose primary purpose is to transform information by processing it in a fashion that is predetermined. Apparently this is useful — people pay for the transformation to be carried out.

The next category consists of the softer services that ride on the infrastructure of computing and telecommunications. The most obvious industries in this category are broadcasting, telephony, and cable television.

The third category consists of what used to be called traditional service industries but that, when you look at them more closely, turn out to be industries purely involved in the transformation of information. These industries include finance, insurance, advertising, entertainment, and various business services. These are symbolic industries that deal with knowledge and transactions. Nearly all these industries deal with number crunching, and usually they deal also with some form of judgment, so that the value-added factor is derived from the expertise and knowledge of the people involved. Investment banking is a good example of such an industry. Other industries in this category simply involve a person being paid for his ability to make a decision of one sort or another. Lawyers and accountants are examples of such industries.

We call the first two of these categories of industry the primary information sector. We use the word primary because it is these industries' primary function to deal with the information base, or the knowledge base, of society. In 1967 in the United States, the primary information industries accounted for 26 per cent of the value-added factors, such as wages, profits, and so on. Now they account for over 30 per cent of the GNP of the United States. In other Western alliance countries the percentage is somewhat less than that, but it is growing rapidly.

When you look at the primary information industries, you realise something important in understanding what the information economy is all about. Many of the information processing functions in the third category of industrial activity (the traditional service industries) are performed routinely, even though these industries are not concerned directly with information processing. For example, a large oil concern has embedded in it a very sophisticated information industry. It has within it not only the informal resources of decision making and number crunching — the managers, secretaries, clerks and so on — but also very well-defined units such as the photocopying room, the data processing centre, and the telephone switching facility. Each of these resources is identical economically to those in the primary information sector.

We realised that we had to invent a name for a further sector in the information economy — this sector that consists of industries whose output has nothing to do with information but that none the less have a very high degree of information input or of information content.

We called this sector the secondary information sector. The secondary information sector consists essentially of overhead — it is the resources of people and machines and buildings in which a company invests in order to be able to conduct its business. Without such investment in information capability, the company would not be able to produce its output.

By 1980 something over 25 per cent of the US GNP originated in the secondary information sector. Now this figure is up to something over 30 per cent. It has grown in proportions that were unimaginable only 25 or 30 years ago.

In some industries 60 per cent to 70 per cent of the cost of the product represents embedded information costs such as research and development, management, accounting, law, computers, telecommunications, and so on.

This means that companies that have nothing to do with producing information as an output have had to become very sophisticated in how they manage their information resources, given that those resources account for a considerable part of the company's operating cost. In the old days, when only 5 per cent or 10 per cent of the output of a company was represented by information costs, it did not matter if they were inefficient in dealing with that resource. The fact that they could achieve only 50 per cent efficiency in the information area would scarcely reflect in their final output prices. But now, with 50 per cent, 60 per cent, 70 per cent, and more in some cases, of the total product cost represented by the information element, very small differences in the organisation's efficiency in dealing with that resource make tremendous differences at the outut price. To the extent that companies are competitive, this makes a difference to their chances of success.

This brings us to the very difficult problem of assessing the productivity of an information resource. It is much simpler to assess the productivity of a non-information resource, and that is why we all spend time concentrating on non-information resources. There is always a natural inclination to follow the line of least resistance, and the line of least resistance for us as business people is to focus on those things that we understand. Typically the things that we understand are also the things that are easy to define and measure. Information is neither easy to define nor easy to measure.

It is clear that information consumes a huge block of resources. Having said that, however, it is much more difficult to go to the next stage, which is to determine with precision and reliability how to allocate one extra unit of information labour, or one extra unit of information capital, to the production process so that there is some measurable increase in output.

Why is it so difficult? Let us go back to the nature of the commodity about which we are talking. When you are dealing with a strictly manufacturing kind of environment, it is very well defined how non-information labour and non-information capital and resources combine into output. Take the simple case of a man building a brick wall 45 feet long and 6 feet high. You know the capacity of the man to build and you know the size of the bricks. It is therefore a very simple operations research problem to combine these factors into a forecast of how many bricks and how much manpower is needed.

With information, such forecasting is never so easy. What is the marginal contribution of a secretary? What is the marginal contribution of a word processing pool? What is the marginal contribution of a manager? In the light of this morning's presentation on lateral thinking, this comment is made all the more dramatic. There is no real way of measuring, in a long-term sense, the productivity of an information worker. The more highly paid they are and the more creative they are supposed to be, the less sense it makes to measure people's productivity in the information field.

Certain very large companies, for example Xerox, have spent a lot of money trying to convince themselves that they understand this problem. Xerox is in the business of improving office productivity, which means that they are in the business of having to convince managers who make purchasing decisions that Xerox understands how to improve productivity, and how managers, secretaries, and organisational units in the secondary information sector get better, smarter and more profitable by the infusion of information capital. To date, our conceptualisation of that problem, and our methodologies for solving it, are quite primitive.

Even if you have accepted the argument — and many people have not — that there is such a thing as an optimal mix of information resources on the one hand and non-information resources on the other, you have enormous problems planning for optimal output, particularly when you also try to take into account the thousands of different types of information capital and information labour resources that are available on the open market. All you can do is make some educated guesses and go by intuition. What usually happens is that decisions are made more according to political factors than anything else. You have probably seen this happening in your organisations.

The information side of the business is usually where the top management are located. Top management usually does a very good job of protecting itself in terms of resources, equipment, and personnel. So in a recession what you usually see is that the information side of the business gets fired less quickly than the other side. My intuition tells me that it should be just the opposite.

Another aspect of this insoluble problem is that the non-information side of businesses is typically very well defined. If you have a factory with machines in it, those machines are very lumpy resources. You either have a machine or you do not have a machine — you cannot have a bit of a machine. You know what your orders are and what your demand is, and so you can very quickly determine what the capacity utilisation of your shop is going to be. As soon as you have determined that, you can immediately determine how many people you need and how much of other resources such as electricity and coal you need. The task is a very well-defined, mechanical one. The transforms between demand, capital, labour, and resources are fixed and well-known.

On the information side however, if I am facing a recession does that mean I need more salesmen or fewer salesmen? Do I need more advertising or less advertising? There are arguments in both directions. If I am a Chief Executive Officer and there is a recession, with everybody competing much harder with one another for the same scarce dollar of demand, then I could argue that I need to increase the information sector and forget about the other side of my business, because the only competitive advantage I have is at the level of information.

At the manufacturing, industrial, technical level, organisations in the same line of business have much the same equipment. If I do not have the funding to bring my technology up to the state of the art then it is a fairly straightforward task to work out what the differential between my prices and everyone else's prices should be.

On the information side, the problem is quite different. I have research and development scientists, but who knows when they will come up with the next breakthrough? I have lawyers, but who knows when they will make a critical mistake because there are not enough of them, or a critical breakthrough because one of them has had some extra time to be able to fix an anti-trust problem, fix a regulatory problem, or fix a competitor, and so overcome the difficulties of the recession? I do not know what the next lawyer on my staff will come up with, or what the penalty is for not having an extra lawyer. These are unknown questions.

My data processing manager may be saying to me, "If you want to be super-competitive, you must look at what your competitor has just bought — an intelligence network. It is not even a computer network, it is an intelligence network. It is composed of people at sites all around the world, bouncing information off satellites. They have a fifth generation program written in PASCAL holding the whole thing together. You must have one otherwise we will be out of business in three years." Do I believe that man? What is the cost of not believing him? What is the cost of believing him?

My advertising manager says to me, "We have to pour more advertising into the marketplace. We have to do it in ever more sophisticated ways. We must have computerised advertising."

As an aside let me give an example of the sort of computerised advertising that is now being used. You hire a movie actor or a well known sports personality and get them to record on tape a series of soothing messages which are then held in a computerised system. The system is able to dial thousands of households at random, or according to programmed parameters. When someone answers the telephone, the soothing voice says, "Hello, I want you to buy my service. I am . . . will you please talk to me?" If Marilyn Monroe got on the phone and said, "Would you please talk to me?" you would at least want to know what was going on.

If the person at the other end says "Yes" the system branches into the next stage of the selling message on tape. If he says "No" the famous person persuades him to talk. Before he knows it, he has spent 15 minutes talking to a totally computerised tape recorder. In the end, the tape asks him to place an order. For example, if it is a solicitation for a magazine, at the end the tape will say, "Well, does that mean I can have your subscription for three years?" The person says, "No, I think three years is too much." "Well, how about two years?" The system is completely interactive. When the person finally says, "Fine. Two years." the tape says, "We will be sending you a bill. When you receive the bill, please pay it promptly." A month from then the same voice will come on the telephone and say, "Hello, remember me? We have sent you your first magazine but you have not yet paid your bill." This is a new industry. It is alive and well in America.

My advertising manager is telling me that I absolutely must invest $25 million in this new information technology. How do I evaluate that?

Remember that the secondary information sector accounts for 30 per cent or more of the GNP in the United States. This is composed largely of people in corporate bureaucracies, people such as the research and development scientists, data processing managers, and advertising managers I have described. The premise of this discussion is that it is from these areas that the business's competitive edge comes. The other corporate resources are well understood and there is very little that can be done about them that is creative. The information side is where the action is, and we do not understand it, other than intuitively.

There are a lot of consultants now running around trying to make sense of how to rationalise the information side of businesses and make some of the processes involved more intelligible to the decision maker.

The basis of my session is that we now have an information economy, that we are now an information society. Rather than just treating that as a piece of sociology, we tried to demonstrate that it is valid in the economic sense. We tried, in a very systematic way, to use the national income and product accounts to measure the information sector. So we defined a sector and measured it. The results are that the primary information industries are about 24 per cent of the GNP and the secondary information industries are about 28 per cent of the GNP. So you can make a case that the majority of economic activity is wrapped up in information.

In terms of the workforce a much simpler case can be made in that over half the workforce now hold white collar information jobs. And the pay cheques of these workers add up to much more than half of the national income, because most such workers are better paid than the blue collar workers.

So it is not just pop sociology that leads us to say that we are an information society. Our forecasts for the different parts of the US national workforce are shown (in thousands) in figure 3.

Figure 3

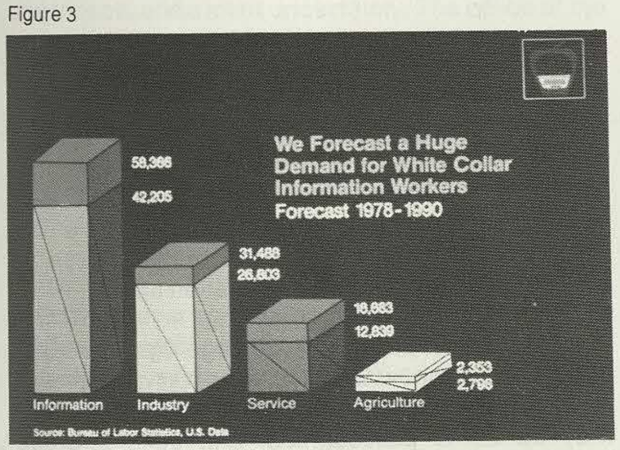

During my work with the OECD, I came across a uniquely European fear — that the microprocessor would abolish all information jobs, and if not all information jobs then certainly all blue collar jobs. There is some validity in that point of view as far as blue collar jobs are concerned. But in the main, what we have seen in the United States, and what we will continue to see over the next decade, is that the information economy has an inexorable, self-sustaining, internal engine of growth. It is almost as if the bureaucracy begets a bureaucracy which begets a bureaucracy and so on. For example, if my company hires a lawyer, your company will hire two. The information workforce manages to sustain incredible growth. As you can see in figure 3, it will grow from 42 million people to 58 million people by the end of the decade, or so we think — we being the US Bureau of Labor Statistics. In fact it can be said that the brightest future for today's children is to become information workers, because that is where the action will be.

I think this will be true also in all European countries. I think that your information workforce will continue to grow, relatively unabated.

This leads one to ask where all this growth in the information sector is leading. Obviously you cannot eat information — you can only eat food. You cannot wear information — you wear clothing. You cannot walk on information — you walk on shoes. You cannot sit on information — you sit on a chair. So the reality of our physical existence has nothing to do with information. The immediate answer is that the information content — in the form of research and development, accounting, data processing, advertising, and so on — is increasing enormously.

Figure 4 shows forecasts of growth rates for a range of occupations. For every occupation there will be an incredible growth in the information workforce.

Figure 5 shows forecasts of growth rates for some of the information industries. There will be some winners and some losers — some companies will go out of business. But to give an example of a success-

Figure 4

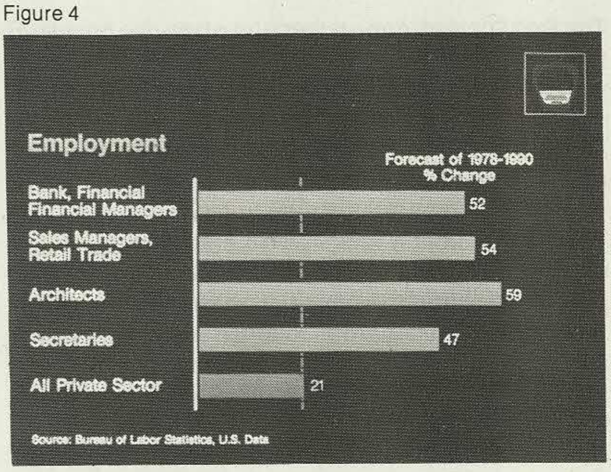

Figure 5

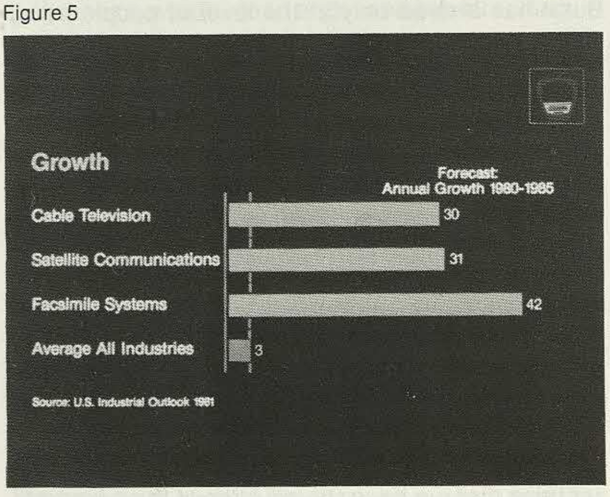

ful company, AT&T is expected to make $670 million more in profits next year than they do this year.

The point is that the information infrastructure, whether at the level of telecommunications, computers, or small off-shoot industries that derive from the information infrastructure, is not going to have a problem There will be some losers within the industry, but for the sector as a whole there will be very substantial growth.

Next I should like to focus on a number of policy considerations that are central to what we need to be thinking about in the information age.

I have brought along a copy of Forbes magazine, which includes the following item.

> "The Bell System is very proud of its concept called telemarketing. Telemarketing for it is a way to seize the switched data network that it has built up and make it do everything for everybody in the world of business in terms of voice, in terms of data, in terms of even its marketing abilities."

The Bell System is spending a lot of money on advertising. It is spending tens of millions of dollars per year talking about the information age and using the words, "We are now in the information age." They are now trying to give some specific and concrete meaning to the concept "information age".

Ten years ago, the concept "information age" was used by Gordon Thompson and about three other people. Now it is trite. When a mass market magazine every week carries two or three messages about the information age, and when companies the size of AT&T and IBM week after week advertise in the popular press to try to give specific lessons about what they mean by the information age, you know that the concept has arrived.

But it has arrived only at the level of people being able to talk about it. It has not arrived at the most fundamental level, which is learning how to make use of the information age in the best way possible.

If the information age is to be taken seriously, it had better address the current concerns of Western nations. I would say that Western nations are obsessed right now with the question of productivity. The information sector is both part of the problem, in the sense that office work is not sufficiently productive, and also part of the solution, in the sense that information makes all industries more productive and competitive.

As far as being part of the problem is concerned, I do not think that we have the courage in the corporate sector to face up to how unproductive, wasteful, foolish, and bizarre some of our procedures are within bureaucracies. Suppose you were able to step back for a second and remove yourself, your ego, and your identity from this question and ask, "How much of what I do and what my group does is truly productive in this world of information?" Or, if that is too personal, step back and say, "How much waste and foolishness have I seen at the level of the information industries?" This is a tough test. I do not mean to be critical but I would say that the information sector within corporate bureaucracy is very wasteful. I think that government bureaucracy is much more wasteful, but we are all used to taking shots at the government — we are not used to taking shots at the corporate sector. People do not know how to improve this situation, although they are trying. They are trying to understand what it means to make people and processes in the world of information more useful. That is probably the greatest issue facing the information age over the next ten years — the issue of understanding our own processes.

As far as being part of the solution is concerned, there is a long list, which you may have heard many times before, of all the wonderful things that information machines can do. Information tools are augmenting our ability to communicate and to process information. From that central assumption flows thousands of examples of things that would be unimaginable, unheard of, preposterous even to dream about, without the information age.

I would say that the information age right now is the most important, and maybe the only, engine of change that we have at the level of industry. There may be some philosophical and cultural breakthroughs that have nothing to do with information. But if we are talking about economics and employment, the information age is the engine of change. We had better understand what that means and pay a lot of attention to it.

Let us look at the role of the information economy in terms of inflation and recession. Slow productivity growth is a major cause of inflation and recession. However, the information economy can ease inflation, create jobs, and increase output. For inflation, the information industry is the only industry where there are price reductions. It is the only industry that is conservative of resources. It is an industry where you can take the information resource and economise on labour.

There are lots of arguments concerning the information economy, but they must be understood much better than they are today. They have to be put forward by people who are thinking not only of selling equipment, but thinking academically. The academic research community has to figure out the linkages between the information economy and inflation.

I alluded to job creation earlier. It is clearly not the case that the information economy destroys jobs overall. On aggregate the information economy creates jobs. In fact that has been one of the problems. It creates so many jobs and so many people want to work in the bureaucracies that we have a secondary problem of what to do with all these people. There is no doubt that, to the extent that the information economy is also responsible for brand-new industries, it is a job creation process. The list of new industries, both technical and service industries, that are derivative of the information age is huge. This was not so of the industrial age, the agricultural age, or the personal service age.

Finally, I would like to discuss the convergence of the industries that make up the primary information sector. As a result of this convergence there is a lot of jockeying going on, a lot of industrial competition and friction, and a lot of uncertainty. The office of the future, for example, is drawing into battle several different kinds of technology that traditionally have had discrete existences — computers, databases, telephone networks, office equipment, cable TV, and satellites. A lot of industries are being thrown together, by necessity, because they are convergent activities.

If the information economy means anything, it means that every single resource that has to do with information will get drawn into a nexus that will operate as one environment, or what Gordon Thompson calls one shared space. To do this properly will be difficult because it involves the violation of habit, tradition, and established industry boundaries. The existing industries become defensive, which causes them to become reactionary and start suing one another. That is the situation we are now facing in the United States. To accomplish a transition from where we are today, with a lot of very healthy, separate industries involved in the information age, to where we need to be tomorrow will require these industries to change very rapidly, to accept a blurring of industry lines and a blurring of function, and to co-operate with each other. All kinds of changes are needed that are not catered for in the law, or currently in the personal habits of the people running the industries.

To take one example, let us look at electronic publishing. In this industry there are several discrete industries that have to start working together co-operatively, or find some way of accommodating each other if they are to succeed. It is clear that electronic publishing is an industry of the future. At present it is already an industry of over a billion dollars in revenue in the United States. The future industry for electronic publishing and the electronic transmission of information, in the home and in the office, has to involve a wide variety of industries — newspapers, publishing, television advertising, cable television networks in the United States and their equivalent in Europe, telephone, computers and electronic components, sales and retailing, banking and other kinds of financial services. These are people who have typically never talked to one another, except tangentially or in terms of some very specialised kind of requirement. These people now have to approach a brand-new industry with a brand-new set of expectations of their behaviour. It is a foregone conclusion that this will not be successful immediately. When these separate industries start communicating there will be friction, misunderstanding, fights, jealousy. There will be technical incompatibilities, standards problems, and so on. In general there will be a period of turmoil.

Yet the information age has to go from here to a point very soon in the future where such an industrial grouping operates like a Swiss watch — reliably, smoothly, with no problems.

This kind of convergence is not only happening between industries, it is probably happening within individual user companies, each facing its own technical future in its own information age. Each organisation has an endowment of information capital — telephones, calculators, computers, and so on. It has an endowment of information workers — managers, secretaries, data processing staff, and computer scientists. And now it has to make some critical decisions about that environment. Somehow the separate elements of the organisation have to agree on a future, probably not by design but by default, in which they can communicate and coexist.

To expect perfect harmony in such circumstances, even within one company undergoing such a transition, is probably asking too much. We will all see foolishness, waste, dead ends, and backtracking, in our organisations. If this is true for one company, it is true for the multitude of companies, it is true for whole industries, it is true for whole nations.

There are instructive historical parallels with the start of the industrial age, which I shall just briefly mention. When the industrial age started, an enormous number of errors were generated. We did all kinds of foolish things, not only technically but also socially and legally. Eventually however everything worked out, and so it will with the information age. We will come to understand problems that we now do not understand, such as, for example, the protection of intellectual property.

The next point I would like to make is again one that operates on several levels. It operates at a national level, it operates at an industry level, and it probably operates at the level of individual organisations. It is what has come to be known as the sunrise-sunset debate. This debate is ultimately not an economic question but a psychological one. At the national level, this debate concerns the fact that any country faces a multitude of investment demands by different industries but has only very fixed resources to allocate to those industries and therefore has to make choices. The choices now are classically divided between those sectors of the economy known as the "sunrise" sectors, which are part of a new age, or between the "sunset" industries, which are industries that have passed their zenith and are now drawing to a close.

If political and psychological forces come together so that the sunset industries, for all kinds of sentimental, psychological and human considerations, are allowed to capture the majority of available investment funds, that is a disaster. Everybody knows this is the case, but there is still a problem in overcoming the human dimension.

What is curious to an American is how obsessed Europeans have become with the notion of displacement and the unretrainability of people in industry. In the United States, and in Japan, there is no such problem. It is simply taken for granted that the sunrise-sunset debate has to come out on the side of sunrise. The short-term dislocations are something that we will have to live with, but they are manageable. The obvious reason for the European attitude is that your unions are stronger than the unions in the US, but I wonder whether there is a deeper answer than that.

Finally, as we enter the information age, the future of all our own jobs and the future of our children's jobs and well-being will very rapidly fall in the realm of information and information handling, which means that we should take a very close look at how well we are training our youth for entry into the information age.

In America, the situation with regard to such training is dismal and depressing. The children who are coming out of high schools in America have a problem with what is referred to as functional literacy. Their ability to comprehend instructions that to us are very ordinary is diminishing rapidly. At the same time, the requirements for them to be able to understand instructions, processes, and how to manipulate information are increasing tremendously.

One example of such a requirement is provided by automated bank teller machines. Banks have discovered something that is intuitively clear to most of us — that all they are is information factories. The cost of running an information factory is very high. So bankers have decided to adopt the supermarket concept of pushing the cost onto the customer. The customer is given a card and made to do the work.

In order to use such a facility, first, you have to have a bank account — and many people, especially poor people, do not have bank accounts. Secondly, you probably have to live in a metropolitan environment. Thirdly, you have to be capable of the fairly sophisticated action required to operate one of these machines. It turns out that some consumers have a real inability to understand the simple instructions that it takes to operate an automated teller machine. The ability to interface with any organisational procedure that is involved with information is diminishing, particularly among children coming out of school.

The fact that the quality of the workforce is degrading very quickly is not a point that has yet been generally recognised in the corporate sector. However, it is something that is actually happening, mostly as a result of the level of schooling and a lack of consciousness of what it means to be in an information age. It seems to me that society should pay a lot of attention to the schooling of children to prepare them for the information age. This does not necessarily mean that they have to learn how to program a computer, although that may well happen. As an aside, you may have read with interest that Apple has proposed legislation to give a computer to every school district in America. There are about 85,000 school districts in the US, and so they propose to give away 85,000 computers, each costing at retail level about $5,000, although their true cost is more like $2,000 each. They want children to have learnt on an Apple so that when they grow up they will buy an Apple.

But I am not talking about computer literacy, I am talking about straight functional literacy — the ability to interact in a complex information society. We do not have functional literacy in our society, and to me that is an alarming message.

I hope that I have given you a sense that the information age is an economic reality. It is not just an advertising slogan or a piece of pop sociology, it is an economic reality. Secondly, I hope that I have helped you to position your jobs and what your organisations are doing within a very rapidly changing context.

INFORMATION ECONOMICS

Gordon Thompson, Bell Northern Research

Gordon Thompson graduated from the University of Toronto in 1947 with a degree in Engineering Physics. He joined the Northern Electric Company, forerunner of Northern Telecom, where he was involved with the design and development of commercial electronic communications equipment and systems.

In 1963 he joined Bell Northern Research and began to question how decisions were made about the new technology. This led to an expanding examination of the relationship between information technology and society. His current position is that of Manager, Communication Studies.

He has published 40 or more papers on the subject of communications and the future, including the well-known paper "Memo from Mercury". He holds 13 Canadian and 11 American patents.

Information technology is quite different from its older antecedent, industrial technology. Because this is so, we cannot expect to know intuitively the best way to use it in terms of producing the maximum benefits. A comparison of the inflation rates, workforce expansion, and unemployment situation experienced in Britain as industrial employment rose from virtually zero to one third of the workforce with what is happening today corroborates this essential difference. During the period from 1765 to 1815, the average inflation rate in Britain was under 1.4 per cent per year, the labour force grew by 40 per cent over the period, and there was what can be regarded as little unemployment, for on occasions rather drastic measures were used to find workers.

Over the past decade, silicon technology has improved its cost/performance ratio in terms of carrying out logical operations by three orders of magnitude. There is still at least another order of magnitude improvement yet to come. The productivity of making cotton rose by a mere 200 times as a result of the industrial revolution. Here we are talking about a ten thousand times improvement. Such a large ratio is simply inconceivable, and cannot be fully appreciated. Consequently, our ability to realize meaningful applications of information technology are very limited. There is simply no precedent for us to build upon.

Since we have not done as well as our forefathers did in applying new technology, and as depending on precedent will not work, a novel approach to discovering meaningful applications of information technology is clearly required. It is hoped that what follows is one such approach.

Innovation is the application of technology. The application of technology produces impacts on the host society, and it is these impacts that characterize the innovation. Just as the physicist has never really seen an electron, and knows it only by its impacts, so we can only know an innovation by its impacts. Innovation, when viewed in terms of the impacts produced on the host society, can be divided into two classes, the *intensive* class, and the *extensive* class.

The intensive class of innovation contains all those very important applications of technology where the major impact is one of intensification of already established processes. The class is characterized by notions of cost saving, efficiency, labour release, and substitution of new methods to do old things. The western idea of productivity-gain describes the major aspect of this class of innovation.

The extensive class of innovations contains all those innovations where the impacts can be described as extending widely across society and deeply throughout time. Here, the essential notions are of labour absorption, creation of new sources of wealth, and the doing of new things. Innovations of this class tend to be non-conservative in the sense that they allow an opening of the otherwise closed economic system, temporarily destroying the conservatism of the "zero-sum game" economy. In a deep and creative sense, this class of innovation produces fundamental increases in mankind's productivity by opening up whole new opportunity areas.

Information technology, it would appear, is being applied as if the extensive class of innovations is inhibited. It is very easy to sell a computer system that saves a corporation money. Unemployment as a result of technological change is no trivial issue. If there was a significant level of extensive innovations of information technology in our society today, our overall economic health would be profoundly better, and unemployment would be a non-issue.

It can therefore be hypothesised that the intensive class of innovations involving information technology is proceeding with alacrity while the extensive class is inhibited from occurring.

One can test such an hypothesis by imagining that it were true, and arguing how the world would appear under that assumption. The result of this is then compared with the real world as it exists. Were we to assume that only the intensive class of innovations involving applications of information technology was functioning, and proceeded to make a list of the new services one might expect, we would quickly see that it corresponds exactly with the lists we see so frequently. Consequently, we have something more than an incorrect or invalid hypothesis. Since there is no proof that there are any extensive innovations at all, we can only state that the hypothesis is a non-trivial one.

As there would be little point of continuing if the extensive class were empty, we shall assume that some such innovations might exist if only we could overcome the constraints that are acting to inhibit their appearance. The challenge is to identify the possible constraints.

Three candidates can be put forth as possible constraining forces acting to inhibit those applications of information technology that could have profound positive economic repercussions. The first stems from the recalcitrant economic behaviour of information, when stripped of any physical embodyment, as an economic good. The second constraint stems from the nature of real language, and the problems machines encounter when dealing with natural languages. The third constraint is the amount of adult learning that is required before massive applications of the technology can go beyond the level of trivial trinkets, given our common lack of understanding of the relationship between technology and people. The constraints are now discussed in order.

What can be measured can be easily talked about. What cannot be measured tends to get left out. Because computers count so efficiently, they have given us the power to measure many things easily. However, they can only count using the metrics we develop. The new technology has made it comfortable to hide behind the old metrics, and so avoid dealing with things for which no metric exists.

This reliance on old metrics makes it easy to sell a computer that increases productivity and efficiency, which are really euphemisms for getting rid of jobs. It is virtually impossible to conceive of, let alone sell, a computer system that creates wealth, that does new things instead of merely old things in new ways, that absorbs labour instead of releasing it, and whose benefits lie beyond our widely accepted, blessed and revered metrics. If it is to ease our present economic malaise, the new technology must do more than merely replace people.

A technology, like a communications medium, can be thought of as having a content. The content of industrial technology is the familiar mass produced, tangible, consumer good. The content of information technology, by extrapolation, must be an intangible or ethereal consumer good.

The miracle of the new technology must be more than a never ending stream of calculators, talking clocks, and other neat toys, all of which are just variants on the tangible consumer goods theme. Our old familiar metrics make this a comfortable vision. But it leads nowhere. It is easier to apply robots to the manufacture of computers than automobiles. Computers are smaller, cleaner and lighter. The new technology brings about the demise of the very jobs created by the adoption of that technology, when it becomes its own content.

Instead of satisfying ourselves with wealth created by *making* the technology, we must seek out ways of *using* the technology that are perceived as creating wealth. For this to happen, we need a new good, and the corresponding new metric. The good and the metric go together, as a matched pair. There is little utility in the good without the metric. The metric is meaningless without the good to meter.

The new good is information, true, but a special kind. It is unembodied, or ethereal information. The technology now allows, for the first time in a massive way, information, free of any physical embodyment, to behave as a private economic good in ways that invite marketplace kinds of transaction. Books and phonograph records behaved in the marketplace more as tangible goods than as some kind of information good. Now that the umbilical cord to the pulp or plastic is shrivelling, we have to deal with the information itself instead of hiding behind its intimate connection to a physical embodyment.

However, as a society, we are completely unprepared for the challenges presented by this new development. The a priori determination of an individual's probable use-value for such an ethereal good is very difficult. We simply do not have the requisite skills or tools.

Massively produced and massively consumed, unembodied information, free of its physical carriers, like books, phonograph records and floppy disks, is the new economic good. It is an intangible good, and like most information, can grow in value with use, a property not shared with the more familiar tangible consumption goods. Inherent in the definition of this new ethereal good are the ideas that it is both difficult to evaluate and cheaper to copy than buy. This last issue must be met from the start if the intrinsic differences between tangible and ethereal goods are to be properly incorporated in the design of a suitable environment for economic-like activity centred around exchanges involving this new good.

These ideas, when put together, define the 'ethereal good' as being intangible, difficult to evaluate, widely produced, widely available for massive consumption, appropriable but not expropriable, and cheaper to copy than buy.

The ethereal good is designed to be the economic atom of a viable, wealth-creating, information society.

The profound differences in common behaviour towards the theft of tangible and ethereal goods attests to the fundamental differences between these two kinds of good. Most people would not describe the taping of a phonograph album on a cassette recorder as being deviant behaviour, but they probably would describe the shop-lifting of that same album as deviant behaviour. If we are to get the real value out of this new information technology, we cannot continue to hide this difference and to force ethereal goods to behave like tangible goods by rigorously applying copyright laws. It is better to dig deeper and seek out the causes of this behaviour anomaly.

The intimate involvement of society at large in the process of evaluation of ethereal goods might be the basis for society's behaviour towards their theft. A song becomes popular as people espouse it, become committed to it, and remember it. Today we have tools for memory, like tape recorders, and committing a song to this form of memory is as much a tribute to that song as memorizing the tune. It is an outward and visible sign of espousal or commitment. It could be argued that since such commitment is the way society contributes to the value of a song, we are not really stealing, but rather are enhancing the song's value when we copy it off the airwaves. Perhaps one is merely taking what is partly his own, if the labour and commitment entailed in making that copy does tend to increase the perceived value of the song.

If this intangible, massively produced and consumed ethereal good of unembodied information is to become a part of an economic system, a meaningful approach to estimating individual use-values must emerge. It is easy to estimate one's use-value for razor blades or bananas. The conventional monetary metric works well for these exchanges. But this is not the case for a screenfull of information. How much is a screenfull worth? The problem of evaluating information has been, and, except for the arrival of this new technology, always will be, the greatest challenge facing any society.

Today, behaviour in our society would suggest that we assess our information, to a large degree, in terms of the ability to attract attention. Mass-media, television, newspapers and radio, demand so much of our attention that their selection rules become the basis for society's rules for assessing information. The things that are shown on these media are selected on the basis of attention grabbing. Is this a reasonable way to decide such issues as nuclear power, acid rain, abortion and the like? The Neilsen Rating is too simple a metric for evaluating a high-technology society's information. For a viable information society to emerge, a more complex metric must be developed, one that better suits the complex needs of its citizens.

To have meaningful utility, the new metric must be of real help in the personal, a priori estimation of use-value that the individual content items will provide upon delivery. Such a metric could possibly be inferenced from usage data for the individual content items, amalgamated with demographic data about the specific users, by some smart algorithms. Both the good and the metric can be designed to encourage synergistic applications of information technology that produce desirable socio-economic impacts.

The determination of whether a good is to behave as a private good or as a public good requires consideration of the distribution means used. This is particularly so as the new technologies multiply the distribution options. Since one person's consumption has no effect on another's access, broadcast radio and television limit their content to behaving as public economic goods. Performing rights are a mechanism designed to address the externalities produced by attempts to force market characteristics on the processes involved in supplying content to these media. However, like any public good, the content of radio and television is generally undernourished.

Two new information media have been developed in recent years. One is a broadcast medium, where the entire library is continually cycled, and the user selects as he would from a smorgasborg meal. This system is called teletext, and is not to be confused with teletex which is a point-to-point messaging system. The second architecture is called videotex. It is not a broadcast system, and more closely resembles a star configuration, with each user being individually connected to a central computer that serves needs on an individual basis.

In so far as the user is concerned, there need be little difference that he can detect between these two architectures, so long as the libraries are small. Two installations could be constructed, using the two different structures, that would appear identical to all but the most discerning user, so long as the content libraries were small. However, the big difference between these architectures is that teletext forces its content to behave as a public economic good because there is simply no opportunity to collect individual usage data, while videotex does allow its content to behave as a private economic good because it can capture all the details of each individual's use of the system's content.

In the case of the videotex architecture, one person's use means there is one less system port available to serve someone else. With videotex, it is a relatively easy matter to deny certain individuals access to any or all the material in the library. These are characteristics associated with the behaviour of private economic goods. Videotex allows, for the first time, unembodied information to behave massively as a private economic good. This allows the possibility of a true information marketplace, one that could be relatively free of unaccounted-for externalities, given proper system design.

A second major difference between these two systems is in the amount of content available to the user. The teletext system must make a trade-off between library size and the length of time the average request takes to fill. This places a real constraint on the size of the library of content. Videotex has no such problem, and so is relatively unlimited in terms of its library shelf size. Offsetting this advantage is the increased capital cost associated with the videotex system.

With today's technology, teletext can be broadcast on a television channel over the air, or fed into the TV cable. In the cable it can be either piggy-backed onto a TV signal, or specially packaged so as to fill a complete TV channel entirely with teletext signals.

Videotex requires the individual circuits which only the telephone plant can provide in any profusion today. This means there is a distinct separation or wall between the two systems.

Teletext, being an inexpensive means of distributing material, is a very desirable technology for content that is accessed frequently. It is, to make an analogy, the bubble pack way of handling information. However, when information is handled in this way, the merchant cannot simply go to the bubble pack rack and easily identify the fast movers. Since the content is behaving as a public economic good, there is simply no record of the demands made by the users. However, the same constraint on display space applies, so it is necessary to develop special techniques to obtain sample usage data if the limited delivery space of teletext is to be wisely deployed.

The cost of running everything on a videotex system is simply prohibitive. Yet, the data that such an approach can collect is the essence of any new metric that would be useful in aiding users establish an a priori estimate of their use-value for specific items of content.

Some combination of videotex and teletext would seem to be very desirable. Such a combination would allow the best features of both to be combined so that the slow moving items of content could be accurately tracked in a videotex environment and the high flyers could operate in a teletext one, with periodic shifts over to the metered videotex environment for accurate statistical data gathering. As the partition between the videotex and teletext environment would be invisible in this combination system, there would be complete freedom to move content back and forth between the two environments, a characteristic that allows optimisation of the cost-benefit ratio of the system.

Were such a combination system to be developed, with carefully designed user protocols, it would allow the content to shift between behaving as a public economic good and a private economic good, without the users being aware of the shifts. For perhaps a mere one per cent of the time, a particularly well used content item might be behaving as a private economic good and so be accurately metered. The remainder of the time, it would be a public economic good. This time-sharing approach allows the spectrum between private and public goods to be filled in, and a good could be one per cent private and ninety-nine per cent public.

The best use of the limited shelf space provided by the teletext delivery strategy can be determined by letting the high flyers behave as public economic goods for most of the time, but sampling their performance occasionally by having them slide over to the private goods mode. The result would be an information system whose basic architecture was determined by overall economic-like considerations instead of simply engineering and cost considerations. The technology to combine the two architectures in this synergistic way is known, but presently is undeveloped for large systems. Such systems will be needed before any truly practical information society can emerge.

When we transfer our conventional ideas of market infrastructure to the information marketplace, we force unembodied information to behave exactly like a tangible good. This is an easy way to run a videotex system, because each time an access is made to a particular information product then a charge can be levied against the user, and revenue sent to the information supplier and the system operator. Unfortunately, recent experience with videotex systems has shown that users experience great difficulty in perceiving sufficient utility to offset the costs of operating the system and adequately rewarding the content suppliers. It is in response to this lack of sufficient economic synergy that the new-good, new-metric approach was developed. The simple objective is to help the prospective user perceive an increasing utility for the system's content. In short, the objective is to stimulate demand.

A marketplace is really a kind of information system, a decision machine, that guides the creative and productive sectors of the society. The new-good, new-metric strategy is far more than just a means of bringing something that is now relatively abundant, and easy to access, under the control of those who would seek mere monetary gain. As the central problem with ethereal goods is most likely to be the detection and development of consensus about the truth, reliability, worth, or simply value, of the divers content items, it is the information system, decision machine, aspect of the marketplace that will be of greatest value. Here is a new mechanism for social synergy. In all probability, the total value of author royalties will be small compared with the value society derives from the reliable assessment of the content.

These concepts need not be applied to an entire society to begin with. They can be applied, in part, to a corporation, or even a smallish information system. The measurement of demand, aggregating and massaging this data, and feeding it back to users to help them arrive at a priori estimates of use value can work in these smaller applications just as well as in the larger ones. The democratic dynamics of the marketplace can be an effective weapon in dealing with the information overload problem. Although our new information technology allows the creation of a marketplace that celebrates the unique characteristics of this new ethereal good and the new metrics that assess it, we have yet to configure the technology so that it encourages economic-like activities that produce sufficient social synergy as to be perceived as new wealth. There is much learning to be done in new combinations of economics and engineering.

That concludes the discussion of the first constraint on the application of information technology. Next, the second constraint, concerning the difficulties that machines have in dealing with natural language, is discussed.

In order to see how language acts as a constraint on the extensive class of information technology innovations, we will examine a technique for estimating the potential impact of any communications innovation, and then do a simple constraint analysis on the factors employed by that tool. If, for example, our estimating tool were to correlate potential impact with the size of the control knobs, what constraint might be acting to limit the size of control knobs? Such is the methodology we shall be using to establish the viability of language as a constraint.

No one has ever seen an electron, a neutron, or any other atomic particle. All that can be directly sensed are the impacts created by the presence of such a particle. A track of bubbles in a cloud chamber tell of the passage of an atomic particle in the same way that a line of footprints tell the hunter of the passage of an animal. We commonly deal with impacts as surrogates for the object itself, particularly when the object is too small, too large or too remote to experience directly. The communications revolutions of the past are candidates for such treatment.

Above all, writing allowed people to record their ideas. Toynbee mentions the recording of lists. Lewis Mumford writes: "Society, as a succession of observers have noted, from Auguste Compte to W. M. Wheeler, is an 'accumulative activity' and the city became the essential organ of that process". Experience was accumulated by means of writing. Writing had a big impact on the societies that adopted the technology. In a phrase, writing eased the access to their stored human experience.

The phonetic alphabet, requiring much less effort to learn and maintain than did earlier forms of writing, further eased the accessing of society's stored knowledge or experience. Books, whether from a scriptorium or Gutenburg's press, also produced this impact. Even the telephone, which allows us easily to seek out and talk with an expert, also eases the access to stored human experience.

Few, if any, significant past communications revolutions fail to make the accessing of stored human experience easier. This was a commonly evidenced impact, and is a likely candidate for the impact of any future communications innovation that seeks to be important.

Conversation presents us with an interesting manifestation of a spatial game. The acoustic space in which the communicants find themselves is occupied by one or the other, and occasionally both. In a well ordered conversation, the occupancies are almost sequential, with only the occasional overlap.

In the cocktail party situation, the game becomes more obvious, and its spatial characteristics are much clearer. Here the dominant talker of a group occupies the attention of the group until he tires, thirsts, or otherwise lacks the stimulation to continue. At this point, another person can easily capture the space, or the attention of the group. Possession of the communication space, as it were, is contested more and more vigorously as the level of alcohol consumption increases. In its well-developed form, such conversational orgies clearly exhibit the spatial-game characteristic of conversation.

The acoustic space in which the communicants find themselves is not the only information space we can identify. There are libraries, disciplines, vistas, and many more. A more useful approach is to consider these as components, or dimensions, of a larger multidimensional manifold which we might call the information space. Such a manifold, or space, has dimensions relating to the physical means that act to couple us together for communications purposes, like the acoustic and the visual ones, which are themselves simple spaces, having fewer dimensions than the whole information space. Language is a dimension of this information space, as are gestures and other cues we use in our communication with others.

For communication to occur between two people, they must share a common information space. Some part of their individual information spaces must be common. They must have at least some shared physical means coupling them, like, for example, an acoustic space. They must have a common linguistic dimension to their individual information spaces if they wish to communicate effectively. In short, the larger the common portion of their individual information spaces that they actively share, the fuller they can communicate. The larger the size of the common information space that they share, the fuller they can communicate.

Past communications revolutions increased the dimensionality of people's information spaces, and opened new dimensions for sharing with other people. If we have read the same lists, we have something in common about which we can communicate. If we can write to each other, we can extend the physical distance over which we can share our information spaces. If we share the same library facilities, we have a rich common resource indeed. The telephone allows us to create a shared acoustic space at a distance with a person of our choosing. Picturephone® did not provide a common visual space that the two communicants could share. Perhaps that is why it did not succeed.

A city increases the size of the common information space that can be shared over and above what a rural life can provide. Time is a shared thing too, ever since standard time, which everyone could share, was widely adopted.

With my word processor, I can share the editorial space of my system with anyone fortunate enough to have similar equipment and software. This allows us to jointly and equally edit, for example, a contract of mutual concern. Everything that happens on my screen happens on his, and vice-versa. It matters not who presses the keys, the effect is the same, for we are fully sharing a common editorial space created by the particular way these two otherwise independent systems are coupled. During this exchange, we would also be connected by telephone, allowing us to discuss, argue about, or whatever, the contents of that common editorial space that we are sharing. We would be well coupled. We would be communicating.

Unfortunately, today's manufacturers of word processing equipment have not discovered the shared space model of interactive communications. Commercial communicating word processors only pass files back and forth. The machines communicate, but not the people.

The shared space idea is really a model of the interactive communications process. The purpose of a model is to explain something in a simple but practical way so that new useful insights become patently obvious. The shared space model of interactive communications does this quite well. There is lots more to squeeze out of this model. Simultaneous shared space voice and data communications, of which the shared editorial space is an example, open the way to many new products and services. These all depend on a good appreciation of the shared space model of interactive communications.

The significant past communications revolutions increased the size of the common information space shared by the communicants. Presumably, any future innovation of communications importance will also increase the size of the common information space shared by the communicants.

Throughout history there has been a special relationship between communications developments and social advances. The Athenians' adoption of the phonetic alphabet fed their hundred golden years. The printing press fostered the Renaissance period in Europe. In cases where communications developments were not entirely causal, they were transformative. Consider the city, about which Mumford says "What began as control ended as communion and rational understanding", suggesting that the city and improved communications were symbiotic in that they were contemporary developments.

Large-scale social developments cannot occur without a wide consensus existing throughout the society. Those developments in communications that we would call revolutionary had important social impacts, otherwise we would not see them as being revolutionary. Because they had these impacts, and since such impacts cannot occur without a wide consensus, it is reasonable to assume that these communications revolutions had the property of aiding the emergence of consensus. All the great communications developments in history eased the discovery and development of consensus.

Before a consensus can become widespread, it must have a beginning. Here is where the new communications developments made their impact. More new ideas were "processed" through society's evaluative mechanisms with each improvement in the communications environment. Where the Inuit have more than twenty words for snow, because snow is so important to them, we have only one word for consensus. Consequently, we are somewhat limited when we try to discuss consensus. To ask one word to range from how the tender beginnings of a new idea are spread through the society, and assessed, right through to how the whole population salutes and accepts an idea is too much. Here, we are particularly interested in what might be described as nascent consensus — the tender beginnings as opposed to large scale acceptance.

In summary, we see that three characteristic impact patterns accompanied the great communications revolutions in the past. First, there was an increase in the ease of accessing stored human experience. Second, the size of the common information space shared by the communicants increased. Third, the discovery and development of consensus was eased.

These three impact patterns can be used as the basis for an evaluative instrument for the assessment of the probable socio-economic significance of a communications innovation before it is in place. The underlying assumption is that if this pattern of impacts has been around for several millenia, it is not likely to change much in the next fifty or so years. Hence, if we produce something that has similar impacts to those produced by past communications revolutions, we will most likely have invented what will later be described as another communications revolution. Although no claim is made for sufficiency, one is put forth for necessity.

A technique for the assessment of the significance of the potential impacts from any given communications innovation can be built on these three dimensions of communications. This technique is described in "An Assessment Methodology for Evaluating Communications Innovations", published in the I.E.E.E. Transactions on Communications, Special Issue on Social Impacts, October 1975. However, we have all that we require of this technique in order to establish the viability of language as a constraint. It is not necessary to know all the details of the assessment technique to appreciate the importance of any constraint that acts on the basic measures of the technique.

Language acts as a constraint on all three of the above characterisations of the impacts of past communications revolutions. To build the ultimate information retrieval system would require some significant advances in artificial intelligence, or the ways in which machines "understand" natural language. People use language in ways that befuddle machines, unless very constricting conditions are placed on the way language is used. The processors that we have in our heads are very different in structure from the computers we now have.

Yehosha Bar Hillel, while all the other experts were making wild and extravagant claims for what computers could do, put it this way, "If you can't teach a dolphin to understand, what hope do you have of making a computer understand?" However, he went on to say that although such a feat was impossible, we could not afford to stop trying! Here is a dilemma. Although a thing is impossible, we must keep on trying! In tribute to his early recognition of this problem, I call it the Bar Hillel Conundrum. Really good machine translations, mechanized telephone operators, universal information retrieval systems and the like may eventually be commonplace, but for now, they can be considered somewhat remote or limited.

Years ago, when Englebart and English were carrying out their intelligence augmentation work at Stanford, Englebart would talk of "soaring through a co-worker's files, choosing this, leaving that . . .". As they got further into the project, the incredible difficulties involved in browsing became evident, and the "soaring" metaphor was dropped. I single out their experience because the files that were being examined were a shared intellectual space, and it was the synergy of this sharing that they sought as the basis for intelligence augmentation. The difficulties they encountered illustrate how the linguistic constraint can apply to the size of the common information space characterisation just as it applies to the information retrieval or ease of access to stored human experience characterisation.

The third characterisation, the ease of discovering and developing consensus is also constrained by language, if one attempts to mechanize the process.

As language is a constraint on all three fundamental dimensions of this measuring tool, then, to the extent that this tool is relevant, language must also be one of the constraints acting to hold back the members of the extensive innovation class. Language now joins the economics of information as a constraint. Let us now turn to the third constraint, the one of perception.

Marshal McLuhan told us that the medium is the message. Without striving too hard, we can find many examples from our television experience where the medium if not completely changed, at least significantly coloured the perception of the events that actually occurred. However, it is not the broadcast media I am interested in exploring here, but rather the familiar telephone.

The telephone has been around for a century, yet it is really a little-understood medium. By means of the telephone we create an acoustic space that can be shared by two people who are many miles apart. With the simplest telephone technology, both parties can talk at the same time. They can fully share the acoustic space created by the device. They need not occupy that space sequentially, they can interrupt, talk simultaneously, share a cry, or whatever. This simple property of the common telephone, the sharing of an acoustic space, is not conceptually understood by most people who use telephones. It is most certainly not given a high priority by those who design the more sophisticated telephone services, like loudspeaking telephones, most conferencing facilities, long distance circuits and satellite systems.

The damage the loss of this shared acoustic space causes is particularly noticeable when a loudspeaking telephone is connected with an ordinary telephone. The person using the ordinary set gets the impression he is not being listened to. Whenever he talks, the sounds from the other end get cut off, and he feels he is talking into nothingness.

The end effect is that one should avoid using the loudspeaking telephone when soliciting. Alternatively, the loudspeaking telephone works extremely well for disciplining tasks, if the recipient has a conventional telephone. After a few calls of this nature, you will just have to place the call, and the recipient will get the message. The medium becomes the message.

Long-distance telephone circuits are troubled by echo that is sufficiently delayed as to be quite apparent and annoying. Here again, technology comes to the rescue, and echo suppressors are installed. As might be expected, echo suppressors also shatter the telephone's shared acoustic space.

The story is told of the elderly lady in Toronto who called her sister in Vancouver to share a cry. The echo suppressors made this an impossibility, and forced them to cry at each other sequentially. They complained. Most users do not complain, because they do not know any better.

The real horror story for unsuspected damage is the satellite with its quarter of a second round trip delay time. This delay is caused by the finite time it takes for the radio signal to get back and forth to the satellite. Only recently have we become aware of how this delay can alter the perception the conversants have of each other's emotional states. Couples that have spent a lot of time using terrestial facilities can have serious difficulties in their relationship, entirely due to the effect of the transmission delay.

Once, on a national television program, I described this phenomenom. Two weeks later I received a letter from a woman in Toronto, whose son had recently gone overseas. She said he had called from France to find out what was the matter with his girl friend, for she had seemed rather hesitant on the telephone. She also had a call from his girl friend, who had found the boy behaving "out of character" during their telephone conversations. Fortunately, the woman had heard the broadcast, and was able to explain about the effect of satellite delay on telephone conversations. Her letter went on to say that the couple were eternally grateful.

Satellite delay can, and has, caused couples to break up. This medium has the power to change a message that was intended to move in the direction of love to go in a totally unexpected and different direction. The medium changes the message.

The effect is so subtle that novice users do not recognise that it is caused by the satellite. In view of the lack of widespread understanding of this phenomenom, it is not surprising that nobody has sued a telephone company for alienation of affection. Once one is aware of the problem, and has a bit of experience with satellite calls, the problem vanishes. Because of the lack of general knowledge of the effect, getting that awareness can be quite painful.

If after one hundred years of familiarity with the telephone we still have problems like these, how can we cope with the new wonders that are flowing from the world's laboratories, and how can we direct such efforts in truly meaningful directions, directions that produce a general rising on Maslow's hierarchy of needs? A great deal of adult learning must be achieved before this can really happen. We simply do not recognize what is happening, and even less can we see the real potential. Our philosophers have simply let us down. The real failure of our educational institutions is that they have not prepared us to ask the relevant questions of the opportunity presented by the technology.

Mercury was the Roman god of communications, commerce and thievery. The ancient Romans were very foresighted in choosing that particular combination as being related enough to put under one god. Our fear is that the technological gift from Mercury will not be appreciated, and we will choose the gifts from one of Mercury's peers, the bellicose god Mars. To avert such an outcome, we must do the learning required to produce an open system economy, one that celebrates abundance as opposed to the historic one that celebrates scarcity.

Of the three constraints outlined above, the economic one is perhaps the most amenable to attack. The economic system we have is not a gift of God, it is a human artifact. It is open to design. To use all this information technology merely to calculate our daily interest earnings is to miss completely the potential of Mercury's gift.

ALTERNATIVE FUTURES

Earl Joseph, Sperry Univac

Earl C. Joseph, obtained a degree in mathematics from the University of Minnesota in 1951 and has been at Sperry Univac since then.

Since 1963, when he was appointed Staff Scientist, he has been researching the future. Previously, his roles included systems manager, project manager, and he has directed, managed and performed the systems design, logic design, programming, and applications of a number of computers.

In his present position, he researches the future and advises management at all levels on future technology, including the design, impact, application, social impact, management of future computers, artificial intelligence, and alternative futures for Univac and society. As a Staff Futurist, *his current forecasting efforts are directed towards futures research, strategic management, future computer design, and long-range planning activities. The activities include the study of alternative futures for microprocessor systems, smart machines, data processing, food and farm automation, defence systems, factory automation, socially desirable futures, economic/financial systems, education, medicine, communications, artificial intelligence systems, knowledge-based systems, expert systems, and much more.*

He holds three computer patents, is one of the creators of Ethnotronic Science, is a creator of a language to describe alternative futures, is the system architect of five major computer systems, has co-authored over 30 books, and has published over 150 papers.

Among his many other activities, Earl Joseph is an Adjunct Professor at the University of Minnesota designing and teaching graduate level courses on the future. He is on the instructional staff of Metropolitan State University and is a (former) Futurist-In-Residence at the Science Museum of Minnesota.

What I am going to do is take you on some trips into alternative futures and show you many developments that are already under way. Everything that I shall talk about is currently at some stage of design, development or research. However, I shall extrapolate these developments into the latter part of this decade and into the next decade and try to show you some of the impacts to expect as we go deeper into the future.

Once upon a time there was a little girl at kindergarten in the big city. One day the little girl came running home from school saying, "Daddy, Daddy, I got an A in class today." Dad was very proud of his little girl and asked her how she had done it. Apparently she had taken a test, which consisted of saying how many legs a horse has. The answer she had given was three. This startled Dad, who asked "How come you got an A?" and the little girl said, "My forecast was the closest."

I tell that story to give you an idea what I shall be talking about. I shall be talking about trends and alternatives.

I do not want to slight children, so I shall redress the balance with another story. Once upon a time, a little boy came home from school. He was a little older than the girl in the previous story. He asked, "Mum, where did I come from?" The mother was quite embarrassed by the question and said, "The stork brought you." The boy knew better than that, but he went on with the next question. "Where did Dad come from?" Apparently the clouds opened up and Dad fell on grandmother's lap. The boy shook his head and went on with the third question. "Where did Grandpa come from?" Apparently he came from another country — a donkey brought him. The boy goes up to his room and starts writing: "After very intensive research on the subject I have been assigned as my theme, I have come to the conclusion that in the last three generations of my family there has not been any normal sex."

My topic is not sex, and it is not normal futures either. I should like to show you some of the contra-intuitive things that we see when we dig deep into the future.

I shall try to answer the question "What will the future be like?" This is a very big question and I can only scratch the surface in the time available. There will be many alternatives additional to those that I shall discuss here.

Do you realise that the future is coming at us at the rate of 60 minutes per hour? I suppose that is obvious, except that when you compare the rate of change at which computers and communications technology are advancing with the rates for the technology of automobiles, steel, and so on, you see that developments in each technology are coming at us at different rates. I shall try to give you an idea about when to expect some of the future impacts that we can now forecast.

One dynamic we will encounter as we look deeper into the future concerns the impact on jobs. For example, as larger portions of information systems applications become automated, as the application parameters become higher level and more powerful, so we will be taking work away from the programmer. However, each time we do this, the tool becomes more useful and has a wider application in society. The result is that, instead of laying off programmers, we hire more. The more we automate the job of programming, the more programmers we are hiring. Of course sometimes we shift the place where they are working. That is an example of why automation creates more jobs than it displaces.

We are in a new information age and we must expect new tools. Most work that people do now involves them in working with their minds with information. In the United States over 70 per cent of the working population have some job where they collect information, process it, or disseminate it. Figure 1 shows a partial list of some new information age tools that have emerged during the last five years. How many of you, five years ago, were using some form of unit with a screen and keyboard in your job? I can see

Figure 1

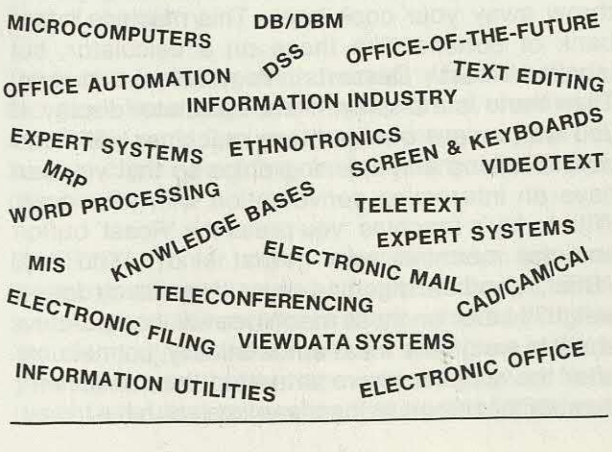

from the show of hands that the answer is about 10 per cent. How many today? About half. That is quite a change in five years.

Advances are not slowing down in our field, they are speeding up. So in the next five-year period we should expect much greater advances than we have seen with word processing and so on. I should like to show you some of the things that will happen, particularly as we begin to add intelligence to our computers and communications.

One way of characterising society is by the key material that is transforming it. For example, we have had the stone age, the bronze age, and the iron age. Our current page in history is the silicon age. Perhaps the next decade will usher in the genetic age, in which genetic material will transform society.

Let us now look at the silicon age. Most of you probably know what VLSI means — Very Large Scale Integration. A typical VLSI device consists of a quarter inch chunk of silicon containing many thousands of circuits. Today, we are capable of putting about ten thousand circuits in such a chip, which is about a hundred thousand circuit elements.

The first process of this technology is to scoop up some sand in a silicon foundry, put it under high pressure, and grow out of it almost pure silicon crystals which are then sliced up into thin wafers. Impurities are then put into the silicon in a multi-step process to build up the circuit elements. Metal is evaporated onto the wafers to interconnect the circuit elements. Then the wafer is broken up into little silicon chips. These chips are then put together into a component, and then a number of components are put together to make something like a computer or the electronics for a communication system.

The process of going from big to little and then building back up to a big system is an added value process — it is the way we have found for making money out of the technology. Today, we have reached a point where we can put into the component sufficient circuitry to build complete machines, and so our building blocks are becoming complete machines. In the next decade we will in many cases be able to use the complete wafer as our component. In other words, our components will consist of collections of machines, including, for example, microsensors so that the component can see like a radar, sonar, or video camera and display information like a flat-screen TV, and electronic communications circuitry so that it can be connected into communications systems.

This gives us a component system or a component institution as a building block. The technology for achieving that has been advancing at a rate of doubling its capability every year for the last two decades. That is a factor of a thousand increase in a decade or a factor of a million increase in two decades. We have knowledge of the science needed to allow progression to continue at that rate for the next couple of decades. I shall show you some of the things we will be doing with that technology.

The VHSIC programme is a new US Government programme to speed up advancement in this technology. This technology has been advancing more rapidly than any other type of technology, and we now perceive it as advancing too slowly. The next step is Ultra Large Scale Integration, and then Super Ultra Large Scale Integration. Then we run out of adjectives and have to refer to systems-on-a-wafer. The next step is the biotechnology step.

So at present our systems use chips. In the next decade they will use complete wafers. Later we will be using biotechnology, and at about the same time three-dimensional componentry will be introduced.

This means that our systems will be provided with new capabilities. During the past five or six years, we have developed the ability for circuitry to talk. We start out by putting the capability into toys such as "Speak and Spell". One of these days we will put it into machines such as the microphone, the automobile, the toaster, the vacuum cleaner, the coffee pot. Then these machines will nag you all day long.

At the present time, hardware for listening is coming out of the laboratory. At the moment the hardware has 50-word, 100-word, 200-word vocabularies. This limited-vocabulary hardware takes words out of context and recognises them in the way previous hardware recognised signals created by a finger pushing a button. So we call them button-pushing languages. I will discuss some discoursing machines later.

That means that from now on our machines will evolve with increasing silicon intelligence. In other words, the circuitry will continue to become more and more complex.

Now I want to turn to the area of artificial intelligence, where we imitate what the mind does. I shall use the word "heuristics" to describe the way the mind goes about discovering something. This includes the inductive and deductive reasoning and discovery processes, thinking and learning, making educated guesses, and fuzzy thinking. We have now reached the stage of putting heuristics in our machines.

In the last decade we have learned more about how the mind works than in all previous history. In the early days of the study of artificial intelligence we thought we should take a bunch of circuits, connect them together, add sensors, and let the resulting systems start out as infants, with no knowledge, and then let them experience and learn to do whatever we wanted them to do. We did not find a way of doing that.

Professor Marvin Minsky at the Massachusetts Institute of Technology, in the late 1960s, said that even if we could learn to do that we should never allow our machines to be like that. When asked why, he said, "Obviously we could never allow a machine to be a teenager in society." He said that what we had to do was to find an architecture into which we could put a lot of knowledge from the very start. In other words, we should build the machine as an adult expert and let it learn from that point forward.

That was the breakthrough, but it has taken from then till now to learn how to do that. I should like to show you the stage we have now reached and then extrapolate a little in the future.

An expert system has a knowledge base, which is an electronic memory with database features that use associative techniques. We are not so much interested in the hardware of the knowledge base, we are more interested in the contents. The contents are the facts or information that an expert uses, together with the heuristic rules for accessing those facts and using them. Expert systems also incorporate heuristic procedures for manipulating the rules and the facts, and then also some programs to put the results into terms that you and I understand.

These systems are based on using commonsense reasoning to do what experts do. But this is the stumbling block. How many of you can tell me what a manager does other than delegate? And what information and heuristics do managers use? Those are tough questions and it takes a long time to find the answers to them. Our real stumbling block is to discover what it is that experts do.

Let me show you an example of a machine that uses primitive intelligence. The machine is a microwave oven, which its advertisers claim enables you to throw away your cook book. This machine has a bank of buttons, like those on a calculator, but labelled 'Roast', 'Desserts', 'Vegetables', and so on. Then there is a display, like a calculator display. If you wait a year of two, these machines will incorporate talking and listening chips so that you can have an interesting conversation with your oven. With today's machine, you press the 'Roast' button and the machine asks "What kind?" You say, "Beef,", and the machine says, "How much does it weigh?" Later on these machines will incorporate a scale to weigh the meat automatically, but not until after the suppliers have saturated the market with dumber machines. In the meantime you have to say, "Ten pounds." The machine then asks, "How well done?" "Medium," "When?" "By 5.30 p.m." "OK," says the machine. Where is the cook book — it is in silicon, either hard-wired or programmed. Such machines are the tip of the iceberg, the first examples of us putting the knowledge of how to do something inside the machine.

This machine does 90 per cent of what an expert does. It asks questions — what, how, and when? Experts usually do not do much more than that. The heuristic consists of a decision tree for answering questions. In order for the machine to be able to do that, knowledge has to be recodified in a different way. When we program a computer, we are recodifying algorithmic or process type knowledge using linear strings of procedure-oriented instructions. When we codify a knowledge base we are using logic rules, the logic of IF, THEN, AND, ELSE. For example, IF roast and IF ten pounds and IF medium and IF by 5.30 p.m. THEN the machine knows what to do.

It turns out that the recodification of knowledge in this form is closely related to programming. The recodification has to be done very precisely. The resulting knowledge base is somewhere between ten and one thousand times more precise and comprehensible than the standard textbook or manual. So what we have here is the beginning of an industry using computers and communications systems to replace the standard textbook or manual, if not in this decade in the next.

I shall now discuss this in more detail by examining particular types of system. First, I shall look at DENDRAL, which is an artificial intelligence program that has been worked on for 16 years at Stanford University. This program is used for doing symbolic chemistry, the form of chemistry that can be thought of as detective story chemistry in the sense that it is concerned with finding out what mixtures are made from. In the last six months DENDRAL has achieved championship status, by which I mean that there is no other tool, human or otherwise, that is as knowledgeable as DENDRAL in its particular subfield of chemistry.

Think of the professional chemist, who now has a tool more knowledgeable than a human chemist to amplify his work. Can you imagine the synergy of the human chemist's mind in conjunction with the machine?

MACSYMA is for doing symbolic-map algebra, including differential and integral calculus. This system also has reached championship level in the last six months. Mathematicians now have a third-generation tool to take laborious work away from them. In the first generation, during the 1950s and 1960s, the large scale computer had an impact on the mathematician. Although these machines took much work away from mathematicians, the machines also allowed the mathematician to achieve things that were impossible without such a machine — for example the design calculations for a jet aeroplane. So, even though we automated a large portion of the mathematician's job, we still hired more mathematicians because they could now take on many new tasks.

In the 1970s, we impacted the mathematician with the hand-held calculator, and the screen-plus-keyboard microcomputer, which made real-time calculation instantly available in support of the mathematician's thinking processes. Again we did not lay off mathematicians, although once in a while we changed their title to programmer.

In the 1980s, the mathematician is being impacted by the knowledge-based system, in many cases more knowledgeable than the mathematician himself. Again we are not laying off mathematicians. In some cases we are hiring more.

In the next decade, we will probably have inference engines, which will be the fourth-generation tool.

Let us look at the architecture of such machines by discussing INTERNIS, which is an artificial intelligence system for carrying out medical diagnosis. The way it works is that a doctor, or a patient, keys in a set of symptoms — for example, headache, rash, and high temperature. INTERNIS then looks in the first part of its knowledge base, which is the symptom catalogue, and finds "headache". The system then makes an association into the second part of its knowledge base, which is a disease catalogue, and finds all possible diseases in this catalogue that produce headaches. Then it goes on to find all such possible diseases that produce rashes, and then all such possible of those diseases that produce a high temperature.

There will be a lot of overlap of such diseases. The system sorts out the overlap by using its heuristic deductive and inductive reasoning ability, making an association in the third part of its knowledge base which consists of a set of questions like those in the microwave oven. However, unlike the microwave oven, it will really know what is the next test that should be made on the patient, taking into account the need to get the most information out of each test, at least cost, and with minimum damage to the patient.

Through this process of using a decision tree of questions and answers, the system ends up with a diagnosis of what is wrong with the patient. Once it has made that diagnosis, it goes into the fourth part of its knowledge base which is a catalogue of prescribed treatments.

That is a typical architecture for an artificial intelligence system.

Think about your particular field and imagine building a knowledge base, a taxonomy of your field, that can be used in the real time of your work and your decision process. You can imagine how your work will be changed as we go deeper into the future.

We are entering the era of ethnotronic machines. Ethnotronics is a coined word like avionics. The definition of ethnotronics is "The science of the relationships that humans and society have with inorganic systems which amplify their mutual capacity for learning, reasoning, decisioning, accessing information/knowledge, and communicating." If I give each of you a standard calculator, you become a new ethnic group that is electronically amplified to do calculations. If later we design and manufacture an expert manager's machine, each of you provided with such a machine will become a new ethnic group electronically amplified to perform the tasks of management. So, ethnotronic machines are people-amplifier appliances with heuristic primitive functions that provide aids for helping us discover, invent, decide, and infer.

As a futurist I am forced to forecast that there is a period in history when even the bureaucrats and politicians will be using these types of machine. I am then forced to forecast beyond that point that eventually we will have intelligent government.

Now that we are nearing 1984 and teleconferencing is here, we have to be very careful about the way we design these machines. Let us go to the mid-point of the next decade and claim that we have a knowledge base for a management expert system. By then the expert management machine will be capable of eavesdropping and listening in to determine when you need some help. It will use as cues the verbalisations that go on and will be able to read the image of the face to determine when the face is registering a puzzled expression. It may also use some silent language, like infra-red signals, to determine when you are getting hot under the collar. If, having integrated all of these cues, it does not understand what you need, it will begin to ask questions, on its own initiative. The machine has only one purpose — to amplify whoever is using it.

Let us assume that it makes a determination of what you need in the way of help and dips into its knowledge base, which is the collective knowledge of many managers in your field. Most of the time it will find the help you need and make you aware of it. Let us assume that once in a while it does not find the help you need. But let us also assume that we have made the telephone intelligent using chip technology so that on its own initiative this device can make a telephone call to bigger knowledge bases, pull out what it perceives you need and bring it back over the telephone line. At the moment, when the telephone bell rings, your intuition tells you that somebody wants to talk to you. In the future, when your expert machine bleeps at you, you will think, "Aha, society has amassed some knowledge that I can use in the real time of what I am doing." Imagine a machine that is amplifying you, that is smarter than you, and that cannot wait to tell you.

Our problem with such machines is the garbage in/ gospel out problem. Incorrect information in the knowledge base creates an enormous problem.

The question is whether such machines will do your work in the future. My answer is "Yes, No, Maybe, and Hopefully". What I have been trying to tell you is that, by using such machines, tomorrow's clerks will be able to do what today's professionals are doing.

In order to understand this we have to know what a professional is. The higher we go up the professional ladder, the closer we get to mechanical regurgitation, so the first point about a professional is that he is a regurgitator. The second is that he is a parasite — he takes from his profession. Whatever he takes he usually delivers, and so the third point is that he is a delivery person.

As the clerk takes over, the professional gets out of the in-line delivery of the profession and loses the parasitic status. What happens then is that many of the professionals add to their profession, or in other words grow it. As a result the profession advances rapidly, becoming much more valuable and useful to society.

In most professions, although not in all, as we automate the profession in that fashion, we require more professionals. So, the more we automate, the more professionals we will require.

So, society has become more and more complex as our computers have become bigger and bigger — our aeroplanes have become bigger and bigger, our governments have become bigger and bigger, and so on. As a result we needed more education, as shown in Figure 2. Now we are educating machines, and we need less and less education of the vocational type to use those machines. That does not mean that we need less education in total — the total amount of education we need will continue to rise, but the education will incorporate some very different types of education, perhaps the type for teaching us how to get along with each other.

Figure 2

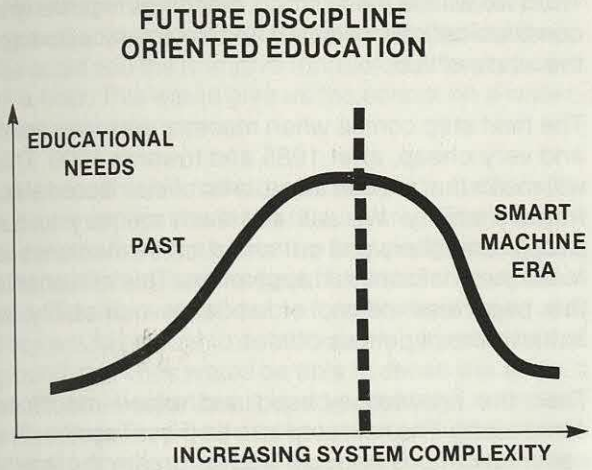

When we look at some alternative future computers, we see smart machines, knowledge bases, expert systems, ethnotronic systems, micromainframes, wafer systems, component machines, and component institutions. In the information age, computers talk, listen, answer, see, and give advice.

Equally, there are considerable developments taking place in communications. Future communication systems include developments such as sensor-based systems, smart communication systems, component communication systems, people-amplifier appliances, database systems, information appliances, knowledge-based systems, expert systems, current-awareness systems, component schools/offices, distributed smart systems, embedded communications, intelligent communications. In communications we are reaching into a period of more opportunities than we have ever had before.

Now I want to talk about machine intelligence. Let us start by looking at home automation where we have not a distributed network of computers but a distributed network of smart machines. A little later on we will have a distributed network of intelligent machines, such as intelligent toasters, vacuum cleaners, refrigerators, and stoves, all interconnected through a communications network.

Let us imagine that this network of expert kitchen machines has a very sensitive, remote ear so that the machines can hear what is going on in the bedroom. There is no need to be afraid of this — you are not afraid to undress in front of a light bulb, and we are talking about just another electronic device. As you wake up the machine senses what is going on, or if it does not understand then it begins to ask questions. So when you stumble bleary-eyed into the kitchen, the coffee pot hands you your cup of coffee just the way you want it and you sit down to the table that has been set out with the breakfast just the way you want it.

Let us assume that you have the latest model of toaster, which has more intelligence than the rest of the machines in the kitchen, and so has been feeling lonely all night. Upon sensing your entrance, it begins to jabber away: "How would you like to have your toast this morning? Would you like it light tan, medium tan, dark brown? Do you want some jelly on it? Shall I ask the refrigerator?" After being nagged by the toaster in that fashion, you spill your cup of coffee and make a mess.

Another machine, which has been hiding away in the closet and has only one purpose, which is to clean, senses the mess you have made, comes running out of the closet, cleans up the mess, and goes back into hiding.

What I have just described is a network of co-operating machines. We can now build such machines, because we can add sensors to them, we can add intelligence, and we can link them into a communications network. That is quite a turning point. Imagine the use of such a network in the office or the factory of the future. Now, let us consider the value system, or the manners, that we have to put in such machines. In the home, the TV dinner of the future will tell you that you have already exceeded your calorie limit, just when you are about to eat it. Clearly, we will have to be very careful about how the machine does this.

The Japanese stumbled on this problem. Some time ago they brought out an oven that talks, but none of the Japanese housewives have bought that machine because it incorporates some wrong words. You can imagine, if there is just one word wrong, how bad these machines become. This requires a whole new study of language and its use. In fact we are beginning to learn more about human values as we enter the period when we have to incorporate value systems in our machines.

In the past, we have had the calculation engine for performing numerical arithmetic, the data manipulation engine for performing information processing, and smart machines with embedded logic. Future types of computer that we can now see emerging are the inference engine, the ethnotronic people amplifier, and robots. So there are three new classes of machine.

There is a long list of trends in computers, or cybernetic type machines. These trends include smart, embedded computers, computer-based machines, microsensor-based systems, knowledge-based systems, AI/expert systems, linked co-operative systems, distributed networks of smart systems and of intelligent systems, convivial machines (by which I mean machines incorporating at the interface between the machine and the human being a lot of hardware dedicated to making the interface friendly, easy to use, and intelligent enough to tell you how to use it), current-awareness systems, and inference engines.

So, in the information age, people amplifiers bring access to society's knowledge in the real time of our action.

As an aside, I would like to point out that love is information — it is a communication. My question is, will the information age be more loving? How many of you can state that your institution this year is more loving than last year? By how much? I claim that if we do not know the answer to that question, and to other questions related to our positive, desirable values, then how can we make the decisions needed to take us closer to desired values?

Let us now look at office automation beyond word processing. Several developments are taking place to screen-and-keyboard devices. If you sit all day long looking at a fixed-position device, you get a pain in the neck. This is being overcome by making the devices portable.

The definition on the face of the screen is not all that it could be, and as a result you get a little eyestrain. This is being overcome by using high-definition video. Another particularly useful development is removal of the need for managers to be able to spell by putting listening and talking chips into the devices. All these developments are totally changing the ground rules for these types of information appliance as we go deeper into the future towards the smart management machine.

I shall now run through past and future developments in office automation. The office of the future started with minicomputer-based word processing systems. Then we added smart, microcomputer-based word processing systems. Then we tied these machines together by means of a communications network. Then we arranged the marriage between data processing and word processing systems. Then we added office automation software, computer mail, and computer conferencing, which is about the stage we have reached today.

What is in development? The first development is the database computer. Next are smart office machines, for example a typewriter that can understand the spoken word and so as a result of you dictating to it type out the words and spell them correctly. Another example of a smart office machine is the electronic filing cabinet incorporating storage, a screen, and a keyboard. Then people-amplifier appliances will be developed, such as the smart management machine, the expert doctor machine, and so on. Then we will tie these smart machines together in a communications network and so totally automate the whole office.

The next step comes when memory gets very small and very cheap, after 1985 and towards 1990. This will mean that we can afford a lot of distributed electronic memory. We will add such memory to our people amplifiers and our smart office machines to make them information appliances. This will usher in the paperless office, or at least our ability to achieve the paperless office.

Then the knowledge-based and expert machines come along. The next step is to tie these together in a communications network and so create the intelligent office.

The next step is to make a component office, an office on a wafer. Today, we can put the same amount of information on a chip as occurs in a fairly large book. Of course, we need a lot of electronics around the chip to be able to access the information. With the VHSIC program that I mentioned earlier, by about 1985 we will have the capability to incorporate direct output visual technology in the chip so that it can display on a flat screen the information it is storing. Ten years from now, we will be able to put somewhere between one hundred and one thousand books full of information in a component of that size. However, by that time we will be using wafers, each of which is equivalent to about a hundred chips. If a chip can hold a hundred books then a component library can hold ten thousand books. Most of the silicon acreage would be unused in such a component library, so we might as well incorporate other capabilities, such as office automation programs, database management systems, management information systems, and so on. We could also include the primitive functions required to provide the mathematician on a chip.

The chip in one of today's calculators contains about a thousand circuits. This mathematician on a chip does quite a lot of work. Imagine when we can incorporate a hundred thousand circuits in a component, or half a million circuits a little after the mid-point of this decade, or a few million by the end of this decade. Imagine the capability of such a device. Maybe ten thousand circuits will enable us to build a smart doctor machine, and half a million maybe a smart management machine.

What would be the primitive functions for the manager on a chip? They would not be Add, Subtract, Multiply, and Divide, as for the calculator on a chip. They would be functions such as Delegate and Communicate, Delegate and Control, Delegate and Plan.

We could add more functions to our office on a wafer. For example, we could take up some of the chip areas with courses, such as are already becoming available on devices that plug into calculators. We could add the primitive functions for the teacher on a chip. This would give us the school on a wafer, or a component school.

Let us look at an initial application for such capability, an application in the leisure area. The component could have the collective knowledge of a group of expert skiers, and so would be installed in a pair of skis. The skis would be able to see where trees are. They would be able to detect where bare patches of ground are. They would be able to sense the snow's condition — whether it is dry, powdery, compact, wet, or icy. They could sense which way the skier's toes are pointing, how the skier is bending his knees as he manoeuvres towards the ski poles, and so on. As you are skiing down the hill, your skis are in constant conversation with you, nagging you to ensure that you never break your neck. This is a real-time, educational system.

Imagine a component school embedded at the interface of just about every machine with which we have contact. With such intelligent, communicating, and discoursing machines, the world around the year 2000 is going to be a very different place in which to live.

Now let us look at what has happened in computers. The computer eras are summarised in Figure 3. What have these eras meant for management? In the 1950s the manager was managing the operation of the tool. In the 1960s, the data processing manager was managing the tool, the data, and the programs. In the 1970s, MIS management was managing all of that, the system, plus information. In the 1980s we are entering the era of information resource management, where managers are managing the total information environment, not only the computer system but also the photocopying room, the communications, and the mail room. In the 1990s, we enter the era of knowledge resource management where we are managing the total knowledge application. In the year 2000 and beyond, we enter the era of intelligence resource management, which is concerned with the automated management of leadership. At least, that is one path into the future. One could develop other future history maps about the impact on management.

Figure 3

COMPUTER ERAS

1950s –	**DAWN OF COMPUTERS & AUTOMATION**
1960s –	**COMPUTER SYSTEMS & DATA PROCESSING**
1970s –	**INFORMATION PROCESSING, DISTRIBUTED & MICROS**
1980s –	**SMART MACHINES, EMBEDDED & AI**
1990s –	**INFERENCE PROCESSING**
2000+–	**GENERAL INTELLIGENCE**

Let us look at this future map now in terms of communications systems and of friendly machines. Let us think about the sort of information utility that is developing now, and the sort of knowledge utility that will develop during the next decade.

In a Japanese Government White Paper last year, it was stated in the introduction that in Japan in the 1990s information will be elevated to a basic need, as basic as food. The United States and Europe are moving just as rapidly into the information age as Japan. However, the Japanese have singled out the industry and the machines for that age. In fact the Japanese have given us a target at which to shoot. The question is, who will be the leader? The Japanese have set in motion the mechanisms through which they think they will become the leader. However, ten years is a long time. Even Europe could become the leader, if it tries. A future is always ten years away until you start to do something about it.

The list of corporations becoming involved in the information age is growing very quickly. In the United States, the largest corporation, AT&T, has just been unleashed. The United States Government is changing the law to allow AT&T to move beyond the provision of a communications highway and telephone terminals into the provision of information services. Many other countries, including Canada, the UK, France, and Japan are building technology for the information utility era.

There are forecasters who forecast that before the end of the 1980s the information utility will add one trillion dollars in annual sales to the communications industry. The artificial intelligence part of that will be one hundred billion dollars. The computer industry itself has yet to reach the one hundred billion dollar level. Imagine an industry that large just in the artificial intelligence area. We have never developed such large forces in such a short period of time. The impact on society will be much greater than anything we have known in history.

But the information utility is only part of this future. There are also the intelligent machines that I have talked about — the ethnotronic people amplifiers, the artificial intelligence, expert systems, and the discoursing machines with which you will be able to have interesting conversations.

This will also bring about the electronic cottage industry, which is already starting to happen in many parts of the world. Today, our designers carry home a device so that we can get more than eight hours a day work from them. Then there are the convivial and symbiotic machines that are developing at present, and the telecommunications substitutes for travel, and telepresent systems such as tele-medicine and tele-education. You are familiar with almost all these developments. You can therefore see that there will be many advances and vast changes.

Now let me switch to another subject, genetics. I will then talk about how we expect genetics technology to be married with computers and communications as time goes on.

In every cell of every living organism there is a program called DNA. Each program uses four basic instructions and has many millions of entries. Physically, DNA is very tiny. It has only been with the very powerful microscopes of the last decade or so that we have been able to see it and read it. We have been reading the programs for a variety of species to find out their instruction sequences and where the switch is to turn them on and off.

One process of life is that the program splits, rather like a zipper, into two copies or templates called RNA messengers. One of the purposes of the RNA messengers is to go outside of the cell and make two new copies of the cell from which it came. Another purpose of the RNA messengers is to be targeted internally to the original cell in order to reprogram what that cell does about genetic stress.

About six years ago, we learned how to cut a DNA program and splice in another part. The result is that we can now splice, edit, amplify, and transcribe genetic codes. This has given rise to the creation of many genetic factories for producing medicines, drugs, industrial chemicals, and repair parts.

To give an example of what is possible with repair parts, about a year ago, some skin was taken from a badly-burned animal and cloned by turning on the living switch over and over again, very rapidly, to produce the cell division process. From one division you get two new cells, from two divisions you get four cells, and from three divisions you get eight cells. From 20 divisions (as you know from binary arithmetic) you get about a million cells. That is a lot of skin. The skin was put on to the badly-burned animal. It was that animal's skin but it was not grown there. It was like a bandage that never had to be taken off and that the body never rejected. Researchers have also learned how to grow bone tissue. Imagine how useful such technology would be for repairing human beings in the future.

Let us now look into the future and imagine that you get your arm chopped off. In the year 2000 that may not be too much of a problem. Technology may be such that you go down to the local pharmacist or drugstore and buy a machine to strap on. The machine turns on the growth process and monitors it so that you grow an arm rather than any other part of the body. Since the whole DNA program is in every cell the process will grow the arm, the wrist, the thumb, and the fingers, and then when it gets to the tips of the fingers, it will switch off.

This process was not discovered by first seeing it in nature. The process was researched in laboratories and then the researchers went back and asked whether the process is a natural one. The answer is that it is a natural one, and so they then went back in history to find examples of it.

One of the areas they studied was corn. In pre-recorded history, before there were many human beings, corn had a single strand of kernels in a husk-like pod. Subsequently, corn was genetically stressed by being given too much food in the form of human manure. The result was that the genetic material of corn had to learn how to cope with that stress of too much food. The DNA code was cut and spliced, usually on a trial and error basis. Mostly of course errors were made, and so it took many generations for corn to determine how to make use of all that food. What it did finally was to build a big surface area that we now know as a cob. Now that we understand that process, we can start to think of making a much bigger and longer cob so that it takes far less energy to produce the nutrition available. We can even start to think about growing more on the cob than just corn. We can think about growing the whole meal — corn, peas, potatoes, carrots, gravy, and beef. There is probably no reason why beef has to be on the hoof.

Why are we interested in all of this from the computer and communications viewpoint? In silicon technology, the line geometries of the interconnecting paths and the size of the transistors, resistors, and capacitors are in the micron range, and we are working with two dimensions. With genetic material, we work with features and geometrics in the nanometre range, and with three dimensions.

Why are we interested in making circuits physically smaller? If we put one circuit in a chip, it costs about a dollar to manufacture it. If we put ten thousand circuits in a chip, we have now reached the point on the learning curve of high production where the cost to build that component is one dollar. However, each circuit costs only one ten-thousandth of a dollar. So one of the reasons for wanting greater miniaturisation is to achieve lower cost per function.

From the component manufacturer's viewpoint, he has to take into account the development cost. Say the chip with ten thousand circuits costs about 10 million dollars to develop, if he makes 10 million, that adds a dollar per chip. So the chip costs a dollar to build and a dollar towards the development cost. The manufacturer sells it for a thousand dollars — until the competition takes a close look at the arithmetic involved.

The computer manufacturer's major problem with computer components is how to charge a million dollars for something that small. He will never succeed, of course.

With genetic material, however, each dimension is a thousand times smaller, and there are three dimensions. If we have ten thousand circuits per chip with silicon technology, the extra dimension of genetic material gives us a hundred thousand circuits, and the greater packing density means we can multiply that figure by a thousand times a thousand times a thousand. That is a big number.

So the first advantage of genetic material is a step function increase in capability. Secondly, once we can cut and splice a genetic code to define it as one bit of memory, we can set in motion cell division, and in 20 divisions we have a self-manufacturing process for a million-bit memory.

So the second reason we are interested in genetic material is its capability for self manufacturing. If we stay with the silicon type of technology, the upper range of development is a factor of a thousand per decade. However, with genetic technology we could, in a decade's time, have made a step function increase of a thousand times a thousand times a thousand times a hundred thousand.

Of course we will have to learn how to marry the genetic material with hardware, because the genetic material will not be able to perform such functions as simple multiplication and addition as fast as we can with existing semiconductor technology. However, with genetic material we should be able to imitate the synaptic junctions of our brain, and so be able to do image processing and inference processing much better. So our problem is to marry the two technologies together.

What stage have we reached with genetic technology? We have learnt how to move electrons and electroactive polymers in enzyme living systems and we know that the switch is in brain tissue, at the synaptic junction. Synaptic junctions do far more than the AND/OR/NOR circuitry of computers. So we have the basic building blocks for making biological, genetic computers. We are at the early stages of self-manufacturing, self-producing, genetic systems, solar cells, computer circuits, and so forth.

A possible future scenario therefore contains biogenetic-live computers, sensors, intelligent parts, worker units, interpreters, bionic systems, and human extenders.

The usual marketing strategy is to copy one's competitors. That is the major strategy worldwide, and it leaves one wondering who is actually innovating these days. I have tried to show you some of the innovations that are on the drawing board at the moment, but if you are just copying the competition then you will not yet be interested in these innovations — your unwritten strategy is to let your competition dictate, plan, design, and decide your future.

In the planning department where I work, we have a new science known as the anticipatory sciences. Figure 4 shows some of the jargon of the anticipatory sciences. At the University of Minnesota, where I teach as an Adjunct Professor, we have had a Masters Degree and a Doctorate in Anticipatory Sciences since the early 1970s. Other educational establishments are now introducing similar qualifications. The study of the future has become real. We do not use crystal balls and tea leaves any longer.

Figure 4

ANTICIPATORY SCIENCES

ALTERNATIVE(S) FUTURES
FORECASTING CONJECTURE TREND ANALYSIS
FORCE ANALYSIS EXTRAPOLATION UNCERTAINTY
SIMULATION LONG-RANGE MODELING
CROSS-IMPACT PLANNING
FUTURES PROGNOSTICS FUTURE HISTORY & MAPS
ALL WIN
POSSIBLE- PROBABLE- PREFERABLE- FUTURES
DELPHI SCENARIOS GOALS PREDICTION
EXPLORATORY- NORMATIVE- VISIONARY

Let me show you a language for decoding alternative futures. The language is based on a co-ordinate system, and is shown in Figure 5. The abscissa is the time dimension and the other axis can be anything you like. For example, it could be your health after the age of 20, your earning power after you retire, your relationship with your spouse, or your relationship with your boyfriend or girlfriend. In other words, the language can be used to describe futures for the mixture of products in your corporation. That is the way the future normally comes, as a whole bunch of mixed up things.

Just about every discipline, science or art, has an adjunct-amplifier language. The mathematician has

Figure 5

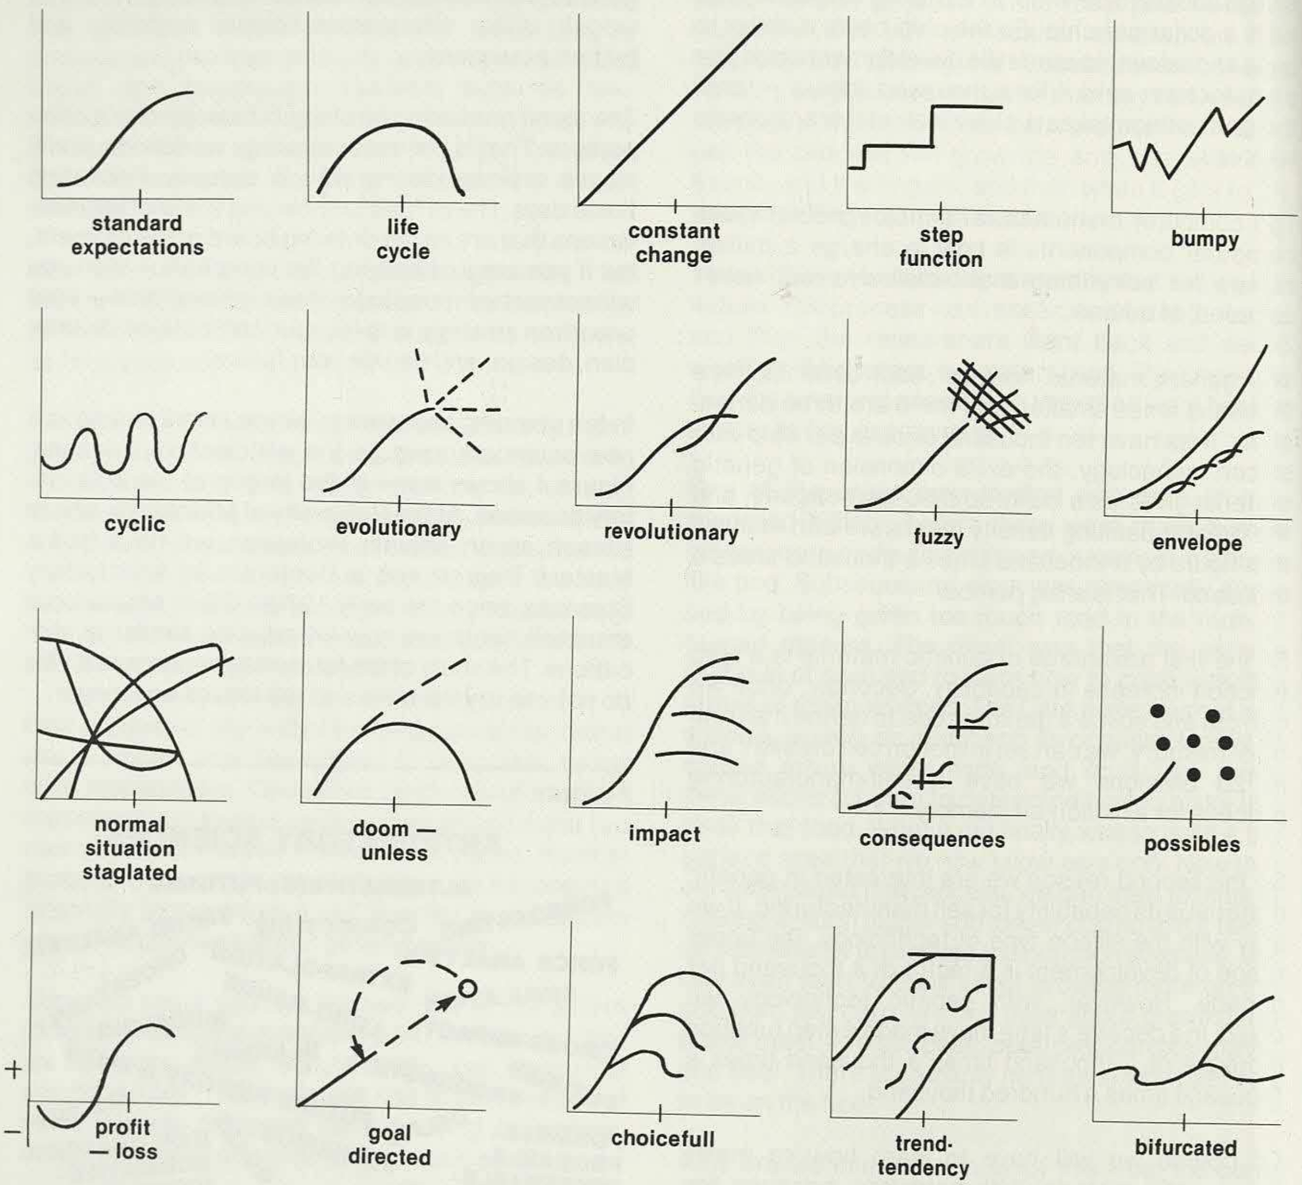

the symbols of Add, Subtract, Multiply, and Divide. The musician has the language of notes.

I discovered this language for alternative futures about five years ago. We are now in the process of defining the vocabulary and the alphabet. Figure 5 shows the first attempt at defining the alphabet.

What I have been trying to tell you is that advances in technology liberate human beings. The blue collar was the labour slave, the white collar the desk slave, the steel collar the robot slave, the silicon collar the intelligent machine slave, and the genetic collar will produce a new species of slave.

THE FACTORY OF THE FUTURE

Raj Reddy, Carnegie-Mellon University

Dr. Reddy is a Professor of Computer Science and Director of The Robotics Institute at Carnegie-Mellon University. He received a Ph.D. degree in computer science from Stanford (1966) and had previously attended the University of Madras, India and the University of New South Wales, Australia.

His research interests in computer science are in the areas of artificial intelligence, man-machine communication and signal understanding systems. In particular, he is working on speech input to computers, visual input to computers, robotics, graphics, distributed sensor networks, and computer architecture. He is the author of over 75 papers and technical reports in these areas.

His current activities in robotics include research towards the factory of the future and autonomous mobile robots capable of operating in hazardous environments such as undersea exploration, nuclear rescue, space manufacturing, and mining.

Dr. Reddy is a Fellow of the Acoustical Society of America and was the program chairman for the International Conference on AI (IJCAI-77) and the general chairman for IJCAI-79.

I was told that most of you are managers or senior vice presidents of information processing in large corporations and that perhaps, 25 years ago, you might have been one of the initial introducers of computers into your corporation. I feel that I am one of you because I was a programmer (applied science representative to be exact) for IBM Corporation in Australia, in 1959. I introduced the 1401 and the 1620 into Australia at that time and, since then, I have moved into the academic environment, but I feel fairly close to your occupations.

My own area of interest in computer science is artificial intelligence. In the area of artificial intelligence we usually look at how one can bring to bear knowledge to solve problems which, when solved by human beings, would be considered intelligent.

This type of activity has been primarily limited to playing games such as chess and backgammon, and activities such as theorem proving, puzzle solving and taking IQ tests. Recently we have been discovering that more and more tasks that we take for granted, such as seeing, hearing and acting, are perhaps a lot more difficult than previously imagined, and in no area does it show up as completely as in the area of robotics.

My talk will be in three parts. The first part will be a general introduction to robotics; the second will be on the factory of the future — what kinds of things we might expect to see and how such factories might be organised; and how can you take a factory and augment it with electronic intelligence so that it might be an intelligent factory. We need to talk about what an intelligent factory might look like. The last part of my talk will, as requested by Butler Cox, be about managerial, organisational and social aspects of this technology.

Let us look at robotics. Historically we have had a view of robotics as an iron knight or as R2-D2 and CP30. Perhaps a more realistic version of robots that we might have seen are systems for aiding the handicapped. This is one area where robotics will be of great help and later, if we have time I will be happy to give a few examples.

Historically, robots have been used as teleoperators, for example in handling highly radioactive material, where an operator with appropriate gloves can control a larger arm, remotely, and perform experiments without being exposed to radioactive materials. These are called teleoperators.

In the 1960s, when I was at the Stanford Research Institute we had a number of activities on robotics. We started by attempting to build a Mars Rover which would be used in exploration of Mars in 1964. We built a robot called Shaky, which could move around in a room and perform simple tasks such as pushing a block. It did not have an arm or a manipulator to do operations, but it could sense its environment and avoid obstacles.

A more realistic robot that you see today is a welding robot. In general, robots today are being used in a wide variety of applications. Welding is one of them; painting and coating is another; others are loading and unloading of machine tools by essentially material handling robots, machining, transfer of materials, and assembly tasks.

The interesting thing about these robots, unlike the science fiction view, is that they are essentially mechanical devices that are programmable. Most of you are familiar with the notion of programming. What distinguishes these systems from classical, hard automation is that these robots are soft automation. Today they can be doing one thing, and tomorrow they can be doing a different thing. So because they are programmable they can be used in a wide variety of tasks without having to reinvest in new capital each time you have to change the factory, or each time the demand for a particular product decreases and a new product has to be introduced. Thus, you do not have the problem of recapitalisation, over and over again. This flexibility and versatility is a direct consequence of the programmability of these devices.

Secondly, they can operate beyond human capabilities. They can work for long hours, in continuous operation, in uncomfortable environments; they do not need lighting or air conditioning, and they can perform hazardous tasks.

Finally, these systems are often capable of performing high-precision tasks, requiring a great degree of accuracy and repeatability that human hands are not capable of.

What are the advantages of using this type of robot? Essentially, the whole technology of programmable automation makes it possible for us to have a single factory in which you could be producing a large number of different kinds of products. You could be producing automobiles today, trucks tomorrow, and tractors the day after, if the factory is organised with the right type of soft automation and programs.

You do not move the machinery around, you do not even change anything; all of the changes are programmed in software, which means that the changeover from one product type to another is essentially automatic. There is no rearranging or restructuring of the factory.

That is not quite possible in many of the existing factories because they were not designed to be flexible manufacturing factories, but in general that is exactly what robotic, flexible automation permits you to do.

Secondly, the use of robots, or in general flexible automation, leads to ease of phasing in product design, modifications to products and changes in products. You do not have to be right the first time around. If you make a mistake in the product design and you notice that some products have gone into the market with a problem, you do not have to throw out either the whole design completely or the machine. Suppose that you programmed the machine tool to machine the product in a certain way. Essentially, most of it was right except for a small part. At present you would have to throw out all the equipment that was designed to make that product and maybe have to re-do the tooling. If not, a person would have at least to re-do the whole prototype from the beginning, in a slow, laborious way.

With a soft robot, maybe one small routine of a few instructions might have to be changed. The changes can be made and the raw material can be cut again, and you can verify whether your product is now up to specification. A classical problem in machine tool shops is where they make two parts that are supposed to mate together. The holes do not line up. Today you have a two-week delay before they machine another part and re-do the whole thing. If you have completely soft automation, where all you have to do is to change a couple of lines of code, you can have immediate turnaround and you can produce a new prototype without the delay that normally ensues. So there is no retooling, just reprogramming.

Improved operating ratios and operating times of the factory are other benefits of the use of robotics. Normally what happens now in a conventional factory is that about 25 to 35 per cent of the total operable time is wasted because there are shift breaks, coffee breaks, lunch breaks and so on. Also, the machine tools and machines are sometimes idle because the parts are not in the right place at the right time, or the person did not act fast enough or in time.

If the whole thing is fully automated or is essentially an unmanned factory, most of those delays are eliminated. This can lead to improved throughput of anything from 20 to 80 per cent compared with an existing factory.

This is what we mean by making a factory intelligent; by adding electronic intelligence, such as computers and distributed systems with sensors to every machine tool in the factory (i.e. augmenting the mechanical intelligence with electronic intelligence) we expect the throughput of the factory to increase by 20 to 80 per cent. On average, it is slightly over 50 per cent in the examples we have seen.

If you take the cost of a factory as anywhere from $20 million to $100 million, then by installing approximately 10 per cent additional electronic equipment or computers and software to this factory, you are able to improve your throughput by, say, 50 to 60 per cent. You will probably be able to pay for the whole addition to this factory within one year. Also, you do not have to build a second factory and increase your capital investment.

The other advantages of robots are that they produce a predictable output. If they are wrong, they are wrong all the time. It is not a question of sometimes you produce a right product, and sometimes you produce a wrong product. The quality is precise and controllable. This is one of the areas in which the Japanese have made great strides. By using a substantial degree of automation, including robots, they are able to produce predictable quality output; and if the quality is not acceptable, you improve it. This is one of the areas where, if you want to compete in the world market, it becomes essential to use automation.

Finally, the systems are able to withstand severe working conditions. One of the reasons that industrial concerns have not prospered so well in tropical climates is because of the heat. You have to air condition the whole factory at great expense so that workers can work comfortably.

There are three different kinds of robots. The robots that are today chiefly used in industry are what we call blind and deaf robots. They are programmable, but they are programmed to do the same task over and over again. If something goes wrong and the environment is not as predictable (for example, the part appears but it is not in the right place), then the robot still tries to do the same thing, so that you might end up with an engine without engine blocks.

The second and third generation robots that we are looking at in the research laboratories are seeing and thinking robots which can be programmed. They are sensor based and adapt to any variability in the situation. The third generation robots are closer to R2-D2 in that they are mobile, can do optical detection and take avoiding action. We at Carnegie-Mellon have about three different mobile robot projects.

Normally, a mechanical manipulator might want to pick up an object and if it does not know where the object is, it is in trouble. By adding a TV camera to the system and processing the image to determine exactly the position, location and orientation of the object, you can direct the manipulator to go to the appropriate place.

Whereas previously the robot was programmed once and repeated the same function for the whole day, the newer robots are being programmed continually. They are, in effect, being dynamically reprogrammed, although it is really not that sophisticated. They are merely interpreting a routine dynamically and, given the co-ordinates, move their manipulators to a different place each time. It leads to some difficult mathematics, but it is not impossible.

With sensory processing you incorporate the knowledge about the types of objects that you want to recognise — pattern recognition — and process the sensory data using this knowledge and model. The total system integration combines the sensory parts to the effective parts so that you have what we call a hand-eye system.

One of the mobile robots being built at Carnegie-Mellon has a number of interesting features. It has two wheels, each being independently steered by a separate microprocessor. If you think of just that one technology, it is a good example of augmenting or replacing mechanical intelligence by electronic intelligence.

Today, a car has a steering wheel with mechanical linkages. If you want to go somewhere, you turn the wheel and you go in a particular direction. The wheels are tied together fairly tightly by mechanical motion, so that certain kinds of movements are not possible. For example, you cannot move sideways and park and, if you are going up a steep hill, you cannot go up in a roller coaster fashion. By eliminating all the mechanical linkages and replacing them by electronic linkages, where the microprocessors are talking to one another and determining that they do not pull in two opposite directions (if they do there is a bug in the program) you can think in terms of many different kinds of motion that were previously inconceivable.

This is not a new idea. I do not know how many of you are familiar with some of the NASA experiments with the so-called fly-by-wire. If you were designing an aircraft in the old days, you would have long wires connecting the cockpit to the rudder and other controls. If any one of the wires were broken you were in big trouble. By replacing the mechanical linkages by electronic signals, by packets if you will, and by having (in the conventional networking sense) redundant paths of communication for these packets, you can now take a direct hit on one side without completely incapacitating the aircraft. This idea of fly-by-wire is another example of replacing mechanical intelligence by electronic intelligence.

Currently, machine tools will do the same thing again and again. They are programmable, but they are not sensor intensive machine tools, so they will do the same operation over and over again. If the tool were to malfunction, it might machine the air, over and over again, without realising that something was wrong. By adding force, pressure, temperature and vibration sensors to a machine tool, and tying them back into the controller and automatically reprogramming the controller dynamically it is possible to overcome the malfunction. We keep coming back to this whole notion of automatic reprogramming. We do not do it routinely in data processing applications, but in many process applications it becomes an important consideration. Whether you can change the program dynamically to cater for changing environmental situations is an example of that.

Now we come to other types of automation. When people talk about robots in factories, less than 10 per cent of the tasks in a factory require a physical robot doing something. If you look at a total factory, about 60 per cent of the people in a factory are blue collar workers. If you look at the tasks they perform, some of them are performing material handling tasks, others are performing tasks that involve assembly; and others are performing inspection tasks. About 30 per cent of the people that are in the manufacturing line are mainly using their sensors to perform inspection tasks. That is one of the most tedious and boring types of task.

An example is printed-circuit board inspection, which is one of the things we are trying to automate at Carnegie-Mellon University. To give an idea of the computational complexity of this task, a printed-circuit board for electronic manufacturing is about 100 square inches in size and you are looking at dimensions which are approximately 5mm, or maybe as small as 2mm. For that type of dimension you have to look at least at 1,000 elements per linear inch, so per square inch you need a million pixels or picture elements. So if you have a printed-circuit board of 100 square inches, you need 100 million numbers.

No computer has that size of primary memory, so of necessity you have to use some kind of paging system to process the information in real time. The computational power required is enormous. Just to look at one element requires anywhere from 50 to 100 instructions. Processing the information requires looking at a large number of these elements and doing edge detection or fault analysis. If you do it in a straightforward way, it is not uncommon to require a thousand operations per pixel. Now you have here 100 million pixels. So all you have to do is multiply those two and you end up with having to perform 100 billion operations on a computer. If you are going to do it economically, with a 1 mip computer, you are still talking about a hundred thousand seconds, or approximately 30 hours of computation.

That is not economically practical. But by using a number of tricks of the trade, such as low-level microcoding of specialised instructions, or using special purpose algorithms that do not have to be accessed, from secondary memory, we have reduced what would normally have taken 30 hours of computation to five minutes on a small computer called PERQ, which is marketed in Europe by ICL. It is a microprogrammable, megabyte engine which costs about $40,000 in single-unit quantities. It has high resolution graphics and Ethernet networking facilities. This whole system is running as an integrated system, and it gives you an idea of the complexity of an inspection task.

What I have just described applies to a two-dimensional inspection task. If you are trying to decide whether a particular part is of the right dimension, let us say a turbine blade or a complex shape, you have to have three-dimensional sensing capability. Again this increases the complexity of the task by another two orders of magnitude from the previous tasks that I talked about. Looking at complex, three-dimensional shapes and determining whether they are the right shape and tolerance turns out to be a very difficult problem.

We are working on about three or four 3-D measurement problems. We use what we call a light striping technique where there are four different cameras and four different light stripe projectors. They look at the stripe at any given point, and the robot or some moving device is moving the whole object up and down. So you essentially get profiles of cross-sections at uniform times. This involves complex triangulation and other types of computation. But it can be done.

Another example of the use of robotics is in the inspection of neon light bulb filaments. These filaments represent about 2 per cent of the total cost of making a bulb. The total shrinkage in a light bulb factory is about 8 per cent, that is 8 per cent of the bulbs that are made are faulty. Of that 8 per cent, 5 per cent is accounted for by bad filaments either because the shape of the glass is wrong and therefore the vacuum seal is not working, or they are double filaments, or one of the filaments is missing, and so on.

This is an example of hard automation. You are producing millions of bulbs very rapidly. People can inspect them, but not as carefully, and they are not able to inspect the shape as precisely because they are not capable of doing that. So we are building a specialised inspection station to do this type of inspection.

I should now like to go on to my next topic which is the shape of the factory of the future: what does it mean; where are we likely to be; and what should we be thinking about?

A good example of what the factory of the future might look like already exists. When I was visiting Japan last year, "The Factory of the Future is Already Here" was the title of a newspaper article describing the CANBAN system of the Toyota Automobile manufacturing plant.

This particular system is really not an example of the use of robots; robots are used, but the really important part of the system is the intelligent use of computers to control the factory. Using this system, Toyota is able to manufacture the same number of automobiles as in the USA and Europe, using about two-thirds of the labour, one half of the floor space, and 15 per cent of the in-process inventory.

Toyota uses classical data processing techniques in the factory. They have tied themselves, their suppliers and their suppliers' suppliers all into the same computer system. All their databases talk to one another and are compatible. So when the production manager decides that he will produce so many cars tomorrow, all the suppliers are immediately notified of the requirements of the parts; and they in turn, depending on their inventory and the supplies to the suppliers, notify their suppliers immediately through computers.

As a result, the parts, 100 per cent inspected, are delivered to the factory floor when they are needed. The system is so integrated that if there is a traffic jam on the Tokyo freeways, a radio message comes in saying that the parts delivery will be delayed by an hour; they shut down the whole factory, take a break, and come back after an hour. Because the in-process inventory is so small, you do not need that much floor space. If you look at machine tools in an automobile factory in the United States, you will find about 200 square feet of space per tool to stack up all the raw material that it needs.

So an example of what it means to have a factory of the future already exists; it means effective control of capital resources by using computers. There is nothing brilliant about it, but it is a systematic and effective use of the technology with which we are already familiar.

This is why when we analyse a factory we work closely with a number of industrial manufacturers in the United States, such as Westinghouse and Digital Equipment Corporation. We almost never spend much time looking at how they can reduce their labour costs by a small percentage. That is important, but it turns out not to be where most of the savings come from. Effective use of their capital, resources and inventory is just as important, if not more so. So we talk about improving the productivity of machines, improving the productivity of blue collar workers, improving the productivity of white collar workers, such as engineers, and improving the productivity of managers.

So it is not simply a question of putting a robot in a factory, because you will not get more than a potential improvement of 10 per cent simply by putting in a robot. It is the managers for whom you need the tools. They have to make decisions in an uncertain environment: a person is sick; the raw material did not arrive; there is a rush job; a machine tool has broken down and has to be maintained. The rules for making these decisions are not written down in any rule book. At present, the managers of a factory make these decisions with incomplete and sometimes inaccurate models of what is going on in the factory. As a result, significant delays and unnecessary wastage of resources (people and machines) occurs.

What tools are needed to increase the management productivity in a factory by, say, a factor of 10? In theory, it is very easy to answer this question, and operations researchers, planners, analysts, and schedulers have been doing this for the last 20 years. Unfortunately their solutions all depend on manual input of raw data, and this leads to the classic garbage in and garbage out problem. One of the reasons that MIS systems have fallen into disrepute is because they depend on somebody to put in the right data so that they can generate management reports. Managers look at these reports and say, "That's not reality. Something is wrong," and they throw it out, without even bothering to look at it.

So one of the most important parts of making effective use of computer technology is to make the environment, whether it is a factory, an office or a financial situation, sensor intensive so that the data that you need is entered into the computer automatically as it is created. This happens to some extent in some of today's point-of-sale terminals used, in supermarkets for inventory control, and so on. It also happens to some extent in banks. But the financial information in conventional manufacturing industries is not as readily available. In a typical manufacturing plant, many different processes are carried out, and the financial information is scattered all over the place, and the databases of the different computers do not talk to one another. As a result, you can never make the right kinds of decisions with that kind of uncertainty about the information you have. So you use your intuition and say, "I have to make a decision today. I'll do it this way".

So, as a first order of approximation the factory of the future is a sensor-intensive factory, where everything that is happening in the factory is monitored and reported by sensors to the computers, with no human intervention.

A consequence of that statement is that if you are building or planning a new factory then, in addition to putting power cables into the factory, you have to put in an information cable, such as Ethernet or some other kind of network. You need some means of directly communicating the status of the tools, the machines, the people and the raw materials into the computers as the data is being created.

Most factories today probably have some kind of clocking-in procedure. I do not know how many of them are automated or connected via computers to an appropriate database, so that by 8.30 or 9 a.m. when work starts, the manager can immediately know which people are not available. Ideally, the manager should have presented to him on a CRT the consequences of those people not being available. Not only the consequences should be presented, but proposed solutions, and how he could rearrange people, giving three or four alternatives, and what the impact of each of those solutions is.

If one can provide a management decision-making tool of that type, it is still an interactive tool, but most of the routine decision making and available alternative options are presented by the computer. That will permit the manager immediately to make decisions rather than waiting for the foremen to report and then running around to see what kind of delays are occurring in what place.

So the first step towards a factory of the future is a sensor-intensive factory in which the whole factory is wired with an information cable.

A second consequence of that statement is that the information cable must be fault-tolerant and fail-soft. It cannot be controlled by a large, central computer which brings the whole factory to a standstill if it breaks down. So one of the important consequences to you and to us in a research environment is how to build distributed systems in which there is no master and no slave. All systems should be capable of taking responsibility as a master and all systems should have distributed databases so that all the information is not concentrated in one place.

This is, in fact, a classic problem in human organisation. Consider an army or navy during a war. If something happens to the commander, somebody else automatically becomes the commander. The procedural rules for transferring command are well known; everybody knows what to do in given situations. So there is nothing magical about organising a distributed system so that it is fail-soft and in which you can literally turn off half the machines at any time. The system might hesitate for a few seconds, but it will continue to run without any stoppage of the system.

The software and hardware tools for designing those kinds of systems do not yet exist. Even programming a system which will dynamically reconsider all of its processes cannot be done with today's technology. For example, if you have a 100,000 line program which is running on five different machines and one of those machines is turned off, the system cannot dynamically reconsider how to use the remaining machines and continue to run. We are working on this type of problem; not for the factory of the future but for the Department of Defense, in a project called Distributed Sensor Networks which is similar to this problem, but for cruise-missile defence of large geographical areas. The research problems that arise are numerous and they are concerned with distributed problem solving, distributed databases, distributed situation displays, distributed architectures, distributed programming languages, etc. We have barely scratched the surface of every one of these problems.

I want now to describe what I call smart sensor technologies. These are sensors that tell the computers what is going on. At Carnegie-Mellon we are designing a direct-drive manipulator. Most robots are designed with gears, which results in imprecise positioning. By using high-torque variable motors which are directly embedded in the joints, you get a direct-drive manipulator. The technique is similar to record players, which used to be gear-driven, but now use direct-drive motors that give greater precision of motion.

Sensors in general result in a large amount of data. If you are looking at a visual scene, it is not uncommon for a TV camera to require a data rate of a gigabit per second. Intelligence gathering satellites are now producing anywhere from 10 billion to a trillion bytes of data every day. So image and sound sensors are huge data producing devices; and 99 per cent of this data goes straight into the NASA archives and is never even looked at. You may well ask why. One of the reasons is that there are no computers today that can process that volume of data and extract the relevant information. So what we need is smart sensor technologies, where sensors are augmented with electronic intelligence. Thus, you would have a microprocessor or even a large processor associated directly with the sensor, which immediately discards a large part of the data so that you receive only the relevant pieces of information in a highly reduced form.

This is a research topic that requires integrating sensor technology, with microprocessor silicon technology. Sometimes it is possible to have both of them on the same chip or wafer, because the sensors are also made out of silicon. Intelligent sensors could also be used in machine tools so that they can be force-adaptive, dimension-adaptive and vibration-adaptive.

The next aspect of the factory of the future that I want to mention is what it means to have a no-inventory factory, a factory that can produce parts on demand. If you can produce any part that is asked for on demand, untouched by human hand, then you can eliminate all inventories. This is what we call an electronic warehouse.

A prime example of the need for an electronic warehouse can be found in the US defence requirements. B52s were built in the 1950s using germanium transistors and vacuum tubes. The technologies have advanced four more generations since then, but we still have warehouses full of these old mechanical, electrical and electronic parts. It is estimated that 80 per cent of them will never be used, but will eventually be sold as surplus for scrap. If you can produce these components, assemblies and systems on demand by having the necessary manufacturing instruction sets in an electronic warehouse, the potential cost savings are enormous.

There is concern in the United States right now about the ability to convert industry to a war-time industry — what we call industrial preparedness. It is estimated that the United States would run out of weapons in less than three months and that it would take three years before industry could be geared-up for war-time production. That is just not acceptable. So there is great concern about how to change the factories, how to design factories in the future so that they could produce peace-time, consumer products but, on demand, could instantaneously be switched to produce some defence-related product. These types of factories are conceivable and possible today. It just needs somebody to sit down and solve the problems and put the systems together. It is a huge systems integration problem.

At Carnegie-Mellon we have a policy of not building anything we can buy. We buy whatever we can, and we design and build what we cannot buy in the market. Recently, we put together an automated swageing cell jointly with Westinghouse. It had two robots, two vision systems and a huge, general forming GFM swageing machine and a furnace, a couple of controllers and about eight different computers controlling all these devices.

The software protocols, and the assumptions made by the software we bought caused chaos. The software assumed that the robot would sit in isolation and work by itself, and that the inspection system would work by itself with no need to communicate with other parts of the system. The controllers for the machine tools and the furnaces used microprocessors, and we had to rip out all the programs and re-write them. The suppliers, of course, did not want to tell us anything about their programs because they said they were proprietary. So, even converting existing tools into an integrated factory of the future is an enormous job. I am sure you understand that, being at the centre of organisations where you have to integrate a diversity of information and worry about all the databases talking to one another.

It is a difficult task, but it is not something that requires a great deal of creativity or new breakthroughs; it is routine programming which, as we know, is a creative activity, but it takes lots of people lots of days or years. But it is possible to do it.

The last aspect of the factory of the future is what you might call white collar robotics. There are several types of white collar workers in a factory whose productivity you might want to think about improving. For example, there are those that design parts, such as design engineers who produce design drawings using CRTs and graphics systems. The drawing is sent to the manufacturing organisation, and they use some kind of graphics tablet to re-input the drawing into a computer, and use that information to create the manufacturing instructions for the machine tools. These are then sent by paper tape to the machine tools.

There is no reason why all of those manual steps have to take place. The design can be converted to a drawing and the drawing redigitised back into an electronic drawing. It is a simple process that most of us know and understand. But for it to happen, two computers have to be connected together. If the two groups have different computers that cannot talk to each other, you have a problem. Or maybe the computers can talk to each other, but there are significant differences in the way in which the same information is represented in the two databases. Translating from one representation to the other could well turn out to be a major and complex problem.

So the first stage of white collar robotics is an integrated system in which you have a functional specification from which you derive the design. The system needs to have access to the accumulated expertise used by an engineer when he designs complex mechanical parts or electronic parts. Usually, there is nothing about the design rules that cannot be captured.

We have an "expert system" project which captures much of the routine design knowledge of an electronic designer. If you want to design a process control computer with a given performance, the system will use a number of cost and performance curves, to suggest half a dozen different architectural diagrams for the process control computer.

The architectural diagram can then be input to another expert system, which converts it into the required logic circuits. The type of 'knowledge' codified in such a system concerns known facts about how to design input interfaces, how the memory interface is designed, what the bus structure is, what the CPU instruction set is, how to implement the instruction set, and so on. This knowledge is usually in the heads of half a dozen people in a corporation or a design group and, if something should happen to them, you are in a lot of trouble. But, in general, there is nothing magical about the knowledge and much of the information is in fact routine. It is only about once every five years that design engineers come up with a novel architectural feature, such as cache and memory hierarchies, or paging and virtual memory. Once you have half a dozen implementations of the idea, you can capture the design features in knowledge libraries, just as you have libraries of subroutines for doing other mundane things. The expert system can then use cost/performance trade-offs to determine which of the half a dozen implementations you require. The input parameters will include the required chip count, or the production volume, or the required power dissipation, and these, and a number of other constraints, will dictate which of those half a dozen choices is appropriate.

So going from a functional specification to design, from design to manufacturing, from manufacturing to production, can all be significantly automated. We are doing a number of these things as part of a design automation project. This type of automation will make it possible to change over rapidly from one kind of production to a completely new product. So instead of stockpiling obsolete germanium transistors and training people in obsolete technologies, you can have a man/machine system that can produce a spare part using modern technology, but which performs the same function.

The second interesting area of white collar robotics is what we call automatic programming from geometry. Suppose you have a mechanical part you want to produce from raw stock. It will need to be machined, swaged, or formed in some other way. Today, machining is an art. A machinist looks at the drawing, looks at the raw material, and begins cutting, removing more and more material at each step. If an NC or CNC machine tool which can be programmed is being used, then a manufacturing engineer writes programs to control the machine. He goes to the machine tool which costs half a million dollars, stops all the other work, puts in his new program, tries to debug it, with a finger on the panic button just in case he misprogrammed it and there is a bug in the program which will destroy the whole machine and put it out of commission for a whole week.

There are ways in which this process can be automated. One way is to capture much of the knowledge of what it takes to produce a special geometric part in a machine tool. In effect, we are automatically generating programs. We call it "programming from example". At present, when you want one of your associates to write a program you call him and say, "Here is the specification of what has to be done. Go and write a program". There is a group of researchers working on how to automate such a programming task — 'what' rather than 'how' programming. You do not give the algorithm, you just give the specification, and the aim is to construct an intelligent programming engine that will convert the specification into a program.

There are already a number of simple examples of this type of process. RPG is a good example. With RPG you can write out an invoice format, the names and other items, and a listing program is created automatically. Unfortunately, if the file also contains other types of information, it too will be listed as an invoice and will appear as nonsense. But you can now think in terms of an intelligent listing program which looks at the extension of the file and says, "This is a text file. This is a drawing. This is a financial statement. This is a table," and will print it in the form appropriate for that data.

These types of concepts already exist, but they are not readily available in most computers today. So automatic programming, or 'what' rather than 'how' programming, from a firm example, whether it is in the factory of the future or the office of the future or financial systems of the future, is very similar for each type of application. But they all require substantial additional research which we have not yet done. Each of them poses a different class of problems.

Automatic programming for factory applications is very difficult because you are dealing with complex geometry. We do not know how to deal with shapes or three-dimensional matrices where you are talking about resolutions of a thousand points to an inch. For a cube of one foot in size, you have 1200 x 1200 x 1200 points, which is close to 1,726,000,000 cubic pixels. Processing that amount of data is a major problem.

Now I should like to say a few words about the use of robots in hazardous environments. At present we are working on three kinds of systems: systems for use in space, systems for nuclear environments, and systems for ocean exploration and exploitation. With the ocean systems we are looking at autonomous, underwater robots for defence applications, for mining, for the inspection and repair of offshore structures, and for salvage and rescue operations in the oceans.

Currently many of the offshore drilling platforms are in shallow water, about 500 feet deep. If you want to drill in a few thousand feet of water, the technologies are non-existent. Exxon is currently experimenting with one deep-water drilling technology, but in general you need autonomous systems that can operate by themselves with sensors rather than current offshore technology.

The same type of technology is useful also for autonomous navigation or for performing specific missions for defence, or for harvesting the oceans. We have an underwater autonomous robot project supported by the Office of Naval Research. The main research effort is in the areas of target identification and landmark detection in the oceans. Other research areas include obstacle avoidance, path planning and navigation in the oceans. On dry land there are traffic signs saying "Turn right" or "This way to Davos". There are no traffic signs on the ocean floor and you have to guide yourself by nature. So navigation is a much harder problem.

We have similar problems with the space systems. We are looking at the concept of a space factotum, which can be used for the collection of garbage, for material handling, construction of space telescopes and antennae, and space rescue missions. In the area of nuclear systems, monitoring a nuclear power station with about 10,000 sensors continuously to detect any abnormalities, and mobile rovers that can inspect and repair in that kind of highly radioactive environment, are also important.

The fundamental techniques needed for all of these systems are essentially the same: sense, think, act and navigate. Each environment poses problems of different kinds: what kinds of sensors to use; how to process the sensed data; what kind of computational power can you have in two cubic feet of space; and so on. But the computer science techniques are basically very similar, no matter which one of the environments you look at.

At Carnegie-Mellon University at the Robotics Institute we have about 100 professionals, probably the largest in the United States, by an order of magnitude. We have about 30 Ph.Ds, about 30 engineers and programmers, and about 40 graduate Ph.D-level students, working on a number of these areas.

What are the implications of some of these technologies for the managerial, organisational and social issues with which many of you are faced in day-to-day problems? Almost all organisations are in the process of decentralising their data processing activities, and this can be viewed both as a threat and as an opportunity. It is a threat because you no longer have control over all the computing in the corporation. The empire cannot get larger, it can only shrink.

But at the same time, in almost any corporation, much of the resident knowledge about how to use computers effectively is within the data processing groups. You need the appropriate foresight to say that it is not simply about payroll and accounting and financial statements. Rather, it is about integrated total corporate information systems for the whole corporation. Much of this is not traditional data processing at all, but is mostly concerned with communications and protocols, electronic mail and electronic publishing, and a whole range of other things that currently you do not do. In general, it is a very different kind of activity to that carried out by most corporate data processing organisations in the past.

We are finding that the role of our computer centre in the university is changing substantially. We plan by 1985 to provide each incoming student with a personal computer with the power of a VAX. When he leaves, he will take it away with him. So we have been asking what is the role of the computer centre and the computer centre director? It turns out that there are a large number of new organisational problems that come up. You need a place where the students can take these 5,000 computers and have them repaired. You need a place where you can train people. You need a place where new software systems can be produced. The role is not very different from the current one, but the plan does change the nature of the services that are provided by the data processing group.

The data processing group needs also to keep abreast of developments in the communications' field. Should you buy a whole transponder or only one-tenth of a transponder? What can you do with a 2 megabit world-wide communication network? What can you do with voice grade lines of 56 kilobits? Currently, AT&T and many of the European PTTs are working closely with CCITT to define a standard that will enable you to use 56 kilobits on your local line without using modems. (Most of the long-distance communication is already digital and is transmitted at 64 kilobits on every voice grade line.) If the local exchange is an electronic switching station, the technologies now exist to enable you to use 64 kilobits without any extra cost. For most corporations that is adequate bandwidth for many of the routine applications. The implication is that you can dial out when you need to, and you do not have to invest in a large private network.

But these alternative means of communication and alternative system designs, alternative fail-soft environments, present extremely interesting issues of choice for management. Other issues, such as incompatible databases, also have to be resolved.

Also, management has to come to terms with the continual rapid technological change, where we are talking about 30 per cent improvement every year in the cost performance; every four years the cost of computing will be one quarter of what it is today for the same performance. Dealing with these types of issues requires a constant monitoring of the technology and making appropriate decisions today for an installation three years from now.

There are also issues concerned with new programming languages. Most of the programs I have talked about today cannot easily be written in languages like Fortran or Cobol, because they are programs that can create other programs. One of the reasons that Lisp is the chosen language for artificial intelligence applications is that the programs and data are interchangeable; they all have the same structure. Also, Lisp programs can generate other programs that can be immediately executed.

Let me give you an example. Suppose you ask a question of a database: how many people of ages 35 to 40 earn between $35,000 and $50,000 in the organisation? Suppose also that the required information is not in the database but has to be computed. You have two options. You can either call a programmer and say, "My boss wants this information. Why don't you write a program and get me the data by the end of the day?" Alternatively you can use a program that can write the required programs; you specify the information required, and the program will create another program which will operate on the database and give you that information. Other languages such as Ada and Pascal are not quite as powerful as Lisp, but they are much better than Fortran and Cobol.

We had a major battle over programming languages with our industrial research sponsors such as Westinghouse. They said, "99 per cent of the people I have in my organisation are Fortran or Cobol programmers. You are talking about programming in all these advanced languages that we don't know anything about. What do we do with our existing people?" The answer is that you have a problem. You should run the old applications in the same way until the machine and the application die. But anything new that you do, for heaven's sake do not program it in Cobol. You ought to be thinking forward, because you do not want to be forever limited in your expression capabilities by some ancient programming language that was invented just as computers were coming into being.

So it is very important for the managers of information processing in a corporation to take the lead and have a better understanding of where to use what language, and how to have interchangeable systems.

Lastly, I will say a few words about the social impact. There are two issues here: one is unemployment and the other is what to do about it. Every time we talk about robotics and automation, a large number of people throw up their hands and say, "That will create a lot of unemployment. Do we really want it?" I am not sure that we have an option, unless you want to build a wall around each country and say, "We don't want advanced technology. We don't want to import or export anything. We want to live with the technology and the society we now have." You could freeze the whole society by putting up walls.

Inevitably, you will be part of the world economy and if you do not embrace automation, someone else will. The Japanese, for example, are bound to use the most advanced tools, so they can be highly productive and produce better quality products at much lower cost. There is no way you can compete unless you also adapt to the same technology.

Then the question is: what to do about the resulting changes in the jobs and the skills required? It is not Toyota that has unemployment today; it is Ford and General Motors that are laying off people because they have not aggressively used computer technology as well as they should.

I have been a consultant for General Motors for a number of years, and one of my major complaints, which they continually ignore, is that their whole computer science department has a total of only 15 people; and half of them left recently because of their poor salary structure. General Motors has difficulty in acquiring and keeping good people. They realise now that they could have done many of the things I have been talking about 10 or 15 years ago. They should have been more aggressive about realising the capabilities of using computers. But they did not. When we talked to them about using sensor-based robotics, which we already had running in 1968, they said, "That's crazy. We don't want to produce small batches of cars. When we design our factories we design them to produce 10 million units". Consequently, they have absolutely no flexibility in their factories. They cannot change rapidly to counter competitive moves. Once they have built their factory they are stuck with using it for three more years if they want to recover their cost. It is an unfortunate situation.

The question is: where will the new jobs be and how will we deal with them? We see three areas. One is the information industry itself. Right now the information industry in the United States provides about four million jobs, which is about 5 per cent of the total workforce. By the end of the century it is expected that the number of jobs in this area will be between 25 and 30 per cent of the total workforce. So we need to begin to train people to take advantage of the new opportunities, and not train more welders and painters.

Secondly, there is no threat of over-production. People say, "If we are ten times more productive than we are today, we may suddenly have a lot of goods that nobody wants". That may be true in any one country, but if you are talking about the world economy, 80 per cent of the population have nothing today. If we try to bring their standard of living up to that of the other 20 per cent, there will be a huge world-market for your goods.

The next objection raised is that people in developing countries do not have money to pay for the goods. That is not necessarily true. Some developing countries such as Saudi Arabia have a lot of money. Others have the capacity and the natural resources. If you look at Japan and India after World War II, or countries in Africa which became independent about the same time, you will find that they all started with nothing. One country is a leading industrial power today, and the other countries are still stumbling along.

The main difference is the skill levels of the people and what they can do. The Japanese were building steel plants, ships and other kinds of industrial equipment, well before the second World War, so even with their entire industrial capacity damaged they were able to build up their economy very quickly. Most of the other developing countries have no skill levels, so the key thing for them is to improve the skill levels of their people. If they do, their productivity will improve, their purchasing power will improve, and the size of the world economy will improve by a factor of 10. That is our current prediction. If that happens, there will be a lot more capacity needed and a lot more jobs throughout the world.

Finally, there is the issue of the reduced working week. It is already a reality in some countries such as France. It may not be long before we are working only 20 hours a week and doing routine jobs. The rest of the time will be spent either in going back to school or in intellectual tasks such as painting, art or music; but it will not be spent on boring work.

COMPUTER FACTORS IN HUMAN SYSTEMS

Philip Kraft, University of New York State

Philip Kraft is Associate Professor of Sociology at the State University of New York at Binghamton, where he teaches in the Graduate School and the School of Advanced Technology. He is a specialist in the sociology of computer occupations.

Mr Kraft's publications include "Programmers and Managers: The Routinization of Computer Programming in the United States" (Springer Verlag-New York), as well as several articles on workers and computerisation. He is currently directing a three-year study of women in computing occupations.

The usual meaning of the word "lag" is the time between the introduction of a new technology and its generalised use. There is also a different and more important kind of lag, which is the period between the introduction of a new production technology and the creation of new forms of social organisation appropriate to it.

The industrial revolution began nearly 300 years ago. It took another hundred years before Smith and Ricardo clearly stated the principles of the market-place — that the price of a commodity is determined, first, by supply and demand and, secondly, by the cost of the factors required to make it, including, above all, labour.

It was yet another hundred years before anyone figured out what that meant. For example, in the United States at the turn of this century, the iron, steel and coal industries were essentially modern in terms of their technology, and yet they were organised as huge cottage industries which happened to be organised under a single roof. Production orders were sent out for bid to what in effect were in-house subcontractors, who consulted with their own work gangs about production prices, division of labour, wages, and the pace of production. Industrialists provided the capital, the physical location of production, and most of the raw materials. Missing from this list were the workman's tools and virtually all of the management as we now understand the term. Although the production technology of these industries was modern, at first it was simply grafted on to an old form of production organisation. This process lasted well into the 20th century and, if you believe some British managers, it is alive and well today. I suppose that last point is open for debate.

Of course the situation changed. Ironically, one of the main figures in this process of change was none other than Charles Babbage. If most people here know of Babbage's historical role in computing, probably fewer know that he also formulated the original principles which still guide modern management theory. Let me read to you a brief excerpt from his "Economy of Machinery and Manufacturers" published exactly 150 years ago.

> "The master manufacturer, by dividing the work to be executed into different processes, each requiring different degrees of skill or force, can purchase exactly that precise quantity of both which is necessary for each process; whereas if the whole work were executed by one workman, that person must possess sufficient skill to perform the most difficult and sufficient strength to execute the most laborious of the operations into which the art is divided."

If the language is quaint, the meaning is clear. Babbage made a fortune from various manufacturing industries and was not a theoretician by occupation. He therefore spoke from first-hand experience. The simple, although revolutionary in his day, point he was making is that human labour is similar to capital, raw materials and so forth. Labour therefore ought to be subject to similar input/output analyses, measurement standards, and control.

Others refined and extended Babbage's insights. The most important extensions were provided by what is now called scientific management, and primarily by the man who has lent his name to that method, Frederick Winslow Taylor. Taylor is even today regarded by many as *the* prototypical American production genius, the man who provided Henry Ford and others with the technical tools responsible for the great increases in American industrial productivity.

In fact Taylor's technical contributions can now be seen as minimal. What he provided that so attracted managers, particularly those confronting strong trade union traditions, was an aggressive ideology for outmanoeuvring workers and their unions. For Taylor, the average industrial worker was altogether too cunning, too knowledgeable and, ironically, too skilled. What was needed were workers "with the constitution of an ox and the disposition of a trained gorilla".

Taylor's method was simple. Production, like the infant in the book of Solomon, was to be split in half. The first half was manual work which was to be made as routine and standard and repetitive as possible. The second half was mind work, which Taylor called planning, and which was to be done by managers and engineers. Production workers would no longer make any decisions about tools, about materials, about pace, or about anything else. They would simply and only carry out tasks defined for them by experts and supervised by managers. Here obviously were the beginnings of the manager as expert or, to put it slightly differently, here were the beginnings of management as a branch of engineering.

Recently, Mr C. de Benedetti, joint managing director of Olivetti, paid homage to Taylor and, by extention, to Babbage. He said:

> "The tendency to analyse the productive process in mechanistic and deterministic terms, and thus to express it in quantifiable, measurable and predictable terms, to give priority to digital quantitative analysis rather than to analogue analysis, is intrinsic to our very method of production. The connection between formalisation, rationalisation and industrial competition is not a casual but a fundamental link. The Taylorisation of the first factories developed as an answer to competition between companies. It is a digitalisation of the production process. At first, it enabled the labour process to be controlled and was the necessary prerequisite to the subsequent mechanisation and automation of the productive process. In this way, Taylorised industries were able to win competition over the putting out system."

In other words, standardisation has preceded, by necessity, automation and certainly computerisation.

The most well-known example of this process is what took place in that prototypical modern industry, the car manufacturing industry. While it is now fashionable to dismiss the car industry as mature and not worth serious attention any more, there are important lessons to be learned from its history.

Car making began as another industry which made new engineering products using very old production systems. We all know that this changed with Henry Ford. In fact, conversely to many people's belief, it did not change because of Ford's assembly line. And in any case, Henry Ford did not invent the assembly line. That honour goes to the meat packing houses of Chicago and Cleveland, although technically I suppose those were disassembly lines. A cow would come down the line, where there would be one person per function. One person would slit all the carcasses, another person would cut off the forequarters, and so on. These workers were paid at different rates depending on the perceived skill required for the task. This production line was at least 20 to 30 years earlier than the Henry Ford conveyor belt.

Ford's place in industrial history is, however, still secure. To him belongs the distinction of transforming the expectations and behaviour of industrial workers. The assembly line obviously had something to do with this, but so did the $5 a day wage, the Ford social work department (which I have to tell you was originally called the sociology department — which is something about which I have mixed feelings), and Ford's private army which protected the company from external threats such as trade unionists and government safety inspectors.

These collective carrots and sticks amounted to a deal which Ford imposed on its employees. Workers would receive substantially higher wages and benefits than those received by other industrial workers. The high real wages would be financed by the above-average profits which the industry enjoyed, partly as a result of the boom in cheap consumer goods (primarily the automobile) and partly because of the oligopolistic nature of the industry.

What Ford's workers gave in return marks the turning point in modern industrial relations. Workers in the United States formally accepted Taylorism as a legitimate form of production organisation. More than that, they accepted management's right to impose Taylorism without question.

Ford and his American successors were now able by right to measure, standardise, and time their workers — in short, to control every aspect of the production process. The 18th century artisan was replaced by the 20th century industrial operative.

I shall give an example. Everyone is familiar with time and motion studies which carefully measure work tasks and for which the car industry is famous. But you may not fully appreciate, because you are probably not in a mature industry, how pervasive such timings have become. I should like to read from one such schedule, in this case not having to do with work at all but having to do with breaks from work, in a UK car assembly plant. It reads as follows. "For fatigue, 1.3 minutes. Sitting down after standing too long, 65 seconds. Trips to the toilet, 1.62 minutes." In short, the acquiescence of US workers to the system of management transformed them for the first time, and quite literally, into human factors of production.

From the car and metal industries, Taylorism spread to virtually every American industrial workplace of any size. My favourite example, precisely because the proprietors have managed so successfuly to hide it, is that model of process control, McDonalds. For years McDonalds has been masquerading as a chain of self-service restaurants. McDonalds actually is an exquisitely designed, brilliantly engineered, ruthlessly efficient, self-contained hamburger factory, peopled by industrial operatives paid at the legal minimum wage. Raw materials are delivered to loading docks at the rear of the factory, from which they are crated and processed on an assembly line until the finished product is delivered to the final user. Of course delivery is made only after the appropriate data about price, volume and time have been entered on to what the customer thinks is a cash register, but what we know to be a remote entry terminal. Use of such terminals not only encourages employee honesty, but also provides immediate inventory control and information on employee productivity, work load fluctuations, and so on.

It is fascinating to study the history of the terminals. They started out as ordinary cash registers and then they became electronic cash registers. Then they became sophisticated electronic cash registers, where you could enter the price of the items sold together with the amount tendered. The machine then automatically calculated the cost, in certain states the sales taxes, and then the amount of change required, which was flashed up on the screen. These cash registers were soon replaced by terminals that had the names of the items instead of numbers on the buttons. These were eventually replaced by terminals that have pictures of the items on the buttons. At this point you can hire people who cannot even read or count, which is very important in the United States because it means that you can hire minimum cost labour. If one of Reagan's proposals, which is to create a dual-tiered minimum wage system incorporating the so-called youth wage, goes into effect then these hamburger factories will be operated by industrial operatives paid at below the legal minimum wage.

It is this pervasive and, more importantly, apparently legitimate structure of management control which has fascinated generations of Europeans. Remember that it was Vladimir Lenin who said, without irony, that "Taylorism plus electricity plus Soviets equals socialism."

For the sake of historical accuracy, we should mention that the Taylorism of Frederick Winslow Taylor also involved a deal between the workers and those who controlled their work. Taylor's deal was as follows. Workers would cede discretion over their work to experts, primarily engineers. The resultant productivity increases would be great, and here is how Taylor's deal differed from latter-day Taylorists, thus permitting owners to share the results of the increased productivity with their workers. Employee participation programmes, by whatever name, are therefore not new to American industry. American managers, however, have had a very poor record in keeping to their end of the deal when the desired productivity changes have occurred or when the market situation has worsened.

This is in contrast to the fabled Japanese management methods, which amount to no more than Taylorism with a vengeance. The Japanese managers do differ from their American counterparts, however, in that they are more likely to honour the Taylorist bargain in bad times as well as good. American managers, by contrast, have been relatively greedy or panicky, or both.

The Taylorisation of industrial manufacturing was the first battle in the struggle to control the modern workplace. But that battle has long since been won be managers and engineers. Today, the process of "the digitalisation of the production process" relies heavily on the computer. Although most of the talk now is of automating conventional production, for example, as Earl Joseph discussed, with robots, the bulk of capital investment in computerisation is probably directed into automating the office. The drive to automate the office has been accompanied by a public relations campaign intended to convince us that clerical productivity is low and that its improvement is to be found in raising the level of capital investment per clerical worker.

This is doubly misleading. First, those aspects of clerical labour such as typing which lent themselves to conventional Taylorist routinisation and rationalisation were long ago subjected to measurement and control. Secondly, the ultimate target of office automation is not clerks but managers.

Let me review these points in turn. Taylor's production technologies, such as typewriters equipped with keystroke counters, were developed very early and for years remained the chief form of engineering control. Ironically, these engineering control techniques do not appear to have been the major form of worker discipline, nor even to have been used widely in the early modern office. The social position of office workers, especially women typists, was so precarious that appeals to middle-class propriety were sufficient to keep most clerks docile. In any case, the unfavourable supply and demand situation of women typists was usually enough to keep them working hard and without complaint. So, although the technology of engineered coercion was well understood and available, the social and market situations of women clerical workers made use of this technology unnecessary and, as some observers claimed, counter-productive.

This is still true. Most offices in the United States are small, and nearly all of these still rely on electric typewriters and filing cabinets, not on word processors and certainly not on local area networks. Female labour in the United States — it gives me no pleasure to say it — is still cheaper than fancy electronics.

As ethereal products , as one of my co-speakers has labelled them, come to replace cars, integrated circuits, and Big Macs as the chief products of American industry, the old equation between clerical wages, productivity, and capital goods is being re-evaulated. You can probably sense that I am struggling here to avoid saying that the United States is becoming a service economy, because the term somehow suggests that Smith and Ricardo's laws work differently in those markets than in more traditional ones.

This is clearly not true, as my McDonalds' example illustrates. Whether we use the phrase service economy, ethereal work, or any other phrase, it is true that in the United States more and more people are producing different sorts of things than they used to. In a peculiar way, industrial history is repeating itself. The cost of this new work has revitalised the efforts of the old Taylorists. The term used by these latter-day scientific managers is 'human factors research'. The particular goal of human factors research has been to apply the principles and methods of Taylorism to work of the mind, to intellectual work.

At first this may seem strange. Taylorism, after all, developed from traditional production and traditional production engineering. Yet human factors research has wholeheartedly embraced the goals and the methods of scientific management. In fact, what distinguishes human factors theory from scientific management is not any major or minor departure from orthodoxy — on the contrary, human factors methods are distinguished by the unrelenting literalness with which it applies Taylorism.

Above all, human factors methods rest on one fundamental claim, which is that mental work no less than manual can be analysed, simplified, monitored, measured, and controlled. Human factors theory has carried Taylorism, quite literally, to its logical conclusions.

The ethereal workers most ripe for this sort of intellectual Taylorism are, appropriately enough, technical specialists. I am sure that you are familiar with efforts in this direction, such as computer aided design and computer aided manufacturing. Similar efforts, to develop structured design methods in programming, reflect management desires to control more closely the work of computer programmers and systems analysts.

Let me give you an example taken from the literature of human factors research in software development. This is from a recent article on the design of text editors. From this we learn that, with the user's hands already in place on the keyboard, it took 1.5 seconds to terminate a search-and-replace task, and 0.4 seconds to reposition the hands on the keyboard when finished, while the average TM (mental response time) was 1.35 seconds. Compare not just the times but the methods and the assumptions with the situation I mentioned earlier for the car assembly factory. There is no essential difference.

There is a certain logic in extending the principles of work standardisation to mind workers, but how does one explain similar attempts to apply human factors principles to managers? There is at least a general consensus about what engineers, technicians, and computer programmers do, but it is not clear that managers do anything.

I do quite a lot of workplace observation, primarily of systems analysts, programmers, and clerks. I feel reasonably confident that I can describe what most of these people do and what most of them produce. The only thing that I can say with any assurance about managers is that they seem to talk a lot. How this sort of activity is measured, analysed, standardised, and controlled is obviously an exciting technical challenge.

It would be appropriate now to find a quote about middle managers from Lenin. I could not, but I did find a quotation,nearly as good, once again from de Benedetti. He said,

> "We have witnessed a shift from hierarchical structure towards a polarisation, the elimination of intermediate groups and a centralisation of information and decisions prejudicial to middle executives. Data processing is the continuation of a story which began with the industrial revolution. Information technology is basically a technology of co-ordination and

control of a labour force of white collar workers which Taylorist organisation does not cover. In this sense, EDP is in fact an organisational technology and, like the organisation of labour itself, has a dual function as a productive force and a control tool for capital."

How exactly does this control tool for capital affect the future of middle managers? Human factors research has made one of its goals, if an implicit one, the ultimate elimination of human factors. These may be described as simply the engineering equivalent of variable costs (or just labour). The elimination of human unpredictability, in other words, ultimately is achieved through the elimination of human workers. If this is not yet possible or practical, as for example in computer programming, the alternative is to reduce the discretion workers have over their work. It is feared that repetitive or at least standardised tasks such as those found in data entry, assembly line operations, and so on, are taken as models. When combined with appropriately designed production technologies they permit the use of less-skilled workers.

Using less-skilled workers reduces labour costs, of course, but not just because unskilled workers are cheaper than skilled. The combination of smart machines and stupid machine-tenders also reduces the need for the services of the middle executives described by de Benedetti.

For example, when telephone operators handled all calls, the ratio of supervisors to operators was as high as 1 to 8. A supervisor would patrol, military fashion, a row of operators and observe their work, their pace, their demeanour and so on. Additional supervisor personnel would listen in on conversations between operators and callers.

As a result of the use of new electronic switching equipment, there are many fewer operators today. Therefore there are fewer supervisors. But there are fewer supervisors relative to the number of remaining operators because the electronic switching equipment also monitors the work pace of operators and informs the remaining supervisors when an operator is not making her quota, taking too long for a call, or just not working. The military patrolling is now done by the machines. The last time I bothered to check the current ratio of supervisors to operators it was somewhere around 1 to 30. This ratio is even less in more advanced switching stations.

In a different sort of industry, in retailing, managers used to keep track of goods sold, goods on hand, goods ordered and, on the basis of these records, calculate inventory needs, authorise purchases, arrange shipping, and so on. Now even a moderately simple software package does all of these things not for one product or even one product line, but for whole organisations. This level of manager, too, is becoming an endangered species.

Granted, by today's standards, these are relatively low-level management jobs. Let us look at something more complicated that affects higher level managers.

When managers are not talking, one of the things they are supposed to be doing is thinking. A typical sort of thinking involves the well known 'what if' exercises which people learn to do in fancy graduate schools of management. Now, thanks to clever software packages with names like VisiCalc and SuperCalc, which cost as much as $200 and run on machines which cost as much as $1,000, the same sort of 'what if' exercise can be performed by anyone who has a definite goal in life.

To give one more example, the CAD/CAM packages that I have seen recently make it absolutely unnecessary to have a manager, supervisor, design engineer or technician. The detailed supervision and ongoing comparison against design and cost specifications are built into the system and simply do not allow the designer or technician to violate them.

This is white collar factory work, but it is still factory work. The white collar factory of the future will resemble the blue collar factory of the future in that the functions of line managers will be built into the technology. In short, as managers oversee the introduction of idiot-proof systems, that is computer-based systems in which employee discretion has been reduced to a minimum, the need for their own services diminishes accordingly. Of course, there are differences in managing managers and managing, for example, machinists or assembly workers. Managers, precisely because their work is so ethereal, require more subtle forms of control. The necessary subtlety is provided by a related management method, generally referred to as human relations. In contrast to scientific management, human relations theory has stressed persuasion, education and careful testing and selection of employees to win employee acceptance of changes in work organisation and production systems.

The recent emergence of what are generally referred to as Japanese management methods, which is actually the re-emergence of industrial democracy, attests to the renewed interest in this approach.

Human relations theory departs from scientific management, and therefore also human factors methods, in its approach to human input. Although also concerned with control, human relations theory allows for some employee discretion as long as high productivity levels are maintained. Human relations theorists claim that treating employees as just another input will inevitably backfire, that is it will promote demoralisation, sabotage, and unionisation. To human relations researchers this suggests the need for more flexible and more subtle methods. The difference in the two approaches is not whether to establish control, but how.

In practice, carefully managed employee participation schemes have been the core of human relations efforts to secure employee co-operation. Currently the best-known approach of this kind in the United States is the so-called 'quality of working life' school. Socio-technical design, promoted by Mumford and others in the UK, is a variant developed specifically to address special problems created by office automation. In the latter case, human relations techniques are employed as part of a broader stategy to prevent unionisation or to circumvent the authority of unions already present.

Human relations methods, which were designed to soften the resistance of production and clerical workers to traditional scientific management as a prelude to changes in production technologies, have been used as a diversionary tactic in two very delicate situations. The first is managing technical specialists whose creativity is the core of their labour and whose work should not, and perhaps cannot, be subjected to intellectual Taylorism. The second is in the control of middle managers whose whole training and self-image resists the idea that their behaviour should be accountable.

It is clear that human relations, which goes under a variety of names such as theory Y/theory X, self-actualisation, and so on, is really not an alternative to scientific management at all, for managers or for anyone else. It is instead a supplementary technique for use on employees, such as managers or technical specialists, who are likely to get their backs up if subjected to the cruder forms of human factors manipulations. It is in essence a holding action, a diversionary tactic to cool down workers whose jobs are scheduled for Taylorisation or elimination.

Human factors techniques do in fact boomerang when they are used in manipulative and dishonest ways against managers or non-managers. There is considerable evidence for this. First, there are the increasingly vocal concerns expressed by US clerical workers about stress, health, and safety issues associated with office automation. Such concerns are now routinely cited as the causes of the rapid growth of clerical unions in the United States, even in workplaces where quality of working life and socio-technical design schemes have been implemented. In fact, clerical unions are not only the fastest growing unions in the United States, they are the only growing unions in the United States — this in an occupation hitherto populated by pliable, docile girls.

In software workplaces, similar opposition has been documented to structured techniques by software specialists. It is true, at least in the United States, that the resistance is muted at present, because up till now skilled software specialists have been able to express dissatisfaction with a given organisation by the simple expedient of finding other employment.

However, the history of all occupations, including those thought to be indispensable, clearly demonstrates that the supply eventually catches up with demand. When this happens the tension latent in trying to standardise the way people think will express itself more fully. When that happens the important questions will be whether technical workers, and even managers, respond to attempts to control them in the same way that office workers are now responding to similar efforts.

I began my talk by distinguishing between two kinds of lag. The first was the time between the appearance of a new technology and its widespread application. The second was the delay between the introduction of a new technology and the emergence of new forms of social organisation appropriate to it.

I should now like to add a third, and perhaps the most important, kind of lag. It is the delay between the emergence of new forms of social organisation and the development of new technology appropriate to them.

Modern scientific management, and thus human factors research, which are essentially the ideological offshoots of classical economics, accepted without question and without reflection two of classical economics' fundamental assumptions. First, that efficiency, productivity, and profitability are essentially interchangeable terms — they constitute an identity. Second, efficiency, and therefore also productivity and profitability, are achieved by substituting capital for labour, and unskilled labour for skilled, including, we may now add, intellectual labour.

These assumptions were easy enough to accept as long as bigger was demonstrably better, crude was cheap, and the Congo, Hong Kong and El Salvador were sources of cheap labour and markets for expensive finished goods. In these circumstances, one could regard as plausible the claims of the then General Motors' chairman, Charles Wilson, who said, with considerable feeling, "I have always believed that what was good for the USA was good for General Motors and vice versa."

That was in 1952, when gasoline was 5 cents a litre and the chief worry of the US car manufacturers was that Americans would run out of garage space to house their shiny new vehicles. The car industry, as I said, is now a mature industry and there are indeed lessons to be learned from it.

One of these lessons is that the old assumptions may not be valid. Perhaps in Charles Wilson's day, General Motors and even Chrysler were efficient and productive car manufacturers. They certainly were profitable. Today there is something obviously wrong with equating the economic health of an entire society with an industry that excels at making 2-ton private motor cars. The equation was indeed correct for Detroit, but it was not necessarily right for the rest of us. In other words, for a while Detroit's technology and its products were good for both GM and the USA. As long as that was true, perhaps there was some justification for dividing people into mind workers and trained gorillas. Today, however, we are beginning to realise that safe, reliable, economical transport from point A to point B is not necessarily the same thing as a private motor car. As a result we are also beginning to question whether it is necessary to create a society populated by alphas and betas — that is, people who work only with their hands and people who work only with their heads, instead of populating our societies with human beings.

But I can hear you say, "These were mature industries. Their mistakes will not be replicated by the high technology industries which will raise society to a new and more advanced state, where we can all achieve self-actualisation." I am not so sure. Car making, steel making and the textile industry were also the high technologies of their days. They, too, were greeted by contemporary pundits as the saviour of the ordinary working man and woman — simultaneously the source of liberation from drudgery and the source of unlimited material prosperity. It has not quite worked out that way.

Each new production technology has been absorbed into the existing social structure, reinforcing existing social divisions rather than transforming them. True enough, the material production of the last century has been staggering. But far from eliminating toil and drudgery, we have created more. If some people have been relieved of monotonous and tiring work, the chances are it is because they have been relieved of their jobs as well. If the job still exists, the people will probably have had to emigrate somewhere else, usually at a lower wage and in an area intolerant of uppity workers and their unions. People who used to have interesting and challenging work like draughtsmen, and even clerks, have seen their work routinised. Modern technology has not eliminated monotony and drudgery, only redistributed it.

This third and most important type of lag that I have described is thus not technological at all, it is political and its implications are profound.

For one thing, the old trick of throwing more resources at the problem will not work much longer. The days of the better black box or the smaller integrated circuit as a cure for deep social inequalities are just about over.

The chief source of new work in the United States is now clerical work. Clerical work is female work. In fact the only expansion in the US labour force for the last several years has been among women. Male employment rates have actually declined, along with the decline in manufacturing and agriculture. This has been true for a long time, but it did not matter because the decline was absorbed by blacks, who do not count. For years the unofficial attitude was that as long as the adult, white male unemployment rate in the United States did not go much above 5 or 6 per cent everything was all right. If the black unemployment rate went up, and usually it was exactly twice the white adult male unemployment rate, that was all right, we could handle it. The situation has now changed because the unemployment rates among white males for the last three or four years have started to climb.

The employment figures released by the US Government at the end of April 1982 indicated that both service and high-technology industries have for the first time also lost employment during the current slump. This is the first time that has ever happened — among the important people, the adult, white males.

The only consistent growth occupation since World War II, has been among janitors and dustmen. Even Japan, which has managed to hoodwink everyone with its mysterious and inscrutable management magic, has experienced precisely the same pattern, although their racism is of a different kind from ours. Industrial employment has declined by about 1.5 per cent a year since 1976. Even as productivity has increased, unfortunately so have the rates of industrial accidents and days lost to injury. If you think all of this is mitigated by life-time employment and accomplished by — to use that wonderful British expression — natural wastage, consider the following. The life-time employment guarantee in Japan works fine if your lifetime does not exceed 55 or 60 years, and if you happen to be one of the 30 per cent of the Japanese labour force employed by industries that offer life-time employment policies. There is less here than meets the eye.

David Noble, the American historian of technology at the Massachusetts Institute of Technology, points out that throughout the industrial revolution every new development in production technology or organisation was greeted as the harbinger of certain economic collapse and misery for the working population. There was, of course, economic collapse and misery often enough, but they could never be attributed to technological development, and usually things got better after they got worse. As a result, says Noble, raising the alarm about the dangers of new technology has come to be like crying wolf — after a while no one believes you and everyone stops listening.

This time, though, the wolf is at the door. It is real and it is particularly vicious. Technological unemployment is a reality now, in spite of the proliferation of ethereal products and in spite of the reassurances of futurists.

What should we do about it? I think we have to at least start by rejecting precisely the assumptions of Smith, Ricardo, Babbage, Taylor, and Lenin. We must understand that Charles Wilson's dictum that what is good for General Motors is good for the USA is no longer true, if it ever was. In just the same way that it is not true that profitable production of private motor cars is the same thing as efficient public transport, it is time to acknowledge that corporate profitability is no longer automatically indentical to public welfare.

We have an extraordinary technology at our disposal. It has been applied with breathtaking imagination to every aspect of our lives. But an entirely different set of social relations than we presently have is required if this technology is not to destroy us. The old economic categories and the technologies appropriate to them, those that preoccupied Taylor and Lenin, worked well for a long time — but their time is passing, and perhaps is already past.

We need new criteria of productivity and efficiency. These must incorporate a commitment to enlarge people-skills and capabilities, and not, as we do in our present system, to diminish them or eliminate them entirely.

If we treat this new technology as we always have, if we attempt to scare off the wolf of technological unemployment and global polarisation with the same old assumptions about displaced workers being absorbed by the new technologies and industries, if we do that, the wolf, instead of going away will get bigger and more vicious, because that is what created the beast in the first place.

IMPLEMENTATION OF INFORMATION TECHNOLOGY

Bengt Rosenberg, A/B SKF

Bengt Rosenberg is presently responsible for the development and co-ordination of methods and standards for systems development, telecommunications, hardware and basic software within the SKF Group. He is also in charge of co-ordination and implementation of office systems.

He joined SKF in 1964 and has held a number of management positions within SKF. His responsibilities have included organisational and administrative development, system development and administrative services.

He graduated in Mechanical Engineering at Chalmers University of Technology in Gothenburg.

Tomorrow, the 75th anniversary of SKF will be celebrated all over the world, but I suspect that most of you do not know much about SKF. I shall therefore introduce SKF and tell you something about our products and our organisation to provide a background for what I shall say later.

In 1907 Sven Wingquist, the founder of SKF, designed the first self-aligning double-row ball-bearing. Later, in 1919, this was developed into a new invention, the two-row self-aligning roller bearing, which permits angular displacement relative to its housing. These are the products on which SKF mainly bases its business.

I will give you a short review of some of SKF's other products and some of the other activities that are going on in SKF. We manufacture mainly high-quality steel. Out of what we produce, about 50 per cent is used internally and 50 per cent goes to the open market. We manufacture the main parts of our machines ourselves.

We are now developing our next generation of steel factories. The objective is to have the factories running 24 hours a day, but manned only eight hours a day. That is in order to make the best use of the investment while also taking account of the fact that workers do not like to work at night.

The main product within SKF, however, is bearings. We make many different kinds of bearing. There are about 10 main types of bearing with more than 20,000 variants of these products.

Part of our philosophy is to build products into other products. For example, in supplying axle boxes for freight wagons, we build all the components into one unit and deliver the complete part.

Now we are adopting the same approach in the car industry. We are manufacturing a complete unit for the wheel. This was first introduced in normal production on the Fiat Panda car, and I think we will see that approach on a lot of cars in the future.

Our business is split between our products as follows — rolling bearings account for about 70 per cent, special steel products about 15 per cent, cutting tools 4 per cent, and other products about 11 per cent.

Our sales in 1981 were 13,570 million Swedish crowns, compared with 12,572 million Swedish crowns in 1980. The number of employees at SKF has decreased by about 3,000 people from 1980 to 1981. During the last decade, we have had a productivity increase of between 8 per cent and 10 per cent per year. This means that from the beginning of the 1970s, we have reduced the number of people by about 20,000. I will come back to this later, but it is completely due to the competition from the Japanese.

We have an objective today that the profits of SKF should be the inflation rate plus 3 per cent. This is a survival objective. If we cannot meet this objective we will not survive in the long term. We have not achieved it yet, but we will have to achieve it one way or the other. This means that we have to reduce the current number of blue collar and white collar employees in future in order to survive.

We currently have 185 companies and 74 factories. Our market is split as follows — about 60 per cent of our sales are in Europe, 22 per cent in North America, 8 per cent in Latin America, and 10 per cent in Asia, Africa and Australia.

Our organisation structure is very much product-oriented. We have four main divisions — the European Bearing Division, the Overseas Bearing Division, the Steel Division, and the Cutting Tool Division. We also have Group Staff units such as Marketing, Manufacturing, Engineering, Product Engineering and Research, Public Affairs, Legal, Personnel and Organisation, and Finance and Information Systems.

I belong to Information Systems, where there are six people who are responsible for the co-ordination of information system functions all over the world. I am mainly concerned with hardware and software methods and standards and, to some extent, the office of the future. One man is responsible for the co-ordination of information systems in manufacturing (mainly within the European Bearing Division). One is responsible for the co-ordination between overseas companies, which involves both manufacturing systems and sales systems. Another is responsible for this aspect of material flow outside manufacturing.

The total number of information systems people within the group is slightly over 1,000. Of these 1,000, 75 per cent are in Europe. Installations range from very small ones with only two or three information systems people to much larger installations, the largest of which has 130 information systems staff.

In 1921, SKF introduced the punched card and used the first punched card machines. In the 1940s, we went over from 40-column punched cards to 80-column punched cards. In 1958, we implemented one of the first computers, RAMAC 305. In 1982, we still have some routines implemented on punched cards, on Olivetti RP60 machines which are 20 years old. SKF feels this is a little old-fashioned now and it is going to change. Today we are completely IBM-oriented regarding hardware, and we have plug compatible peripherals from a number of other vendors.

Now that I have given a short overview of the organisation, I will discuss our information systems in more detail. In 1971, about 25 people met at an hotel in southern Sweden. We held a relaxed discussion on how our information systems may look in the future, with common systems, common developments, and so on. I will follow three lines of development from this point in time.

The first line is the MASCOS project. This project was intended for our sales companies all over the world — MASCOS stands for Modular Automated Sales Company Systems. The project started in Madrid, where the local company provided the financial support for the first generation of the system to be developed, which took from 1972 to the middle of 1975. Once the first generation of the system was implemented, the project was reorganised and located in Brussels, where work was started on the next generation of systems. The first generation was based on an IBM370/135, but the second generation was based on IBM System 3.

There were a number of subsequent development steps and, in 1979, the whole system was converted to the IBM System 34.

So there are three generations of the same system, incorporating improvements and adjustments to the new hardware and basic software available at the time. Today the same system is implemented in 27 companies on 25 computers.

The role of this system covers three areas — marketing, material flow, and finance. It includes the normal functions. Material flow includes customer orders, back orders, supplier orders, invoicing, and inventory management. Finance includes accounts receivable, general ledger, inventory accounting, and inventory and sales statistics. Marketing includes sales statistics, budgeting, sales support, and GHQ reporting (the sending of information to group headquarters).

The objectives of the system were to reduce costs, to improve effectiveness, to give better service, to reduce stock, to increase control, to provide quick information, and to provide better communication. These are very general objectives, about the same as you meet in most such projects. But they are still important.

There is one central MASCOS group which has project management responsibility for development, maintenance and implementation support of all the systems around the world. The implementation of a new system is carried out to a very fixed schedule that covers about one year. This schedule is a detailed plan that all companies implementing the system have to follow precisely. Nearly all of them have been successful in doing so.

The project today has a steering committee. The members of this committee are from some of the more important users of the system. They decide what the budget for the project should be and what improvements to the system should be carried out.

I think most of the companies involved are grateful for these systems. There are many reasons why it has been such a successful project. It was of a reasonable size and had a good project manager, which was particularly important in the early stage of development. It had reasonable top management support. It was also very business oriented, which was something that the company needed. The result has been an extensive saving in administrative costs. The number of administrative people has been greatly reduced in some of these companies. One of the important objectives was to move people over from administration and material flow activities to marketing and sales activities. This has been successful.

The system has been implemented by all kinds of companies. You must realise that some of these companies have had no previous experience of data processing at all. Some of them had not even seen a computer before. So the system is implemented all over the world in very different kinds of environment, and still it works very well.

The maintenance and improvement costs are very reasonable. Each company pays about £10,000 per year for maintenance of the system. Improvements to the system are handled in a slightly different way. It should be mentioned that we are now planning for the fourth generation of this system.

A factor that we have found particularly important is that managers who move from one company to another — and there is a lot of mobility of managers between the sales companies — will meet the same administrative system wherever they go. When a manager moves to a new company he does not have to bother about developing a system — the one he knows is already implemented there. He may have to request a report that is new to his new company but that he was used to receiving previously. Otherwise he will meet the same administrative environment and get the same support in any of the companies.

On the other hand, it is not reasonable that every manager coming to a sales company should be permitted to change the administrative systems. It is just a waste of time in many cases. The value of coordinated use of the same system is very important.

These systems are long-range investments. Managers must be patient and not give up before the results appear. In the beginning, there were not many companies that actually supported the implementation of such a joint system, but still this approach received top management support. That was the reason why it was possible to carry out the project. Without such top management support, the project would have been stopped. The top management believed in the idea and could see its long-range consequences.

This way of concentrating resources on a few tasks gives a very good result, even if it does make cost/benefit analysis difficult. At the beginning of the project, cost/benefit analysis was nearly impossible. Now we can see the motorway effect of the project, with a lot of companies wanting to join in and use this system, even non-sales companies.

But to be successful is not mainly a technical problem, it is a management problem. The technical problems can be solved in one way or another, but if the management does not support it then it will never be successful.

The second line of development to have started from that meeting in 1971 is the Information Systems Board. This Board was formed to provide co-ordination of information systems within Europe. When it was formed in 1971 it consisted mostly of the finance managers of the companies who formed the Board. These were the people responsible for the development and operation of information systems in the companies. Under them were those who were directly responsible for information systems activities — the information systems managers. These managers formed another group called Information Systems Managers Meeting. Since 1972, both these Boards have met regularly, three to four times a year. They have been very important for the co-ordination of information systems activities in Europe.

Today the situation is such that the support within each organisation for information systems development and operation is so important that I think one of the top managers in each company should be completely and solely responsible for information systems development in that company. Companies should take the same approach as society in general — they should regard information systems as a part of the information society, but in this case the information society is the company. It is impossible for a manager in a big company to share responsibility between finance and information systems. It is a full-time job to co-ordinate the information systems and to be the link between the company's business activities and the technique, development, and operation of the information systems.

Of course, in 1971, we tried to do what most companies tried to do around that time. We tried to build a complete system that would deal with everything. A project was started in Paris, with a contribution from all the European Bearing Division companies. The project was named SKF Total Integrated Computer System for Customer Service System. This project lasted for 18 months, but then it was stopped.

One can say that it failed for a number of reasons. First, people at that time were not clearly motivated to carry out common projects. Second, the hardware and the basic software were not the same in all the companies — each company had about 10 or 15 years of systems development behind it and could not quickly convert to something that was common.

Third, new projects arose that were more important for the future. So this project was stopped, but it was an important introduction for us to the development of common or joint systems.

The next project was called GFSS (Global Forecasting and Supply System).

As I said earlier, in the 1960s competition from the Japanese in the steel industry increased tremendously. In 1967, SKF for the first time introduced a group organisation and the group headquarters in Gothenburg was formed. Another important factor in that context was the introduction of English as the group language. One problem that we had during the 1960s with the Customer Service System was that people in separate companies found it difficult to communicate with each other.

One consequence of this group organisation was that, in 1972, it was decided to restructure the manufacturing of products in Europe. In 1969, all companies were responsible for manufacturing their own products for their own markets, which meant that each company manufactured between 40,000 and 50,000 products for their own markets. It was decided that the number of variants should be reduced to about 20,000 and that the manufacturing should be distributed among the five factories in such a way that each product was manufactured in only one factory. The project was planned to start in 1973 and end in 1978. Actually it ended according to plan, with the result that the number of products was reduced to about half.

This also required completely new information systems and so represented a new era from the information systems viewpoint. We formed a central database to hold all the data (other than the technical data for manufacturing) for all SKF's products. This data included details of where the products were manufactured, and how they should be transferred from one factory to another.

A special co-ordination centre was created in Brussels, named the Forecasting and Supply Office. This centre was given world-wide responsibility for sales forecasting, stock control, manufacturing control, and delivery control. This meant that all information regarding stocks, sales, manufacturing, and market forecasts had to be collected and co-ordinated at this centre. From an information systems point of view, it was important to support the companies all the time.

The first step regarding communications was that we leased lines to all the companies involved. But soon we found that the lines were not reliable enough given the high volumes of information to be transferred between the companies. Monthly transfers of data could last for five or six hours between each company. This caused us to reconsider the structure of the entire group.

When a company first grows to a very large size it becomes very difficult to handle. The usual response to this difficulty is to break the company down into a number of smaller, autonomous units that go their own way. It is then found, however, that the smaller units have lost the advantages of scale and weight. So the third step is to create a central administrative co-ordination body in order to re-create the advantage of scale but still retain the flexibility advantages of having a number of smaller business units.

This third step is the philosophy we have at SKF. It is a particularly important philosophy for information systems development. The essential point of the philosophy is that we maintain flexibility but still have everyone working in the same direction. In parenthesis, perhaps it could be said that the difference between such a company and many American companies is that in the latter when you lose profit you are fired, whereas in our kind of company you are fired only for not following the centrally drawn up guidelines. SKF is acting as a group of companies, and that is the main objective.

This philosophy changed our ability to co-ordinate the activities within the group, because all functions within SKF are acting in the same way.

As I was saying earlier, from a technical point of view we found that leased lines were not adequate, and so in 1976 we started to specify the requirements for a more advanced network. We worked on these requirements for more than a year. Then we took another six months to look at what the vendors had to offer. In October 1977, it was decided that we should use IBM's SNA network.

This network was implemented in two steps — the first in 1979 and the second in 1980. The first step consisted of dividing the lines into two channels so that we could have interactive traffic on one channel, and batch or file transmission on the other channel. The second step enabled us to mix completely interactive traffic with file transmission. It also enabled all terminals to reach all facilities within the network — every terminal connected to the network can reach applications in any of the computer centres.

We have now defined six functions for our network, the first four of which are available to end users now. The first function is file transfer, which is the transfer of information from one computer centre to another. This is handled automatically — we have functions that are more advanced than standard SNA functions. Our system automatically queues file transfers and then makes the transfers to the right computer centres and sends a message to the receiver to inform him that the transfer has taken place.

The second function is job access. This consists of remote job entry, with any RJE terminal able to reach any application in any of the computer centres. For example, we have programs running in one computer centre that are used by all RJE stations in the group.

The third function is a message transfer system. It is a very simple system that enables messages to be entered at any of the 3270 terminals and sent, via internal distribution, to any person or printer in the system. This system is much used today.

The fourth function is transaction access, which consists of the normal on-line transactions access to applications.

The two functions for the future are remote-program access and down-stream access. The object of remote-program access is to enable one program in one computer centre to access another program in another computer centre automatically. For example, if you have a query program that picks up some data from the database in the computer centre to which you are attached, this program can automatically access another program in another computer centre in order to pick up additional data to complete the screen of information you need. This function is not used yet to any great extent, but we understand that it will be an important function in the future.

The down-stream access function has to do with the philosophy of hardware and software. I do not think it is reasonable to introduce distributed computer systems if you cannot run them more or less without any operators or technical people. So the requirement in our system is to be able to load programs, find errors in them, start a batch operation, and so on, all from the distributed computers. The normal situation in many installations is that you run interactively during the day time and then you have one person who stays after the others to run the batch system. This requires the computer centre to be manned 16 or 24 hours a day, and means that it is not reasonable to distribute computers to any extent because what you gain in flexibility is off-set by increased costs.

I will take this network as an example of the importance of information systems people themselves having complete responsibility for the system. It is pointless asking the users whether they think they need this sort of advanced system. Most users cannot see and cannot imagine how such a system could be used in the future. I think that the information systems people must take full responsibility for all technical questions, which means that they have to set policies and standards, decide what software should be available, decide what communications should be available, and so on. Such questions are their internal responsibility, and they have almost nothing to do with applications.

The important point is for the information systems people to be broad-minded and to have foresight. The key is for them to be able to see what the users will need in the future and what the technology can provide.

Based on the network and the requirement for faster order handling between companies, in 1979 we started a project called International Customer Service System. In each company there is an International Marketing department which takes the products from manufacturing and distributes them abroad. Each company also has a Domestic Marketing department which receives products distributed by the other companies' International Marketing departments. This system uses the network to fulfil the requirements of communication between these two types of organisational unit. The intention is to have five databases, one for each company, and to arrange for those databases to be accessible from all the companies. Each company should be able to access each database in order to send orders, enquire about order status, receive order acknowledgements, and so on. The system is shown in figure 1.

Figure 1

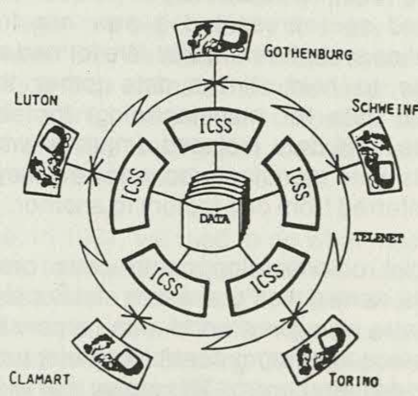

This project was initiated with a feasibility study in 1977, was designed in 1978, and will be implemented during 1982 and the beginning of 1983. It is a giant project, costing in total about £5 million. Members of the project team are taken from all five companies, and we have added a number of consultants to the team. The system involves 500,000 lines of code. It will take about 200 man years (60 man years from the central development team, and 140 man years from the five companies).

This project has been a little more difficult to carry out than the others, because each of the five companies is very strong. Each has its own ideas about how a project should be carried out and what systems it wants to develop and implement. But I think the attitude has changed now. We can see that the companies are working together much better and they understand one another better. The implementation of the ICSS system has improved the technique and knowledge level in the companies.

Now I shall discuss some points that I think should be considered when developing common systems. The first point is that the system must relate to business operations. In the ICSS system there is a steering committee which consists of the material flow managers from each company. This provides deep involvement from the user side.

The second point is that such systems should have an umbrella concept, which means that they should not deal with small detailed areas but cover some complete area.

Third, such a system must be functionally co-ordinated, which means that there must at some level in the group be co-ordination of all the companies from the manufacturing point of view, from the finance point of view, from the material flow point of view, and so on.

Fourth, the life of a system should be long, between seven to ten years.

Finally, strict investment thinking is required in the systems development area.

The experiences we have gained from investing in common systems are as follows. First, the lead time for such systems is often longer than that for local systems development. This is as a result of the added difficulty in co-ordinating and organising the project from the beginning, collecting the right people, creating a steering committee, and so on. Also, there is a lot of discussion between the companies because they have to understand the system and provide their input.

Second, there is much better exploitation of experience and expertise in developing common systems.

Third, such systems require modularisation and step-by-step implementation.

Fourth, such systems require the setting up of permanent support and maintenance units.

Finally, the common systems approach results in low cost for development, implementation, and maintenance per company.

We have also learned that there are three main areas that it is important to co-ordinate — the hardware and software, the databases, the communication, and the applications. Each must be developed in parallel otherwise the complete development does not work.

Let us look at the cost of information systems in SKF. The cost for our European installations has undergone an average increase of 14.5 per cent per year over the last 10 years. I think that top management, at least during the last 10 years, have thought that this increase is too high.

The distribution of these costs is 50 per cent personnel, 31 per cent hardware, 5 per cent software, 4 per cent communication, and 11 per cent other costs. Since 1979 the share of cost accounted for by software has increased from 2 per cent to 5 per cent, and it is still increasing rapidly. Personnel costs have been stable at about 50 per cent during the last 10-year period.

In 1980 and 1981 there was a lot of talk about the office of the future, and in particular how the investment per office worker compared unfavourably with that for factory workers and agricultural workers. I have taken the information system costs for one of our companies and divided it by the number of office employees. Our annual cost per office employee is made up as follows (the figures are given in Swedish crowns) — 18,300kr for personnel, 11,300kr for hardware, 1,900kr for software, and 5,100kr for other costs. This gives a total figure of about 36,000kr per employee.

Another factor is the investment in the development of systems, which works out at about 15,000kr per office employee.

If you accumulate these costs over a 10-year period, you will find that our systems development investment is about 98,000kr per office employee. If we forget about inflation, replacement costs, and so on, and add together the software, hardware, and applications costs, we have invested about 130,000kr per office employee. If you think that the office of the future is in the future, you are completely wrong. Office automation has been around for 20 years. I very much dislike the way consultants and vendors ignore the investments of the last 20 years.

Now let us look at the future. Figure 2 (overleaf) shows the cost of my work station, which is a 3270 terminal. The figures shown include a number of items that I do not actually have at the moment but that I feel I need. The total investment is 67,000kr per employee, and I wonder how many organisations can afford that much.

Figure 2

THE COST OF AN ADVANCED WORKSTATION

		PER MONTH	INVEST-MENT
Hardware Costs:			
* Colour Display Terminal		1000	36000
* Simple Matrix Printer		250	8000
* Share of (8 users):			
- Control Unit		500	14000
- Colour Printer or Plotter		200	4500
- High Quality Printer		200	4500
SUBTOTAL	SEK	2150	67000
	GBP	215	6700
Operation Costs on Central Computer:			
* Text Processing			
* Graphics			
* Decision Support Systems			
SUBTOTAL	SEK	3000	
	GBP	300	
TOTAL	SEK	5150	67000
	GBP	515	6700

Next, I shall show you some figures taken from official statistics for the private sector of Sweden. These statistics, shown in figure 3, are concerned with salaried employees within the private sector of Sweden. I have extracted the figures for three areas that I think have been most affected by computerisation — manufacturing planning, order processing, and finance control. You can see that the number of people has decreased by about 1.7 per cent per year, which is a very slow decrease.

Figure 3

NUMBER OF OFFICE EMPLOYEES WITHIN
- MANUFACTURING PLANNING
- ORDER PROCESSING
- FINANCE CONTROL
AS % OF TOTAL NUMBER OF OFFICE EMPLOYEES

TOTAL NUMBER OF OFFICE EMPLOYEES
1975: 320.000
1981: 319.000

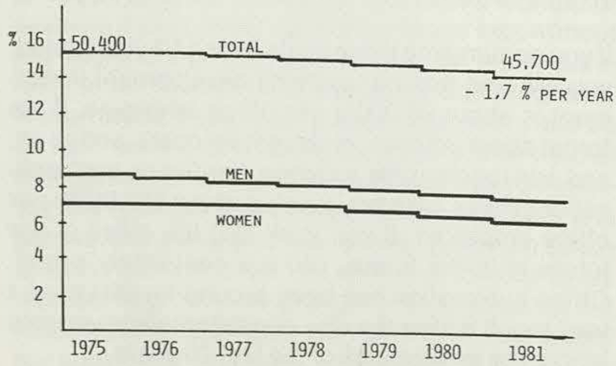

Figure 4 shows the trends for various position levels. Position levels 3 to 5 are the most qualified workers, and 6 to 8 the less qualified workers. You can see that there has been a decrease of 3.5 per cent for levels 6 to 8, the less qualified workers. However, the number of people doing the more qualified work has increased by 1.7 per cent. This structural

Figure 4

NUMBER OF OFFICE EMPLOYEES WITHIN
- MANUFACTURING PLANNING
- ORDER PROCESSING
- FINANCE CONTROL
AS % OF TOTAL NUMBER OF OFFICE EMPLOYEES

TOTAL NUMBER OF OFFICE EMPLOYEES
1975: 320.000
1981: 319.000

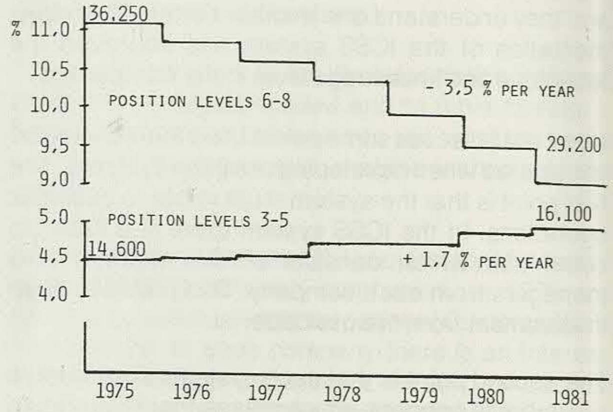

change to the work force is due to the introduction of computer systems.

I shall now look at how many of these office employees are information systems people. During the period 1971 to 1981 the number of people working in systems development has increased from 5,000 to 6,900, which is about 5 per cent per year. The number of people working in systems operation has increased from 2,100 to 3,100. The number of data entry people has decreased from 3,800 to 3,000. But in total, the number of data processing people as a proportion of the total workforce has remained nearly constant; it increased from 3.1 per cent to 4.1 per cent — a very slow change.

There has been little change also in the proportion of those people who are women, except in the systems operation area where the proportion of women has increased from 11 per cent to 25 per cent.

In contrast to what some of the other speakers have said, I think that changes in the work force are happening very slowly and are completely dependent on factors other than the introduction of information systems. Each company is more or less self-contained in this respect, with increases or decreases in the various sectors of the work force depending upon the total environment of that company, what requirements they have, on their profitability, how successful they are, and so on.

During the 1980s I think that office system products will continue to be developed and improved at a rapid rate. Their implementation, however, will take place at a slower rate, and their introduction will not significantly reduce the number of people employed in offices.

LIMITS TO GROWTH

Louis Pouzin, INRIA

Louis Pouzin is Director of Pilot Projects at INRIA. He has managed many large software projects including CTSS as part of project MAC at MIT and Meteor, a real-time operating system for the French weather bureau. He joined INRIA in 1972 as Director of Cyclades, an experimental computer network linking universities and research centres in France.

In my talk I shall not try to be too subtle. I shall draw attention to a number of major obstacles to the development and application of technology. I am not talking about the development of technology in laboratories — raw technology is developing so fast now that we are overwhelmed by it. I am talking about obstacles to the application of this technology. There are seven obstacles, which I call the seven curses.

The first is the human interface. This subject has been abandoned for years in the data processing and communication world. We assume that people should be able to use computer systems but we forget that the concepts used in computer systems are completely different from those with which humans are used to dealing. Computer concepts are abstract and rigid, and they evolve from one into another by means of well-defined processes. Human concepts are totally uncontrollable. We can try to make models of them, but usually we manage to approximate to something that is convenient but does not accurately represent the human process.

I will give some examples to illustrate what I have in mind when I say that we should try to improve the human interface. Suppose you have to attend a meeting of important people and provide a report about the activities in your division. You probably will come up with a number of papers, such as activity reports, containing tables of figures in rows and columns. Having been given these papers, someone in the assembly will put a question such as, "What proportion of the money is going to XYZ group?" XYZ group may be one of the affiliate companies, or the parent company, or perhaps just a company on whose Board the questioner sits. Another such question might be, "How much did we spend on that contract from the beginning?" Such questions are not silly questions. They are natural questions, but it is very difficult to anticipate all of them.

So typically what we do when we come to a meeting of that sort is to have a collection of papers that are produced on a regular basis, and then in our minds we carry a number of typical figures that are not too hard to remember and do not change very often or very quickly. For questions requiring other information we are without answers. That does not mean that we do not have the answers anywhere. We probably have them in our files in our desk, but it would take perhaps 15 minutes to get them.

The way we would come up with the answers if we were given the time would be to leaf through our files, knowing which files to search for the figures, write down the numbers on a piece of paper, make some simple calculations, perhaps with the help of a hand calculator, and within say 15 minutes we would have the answers. But that is unacceptable in a Board of Directors' meeting. It takes too long and no one would accept it. So for some questions it looks as though we have no answers and we look silly.

Perhaps we should have terminals in such meetings so that we could put questions to the computer and get the answers immediately. Let us look at the kind of tools that are available to do this. Typically, we would come across some sort of system that uses commands and statements such as shown in figure 1.

You can open the manual for any of these kinds of system and you will find such statements. Vendors

Figure 1

```
#  DIS ROWS 1 TO 3 OF SCALES

#  MAKE TABLE END_ARMA
   FROM ROWS 17 TO 21 OF ARMA

#  TYPE MEAN OF COL 7 OF ARMA
   WHERE COL 4 = SHAPE
```

call that natural language! I do not wish to say that it is not natural. Chinese is natural for the Chinese, but not for me. The trouble with using the term 'natural language' is that it does not define for whom the language is natural. A language may certainly be natural for the people who use it all the time, and you can probably become fluent in the language if you use it every day for six months. The trouble is that for the sort of questions we have been discussing the language would be used for only about half an hour every third month, and so we cannot be fluent in that language. To us, it is not natural language.

Today's systems, typically produce output such as the table at the top of figure 2. Such information is certainly correct, but nobody can understand it. What people do understand is information such as the graphs shown at the bottom of figure 2. Such graphs may not be accurate, but they do not need to be — they only need to show trends. They need to contrast events and make proportions visible, possibly using colour. That is what people like — that is what they understand.

Fig. 2:

	75	76	77	78	79	80	
KALI	5	73	322	680	712	530	2322
TORA		18	410	615	440	510	1993
SITAR		5	47	205	380	540	1174
	5	93	779	1500	1532	1580	

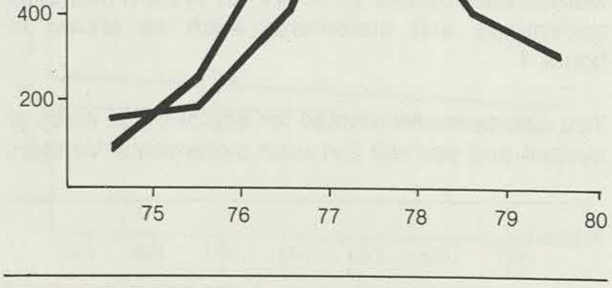

Could we not find a way to get computers to do just that? We know how to do it manually, except that the process is too slow at the time we need it. Let us try to imagine the kind of tool that we need to produce the sort of information that people can understand.

Suppose instead that we have a terminal that has a display, a microphone, and some kind of pointing device so that you can indicate the material you want. Suppose that the computer can understand your voice giving simple instructions such as "Next page", "Next year", "Stop", "Go further", and so on — say a vocabulary of about a hundred words. We know that there exists on the market a very cheap system today that can understand a few hundred words. The system might need some training to recognise individual voices, but if you can give the system a couple of hours notice that you are going to be using it then the training can be carried out.

Such a system would be able to present you with your files, and without having to bring several pounds of paper to the meeting you could leaf through them by saying "Turn the page". We are asking the computer to carry out a basic, slave job, understanding a few sentences and carrying out the simple instructions correctly. For example, you may want to pick out particular figures from your files. You would point with the pointing device and say "Keep it, because I need it later". That way you could build a list of figures and then instruct the system to total all the figures, or carry out a subtraction, or calculate the average, or whatever you like.

The technology makes this sort of system feasible. But building such a system is not a very glamorous job. To do such a job one clearly has to be very good at understanding the human interface, but in terms of computing science the job is very trivial. I suppose that is the reason why not many people get interested in doing that job.

In terms of handling the human interface, I think we are still at the point we were 20 years ago when Fortran was invented. Fortran was simply a way to call subroutines and to organise the order in which you call them. When you look at the so-called natural languages of today, you can see that they are nothing more than a mapping of subroutine calls. There are little differences in syntax. You do not have to remember commas or parentheses, but you do have to remember key words. We have changed the decor, the way languages look, but in fact they are effectively just the same as Fortran. I think we could improve on that.

Sometimes we are deluded by the apparent ease with which we already do things. Most people carry a diary in their pocket. An example of the sort of information that is put into such diaries is shown in figure 3. A lot of assumptions are made in using such a notebook. For example, you know who Robert is, and so you do not have to write in the full name and address. You know that '25F' means travel expenses that cost 25 francs. You know that 'dentist' means that you have an appointment with the dentist.

You put in just the information you need, referring to meetings, travelling, and personal matters. Even though we may be a businessman during the day and a family member during the evening, we are the same person and so it is convenient to keep all this information using a single tool.

Fig. 3:

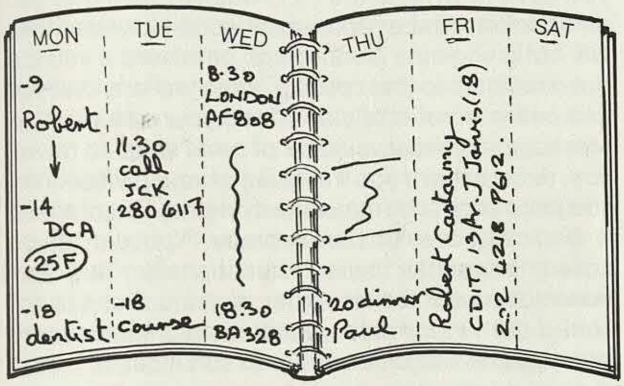

The diary is extremely easy to fill in and to use, but automating it on a computer would be extremely hard. Do you know of any computer system that fits in your pocket, that uses no electricity, that you can read wherever you are, and is also a file? I also use my diary to carry pieces of paper in.

I know of no computer system that could do all that. I think the pocket diary will remain for years as a very competitive tool that is very hard to beat, even though it does contain a lot of information that perhaps could be handled by a computer.

The pocket notebook is a typical example of something that is hard to couple with a computing system, even though it may look easy to do so. Some people think that in the future we will use less and less paper. There is talk of the office of the future being a paperless office. That is not my belief. My belief is that the future very much involves paper.

The best display is the A4 piece of paper. It is a display you can fold and put in your pocket, and it requires no plug into the wall. Since paper is so convenient to use, I believe that any system has to provide for paper input and for paper output. I do not mean that we have to use paper for storing information. Paper-based stores are very bulky and make information difficult to access. Nor do I mean that we should carry paper over long distances, which uses up a lot of energy. It is simply that paper is an excellent medium for people to read, and so any system that has a human orientation has to be able to handle paper.

One approach that specialists use to try to get computers to help people is known as artificial intelligence. At the moment there is a surge of interest in artificial intelligence, due perhaps to the announcement by the Japanese of fifth generation computers incorporating expert systems. These expert systems model the human expert. They have the ability to process knowledge, by which I mean an accumulation of facts, rules, or relations. They do not however have any understanding of what they are doing — the most difficult thing of all to model is common sense.

Marvin Minsky, one of the high priests of artificial intelligence, once said, "You can tell a computer that a bird flies. How about a dead bird? You can tell a computer that a dead bird doesn't fly. But how about a toy bird? You can tell a computer that a toy bird doesn't fly. How about a bird on which you put 20 kilos? A bird on which you put 20 kilos doesn't fly. All that is obvious for people, but you cannot imagine the thousands, or even millions, of rules needed to explain to a computer in which condition a bird can fly." So common sense is very difficult to model.

I believe that expert systems are useful in medicine, in financial investment, in any application that requires a lot of knowledge but not much understanding of it.

Another approach to using computers that is now emerging is to try and have the computer carry out the chores that have been performed up to now by people, but leave matters of initiative to people. A good example of this approach is VisiCalc. When you develop a budget, you allocate money to various activities, to salaries, to investment, and so on, and then you have to add a percentage for inflation and taxes. Then you usually want to change it. All this amounts to a lot of work and involves a lot of calculation. VisiCalc saves you having to do all this work and calculation with a pencil and eraser. It is effectively an automated pencil and eraser plus a hand calculator, and it works very well. It can be used to handle any calculation or manipulation involving figures, and is particularly useful where one is working on relationships between figures. VisiCalc has been imitated by so many manufacturers and software houses that you can now find that sort of software on almost any minicomputer or microcomputer.

This approach does not try to model the mind at all. It is not pretentious. It is simply helping people to do quickly what otherwise they would have to spend a long time doing manually. So perhaps that is what we need instead of — or maybe not instead of, but in addition to — expert systems. Most people do not handle very many facts — they spend their time handling quite ordinary chores.

So, my first curse, the problem of the human interface, is a major handicap that we will have to face up to during the next 10 or 15 years.

The next curse is data communications. Today we are told by specialists, software houses, and suppliers that there is no problem any more with data communication. You can have anything you like — you can have leased circuits, telephone circuits, packet networks, satellites, cables, or whatever. What you have to remember, however, is that when you want to access a computer you usually have to use the telephone system.

Every industrialised country has a telephone system, and most of these systems are 50 to 80 years old. Attempts are made to upgrade these systems, new exchanges are installed, and so on, but mostly you have to use very old electromechanical exchanges that do not work very well. That creates a lot of non-trivial difficulties. I will give some examples which I experience almost every day.

Suppose you dial a number on a data network. Typically what you would get is the busy signal, or it would ring and ring with no answer. Or you might get a recording that says, "The number you dialled has been disconnected. Please check the phone directory." Of course it is not true that the number has been disconnected. Either the system has got mixed up, or perhaps it is designed to tell people they have dialled the wrong number when the system gets congested, so that by the time they have checked the number the system has had time to recover.

Sometimes when you dial, instead of getting the normal welcome banner on your terminal, you get all kinds of garbage, because somehow a spurious character has got on to the line and the computer thinks that you are a different kind of terminal. Or perhaps you get your connection but then as soon as you get it you are disconnected and you have to try again.

According to published data network characteristics, opening a connection takes 200 milliseconds at the most. In practice, as a result of various occurrences such as outlined above, it actually takes anything from 2 to 5 minutes. Having got a connection, you have available to you 30-character-per-second transmission which is being gradually upgraded — in some places you can now have 1,200 bits per second transmission.

Of course you can have leased access to the transmission facility, but that is only for people who want to make connections from within their company. If you want to use the facility from you home, or to use it only occasionally, you cannot afford to have a leased connection to the data network.

Suppose that you have the connection and it is working. You might believe that the trouble is over. But it is not. Suppose you are in France and you want to receive mail over the data network from Switzerland, or from Germany, or from the UK. Because there is no reverse charging mechanism, you are prevented from getting that mail. International calls in data networks have to be placed by people who have a subscription. In order to have a subscription, you have to write to the PTT, who will then give you an account number and a password. However, they will not give you a number and password if you are not a resident in that country, which means that you can call internationally only from your own country. Message systems would be of great value to travellers, except that if you travel out of your own country you just cannot use message systems. This problem is basically one of bureaucracy. You can place collect telephone calls internationally. That has been sorted out for a number of years. But it is not sorted out yet for data communication, and it may take another five or ten years to sort it out.

Suppose you want to log-in to a system. You could well get 'LLOGOIGNIN' as a result of your characters being echoed. If you always had just one echo that would be no problem, but in fact you sometimes get no echo and sometimes two echoes — one echo from the local teleconcentrator and another echo from the host. When you have two echoes you have to find a way to turn off one echo. Of course, this is feasible — the problem is in knowing when you need to do it.

Finally, assume you have succeeded in having only one echo. The system says "Password". I cannot use 'echo' as a password because, if you ask for the echo, that is precisely what you get. The data network does not understand that you are typing in a password — it ignores the semantics of the conversation you are having with the computer. So if you do not want your password to be shown, you have to turn the echo back off again, on again, and so on. Then all of a sudden you get a message, "Please log in," which means "You've been disconnected".

Or you may get the message "address not found". Typically every week or every two weeks when I dial a computer in the United States through Transpac, Tymnet, or Telenet, I get that message. I know perfectly well that the address exists. This is a typical ploy in computing systems — when they get into trouble they send back silly diagnostic messages. We tend to believe those messages. We should never believe a diagnostic message. They make no sense, other than to tell us that the system does not work.

We cannot expect any improvement in the difficulties with communications. New software is always being installed, but in fact software that has been upgraded is usually worse than the old software, because it has new bugs. I think there is no possibility of improvement over the next 10 years because we will have more and more data networks, more and more customers attached to them, and more and more procedures available, which means more new software. These systems will become stable only when they become obsolete, 10 or 15 years from now.

My next point about the data communications curse concerns the term 'public data network'. I do not know where the word 'public' comes in, because I have never seen a public data terminal anywhere. It would be quite easy to install a public terminal and very useful in many places. For example, when you check out of a big hotel, you often have to queue up for 20 minutes. I would much prefer to learn how to use a terminal so that, instead of queueing up, I could check out through a terminal. Perhaps I would have to put a credit card or passport into it, but I would accept that. To me, avoiding having to queue up for 20 minutes would make it worth learning to use a terminal.

I think this situation will worsen in the future. The salaries of people such as hotel receptionists are increasing, and so organisations employ fewer and fewer such people. Also, the people are getting more and more stupid. As a result, the customer has to wait longer and longer. Ten or 15 years ago, when you went to a Hertz or Avis rent-a-car desk then immediately several girls would appear, all smiling and saying, "What would you like, sir?" Now you find one harassed girl who says, "Please wait", and you have to wait 15 minutes or half an hour to get a car.

So I should like to see public terminals available for simple tasks like renting a car, checking out of a hotel, booking a flight, and receiving messages.

The third curse is technical assistance. Whenever we use systems that are not simple, we need help. Even the telephone system is not always that trivial to use. If you want to make a special telephone call to somewhere out of the country, then you can use a public booth on the street. The telephone system is public in the sense that you can find public terminals almost anywhere — perhaps not in the woods, but certainly in the street, in restaurants, in hotels, and in airports. So the telephone system is public, and it provides a single number that you can dial to get help.

Have you ever tried to get assistance on a public data network? One way you can try to get assistance is to type "Help". The chances are that the system will reply, "Help not found, please try again," or "Address not found," or some other kind of nonsense. If you know the phone number and it is during business hours, between 9 a.m. on Monday and 5 p.m. on Friday, then perhaps you could get assistance. But if you are using a number of data networks for international calls, and you tell them that you are having problems, they will say such things as "We talked to Telenet, but they do nothing." If you ask someone from Telenet, he says, "Well, that's Transpac of course. We know that." If you talk to Transpac, they say, "It's Tymnet," or "It's Datapac." It is always the other network's fault, and the result is you get no help.

This is not new. It has always been the case with the telephone system. International networks are not supervised, they are just pieced together. There never has been any international help. The system has relied on the fact that the telephone is sufficiently simple that if you get a bad connection you can try again and perhaps get a good one. If you really cannot get a good connection, you call an operator and say, "Listen. Can you hear that? Give me a better connection." In a sense the system sorts itself out, because it is obvious enough for people to understand why it does not work. But in data networks it is very difficult technically to find out what is not working. In data networks, faults can be very hard to track down. The systems are not built in order to help in tracking down faults. They are built with boundaries around them like separate countries. Each network is well guarded and supervised internally, but there is nobody in charge of the whole thing. This will continue to be the case for the next 10 or 15 years, because that is the way the world is organised. Each country has an authority — a common carrier, a PTT, or whatever — which is more or less a monopoly in some way related to the government of the country. There is currently no future for an international body in charge of data communications. So we will have to continue to live in the kind of world in which no technical assistance is provided.

This leads one to imagine a new kind of business. In the world of travelling, if you want to go on vacation in some place where it is nice and sunny, but you know nobody there and you do not know anything about the hotels or the prices, then you go to see a travel agent who will talk to the various airlines, the hotels, and so on, and make the necessary arrangements to give you what you want. To give another example, if you want to send potatoes to the Russians, but you do not know Russian and you know nobody in Russia, again you will use an agent to handle the transportation, including making arrangements with customs, the railways, and so on.

What we need in the data communications world is a similar kind of agent. We need an organisation with a number of characteristics. First, it would have offices in many countries, whether the local offices are wholly owned subsidiaries or affiliates. Second, they would be experts in tariffs. Tariffs are a complete maze. They are described in very thick books, and nobody can understand them. However, it is possible to use tariffs to your advantage, as long as you know what you are doing. So the agent will help customers in properly using tariffs.

Third, such agents could offer packaged services. They would buy communications from various common carriers and PTTs, perhaps also buying the computing time that is needed to convert files or produce documents, and provide whatever is required as a complete package. Such an agent would send you one bill in your own currency, so that you do not have the problem of having to handle bills from many places and in different currencies.

Fourth, they could also try to aggregate traffic. Many communications systems offer discount rates if you transport more than certain volumes of traffic. These discounts can be quite substantial, in some cases up to 60 per cent. These agents could perhaps buy wholesale and then sell so that they make their revenues on these discounts, rather than charging customers directly for the service.

I think that this kind of business could be extremely useful in the data communications world, but it has yet to be created.

My fourth curse is standards. Standards are, like the flag and motherhood, something everyone has to be in favour of. They are something with which nobody can disagree, except that in practice they are difficult to achieve. Standards tend to change over time, they evolve with technology, and they are defined by various countries and PTTs.

We tend to forget that even at the most basic level of data communication we do not have all the standards we need. If you try to call a computer in the United States from your own acoustic coupler terminal in Europe, it will not work, because the Bell standards in the United States and Canada are different from the CCITT standards. They do not use the same frequencies for modems. Yet this is a situation that I think could be avoided. Anyone who looks carefully at a present-day terminal will find a lot of very tiny switches, perhaps hidden behind a small plate. These switches provide the terminal with a lot of options. For example, the terminal can work at 300, 600, or 1,200 bits per second, it can work half duplex or full duplex, it can provide carriage return plus line feed, and so on. The terminal manufacturers have known for years that the environment in which the terminal will be used is extremely diverse, so they have anticipated that situation and put in such option switches. However, nobody has put in the Bell/CCITT switch yet. All that is required is one more switch, to enable the terminal to use different frequencies. I am not sure why this has not been done. Perhaps there has been some opposition from the PTTs, or perhaps nobody really wants traffic to develop over the Atlantic. But clearly such an option is technically feasible. Modems in a way are an easy target for standards. Modems are something that everybody can understand — either they work together or they do not.

In the area of higher-level standards, perhaps the only standard that is well accepted throughout the world, and that works almost everywhere, is V.24. This standard provides a way to hook two machines together. I believe that one reason why this standard has been accepted and is used almost everywhere is that it does almost nothing. It is simply a means of plugging wires together.

Now we have higher level protocols such as X.25, which is much more complex than V.24. We now have X.25 versions 1, 2, and 3, and revised versions 1, 2, and 3. Every three of four years X.25 is revised, with new options. You might believe that finally we will get a version that applies everywhere. This is not so. X.25 is rather like an expensive restaurant menu. If you go to the Tour d'Argent restaurant in Paris, you find a long menu, but nobody is supposed to take the whole list. Now we have Transpac X.25, Datapac X.25, an IBM X.25, and every one is different. Perhaps the different versions could be made compatible. They have options and there are ways of tuning them. But typically to make two versions of X.25 work together requires several months of work for people to understand one another, to decide what options to use, to plan tests, and then to find the end product that works.

So, at the moment, we have new standards, but they are so rich, they have so many options, that they are standards only in name.

In the past we tended to use the word 'standards' to describe the means of getting things to work together. Now the word means nothing more than that the 'standard' has a rubber stamp from a standards organisation — it is not a means of getting different pieces of equipment to work together.

Even though we can criticise the state of the art of standardisation in data communication, it is still a paradise compared with the computing world. Suppose you take a number of manuals from different manufacturers and look up the command languages — the basic languages that users are offered once they have finally succeeded in logging in.

In order to delete a file you will find words such as delete, kill, cancel, remove, yank, erase, flush. All the manufacturers use different words. You might decide that you can accept that, on the basis that it is like learning different foreign languages. The problem is that when two manufacturers use the same word it does not mean the same thing. For example, if you type "delete" in order to delete a file, you may find that it deletes your process and logs you out. This makes the languages impossible to learn, because as soon as you move to a different system you not only have to forget what you knew but you also have to control your instinct.

In order to print out a file, a range of words are used, such as list, print, type, display, show, read, write. I even discovered in the very same system that to get information printed out on a typewriter, the command 'read' was used in one context while the command 'write' was used in another. Technicians use these words depending on their mood. There is absolutely no standard in such matters.

So the computing world is still a virgin in terms of standardisation. Attempts have been made for the past three years within ISO to devise a standard to support a general approach to building systems, but the point has not yet been reached of defining standards for user interfaces. Current work is aimed at trying to standardise concepts, not the way the concepts are made visible to the users.

The fifth curse, or obstacle, is reliability. If you take a subway train, or place a telephone call, or use a washing machine, you expect them to work. But if you use a computer you feel pretty lucky when it works, because often it does not. It is reasonably accepted everywhere that computers cannot be relied on to work.

We have known for 15 years how to build systems that are reliable. Real-time systems are reliable — they are built to be reliable. Such systems have duplicate databases, with duplicate access to the databases from duplicate hosts, duplicate access from the hosts to a communications system, and the terminals have duplicate access to the communications system. That is what reliable systems are, and they are probably available 99.98 per cent of the time.

Most of the systems that are available today through communications systems, however, have a single host, a single file system, a single access to a single network, with a single access to each terminal. If any of these elements goes down, the service is not available.

If you talk to the person who is running the network, he will say that the network's up time is over 99 per cent. If you talk to the person running the system, he will say that the system's up time is over 98 per cent. But the whole thing put together achieves no better than 80 per cent reliability. At times, the system may technically be working, but it is also congested. You may get a message saying, "All ports are busy. Please try later." This means there is no service. Or you find that when you try to dial the network you cannot. You may get a message such as, "Our offices are closed. Please call on Monday morning". This means there is no service.

So we cannot consider these systems reliable. They are totally unacceptable for ordinary life. They work much below accepted standards.

It is not that we do not know how to make reliable systems, it is just that we do not do it. Whether we should put the blame on to manufacturers, on to communications carriers, or on to users I am not sure. Society as a whole seems to accept that computers should not be reliable. Perhaps we do not want computers to be reliable. Perhaps we would feel threatened if they were too reliable, too safe, and too correct. Perhaps there is resistance to pay the price necessary to achieve reliability.

The sixth curse is regulations. This again is a frequent target for criticism. Anything that involves communications services is more or less controlled by state regulations, whether they are enforced through a state monopoly or a judicial body.

We should not be upset about regulations in the sense that they are simply another constraint on business. We are constrained by regulations, by standards, by laws, by the unions and, by financial restraints. Businesses are constrained in all sorts of ways and regulations are just another constraint.

There are however some important differences in what is meant by 'regulations'. For example, in the United States, regulations are rules, made by a federal body called the Federal Communication Commission, that apply to the carriers. So the companies that offer communications systems are regulated by the federal government. They cannot introduce whatever service they like at any price they like. They have to submit an application to supply a well-defined service at a well-defined price. In Europe, regulations are rules made by the PTTs that apply to citizens. These regulations have nothing to do with the regulations found in the US. In any European country, the regulations that apply to data communications, or to any kind of communications, usually amount to the fact that citizens are forbidden to do almost anything, while the PTTs can do what they like, with no responsibility and no liability of any sort.

It is a little worrying, in the sense that the industrial world has tended to harmonise itself over the years, that in the area of data communications regulations these two worlds are totally incompatible and have a totally different philosophy.

In Europe we have been trained to live with the regulations made by the various state authorities, but at the moment we are in a situation where nothing is clear. Typically the mail system, the telephone system, telex, telegrams, and so on, were all a state monopoly. Then came data communications and for a time it was not too clear what was allowed and what was not. For example, it took some years for the PTTs to decide whether they would allow a service bureau to share leased circuits between several customers. They finally agreed that it was acceptable as long as a surcharge was paid.

Now we are evolving rapidly into a world where we have all kinds of new services such as videotex, electronic mail, digital voice, and so on. The boundary between such services and computing services is fading away. So the question is, if there is to be a monopoly, what services should it cover.

Every country at the moment is worrying about this issue. Many countries are saying that the monopoly will be limited to some services, and that for the rest of the services there will be no monopoly. But I find that difficult to believe. In the UK, at the time when there was still a state monopoly, the Post Office (now known as British Telecom) embarked on building a set of big computing systems intended for a videotex service called Prestel. The Post Office originally thought they would have de facto monopoly in this area because they were able to invest so much money so quickly that no other organisation would be able to compete. But then they were demonopolised and a number of people complained about the Prestel situation. So now British Telecom is changing its mind and saying that it will still offer videotex services but it will no longer prevent anyone else from doing so, and it is willing to switch calls to a private videotex host. In Germany and France the situation is a little different.

So every country now is trying to find a way to decide about monopolising or regulating these new services. What is worrying is that this will take a long time, because it is not a technical issue. It is a matter of people talking to one another, trying to decide through commissions and committees what they should do. Typically that sort of process takes years.

Once agreement has been reached in every separate country, it will be discovered that each country has taken so many steps that are different from those taken in the other countries that the whole business will have to be harmonised at a European level, which will take another five or ten years.

The real problem for these new services is that there is such a lack of long-term stability in terms of regulations that it is very difficult for investors to put money into development of the services. Investors are afraid that their money will be lost as a result of the suppliers not being allowed to develop their business the way they would like or to offer the services that from a commercial point of view it is rational to offer. For example, it would be rational for a service bureau that already offers message systems to upgrade into voice message systems. But it is currently impossible to see whether voice message systems will be treated as being different from telephony or not. So there are some nasty problems and it is difficult to make decisions. In Europe a lot of companies are therefore just staying where they are, waiting to see what the PTT is going to do, and what business, if any, may be left for them to take advantage of.

The final curse is employment. Anyone who has leafed through sociological magazines will have seen charts predicting the changes in employment over the next 10, 15, or 20 years. Most predictions show that the personnel assigned to production will constantly decrease as a percentage of the workforce, because there will be increased automation in factories. The number of secretaries will perhaps slightly increase, because automation will create new businesses and so there will be more offices and more people working in offices. The proportion of people assigned to management will not change very much, because managers can increase their effectiveness and scope by using new tools and, anyway, there is a limit in any company to the number of managers there can be without creating promotion difficulties. There will, however, be more marketing and sales people, because it is generally believed that business will expand worldwide.

My first point about such predictions is that they are only concerned with people who are employed — the total always adds up to 100 per cent. The unemployed never show up in the figures and so the figures have absolutely no social meaning. Perhaps they have some financial meaning, but they certainly mean nothing in terms of people and activities.

We should be worried about this situation, because the office is the last ghetto in which we have been able to continue being inefficient successfully. The people on farms were shifted out to factories when we introduced machinery on the land. When we automated factories, we shifted people into offices. Now we are introducing office automation. Some people, like the previous speaker, claim that office automation has been used for 20 years, but there really are new developments taking place now as a result of there being so many vendors and so many technological pressures to put tools into offices. We should be wondering where we are going to shift the office workers to. There is no place for them to go. Should they go back to the farms, or should they be turned into guides who show crowds around the office of the future? I do not know the answer, but we will have to find some way of using these people.

If you go to an exhibition on office systems you can see examples of the futuristic 21st century office work station, with a display, a microphone, a magic tablet, a mouse, a keyboard, all integrated into it and all interconnected via local networks to other office work stations. I do not wish to suggest that such systems will not work, I am sure they will. And I really believe that they are useful. What I question is whether these kinds of system are really adapted to ordinary people in ordinary offices today.

In today's office, we find a complex of many activities: coffee, telex, meetings, social conversations, and so on. It is not at all like the factory environment where people work on production lines. If you visit a car factory, it is hell. I do not understand how people can live in it. But if you visit an office it is very friendly and warm. People chat for hours because they have nothing else to do. They can have a party in the office, which is much more convenient than having it at home. There is a boss somewhere, but he is so afraid of talking to the employees that he stays in his office. I do not think that office workers will accept a drastic change in their life, because their life is very pleasant. I think the office is something that will be extremely hard to change.

In a way, technology could be used as the excuse for changing the office. It could be the excuse for management to say, "Modern technology means we must change." Of course most people would agree with that, but then they would ask, "What should we do about modern technology, if it is not adapted to us?" So technology will probably be the excuse for both sides to argue about what should be done. Technology will cause delay.

When we try to plot employment over the next 10 to 15 years, we should expect a slight increase in population. Most countries, especially if they are reasonably industrialised, have a population that increases slightly. We can certainly expect the proportion of production people to diminish, as predicted. And we can realistically expect the proportion of marketing and sales people to increase. But I wonder if we really can expect the proportion of secretaries to increase. What I feel is that it will be as difficult to get a good secretary 10 years from now as it is to get domestic help. If you live in Africa, it is quite easy to get domestic help — you can even have a whole family to help in your house, and it is cheap. But if you try to get help in France, in Germany, or in the UK, first, you cannot find it, and secondly, if you do find it, it is very expensive. In such countries nobody wants to help in the home any more, nobody wants to be a servant any more — people prefer to be unemployed. In the future I think that many of the women who might be secretaries will prefer to be unemployed. The only way to overcome this problem is to do as we do at home today. We have tools to clean the ceilings and the floors, to wash the dishes, and so on, because nobody else will do it. In the same way, people will have to become their own secretaries, supported by new tools.

So my guess is that the chart predicting employment for office workers will not even be filled up. The secretary part of the percentage will decrease because nobody will be able or willing to do the job. The difference will be made up by unemployed people. So unemployment among office workers will increase for two reasons. First, you cannot turn production workers and factory workers into office workers quickly, especially if you have introduced technology into the office. They may need so much retraining that they will just give up completely. Second, secretaries and office clerks will give up too. We may have something like 20 per cent unemployment, but perhaps a good half of it will be voluntary unemployment. People will refuse to take jobs that are available.

No society in the past has lived for a long time with such a high rate of unemployment. Perhaps it is feasible in financial terms, but there is real trouble in cultural terms. We have been trained to think that working is good and not working is bad, that being unemployed is shameful. Such attitudes are difficult to change. We may change them over 40 years, but certainly not over the next 10 years.

So we should expect a lot of trouble in social matters as a result of our inability to retrain people or change our cultural values quickly enough to keep up with technology.

CONFERENCE CONCLUSION

David Butler, Butler Cox & Partners Limited

I believe that this conference has really been trying to tackle a very important problem using a proven methodology. The proven methodology is to identify one's problems, to find a frame of reference for those problems, then to look for tools to provide a solution, and finally to look for solutions. Edward de Bono stressed early in the conference the importance of the frame of reference, the importance of the way in which one looks at the problem, and I suppose that we have been struggling for some insight on the scale of Galileo, Newton, Einstein, or Genghis Khan.

If we look at the problems which have surfaced during the conference, and some which we have referred to only en passant, we certainly have the problem of finite resources — a fixed amount of money, of labour, of equipment. We have hinted from time to time that perhaps we simply have too many people, although radical solutions to that problem are frequently in our minds.

We certainly also have the problem of technology absorption; in other words how fast we can take the technology which is coming out of the laboratories and put it into a useful application. We have problems of economic imbalance, one being the north/south imbalance which we have discussed quite a lot in this conference. But we also have imbalances within our own society — the comparatively wealthy worker and the comparatively poverty-stricken non-worker.

We have political schisms, again on a global scale, east versus west, and also within our own societies on the bases of class, sex, age and race. Those are some fairly daunting problems which we have identified.

So far as the tools to solve them are concerned, we have certainly talked a great deal about information technology, and I have no doubt that it has a role to play. I made a note of some of the other techniques that we discussed: genetic engineering; molecular biology, linked in some ways with information technology; nuclear engineering was discussed; I think we touched on ocean engineering, and if we did not we should have done. I should add that these tools are in no special order.

This particular conference made me wonder what would be the Butler Cox Po Foundation. You remember that Edward de Bono said that the Po aeroplane lands upside down and the Po motor car has square wheels. I am wondering, looking at the Foundation and the wishes of its members, what would the Po Foundation look like.

Apart from the fact that the annual subscription would be zero, which I will pass over very quickly, what would it be like? I wondered whether it would be trying to build more links between the different sciences and technologies we have heard about. If you look at them, they all have a part to play. All the problems are connected in some way, yet with the possible tools which we are trying to bring to bear on them we follow our own blinkered line. Presumably people from other disciplines are doing the same. I wonder whether in some ways we ought to think about trying to broaden our vision on such matters. Maybe one day we should be daring enough to have a conference to which we would invite speakers from all these different disciplines to tell us how they see their science or technology creating the possibility of solutions to some of the problems.

If we did hold a conference like that, we would have to get everybody to agree in advance to expect from it something rather different from our normal conferences. In terms of what you, the members, usually get from our conferences, it might turn out to be a total disaster and it would be a risky thing to do. But maybe it is something we ought to consider.

We now have an international advisory board of members in existence, so that if you have any ideas about what the Po Foundation would look like, you have a mechanism for voicing them and we would be delighted to hear from you.

I should like to thank all our speakers who have come this week and presented us with very stimulating and rewarding presentations. I should like to thank the translation team who I think have done an excellent job. Sometimes when speakers are carried away by enthusiasm it is not the easiest thing in the world for the translators to keep up. We all understand that, and I think they have done a great job for us.

I should like to thank you, the delegates, for attending this conference. I hope that you have benefited from it and enjoyed it. I hope that you will have a safe journey back to your homes. Thank you very much.

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

Belgium & The Netherlands
SA Butler Cox NV
Avenue Louise – 479 – Louizalaan,
Bte – 47 – Bus,
Bruxelles 1050 Brussel
☎ (02) 647 15 53, Telex 61963 BUTCOX

France
La Fondation Butler Cox
Tour Akzo, 164 Rue Ambroise Croizat,
93204 St Denis-Cedex 1, France
☎ (1) 820.61.64, Telex 610789 ASFRA

United States of America
Butler Cox & Partners Limited
216 Cooper Center, Pennsauken, New Jersey 08109, USA
☎ (609) 665 3210

Switzerland and Germany
Butler Cox & Partners Limited
Morley House, 26-30 Holborn Viaduct, London EC1A 2BP
☎ (London) 583 9381, (Zurich) 302 0848

Italy
Sisdoconsult
20123 Milano – Via Caradosso 7 – Italy
☎ 86.53.55 / 87.62.27, Telex 311250 PPF MI

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