Transcript

Management Conference

New York
May 14 & 15 1979

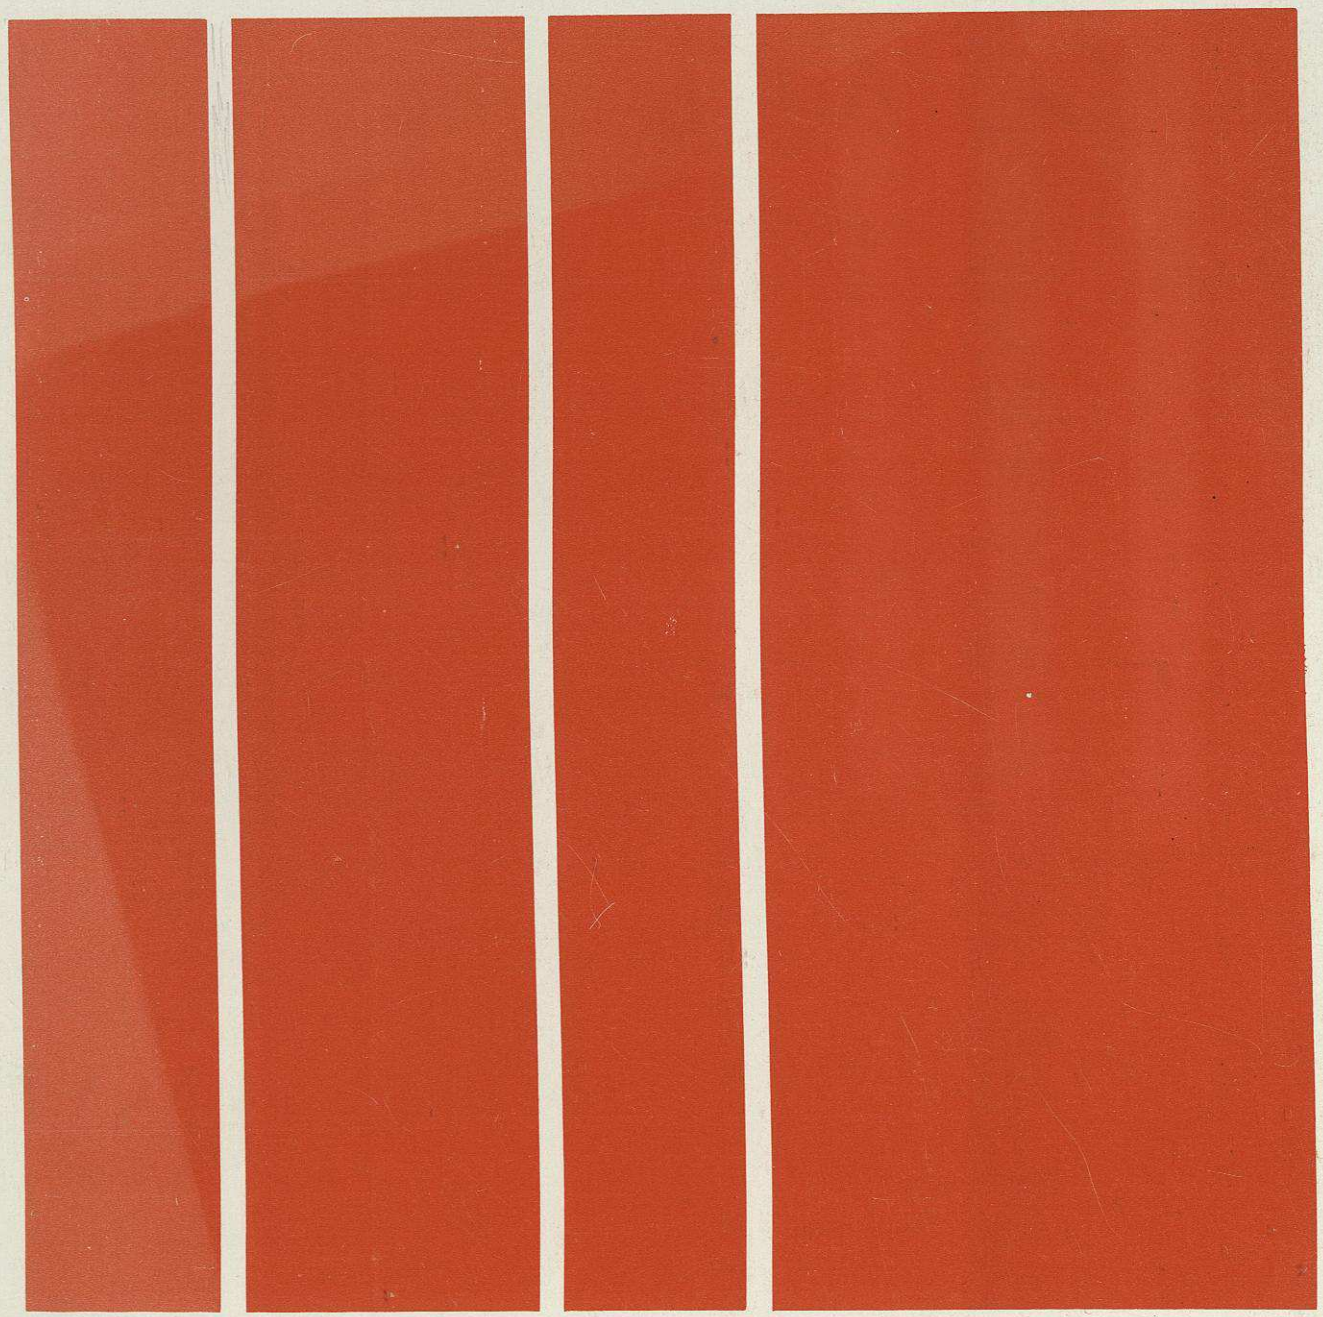

The Butler Cox Foundation

Butler Cox Foundation

MANAGEMENT CONFERENCE TRANSCRIPT

New York, May 14 and 15, 1979

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

CONTENTS

LIST OF DELEGATES

AKZO	Mr A Zijl
BICC	Mr P M Burman
BL CARS	Mr M D P Colin Mr P C Marchant
BOC	Mr A Bargioni Mr K A Grant Ms K Paterson Mr J Reed Mr R H Tester
BRITISH AIRWAYS	Mr A L Jacobs
BURMAH OIL	Mr P Groves Mr G Kaul Mr D McDermott Mr C Trafford
CADBURY SCHWEPPES	Mr S M Ross
CALOR GROUP	Mr C R Gilbert
COURAGE	Dr C E Montagnon
DEBENHAMS	Mr R Hodge Mr I Stoneman
DISTILLERS	Dr A Stevenson
ENGLISH CHINA CLAYS	Mr M H Keat Mr H A Patrick
GRAND METROPOLITAN GROUP	Mr L F Green
IBM	Mr A Head Mr H Mitchell
ITT BUSINESS SYSTEMS	Mr P I Rafield
METAL BOX	Mr T Allen Mr P J Bennett Mr J Dutton
PLESSEY	Mr D A T Rayfield
RECKITT & COLMAN	Mr R E Ferris Mr J Fields Mr G Gregory

ROWNTREE MACKINTOSH	Mr D J Fearnley Mr R M Pontefract
SAMUEL MONTAGU	Mr R W Chandler
SHELL	Mr T A Tobin
TURNER & NEWALL	Mr G Cook Mr J Hill
UNILEVER	Mr T Baker Mr D G Logan Mr B G Maudsley
WEIR GROUP	Mr R G Codd
Speakers	Mr H B Becker Mr D Butler Ms R Goldfield Mr J A Gosden Mr J B Holden Mr J M Ireland Mr J L Jones Mr R B McClelland Mr A Mallia Mr C H Reynolds Mr S Rowe Mr G G Specker Mr N Sullivan
Butler Cox & Partners Staff and Associates	Mr R Bellini (SisdoConsult) Mr G E Cox Mr D Flint Mr E Goldblum Mr R Koffler Mr K Kozarsky Mr M P Ray Dr H C Zedlitz (Akzo Systems)

CONFERENCE OPENING

David Butler
Chairman, Butler Cox & Partners Limited

BUTLER: Gentlemen, good morning. May I welcome you to this, the only tropical rain forest in North America.

May I particularly welcome those of you who have not attended a Conference of the Butler Cox Foundation before. The list of companies is, I think, too long to be read out.

However, I would like to mention one company represented here for the first time — Akzo — represented by two people with different portfolios, which I will explain. First Tom Zijl, seated on the right, and next to him Chris Zedlitz. Many of you may remember that Chris was a speaker at the Stratford Conference.

The presence of Akzo here is important for two reasons. First, because they are our second member from continental Europe and our first member of the Foundation from the Netherlands, and we welcome them on that score.

Secondly, as I have already discussed with one or two of you, Butler Cox and Partners has reached an agreement with a newly established subsidiary of Akzo which is to be known as Akzo Systems B.V., which will be based in the Netherlands but will also work in Belgium, Germany, Switzerland and Austria and which will represent the Foundation in those countries.

Chris Zedlitz will have the job of looking after the members, the member organisations who joined the Foundation in those countries. In the future we intend to make absolutely certain that we integrate him and his team into the management of the Foundation in an effective way, for the benefit of members in those countries and, perhaps most important of all, from the point of view of the existing members. Plans are well advanced for Chris and his team to make an injection of expertise and knowledge from those countries into the report projects which we carry out. As time goes on the reports will become increasingly international in scope and increasingly representative of an international approach to the problems.

I am sure that during the course of the next couple of days both of our friends from Akzo will have a chance to get to know you and talk to you, and to absorb the atmosphere which I think is very distinctive about the Foundation.

There is one change in the Agenda as published, and it is not for me to say whether it is a change for the better or for the worse. Session A — THE KEY FACTORS LIKELY TO AFFECT THE GROWTH OF THE AUTOMATED OFFICE IN EUROPE — we decided could more conveniently be given by me than Tony Gunton, so that is what is going to happen; I am going to give Session A. Those of you who are members of the Gunton Appreciation Society — or 'GAS' — will have to wait for another occasion to express your approval.

It now gives me great pleasure to introduce the first speaker of the morning, David Butler.

THE KEY FACTORS LIKELY TO AFFECT THE GROWTH OF THE AUTOMATED OFFICE IN EUROPE

David Butler
Chairman, Butler Cox & Partners Limited

David Butler is Chairman of Butler Cox & Partners Limited, and is responsible for all the Company's research-based services including the Butler Cox Foundation.

He has written numerous articles on computing and allied topics for technical management and national publications including The Times, The Guardian, Management Today; and has lectured for the British Computer Society, the Foundation for Business Responsibilities, the Institute of Directors and the Industrial Society. In 1977 he was awarded the first prize in the National Computing Centre/Computing essay competition for a paper on computers and communications. In 1978 he won joint first prize in a competition arranged by Dataskil and Computer Weekly for a paper on the 'Information Society'. He is the author of 'The Convergence of Technologies', published by the Butler Cox Foundation.

BUTLER: What I want to do in this session is to try to paint a broad picture of the requirements for office automation and to try to identify some of the factors which will determine the rate of growth of office automation in Europe.

The words "in Europe" are included in the title with great deliberation. During the course of the visits which follow this conference we are going to see some examples of advanced office applications in different American companies, but we will all of us always have to bear in mind the very different environment which prevails in Europe and the very different rules which are likely to be applied in the development of office automation in Europe. If one looks only at the difference, for example, in the telecommunications regulatory situation, then that in itself is a major difference between the environment here and the environment back home.

First, I want to introduce a model of a typical, theoretical or conceptual company and to try to define what role within that company office automation might play; second, to try to identify some of the lessons which have been learned by our experience to date in the area of office systems; third, to look at some of the external forces which will influence the speed at which we can advance towards automated office systems. In particular, there are three which I think are worthy of consideration: — the role of the PTTs; — the role of the competitive market for the supply of information systems and products; — the never-ending quest for more sensible standards in these areas and the different expectations which we can have of contributions to be made to the debate on standards.

Finally, I want to turn to six strategic issues which really should dominate our thinking about office automation over the next four or five years.

That is not to say — let me say it now and I will repeat and emphasise it when we come to this point — that every company has to have fully developed and fully adequate policies in each of the six strategic areas which I will mention. I think that one of the dangers in the current situation is that one puts too much of a counsel of perfection; we could easily get ourselves into the situation where we are doing so much planning that we never actually have time to take any action. That is perhaps a lesson which we should have learned from the past. But there are certainly six areas in which we should be refining and developing our thinking as fast as we can.

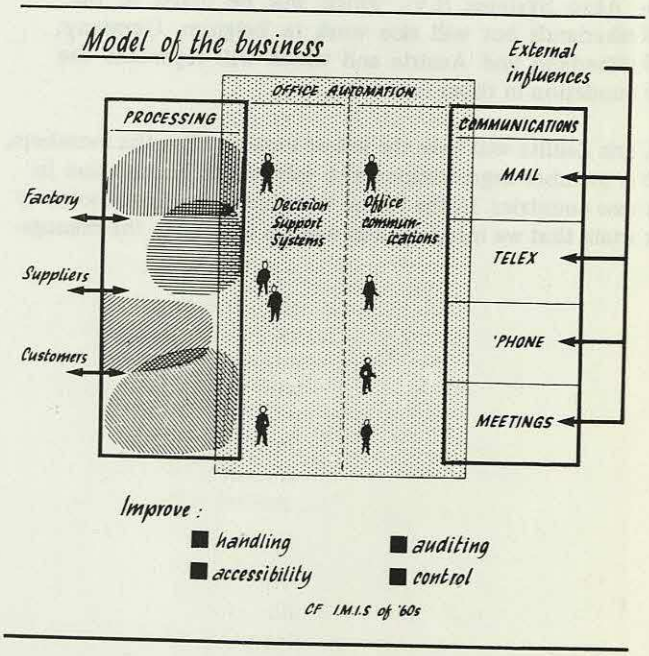

Let me begin with the broadest possible model of a company and its information systems and processes. What we have on the left of this slide is a series of processes connected with the factory production process, with relations with suppliers, with relations with customers. If the outlines of those systems look amorphous and blobby, then that is not an accident of design in the visual, they are intended to look that way.

What we also have are communications systems which, by and large, neither connect particularly well with each other nor with the processes which they are meant to serve. Those communication methods may include conventional things such as the mail, telex, telephones, and meetings; and, as time goes on, doubtless more advanced methods of communication.

Because these methods are neither integrated one to another, nor particularly well connected to the processes which they are intended to serve, what we see as a characteristic of our current situation of where we are today is a very limited degree of connection between the systems and a great deal of making it up as we go along. The role of these chaps here is not actually refining Brian Gladwyn's (Spillers) flour, as you might think, but doing all the ad hoc linking of these communications systems with the processes of the organisation! Basically, it is the ability of people to stitch these systems together and to link the communications system of the company with the processes of the company which keeps the whole organisation ticking over right now. In terms of systems, there is an aching void there. In a sense perhaps one could argue that it is not desirable that that aching void should ever wholly be removed since the people in that gap derive a good deal of their satisfaction from making these linkages, often in ingenious and satisfying ways.

So, what is the role of office automation? It is a grey area — and again, that is by design. But basically there seem to be two parts to the problem: the office communications system, which is what I principally want to talk about today; the way that office communications can be improved both to improve the degree of connectivity between these systems, and also to make them easier to link into these processes, but also decision support systems which we have been talking about for a long time; systems that help managers to take better decisions and present more accurate, up-to-date and decision related information, but which so far I think it would be true to say have been distinguished more by description than by availability.

What are we trying to improve? We are trying to improve our sheer handling ability, our sheer ability to cope with data, with voice, with text. The accessibility of information, trying to reduce the extent to which filing systems continue to be the classical 'black hole' into which everything can go, but by definition nothing ever comes out. The auditing of the system, trying to get a better fix on the total cost of this process, which again right now is something which none of us really understands how to do. And control.

When I think of that range of activities, it does remind me an awful lot of the past. It reminds me of the 1960s when we used to talk about the integrated management information system as though it were the touchstone of everything in the field of conventional data processing. I look around and I see some of the faces in this room — albeit older, sadder and wiser — that were involved in the debate about the integrated management information system, and I think to myself, "Perhaps, darling, they're playing our tune again. And the last time we danced all night, we broke a leg, didn't we?" So some note of caution ought to creep into our thinking lest we make that particular mistake again.

Lessons

- *The techology is here*
- *Build on present systems*
- *Look for a platform*
- *Tackle internal systems first*

What are the lessons, in fact, that we can learn from the past? First, although technology is advancing, and advancing at a speed which is sometimes very frightening, when you actually look at what you would like to do in the office, it is amazing how often it turns out that the technology required to do it is already here, and has possibly been here for five or seven years. Any idea that to make real progress in the field of office automation we are waiting for the research and development laboratories to turn out cheaper or better or wilder products, I think really does not bear very much examination. What we do not have, as yet, is the ability to take the technology that we have, for example, in the existing generation of word processors, or the existing generation of viewdata systems, and see precisely how we can apply that in ways which are self-evidently worth the money and effort of doing so. That is a point to which I will return.

Second, people are not saying right now — and in my view never are going to say — "Let's have office automation. Let's try to build it up from scratch." Basically what we are going to do is to build on the existing systems, starting from where we are now rather than from a green field vision of the future. There are one or two examples of how people are beginning to do this and how it turns out that it can provide a very effective entry point to office automation. For example, Massey Ferguson, in the United Kingdom, started using an ordinary data processing time sharing system which happened to have a mail box facility built into it, using ordinary data terminals, and have found that riding on the back of an existing and justified time sharing application they have now developed an extremely cost-effective approach to message switching.

Similarly, another company in Britain whose plans have not yet been made public is seriously considering the installation of an in-house viewdata system, simply to replace an existing printed mechanism for distributing corporate information around the company and telling people about some of the things that are happening in the company before they read about them in the press. So again, a good opportunity to build something on an existing system more cost-effectively than had been done so far.

Both those companies have achieved what one might call 'looking for a platform' on which to construct some advance towards office automation. It really does seem to be one of the key things that one has to do — to find such a platform — if only in order to avoid two problems: first, the rather intangible nature of the benefits which are likely to be delivered; and second, to overcome the threshold problem of 'how do I get from the situation where I am now to having enough of these terminals — whatever they are — installed

reasonably to expect to get good utilisation and effective use of them on a company basis?'

The fourth lesson seems to be to tackle the internal systems before attempting to tackle those outside the scope of the individual company. I think that there are a number of reasons for that. First, because it is the internal systems which determine how fast you can respond to outside pressures, orders, price changes and so forth. Second, because you are likely to waste less of your effort if you concentrate on the internal problems — you are going to waste some of your effort anyway — rather than try to tackle the external problems where changes in regulatory policy, PTT policy or whatever can have a very serious impact on what you are trying to do.

Clearly, the picture emerging from my first slide and the need to absorb these four lessons from the development history of the past represent a fairly sizeable agenda for any management services department. Not surprisingly, if we dwell too long on that, it probably becomes a little depressing.

External forces

- Role of the PTT
- The competitive market
- Who sets the standards?

But, on top of that, there are also external forces which need to be taken into account in trying to frame a policy in this area. Typically, I suppose, coming off the top of most people's worry list would be these three: the role of the PTT; the competitive market and what evolution in that market place is likely to do for the end user, the customer; and finally, who sets the standards by which these systems are going to talk to each other. Those are the three external forces that I should like to talk about.

The PTTs

- Conflict with IBM?
- Promotion or protection?
- The honest broker?

First, let us look briefly at the PTTs and their role in helping or hindering the user to advance towards automation in the office. The first point which has become a rather fashionable debating point right now is the role of the PTTs vis-à-vis the large computer manufacturers, and particularly vis-à-vis IBM. In the past two or three years, one has seen a good deal written and said about how head-on confrontation between the PTTs and IBM is inevitable. Partly that speculation is triggered off by regulatory conflicts here, particularly between IBM and AT&T. But from what we in the Butler Cox Foundation have been able to observe of relations between IBM and the PTTs in Europe so far, it does seem as if this could go down in history as 'the battle that never was', in the sense that there appear to be very deep-rooted feelings, both within the PTTs and within IBM, that in a market which is going to expand as fast as this whole area of information systems in the office, there is scope for the PTTs and IBM to co-exist, and that if confrontation can be avoided then it must be.

Certainly, in so far as one can see examples of IBM's direct dealings with the PTTs and their contributions to government debates about the role of the PTT and so forth, it does seem to be IBM's policy right now, for reasons that seem understandable and valid, to be supportive of the PTTs rather than the reverse. Therefore I believe that one has to treat with a good deal of caution forecasts about these head-on clashes between the PTT and IBM. At the same time, you cannot expect either organisation to make life completely convenient and easy for the other. For example, if one looks at packet switching networks, it is clear that IBM is going to do its best to support the CCITT initiatives in this area, and particularly the X.25 protocol, but only at the level of a link protocol; so that although it is not impossible for IBM to use the packet switch networks of the near future, it is not particularly economic for them to do so either. In a sense they will be paying both for the IBM component of the software and also for the PTT's component in the software.

So there will be inconveniences. There will be points at which the computer manufacturers decide that they want to try to screw a little bit more out the the PTTs, and vice versa. The next niggle — and it will probably be quite a niggle — will be the question of tariff structures for leased lines after packet switch networks become more universally available within Europe as a whole. The computer manufacturers will doubtless have their contribution to make to that debate.

But by and large the message that is coming from our research right now is that if you are waking up at night, worrying about the conflict between IBM and your PTTs, then probably you are over-reacting and probably they are determined on peaceful coexistence.

A more serious problem for all the PTTs in Europe, particularly for those which have already gone some way down the path towards the liberalisation of their market, is the question of whether they are in business to promote the use of the public switch telephone networks in Europe to the greatest extent, or whether they are in business to nurture and develop the indigenous telecommunications industries within their own countries.

Policies vary rather widely on this. In Britain, we have been surprised at the speed and apparent facility with which, for example, IBM has secured permission to do some extremely interesting things from the point of view of office automation, with its new 1750 voice switch. I am thinking of the teleprocessing line handling protocols and the like, which seem to us to open a number of very interesting doors to IBM.

On the other hand, in the Netherlands you had a situation where, even with the previous generation of switch — the 3750 — the Dutch PTT was simply not interested in issuing IBM with permission to connect that switch to the public

switch network. So there is a diversity in the approaches being adopted.

But the situation in all countries is fundamentally the same, that if the PTTs were to say, "We are in the business of promoting the use of bandwidth, and the more bandwidth we can sell, the better it is for us. Therefore we're going to take the most liberal view that we can about what can be connected to the network. If a device comes into the market which can increase the use of bandwidth, we don't care whether it was made in our own country, or in America, or in Japan, or in Taiwan," then the situation would be changed very rapidly, and disastrously some would say, compared with what it is now. There can be few things that would imperil Europe's telecommunications industry so much as a rapid and forthright acceptance by the PTTs of the policies which are recommended to them by many, many experts in telecommunications; because the European industry simply is not geared up at the moment to deal with an open, competitive, interconnect market.

This is not to say that no progress at all can be made in that sphere. The way that the replacement for the telex protocol is being developed now — the Teletext protocol — is a good example of the way in which the PTTs can make possible the advent both of pretty sophisticated equipment for one section of the market and simply much better based products at another level in improved telex services which are still, in many cases, the work force of communications in many organisations.

There must, however, be a role to which few of the PTTs are currently addressing themselves with any great vigour — and that is the role of resolving incompatibility between different devices which people might wish to connect to the network; in other words, regarding the public switched network as something which is capable of developing its own intelligence and capable of resolving incompatibilities between terminals. This is something which most PTTs have been very cautious about committing themselves to, certainly dipping one toe a few millimetres into the water of packet switching before trying to go very much further.

It seems to be a role for which the PTTs are well qualified. After all, they are the only people with their own networks already fully developed and in place, and with access to enough capital to develop those networks in the way that users might wish to see them developed. There are already examples of private enterprise companies moving into these areas, like Codex for example, moving up from the supply of ordinary old modems and into the interconnect business in a fairly big way. It might well be as the honest broker of connectivity between otherwise incompatible terminals, the friendly nationalised offerer of network coherence, that perhaps one of the most decisive roles of the PTTs might be seen.

As far as the market is concerned, I believe that we are seeing structural changes in the market place affecting office communication systems of a very fundamental nature, and one which the user really has to understand if his purchasing policy is to be anything like intelligent. We have people who are capable of doing the whole systems assembly job, the packaging of total systems required to handle all kinds of communication, of whatever type, within a large organisation. IBM, of course, springs to mind as the prime example in

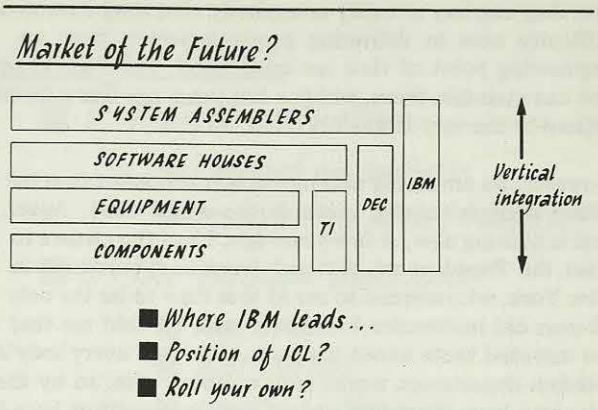

that category. I suppose that one would also say that for certain types of communication and certain types of company, Satellite Business Systems Inc. is also attempting to put itself into that role.

There are software houses which are capable of providing components, equipment and software, but undoubtedly leaving some parts of the system assembly process, either to other companies or to the user himself. Then there are equipment and component manufacturers. We have chosen as an example of the former category, Digital, and of the latter category, Texas Instruments.

What is interesting about that picture is that all the way up the line it seems that people have ambitions to get into the next category up. People have ambitions, if they are capable of delivering software systems, to get into the business of total system delivery; and if they are capable of providing components and pieces of discrete equipment, they also seem to be interested in getting into the business of providing complete software systems. So one is seeing, with this process of vertical integration, a good deal of confusion in this market.

There are a few observations that one can make about it, based on what is happening right now. Certainly we cannot expect to see too many of the companies that can offer a broad range service delivering a total system. Indeed, if one looks ahead ten years from now, it is probably not imaginable that there could be more than two or three companies in that business; and as a matter of chance, or maybe planning, all three of those companies are represented in this room today.

The second point is that the cost of entry at the bottom of this inverted pyramid seems to be declining fast enough to make the market a rather confused one. I said in joke quite recently that my Uncle Harry who runs a secondhand furniture shop in Camden Town is thinking of going into communicating word processors! The cost of entry really does seem to be dropping and the ability of people to buy components and stitch them into some kind of half way reasonable product and sell them to people and get started amazes me.

Last year, we had as part of a consultancy project the task of touring round very quickly all the small computer manufacturers in Britain and the biggest change that I see, talking to them, compared with five years ago, is that because of the improvements in the reliability of components

that they can buy as easily as anybody else, they have no real difficulty now in delivering products which from an engineering point of view are quite good. They work, and you can maintain them, which a few years ago was a facility limited to the very large.

However, we are seeing also in this whole market that the added value is moving upwards the whole time. Again, that is nothing new. A few years ago, I had the chance to meet the President of National Semiconductor here in New York, who seemed to me at that time to be the only 15-year old millionaire I had ever met. He told me that the essential truth about his business is that everybody's research department works for everybody else, so by the time you have a product out all your competitors have it very, very fast indeed. The added value seems to be moving up this way.

I do not know how many of you saw something which I am sure my friends from IBM will forgive me for quoting, but it did seem to me to be a really interesting piece of good news/bad news tactics. If you wanted a 4300, which was announced on January 30 in Europe, you had to get your order in by March 5 — which did not leave too much time for detailed appraisal. I have been waiting for the other shoe ever since, and it came last week in *Computer World*, which rightly or wrongly — and it may be wrongly — estimated that the cost of software to a 4300 user two years down the pike would be twice the cost of hardware. That is the second shoe. But certainly the added value in this market is moving up the market and people are trying to pursue it.

The inability — not just with any one company but the complete inability within the market as a whole — to get any kind of fix on the realities of price demand elasticity has created a very interesting situation. If anybody could do a really effective consultancy project and come up with some real wisdom on price demand elasticity in the market for medium and large computer systems, I believe that they could name their own price; because what it must be costing the computer manufacturers to under-deliver in the way that they are currently doing must be absolutely prodigious.

But the result of that under-delivery — and I think the backlog for IBM processors alone now runs into many tens, or even hundreds of thousands of systems — must be to take away a big, black cloud from over the heads of the Amdahl family and a lot of other people as well. There really do seem to be extremely good auspices right now for the plug compatible manufacturers, stemming from the rather gross underestimates of the size of the mainframe market over the last five years or so, and presumably extending into the future.

So one of the implications of this for the average user is that if he is dependent on external sources of supply in a mainly IBM environment, he is probably a bit more secure right now than he was two or three years ago.

The other question which it raises — we have said 'the position of ICL?' but it also applies to other computer manufacturers — is that when there is so much choice available within the area of IBM compatibility, what is the merit of IBM non-compatibility, which was a deliberate market strategy pursued by some companies in the past? If there are sufficient choices available within IBM compatibility, you have to have some other, powerful reason for wanting to stay non-compatible.

I have no private inside information on this, but if I were in the board room at Putney Bridge right now I would expect to see written on the wall, "Come back System 4. All is forgiven."

Finally, on this end of the market, I would say that we are going to see an increasing tendency for users to want to buy in their own components and start devising systems for themselves. When the added value is at the level that we have all been tackling anyway for many years, then the tendency to keep some of that value in-house will be an attractive one. I am looking forward with a good deal of interest to the result of our research project on the use of microprocessors in information systems, both in Europe and here in the States, to see how far this tendency of do-it-yourself systems is beginning to manifest itself right now.

I think that one has to be careful not to expect too much from that area. From what we have seen so far, there do not seem to be any opportunities for great cost saving in that area, because what you save on the hardware you will probably spend in increased software development costs. But there do seem at least to be opportunities for companies to get a product which is much more customised to their requirement and much more susceptible to improvement and enhanced as far as they are concerned in that way. But do not be misled into thinking that there are major cost savings to be had in that area, because the evidence that we have so far suggests that there are not.

Standards

- Key role of established suppliers
- CCITT
  - X series recommendations
  - Effect on FAX market
  - Teletex by 1981?
- Plugging the 'high level' gap

Turning to the vexed question of standards, office automation cannot be seen, as I said before, as something which we are doing from a clean sheet. We have to build it on the basis of the systems that we already have. Therefore, there is an important role for the established suppliers, with their established products and also with the new products which they want to bring out to preserve the loyalty to their products which their customers already have. In so far as there is a determination of roles to be made here, there is no doubt in our minds that it will be the established companies that will be the pace setters in this area of standards, simply because you all need to preserve the investment that you already have in the equipment of the past.

Similarly, it is possible — and again this is something to which expert opinion over the past two years has contributed

in a slightly unsatisfactory way — to exaggerate the role which has been achieved and is likely to be achieved, and indeed can be achieved, by international standards making bodies like the CCITT. If you were to read the magazines, I think that you would have the impression that the CCITT is working right now as a kind of sovereign authority, laying down standards which are going to determine the shape of the communications business for the next hundred years. If you think about it, this is far from the truth, simply because of the limits on what international standards making bodies can deliver.

CCITT, in delivering the X series of recommendations, for example, delivers connectivity between devices, but it does not deliver coherence. It does not deliver the ability for systems to have real conversations. It simply delivers low level protocols for them to communicate one with another. That is probably a sensible limit right now to the ambitions of international standards making bodies of that sort. Sure, there are discussions going on about more elaborate protocols, but that will take some time.

That is not to say that it is a waste of time, because one has seen, for example, how the fax market has been changed simply by the establishment of connectivity by the CCITT. We expect to see the Teletex standard for communicating word processors — the super telex standard — approved by 1981, and making a valuable contribution to the ability of devices to communicate one with another. But if you compare what you are likely to be able to do then with a machine talking to a machine, or a machine talking to a human, with what you can do on the telephone right now, simply because you have human to human with existing human protocols for conversation, then I think you can see that we are a very long way from any situation in which the international standards making bodies can determine the kind of user level discourse that systems can have, one with another. Just a word of caution on that.

There is still this 'high level' gap — the user protocol gap — which has not yet been filled by the international standards making bodies, and which in a sense you could say is more likely to be filled de facto by the established suppliers of equipment, with the international standards making bodies being to some extent obliged to recognise the reality of those de facto standards in the future.

Let us turn to the six strategic issues which I think should be high on the list of priorities for management services directors thinking about office automation in the next few years.

First, project evaluation. I believe that we have a real problem in that area, and one which we are not in sight of solving adequately right now; and one which nevertheless we have some pointers towards an improved area.

We have not yet finalised it, but we are planning at the next Foundation conference to have as a speaker Paul Strassman, of Xerox Corporation, who is a very distinguished advocate of the point of view that we are not assessing properly the impact of systems on our business as a whole. The rather crude measurements which probably everybody in this room has heard from our research before, that standard data processing applications seem to account for something like 2% on average right across the board

Strategic Issues

- 1. Project evaluation
  - impact on business
  - measuring costs/benefits
- 2. The human dimension
  - hygiene
  - motivation
  - involving the user
- 3. Planning
  - facing convergence
  - M/S in the corporate process
- 4. Foundations
  - role of the network
  - the database approach
  - skills and experience
- 5. Timing
  - competitive market
  - pace of change
- 6. Technology
  - what can it really do?
  - how does it fit?
  - selection policy.

of all industries of corporate expenditure, but information systems of all kinds, handling voice data and text, account for 20%. We are still not coming to grips with that.

I think that we are also not getting ourselves out of the box that we have made for ourselves on evaluation of office automation projects. I am aware that the view that I am about to express probably is not a mainstream Foundation view, so I had better preface it by saying that it is a personal view rather than a researched view. It is also a somewhat contentious view, which is unusual, coming from me. It is that most organisations that we talk to believe that they have a pretty hard-nosed approach to evaluating office information systems. In other words, they believe that unless they can see a clear money return on what they are trying to do, a clear investment justification, then they will have a hard job selling that concept to their management.

In the first place, I am a little doubtful whether this happens quite as often as we like to think it does, because like most of you I have seen products and systems sold, particularly to top management, on much less convincing grounds than that. But second, I am not even sure whether it is desirable,

because in my view, if you think about it honestly and step back from it for a moment, what you are really trying to do on the basic old cost displacement argument that 'if we do it this way we can save money', is taking one view of a notional future which might exist with the system, and another view of a notional future which might exist without the system, and you are subtracting one from the other. And if you take one fantasy figure from another fantasy figure, you will end up with fantasy. Is this a satisfactory basis for an investment decision?
I am actually reluctant to be too negative about that, because if I believe that — and I feel increasingly that I do — I ought to have something to put in its place. Apart from the mindless zeal of technology for its own sake which none of us wants to advocate, I am not too sure what that something is. But I do think that we have a big problem in that area. Again, one of the reports on which the Foundation research team are working right now is investment in management information systems, and at least we will have a better idea of what the actual patterns of behaviour are.

But I think that we are reaching the end of the road in which cost displacement was the main or only means of justifying investment in information processing systems.

The second area of fundamental concern here is what I have called the human dimension, what happens to all these people here in the picture that we are painting for the future. First, I have used the word 'hygiene' in the standard Hertzberg sense of simply looking for opportunities to use the technology to clear up some of the messes that we have created. I think that an awful lot of the jobs that people do in offices are routine, repetitive, soul-destroying and so forth. I am sure that there are opportunities to improve the quality of those jobs, making more task-oriented and so on. That is one aspect, just using the technology to solve some of the messes that we have created without the technology.

But the other big area is the whole area of motivation of people to undertake office automation and how, if we think that it is a good thing, we can stimulate them and entice them to do so. There are three tasks which we have to undertake if we are to fulfil a responsible role in that area. The first sounds obvious, but I think that it will be much more difficult to do than any of us imagines right now. It is keeping space for rational solutions to these problems at a national, economic level. At a national economic level in the United States, in Europe, in Japan, in all countries, there will be pressure for these problems to be solved at the level of political sloganising rather than rational thought. One is already beginning to see the seeds of that sort of slogan. I imagine that most people in this room, certainly most of the ones who come from Britain, have seen the film made for BBC television last year entitled "Now the Chips are Down", and have seen the reaction of Mike Cooley of the Engineering Union, arguing that people define themselves by their jobs. If you ask somebody, "What are you?" he says, "I'm a plumber", "I'm a dentist", or "I'm an order clerk", and that you cannot take away from them the dignity of those skills without having a careful look at what you are going to put in their place.

Most people would go along with Mike Cooley's analysis as far as that goes, but it does seem to me that it is an inadequate analysis of the real situation. What makes me suspicious of it is the simple fact that if many people did not hate their work, what would Mike Cooley do then? He needs them to hate their work, doesn't he? If we could restructure their jobs so that they got up in the morning and thought, "I like going to work. I get on well with my boss. I think I get a reasonable reward for what I do," old Mike would be finished. He would have to find another job as well.

It is worth thinking about that. I said that there were three areas. The first is keeping space for rational solutions to these problems. The second is that I think that it is absolutely essential, if we are to get a proper analysis of the motivation for office automation, to decouple the question of production from the question of effort. We tend to think that if work does not hurt, then it is not real work. I know that my father told me that his aim in life, to which at his present advanced age he has not succeeded, is an occupation which does not involve lifting heavy weights. Those of you who arrived at Kennedy Airport yesterday will realise that you have not reached that goal in life either. But what I am trying to say is that we really do have to decouple the effort put into work from the results achieved. If it were possible for one person, working for five minutes a year, to generate enough wealth to keep the United States going for that year, then you have to ask yourself, "Why not?"

It does leave you with another problem: what everybody else does for the whole of the time. But that is another problem, and decoupling the question of input from output in the work environment is a first prerequisite from any kind of rational analysis of what we are trying to do with office automation. I think that will be a subject to which we will return, time and time again, over the next few years.

The third thing which I think we have to do — if it is not taken too far — is that we have to emphasise the more positive aspects of office automation and what it can do to make people's lives and work more enjoyable, because for every one who is trying to do that, there will always be ten people saying that it is leading us down the road to mass unemployment, total exploitation, and the mindless morons of George Orwell's novel, '1984'.

Finally, whatever we do we have to involve the user of the end system in it, which we have not been particularly good at in the past.

The third strategic issue which needs to be in the mind of every management services director is planning, first how he is going to face the reality of the convergence of technologies in this area; and, second, how we are going to define management services in the corporate process as a whole. Certainly, any attempt to organise management services to cope with convergence in the short term is likely to lead to problems and resistance, and therefore I for one have stopped simply giving people a counsel of perfection and saying, "Make sure that all these office automation functions come under unified management." In many cases, it simply is not possible to do that in the short term. But what one must do is to look at management services, even if it is in a relatively fragmented and in some cases uncoordinated state right now, to see how it fits in with company planning as a whole.

This is a difficult problem to resolve. We find in our consultancy practice that it is a difficult problem to resolve; and, if I may say so, those of us who work in management

services do not make it any easier to resolve, since recognising the fragmented nature of our skills, we are nevertheless ready, willing, eager — and indeed, unstoppable — when given an opportunity to generate a fragmented solution to a barely defined problem in the mind of a user, who is really not sure whether he is talking to the right person anyway. So I think that prudence is required in that area.

I know I am building up a rather formidable list of these strategic issues. I did say right at the beginning that we do not expect every management services director to have his policies completely worked out in these areas, but we do think that they are areas in which he needs to be thinking right now.

There are three basic building blocks or foundation stones on which the policy for office automation is likely to be built.

The role of the network. Virtually everything that we are striking now in the attempts by companies which are a little bit further down the highway in these areas brings us back time and time again to the role of the network. What sort of intelligence does one want in a network? What sort of functions does it make sense to have in the network? What sort does it make sense to have still under centralised control, either through a mainframe based system or a terminal based system in the network?

Certainly, the experiments going on — some of which will be reported to us this week — with intelligent networks and the offerings which are being brought to the market right now deserve the closest scrutiny and the closest assessment for their implications for the future.

Second, what we have called the database approach, but it requires some explanation if we are to say what we mean. It is that the approach adopted by the data processing community at large to the question of database management systems turns out, on inspection, to be a much better approach than any of us gave it credit for in the past.

I wonder whether your experience is the same as mine. I remember a few years ago, going to give a talk at the British Broadcasting Corporation, at an internal conference that they organised on database management systems. I gave my talk, which was kindly received; and I happened to find myself, by chance, in the men's room with the management services director. Knowing that they had not yet committed themselves to a database management system, I said to him, "What do you propose to do about database?" He actually walked the length of the men's room, checking that all the cubicles were empty, before he said to me, "Nothing!"

I think that rings a bell with all of us, because we have thought about database in the past as something which, in the main, was wished on reluctant managers by enthusiastic technical experts within the organisation. Those of you who have had a chance to study the results of our survey of user experience with database management systems will know that the results turned out to be quite surprising. The results turned out to be that the users of database management systems feel that they have, in the main, had a very good return on the effort investment that they have made in DBMS. What is really surprising to me, looking at those results, is that they actually think that they have had more out of the systems than they expected when they first installed them. That is interesting because it runs counter to the accepted knowledge and wisdom of the database using community.

So what we are really saying here is that, although persistence and a fairly long term view may be needed for office automation, particularly for the development of intelligent networks, there is precedent in the use of database management systems for believing that technical complexity alone does not stop you getting that return if your investment is a reasonable one and it is pursued with a fair degree of diligence.

The third building block on which these premises need to be established is skills and experience, which we do not have right now — a kind of substitute for the track record of knowledge and experience that we already have in data processing but which we do not have in networking or office automation.

The fourth strategic issue is concerned with the timing of these developments. All the products which we have talked about in the Foundation reports in the past and which we are talking about right now really do seem to be pretty short term ones. We have seen designs for multi-function work stations which will be on the market place in the next year or two. We have seen designs for communicating word processing equipment which is well able to implement the Teletex standard, if it is published in 1981 as we expect it to be. But it is thinking through the implications of that rather confused competitive market situation which we described which is difficult, and particularly thinking about the pace of change and the rate at which products are going to obsolete themselves. Again, it would be extremely easy to reach a situation where, because the next round of technology which may be only a year down the pipe is looking so attractive, we cannot really bring ourselves to buy anything that is currently available.

Finally, what the technology really can do for us. All the forecasts in the world about how the density of semi-conductor components will change, or how magnetic storage media will change, are valueless unless we can see ways in which the technology will actually be used. How does it fit with the systems that we already have and how does it fit with the people that we already have? since not even office automation will transform the nature of those people. Finally, at a more detailed level, what our selection policy is going to be; what sort of manufacturers we are going to choose to get in bed with; and who we think we can trust the most to provide the kind of support that we want.

Ladies and gentlemen, I have tried to present a fairly broad picture, starting with a theoretical model of any company and its requirements, and then concentrating on the office communications aspects as being the ones in which the salvation of the individual user is most likely to lie within his own hands. We have looked at some of the external forces — the main ones I think: the PTTs, the suppliers, the standard makers, who will seek to influence the user community to go down one path or another. We have looked at six strategic issues which our research suggests should be fairly near the top of the management services director's worry list, when he concerns himself with the future of office automation.

But all of that is really meant as a kind of framework within which we can fit the other sessions that we shall hear at this conference on Office Automation and Distributed Processing, particularly the experience of some of the pioneering companies that we will be visiting during the rest of the week. It will be interesting at the end of this week to look back to see whether any of the strategic issues which we have identified require revision or deletion, or whether more need to be added. We will be doing that during the course of the week and reporting back to you, through Foundation research reports in the future.

DEVELOPMENTS IN EQUIPMENT AND FACILITIES CONCERNED WITH THE AUTOMATED OFFICE -1

Jeffrey Holden
Computer Corporation of America

Jeffrey Holden is the Director of the Communication Technology Division of the Computer Corporation of America. CCA are a high technology software firm specialising in advanced software systems.

Before joining CCA Mr Holden spent several years in the communications technology field with COMTEN Inc. in marketing and as product technical support manager.

While a Captain in the Army, Mr Holden served as an evaluator of computer systems hardware and software for the Department of the Army Applications. His formal education includes a BS in Engineering Management and an MS in Management Information Systems.

RAY: During the remainder of the morning we have three forty minute sessions from major suppliers of communications and office systems in the United States. Two sessions are concerned with networks — from Computer Corporation of America and from Tymnet and one from Wang, who are well known to us in the United Kingdom and Europe.

First, I would like to introduce to you Jeffrey Holden, who is the director of the Communication Technology Division of the Computer Corporation of America. I guess you will all be relieved that he does not look like his photograph in the Agenda!

Jeff is going to describe the Comet message system and explain some of the applications which are currently running on that system.

HOLDEN: I want you to keep in mind my comments with respect to the fact that I represent a vendor, but I think we have some knowledge that we can convey to you, and some help as you embark on studying and hopefully implementing office automation.

As you know, I represent a firm which is a supplier of office automation products or a product. I would like to zero in on a specific term and that is a computer based message system — that is a precise three or four word definition of the product that we provide. The literature shows us that computer based message systems are recognised to be a key element of the office of the future, and in fact our viewpoint is obviously that it is the most important element.

We know this to be true from our own experience.

The theme is how and when will office automation happen? Well, I would like to delay that for a couple of minutes and explore what is office automation?

I suspect some of you are saying as you sit there — what do I mean, office automation? My boss thinks I'm doing a good job, the people that work for me seem very happy, so I am really not sure why office automation. Well in your case the answer comes out to be the same — productivity — because if you look at it from that viewpoint, that means if we provide you with the means to achieve an automated office, and they are more productive, you ought to be able to do more tasks in that given time, and handle more things in a given day, week or year.

Improved productivity is the promise of office automation. On the other hand though, if you have less controllable day to day work, I suspect you have what I term a "hassle factor". I think we all have a certain amount of hassle, and I hope that term isn't too colloquial. By hassle I mean being bothered, being interrupted, being out of control to some extent — spending a good deal of the day not in control of the activities you are supposedly carrying on in finding yourself at 4.30 or 5.00 trying to jam in the three or four things you intended to get done, because you have been interrupted all day.

I think we can all relate to interruption. We have the experience of having a nice dinner at home; phone rings, and we have been interrupted. And even in the office this scenario is something we have all experienced. We have spent three or four days trying to get 15 minutes to see our boss or some higher level executive; you get in there, sit down and get one sentence out when the telephone rings.

In the United States we have another phenomenon — the multi-button phone; you may have that too. I think this is kind of interesting. You get in there, you say one sentence, the telephone rings, he gets to talking for two or three minutes and the second line rings — "Wait a minute Johnny, I'll be back to you in a minute" — and we allow this, we accept it and we continue to try and operate this way. It doesn't take a psychologist to determine that we work best when we have uninterrupted time to concentrate on the tasks at hand.

Office automation brings up another phrase or buzz word phrase which I think is being bandied around now in the field, and that is the "Saturday Syndrome". I am sure you have all had occasion to have to go in the office on a Saturday morning to catch up on some tasks which have

fallen behind, and you say to yourself, "I'll go in for three hours — 9.00 am to noon — and I'll get caught up. You go to the office and turn on the lights and you look up and it's ten past ten and you have done the three things you have come to work to do. And you say, "I didn't realise I'd get those things done so fast," so you do some more and you leave about 11.30 am, and you buy your wife some flowers and you go home and you feel really good.

Well how would you like to feel like that every day? Because that's the way you can. Why office automation? Because our experience and other's shows that today's executives are over-worked, short of time, constantly hassled, and office automation provides relief; it makes the manager more efficient by organising his or her communications and allowing that manager to be master of his communication in his own time. I and others here can testify that there is a better way to live.

But still, there are economic reasons for office automation as well, and we will be exploring more of these throughout the day, and you will for all week for that matter.

I would like to point out, perhaps an obvious point, but one which I think is worth emphasising. For several decades now we have invested tremendous amounts into the auto manufacture, into textiles, into steel mills, to make workers more productive. Yet we have done very little to make the office worker more productive. Investment is very small; in fact, related to inflation it is very backward. So maybe it's time we did some investment for the office worker, and that's why all of this office automation is buzzing now, I think.

At this point I think it would be helpful to talk about our view — mine anyway — on what office automation is.

A great deal of office automation centres on improving communications. After all, it's the resource the office uses to effect action upon particularly managed activities, i.e. the aeroplane in flight — a very automated office, we hope.

A hundred years ago we really had only one way to communicate — face-to-face. Then the telephone was invented, and you will note that it didn't replace mail. I suspect that you will all have pretty full 'in' baskets each morning as I do, full with traditional mail. In the same way we are now getting office automation tools, and we have seen in particular computer based message systems that are able to supplement the mail, telephone and face-to-face meetings. But unlike the telephone, these new systems are several things. They are facsimile equipment, word processing, especially communicating word processing devices and, of course, computer based message systems.

I am particularly interested in computer based message systems and I would like to explore a little more about what they are.

A computer based message system is a special form of electronic mail. It's special because it allows the user — the user being the manager or everybody in the office — to access his or her messages, at his or her convenience, and to dispose of them electronically and to file or to pass them along as he or she sees fit. It leaves a perfect data trail and eliminates or at least reduces the need for paper. What are the advantages? Well very simply, I can sit down and handle some forty to fifty pieces of routine correspondence in as little as an hour. I know that to be true from my own personal experience and we know it to be true from our users. Now this compares with the ability to make some five or six telephone calls in that time. How many of your phone calls actually get to that person you intended to call on the first attempt? Often you place a call, you go to lunch, the other guy will place the call back to you. You will try to get him back, about the fourth or fifth time you might get that phone call through, and with some luck you might remember why you placed the call in the first place!

Another comparison; thirty to fifty pieces of correspondence using a computer based message system, five or six telephone calls and how many meetings? Maybe one or two face-to-face meetings, and how many letters can you write in an hour? Well, certainly not thirty to fifty, unless they happen to be standard letters. But computer message systems also allow the manager to make use of all of this non simultaneous time to solve non simultaneous problems, instead of using real time tools such as the telephone to provide those answers. It also helps in the organisation of that manager; it presents problems in an organised manner, so that you can select the important ones, and handle those that are relatively important, and do this during a specifically chosen time. You choose the time.

We could go on for a long time with the 'what is the computer based message system?', and I think it would be more of a point now to leave that until the question and answer period, and go back to the specific topic today — How and When Office Automation?

Well, how office automation will come to be. I don't think there is one answer, because as I pointed out just a few minutes ago, office automation is really several things. One scenario which we believe and I believe in is that it will come from a squeeze play, that of attaching the unstructured communication within an office and attaching a structured communication within an office. Now this squeeze play then can relate back to the office automation types of things. Word processing and facsimile certainly attack the structured communication, and they are doing it quite well, and I would like to go through an example of a structured communication, then talk about the advantage and disadvantage of attacking from that viewpoint, and I am sure Wang will be presenting even more detail here. I look forward to that.

But I am sure we all had to fill out or at least sign a travel request form to get here. That's an example of formal communication. This can be very easily handled by a word processing device of course or by that type of architecture. And it has some interesting advantages. It certainly is very measurable; you can sit down and analyse that form, what steps are taken, who has to fill out what, where does it go, and it is a nice area for vendors to attack. Wang can look at the application for a few minutes and develop a proposal, show the steps that they have taken away from the manual cycle and now provide you with an opportunity to say to the boss, "We can do this now with these devices and we are going to save about 2,000 dollars a month, and therefore pay for the equipment in x months". A very clear and nice advantage.

But there is another side to the coin. I would like you to

realise that such formal means of communication most readily relate to the clerical and not to the managerial parts of the office. I personally see very little of those forms, except to initial my request, and to initial the request of the people who work for me. I typically, when I am going to take a trip, type a very very quick message, "Eunice, I need to go to New York on x date; please get the reservations at this hotel. I don't need a car" — done. I am not the world's best typist, but she is able to decipher what I mean and can fill out the forms, and the next thing I know I have to initial this form. So what it really does is to reduce the clerical elements of the office.

Now studies show us that top management spends 80% of their time communicating, and middle management spends some 60%. It is important too, that the great majority of this communication is unstructured. For example, a very simple unstructured communication form is, "I would like to meet with my boss next Monday or Friday at 9.30." Well, how am I going to do that? I don't know where he is and we both travel quite a bit, so it's unlikely that I'll really see him. I'm likely to have my secretary call his, so I have involved two other people. And she, Tom's secretary, is probably going to interrupt them at some point and say, "Can you meet with Jeff this Friday at 9.30?", and he is going to say, "No I can't do that. See if 10.30 is O.K.", and it's going to go back the other way.

All I want to do is see him for fifteen minutes. How can I arrange it quickly? Well, I send him a quick message. "Tom, can I see you on Friday at 9.30?" He is going to read that; he's going to say "Yes", he's going to say "No, but I can at 11.30," or he is going to say, "I can't possibly meet you for another two weeks", or simply "O.K." — it's done.

Another type of unstructured communication is a problem, and if we have a problem with our system (which we do once in a while) it's likely that within a short amount of time someone is going to ask what was the nature of the crash and what are we doing to prevent that from happening again. Now I could wait for the weekly report, a week from now, but I know that Tom's going to ask me sooner than that so I am usually going to ask Ross sooner than the weekly report. Now what would I do without Comet or some other computer based message system? I would telephone him and not catch up with him for a day or so, trying to telephone him, or else walk down to his office and interrupt his meeting. It is a lot simpler to send him a message from my office directly — "What happened?" And it's very simple for us to respond in a sentence or two.

You can't measure that; the disadvantage is, you can't measure all that in formal communication.

So in order to attack the unstructured communication we need to build up some case studies, because I don't expect your accountants to just accept computer based message systems without some proof of the benefits.

How to go about implementing office automation and, more specifically, I would like you to think about how to implement this very important computer based message system type of office automation. What I would like to relate are typical scenarios within the United States for proceeding towards implementation. But I would like to relate these and enter some editorial comment because I see things that are being done somewhat wrong, and because I see new developments coming along which will help us.

The typical scenario says that the corporation or organisation establishes a task force to investigate the needs and uses of electronic mail or office automation or, more specifically, computer based message systems. (Incidentally on that point, in the United States we are told by consultant groups that some 450 of the fortune 500 companies have established these task forces, and if that has happened it's pretty much on it's way.)

The task force makes a recommendation for a trial test with some off the shelf hardware/software — such as Tymnet or Comet or others — then management agrees to a trial system. Using some group of people to test the value of these things and specifically the value of, in our case, computer based message systems. It's here where I question what is happening. What is happening is that they are going out and asking various people if they would like to be a part of this experiment or this trial. And typically they get a group of people who have never communicated before, have no need to communicate, but boy do they love to tinker around. So what happens is that they tinker around using this electronic mail or computer based message system — but what have you got for meat for judgment? Nothing, really nothing.

I recommend that you isolate a group which has a purpose in communicating, which does its job by communicating. Ideally, you will also totally immerse them in this product or these products. That way you have a real case study to see what happened. They get all their work done by 3 o'clock, and therefore you know they have more spare time since they have tried this office automation or perhaps they are beating their schedules for the first time in their history. But only if you totally immerse the group of people are you going to be able to see that change. If you are going to require them to double-do their communication, i.e. on the office automation tools and in their traditional roles, you are going to slow them down as opposed to improving their productivity. So I would like to caution you along those lines.

Then what happens ? The trial period concludes they may at that point, as the scenario goes, now see that they need to get a real live group of people that communicate every day, and establish these measurements to test the real value of them. But soon the trial period is going to give way to a decision time. The free users face a choice; they loved the system when it was free, but will they love it when they must pay for it? You will have to decide at that point to put money in the budget and then all the questions will come up, like "Can it be cost justified?". I was pleased to hear that at least one individual from the Foundation feels that maybe that's not so important as it was.

This is the point where most of the experiments are in the United States today. The crucial time is coming up to settle who has authority? Who has budget? Who is going to control this implementation, and when are we going to start?

Note this though, as a further emphasis of a problem of working up a cost justification. All the time this is going on — the other costs, the internal mail, the telephones — they are not usually costed back to the using department. Often they are made up or they are part of a corporate overhead and are billed according to some sort of formula, i.e. number

of employees per square feet, but the office automation tools are going to be very visible at this time. Those costs are going to come in every month and users in the using department are not going to see any relief from that other corporate overhead pool. Therefore, it is going to appear very expensive. Notice I say going to appear very expensive.

Certainly the interest growth is fantastic right now, although the level of real expenditure in computer based message systems is growing at a slower rate at this point in time.

We believe, and others believe, that the resulting momentum from the market pull and push will be quite astounding.

Before all this happens we have to realise some of the motivations behind the whole situation. Those executives, those managers, are just like us; they have all those behavioural hangups. When a manager is faced with the decision of doing what is good for his organisation, or what's good for himself, generally he is going to be a bit selfish, and he will choose a solution that is going to be a little bit easier as opposed to harder.

Management is judged on many things and not, as of yet, is it being judged on the use of office automation equipment to become more productive, but maybe it should start there. Managers get ahead in their organisations by being good politicians and by keeping their programmes on course, by controlling things. So all in all, I think it goes back to the bottom line. Companies are in the same relative business — and that is to make money. And if an organisation or a group of companies sees one of the industry members being innovative and suspects, or can prove, that that company is running better — leaner, able to grab more of the market, perhaps showing higher profits, hiring more capable executives — then I am sure that will be the push necessary for many of us.

And there is a push necessary to us at this point, especially in the area of attacking the unstructured communication. It's going to take bold, innovative work, and shakers and movers and believers within the organisation. There is one other motivation, which I think you may appreciate, and that is the realisation that the use of these tools may help personal advancement. I believe it will.

So this is where we stand. Computer based message systems are certainly powerful. There is no question that those who have observed and studied them agree that they provide the ultimate in immediate message transfer, and they are available today.

In summary, the 'how' question is, I think, answered by whether you are going to agree with me that a squeeze play is going on. You may choose which end of the spectrum to start; that is your choice, it is all of choice. But even our own organisation has recently announced a word processing comparability feature with our computer based message system, because we see the technology allowing us to move closer and closer toward the structured communication as well as the unstructured.

So the remaining question is when? When will all this happen? Well certainly it is now, for the Computer Corporation of America. It's now, for Digital Equipment Corporation, and it's now for a few others — Citibank and perhaps others. But when will the mass buying really begin? If I can get this one right, I'll be a millionaire, and perhaps we all will! I don't think I can answer that for you or for your organisation, and a better refinement of the question is, "When will your organisation involve itself with office automation? Now, or later?" You may already have begun. If later, how much later? After a prime competitor does it, and you see an example of what can be achieved, or when two or three of the giants have everything lined up?

We are talking about change — real change — and I guess it is somewhat scary. But I think the problems are hardly economic — I think that is really an excuse because, with just a little study and a little paperwork, the pay off is obvious. If you can get or you can identify a group or a division that you can totally immerse in office automation, you can test the pay off quite easily. Another way to test it is to take a group and increase the workload and at the same time offer the group management new office automation solutions such as fax, computer based message systems, teleconferencing, and find out what the executive would do.

We were under the impression that there was a great deal of sensitivity to age of an individual, relative to his acceptance of a terminal or a keyboard, and therefore office automation in the form of computer based message systems. We found that not to be the sensitive point, but in fact the sensitive point is how long that manager has been in his present job. If he is new to the job he's likely to be taken with the system — he wants to be innovative, he wants to make his mark, he wants to get things under control. But if he has been in that job for three, four or five years he has probably worked out all the wrinkles he's got into a routine; he plays golf on Wednesday afternoon and man, he doesn't want to be upset. He likes what he's got now.

But take that group and step up the demands that you place on them and a couple of things are likely to happen. First the manager will try to step up the things he has always done. Have more meetings, get into work a little bit earlier, give up that golf on Wednesday afternoon, stay a little later. But if you have significantly increased the demand, such as saying you are now going to manage nine people, not seven, that's quite a work load, and although it's going to meet with some success by stepping up his activity and all his people's activity, he is going to have to find innovative solutions, and it's at this point he is likely to take on some of the office automation that he's been avoiding like the plague.

But beyond that, there are problems. I don't think the real problem is cost justification. But there are real problems in the behaviour and bureaucracy around the sizeable organisations that we are dealing with. And by bureaucracy, I don't really mean to be specific to government. In fact, in our experience to date, governments are showing the greatest amounts of interest. I think this is partly because they have realised that to control things — and there are a lot of things in government that ought to be controlled — they have to get a handle on communication.

But bureaucracy, meaning the size of the wheel involved. Take the example of Computer Corporation of America and Digital Equipment Corporation. DEC is an investor of Comet or they have partly funded Comet, and I guess you should know that. But both of us have determined that we will be totally immersed in office automation. Well, we are an organisation of just over a hundred people, and for us to

effect that change was quite easy. It took us, literally, a couple of weeks. Comet existed within a division, and about a year ago the president of the company decided that it was good for all divisions, and we talked to all the other divisional directors and we effected the change. At least we agreed the intention within a week, and it took only thirty or sixty days to convert a lot of work that was being done by internal phone calls and internal memos over to a computer based message system. Well, on the other side of the coin you take Digital Equipment Corporation. Their intentions are very similar. They have all the belief in the world, and yet it will be another three years, I'm told, before they will be able to effect that total change-over. I don't mean to say that in three years that will be the end of their office automation programme. For a significant first step will have been taken, and that is to get most of their organisation running using a computer based message system. And they will also supplement that with facsimile equipment and word processing to a greater extent.

So when office automation will happen is partly dependent upon the size of your organisation. It's a big wheel, and to get that wheel moving is going to take a lot more for the larger firms. But there is another dependency, and that I think is relative to where the effort starts. If you can get some group of top management of a division or even the entire executive office of that corporation and work it down, it's more likely to go faster than coming in at the side and trying to spread the word up and down.

But the level issue is important; it's an old sales axiom, and I caution you on it. You might want to test your legs in not as high a position. But when you are ready to really recommend full implementation, if you can start it at or near the top, not necessarily of the entire organisation but at least of some profit and loss centre like a division, then it's more likely to happen rapidly. Our experience of our users has shown us that the real successes have followed this scenario.

I would summarise for you a little bit about what I see as office automation tools available. I've tried to tease you about some of the pay offs, and I want you to realise the squeeze play that's going on. The pressure is building from several fronts, especially from the word processing front, and especially from the packet switched networks and communications end of the market, and I want to also make you realise it's a big change. I can testify to that. It's a big change which brings on some very interesting results, some of them not all that good too. But most of them are very, very good.

At this point I would like to welcome your questions on any part of the discussion, or things I haven't touched on.

QUESTION: Have you any advice on selecting a test group?

HOLDEN: Well, I think there are a couple of hints there. One is that if you can bring on a group, or identify a group, which tends to be more technological, their resistance to that thing called the computer or that thing called the terminal, is less, and you're going to reduce that hurdle a little bit. And it is a hurdle. Also I think a group which has a natural geographic spread is helpful. A lot of telephone hassle comes from geographic problems. For me to talk to west coast offices is quite a hassle. There is not that much time available in a day to catch one another.

Obviously you would like to get a group which is willing. I don't think that anything works nearly as well if you have to force it, so you want to get some willingness from your people.

QUESTION: You have spoken about data and text. Can you tell us how voice messages fit into your plans?

HOLDEN: First of all, it's in our plans — there is no doubt about that. I tried to take a serious look at my own use of our Comet system, and to see how many of those messages I would speak as opposed to type, how many I would listen to as opposed to see on a TV. I think there is less advantage than might be supposed, but it would certainly be a great marketing situation to be able to walk into this chief executive's office and say, "Yes, we are going to get this computer based message system — oh no, you don't have to type anything — listen, and you can talk. And it is in CCA's plan within a three to five year time scale. There are technological problems there, not insurmountable, and there are to some extent cost problems. The issue of, for example, studying the potential of voice on a packet switched network — that has been addressed by the US Department of Defence and we have been a party to that research, and we hope to implement a capability in two or three stages.

Another area, of course, is the individual terminal devices that might be presented. Not being a supplier of terminals, we have to wait for certain devices.

QUESTION: I think that some of us would appreciate a brief rundown in everyday terms of the sort of functions that a computer based message system might provide?

HOLDEN: First thing I think you should note is that I almost never get an internal telephone call, and I mean this really sincerely. I welcome you to visit our facilities. I bet I don't get more than one internal phone call a day.

A year ago the average was eight per employee, and the average today is two — two internal phone calls a day. So it has dramatically reduced the interruption factor with internal phone calls.

We schedule meetings using Comet. We use Comet as a broadcast tool for such things as job announcements, parties, corporate functions. We use Comet to do all of our weekly activity reporting now. That used to be done by a memo and it has now all been put on Comet. We do travel requests or we initiate travel requests on Comet, and we intend to eliminate a lot of the paperwork there. We do a lot of thinking on Comet. My boss is a tremendous user in this way; he will send out a message to myself or two or three others which says, "What do you think of this idea?" We then think about the topic and we get a lot of pre-thinking done so that by the time we meet in three weeks, our meetings tend to be very productive.

QUESTION: Does it do any filing?

HOLDEN: We recently had an audit and I set this up to see what would happen — we have, of course, a standard set of contracts and documents and all that. But we really support our Comet users using Comet, and we keep a file of all the contract activities in there, and happenings of people coming on and coming off the system, upgrading so on and so forth.

When we got the audit, one thing they did of course was to do a test of a Government contract and a test of a commercial contract, and I said to the audit team, "Would you allow me to persuade you that Comet itself contains all the documentation you need and is a valid audit?" I spent about 1½ hours at the terminal with them, and they were convinced. They requested that we merely show some back-up paper to certain things, but in effect we were able to use Comet to convince them that it was a viable means of filing.

RAY: If I may, I'd like to stop question time now. I'm sure that a number of you have got questions you would like Jeff to answer outside this session, and I know he will be pleased to do so.

Finally, I would like you to join me in thanking Jeff for his presentation, and particularly for his words of advice on experience that his organisation has had in using the system.

COMMUNICATING WORD PROCESSORS AND VALUE-ADDED CARRIERS

Neil Sullivan
TYMNET Inc.

Neil Sullivan is Product Manager at TYMNET Inc. Prior to joining TYMNET he spent eleven years with TYMSHARE, the parent company, as Manager of Engineering Software and Manager of Systems Support. He had previously worked for Control Data Corporation developing method switching systems and data acquisition systems.

Mr Sullivan has a BS from New Jersey Institute of Technology. He is currently Vice Chairman of the Santa Clara Valley Chapter of the Institute of Electrical and Electronic Engineers.

RAY: I would now like to introduce our next speaker, the second of the two speakers talking about computer-based message networks operating in the States — Neil Sullivan, from TYMNET.

A lot of Foundation research — and if you remember Brian Cartwright of the BPO continued this at the last conference — has talked about the importance of communicating word processors as the way into office automation and as a gateway to facilities on public systems. Neil will concentrate on the sort of services that are available to users and which his organisation will be providing to users, which can be accessed through that sort of terminal.

SULLIVAN: Jeffrey Holden provided us with a lot of examples within his organisation of how message switching and automated message retrieval on computer helped him. One thing I might say is that within our organisation, although we do use an awful lot of it, I find that to some extent what happens after a long period of time is that the phone calls start to catch up too. So now you have two streams of things coming at you. One is the sequential messages that are stored on the computer file which you can clear at will. The other is the group of customer contacts and outside contacts that are steadily coming in also. So sometimes it can be a double-edged sword.

In the articles and reports that are written concerning automation, the area of telecommunications is always a key item. Telecommunications is important for communicating word processors, electronic mail, access to centralised databases, things like order entry, record management, reporting systems, commercial data access to databases. There are many databases that are available now through data networks. These include everything from metals and medicine to breeding horses.

These are available through groups like Lockheed, SDC, Dow Jones, the National Library of Medicine, Batelle, and New York Times. There is a wealth of information that is available through the data transmission networks that are in existence today. There is also access to centralised computers through these networks for modelling, problem solving and, lastly, centralised document preparation. There are two major approaches to document preparation:

- One is the central approach where a large mainframe is used and terminals have less intelligence;
- And the other is to use intelligent terminals. There is a variety of techniques in between these two.

During this presentation I will discuss three customers of TYMNET and how they use the public data network to increase their profits and reduce costs. I will briefly discuss TYMNET now. I do not intend this to be any sort of sales pitch; most of the things that will be covered are also applicable to our competition. But I should like to give you an idea of what the network is and what the costs are within the network.

TYMNET is the oldest and largest of the packet networks.

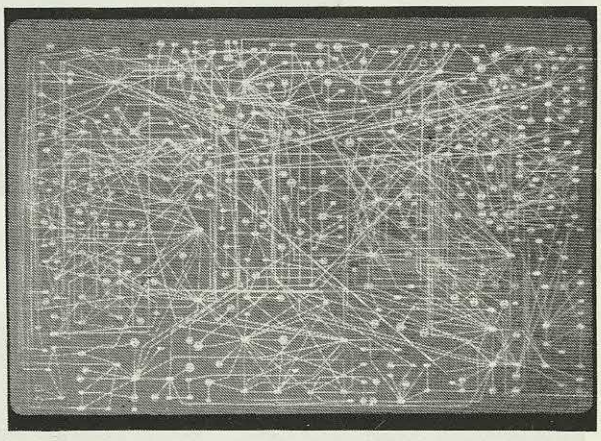

That is a picture of our network. It was designed by a spider! The network contains 410 nodes and has about 170 computers on it.

It can be accessed from about 168 locations in the United States and 22 other countries. It can be accessed from

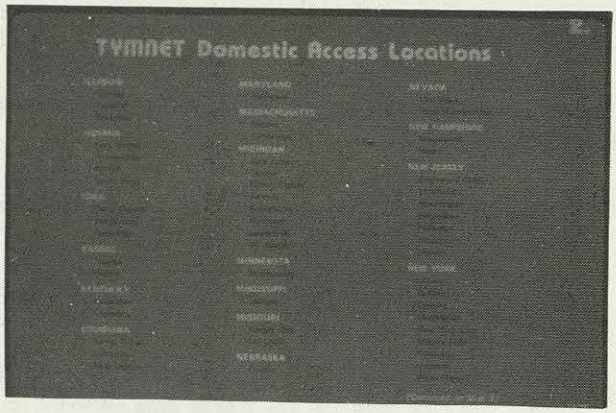

Canada through the Datapac network, from 59 cities in Canada.

Internationally, the networks TYMNET and TELENET are connected by way of the IRCs (International Record Carriers). We are connected to 22 countries by the International Record Carriers.

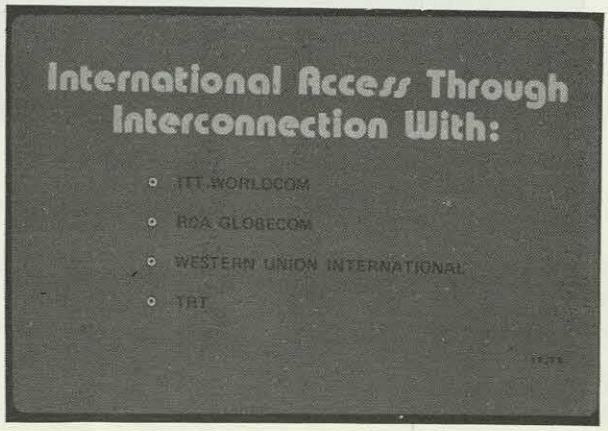

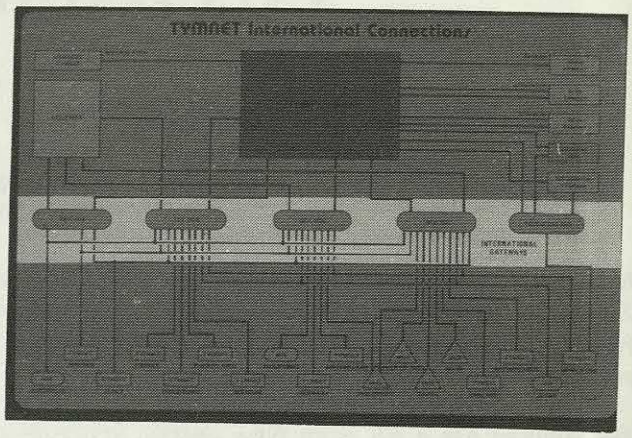

To use TYMNET the connect time that we are looking at is $1 an hour. (The prices for our competitor, TELENET, are very closely the same; in fact it is almost on an application by application basis as to who would win in any particular job).

If you have a host in the United States, you can access it from any high density location within the United States,

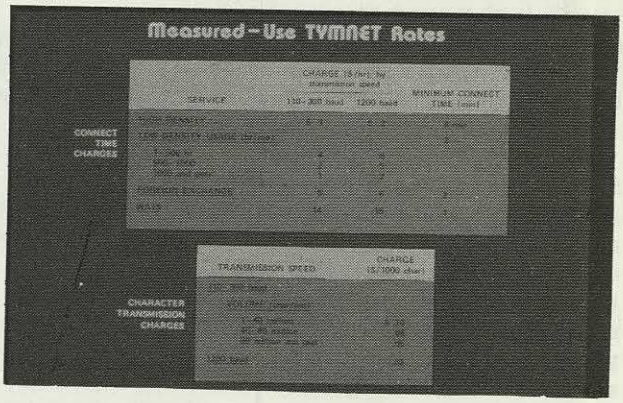

high density locations being the major cities, for $1 an hour. If it is a low density location, the cost is $4 an hour, with a declining scale down to $1 an hour. This is for the normal 110 to 300 baud terminal.

If you use the foreign exchange lines it is $5 an hour, and if you come in on WATS service it is $14 an hour.

The high density/low density combination, that is the combination of areas where we have direct local service, encompasses about half the population of the United States. Neither we nor our competition will cover places like West Sweetgrass, Montana, or any of the real outlying areas. It is a service which is mainly within the metropolitan areas of the United States.

In addition to the connect charge, there is a transmission charge of 10 cents per thousand characters, and that is on a declining scale down to 5 cents per thousand characters.

All of these charges are related to the host computer so the user, in this case our customer, is the host computer. We

charge him for access to his computer. This is different in Europe, where the actual user who is placing the call is charged for the services.

So these quantity discounts that are shown here are quantities that are applicable to the host. In the case of the host computer in Los Angeles that has 500 hours of service from various low density locations within the United States, all of these would be added together to find out if he had exceeded the 500-hour mark and was going down to $2 an hour, or had exceeded the 1,000-hour mark and was going down to $1 an hour.

If you have a lot of traffic you can also choose an option of coming in on 'port pricing', which means that if you pay for one port connecting your host computer to the network,

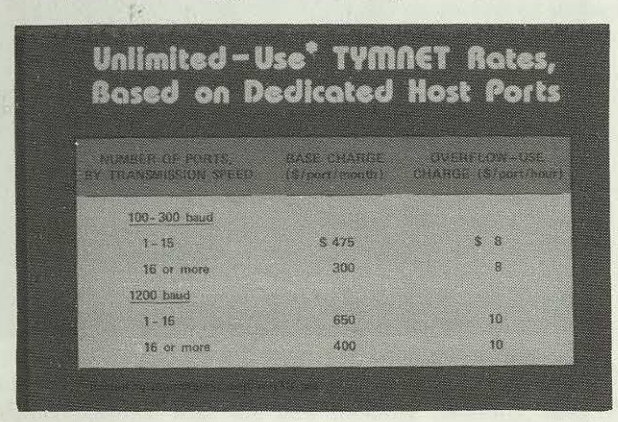

you pay $475 a month for that and you can shove as much traffic through it as you can possibly get in there; it is a single price. You can have more than one port. You can guess at how many you want. There are programs that will estimate how many ports you need.

Both TYMNET and our competition, TELENET, have been experiencing very rapid growth; we have been growing at about 80% a year. The public networks are especially useful for any places that are very geographically dispersed.

Let us look at the benefits for using any data network like TELENET or TYMNET:

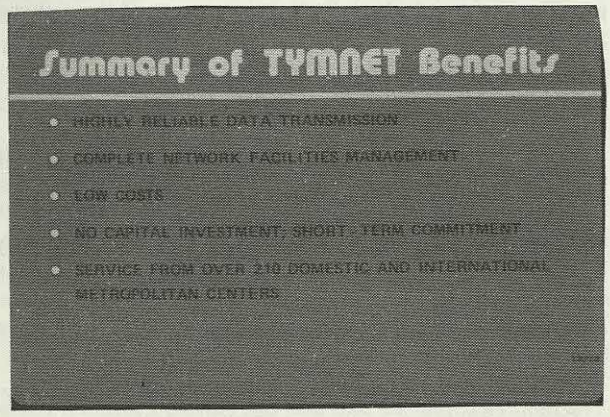

— *Highly reliable data transmission.* Between the computer nodes of the network the packets that are flowing between them are check summed and the probability of error is extremely low. It is calculated that if you are running on a normal line that is running at about 2400 bits per second, the probability of an error in that wire would be 1 bit in 20 years, with the check summing techniques that we use. I guess that on a 4800 baud line that would be 1 bit in 10 years.

— *Complete network management facilities.* That is the carrier manages the network. You do not have to worry about lines and modems and all of the rest of the things connected with those. Low costs. Again, the average user going from California to New York City accessing a host would pay $1 an hour. He would pay about $2.80 an hour in character transmission. So he would have a total bill of about $3.80 an hour.

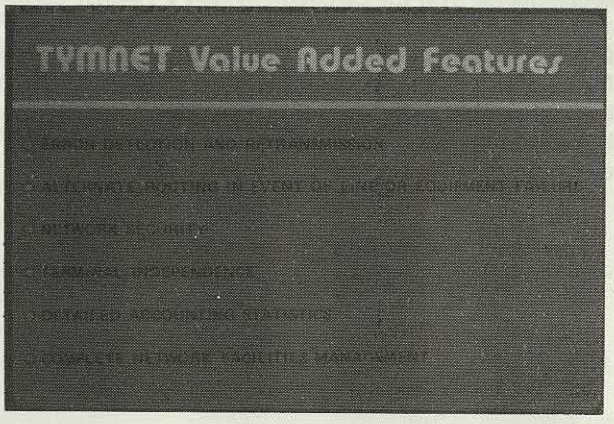

In the case of line failure the call can be rerouted through the rest of the network; just as you can in a Bell system, you can dial again and get a new routing through the network. Network Security. There are passwords, non-printing passwords that are used to prevent access by other people than those that have valid access to the network. So you may have a user name such as Jones to enter the network, and you may have some password such as Ulan Batoor or something like that that it does not print, it does not show up any place; therefore it provides you with a degree of security that nobody is going to use your user name to get into the network.

— *Terminal independence* is something that these types of networks provide. As a host computer you do not have to be prepared to handle a whole bunch of different types of terminals. You can have all of the terminals mapped into one single type of terminal. These can be things like an EBCDIC terminal, or an ASCII terminal, a correspondence code terminal — whichever your computer will handle the best, it can be mapped into that kind of a terminal.

— *Detailed accounting statistics* is a very interesting point. We actually have some hosts on a network where that is one of their primary reasons for being there. Since the network will provide you with a session by session detail of where the user came from and extremely complete accounting details. It will show you the date, the time, the type of terminal the person used, the node that it came from which tells you the location and the country that the person came from, before he came in, how

long he was in, how many characters he typed in, how many characters were sent out to him, and the name of the individual. So now you have not only an indication of how much data or how much traffic is coming in to you, but also its geographical locations within the United States.

This may not be available to you if you are just receiving telephone calls.

Within our company we use electronic mail quite extensively. One of the areas that I use it most extensively for is in communication with the various PTTs abroad. Using the network itself and the mail system within the network, we communicate with the 22 countries that TYMNET currently connects to. Since we have to supply the accounting data and the process of routing which is done through our supervisors for all of these countries, a great deal of communication is necessary; and because of the time distance and the costs involved, phone calls are not really that useful. So we use an electronic mail system.

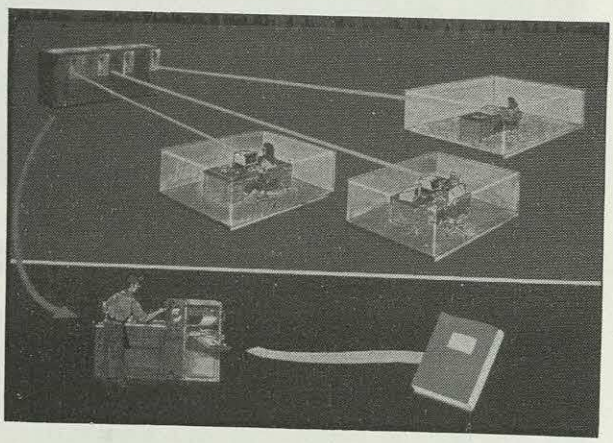

It is extremely handy to be able to come in and just sit down at a terminal, log into the terminal and have all of the messages that are waiting for you printed out, and then be able to handle them one at a time rather than being interrupted all the time.

It is interesting that this message system is used by the PTTs for talking with us, even though they essentially do not like it to be used, and in fact restrict its use for their own customers who would like to communicate with the United States.

The other parts about the message system that are quite nice are the ability to get copy out on terminals that do not look like Qyx machines, that do give you upper and lower case, and that produce a letter quality copy. Also the ability to store things away easily on files. What I suually do with my mail is to print out a copy of it, then I will store another copy away on a very large file. Then once a month I will run a computer listing of the whole stack of mail and file it away, and just to reduce the amount of storage I get rid of the file. But this provides me with a printed copy, which I frankly find that I do not use very much. Usually the time that I do have this stuff on file is a long enough period so that my necessity to go and look at it is within that period of time, and I can just log into the computer and use the editor program to extract those messages that have to deal with whatever subject I am interested in.

The first example user that I am going to talk about is Chrysler Corporation. They developed a system which was called the MOPAR parts connection and it links about a thousand users. It has the status on 150,000 parts. This was a replacement of a telephone system that they had been using. Their normal procedure had been to call one of the 19 parts depots in the United States and place an order for parts. Then that parts dealer would sit down at a terminal and call a central computer in Lansing, Michigan. Now they have distributed terminals to a thousand of their dealers, and the dealers themselves log in to their local nodes in TYMNET and access the computer.

The major advantage that they have found with this procedure is that, by removing the middle man, they have reduced the errors in the system substantially, they claim almost to zero. The system also allows these dealers to log in and check on parts that they have ordered. The system will tell them not only when the order will come through, but from which depot it will be despatched, and therefore the dealer can give the customer a much better estimate as to when his car will actually be fixed.

A large dealer finds that he must use the system about 10 times per day. This means savings of about one hour per day for a dealer of that size.

Another function that the system has is that it automatically updates part numbers. Part numbers change every once in a while, and the dealer is not always aware of these changes. So within this system, when a part number is changed, it can be mapped against the new part number and the dealer informed about the new part number.

You will notice that in this application only 30-character per second terminals were used. It is not necessary to use terminals that are synchronous, running at 2400 or 4800 baud. For a wide variety of these types of applications the 30-character per second terminal is sufficient. It is a lot less expensive and the modems are a lot less expensive. You are not bothered with the problems of leased lines. So here we have a very simple terminal in use.

The next example we had was the case of Florafax. Florafax is the third largest floral delivery service in the United States. They are like Teleflora or Floral Telegraph Delivery. They are in the same type of business. They are sending floral orders to people. They have 12,000 members in the United States and they run about 30 million in sales. In 1977 they began to introduce Texas Instruments (TI) terminals into their florist shops. The TI terminal that they used was the

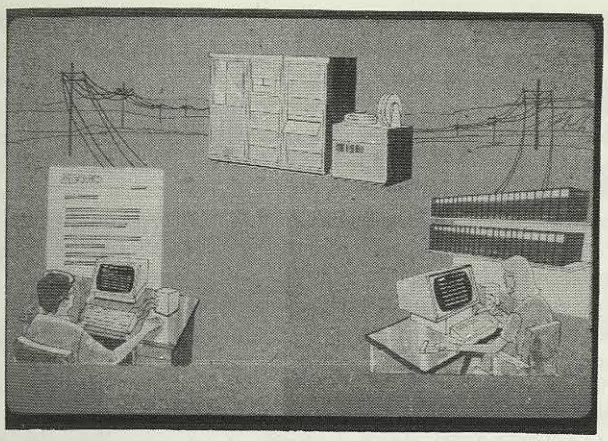

bubble memory terminal, having 10,000 characters of storage in the terminal.

The user sat down and entered the floral orders in an interactive mode, talking to the terminal and answering the terminal's questions. Then he dialled into the network and transmitted the floral orders to the message switched system.

That type of terminal is really quite interesting in that the use of the intelligent terminal cut down substantially the amount of connect time on the network. Also the network could be used for a number of other purposes. By typing another application name, the computer then sent and down line loaded the memory within the bubble memory terminal with the new memory contents, such that it was asking a different set of questions. Essentially he was reloading a new program into the terminal. So now this terminal could be used for many different things. It represents a balance between having a terminal that is more expensive, that has a floppy disc or other type of storage device on it, and a terminal which would be totally transparent as in the Chrysler case, in which the computer was asking all the questions. So here you have a balance between the two.

The terminal is smart, it can ask questions, you can do corrections, you can do anything in it; but it does not have rotating memory or other storage devices that have a higher chance of failing. Instead, the program is loaded through the network into the terminal, is executed, and you can continue to execute the same program over and over again until you have need for one of the other applications that are on there.

Using this system, last year they handled about 90,000 orders during the two weeks before Mothers' Day. I have not checked, but when I left last week they were handling 30,000 orders a day in the period before Mothers' Day. That is the real crunch period in the floral industry.

What does it cost the user to use a terminal like this? To use this TI terminal, Florafax charges their users $79 a month. In addition, they charge them 25 cents per order. This was a considerable saving over the previous amount of 91 cents per order that they were paying on average for telephone calls. So it went from 91 cents down to 25 cents.

In addition, the system had a lot of other advantages. After placing the telephone call, the florist used to have to sit down and fill out a bunch of forms and send them into Florafax so that they could settle between florists. If you sent out $200 worth of orders and you received $100 worth of orders, you had a net imbalance of $100 that you had to send to Florafax and Florafax billed you for it. So all of this paper work had to be filled out and sent in. Now, using the message switch, a copy of all of the orders comes off on a mag tape, it is sent to Florafax and they process it on their computer. Using that mag tape, they can then automatically come out with a settlement between the various florists. So it reduced the amount of paper work and it reduced the number of errors.

So there we have two separate types of message switching that was being done, one slightly more complex than the other. If the florist does not dial in to receive his messages, the computer will automatically dial his terminal after a certain period of time has elapsed. So it does have the ability to call out and get to the terminal.

An interesting point that occurs here is the fact that this terminal is being down line loaded, so we do have information that is being passed through the network to load a RAM in the terminal. Unfortunately, there are a lot of networks that we see now being built that do not have transparent modes. Networks should have the ability to go into a mode in which they are totally transparent, in which nothing is done, in which there are no special characters. You cannot easily do things like plotting, numerical control, typesetting and a number of other functions if you do not have transparency to 8 bits. This idea of reserving characters or reserving the parity bit is absolutely absurd. In every network of this type there should be a transparent mode that allows you to do things like plotting or any of the myriad of tasks like that which do require a full 8 bits of data, or at least are made a great deal simpler by a full 8 bits of data.

One of our users is the Augmented Research Centre which uses the Augment terminal. This was developed by SRI. It contains some unique approaches to document handling. It uses a centralised database based on a PDP-10. I am interested in this particular development because, in 1965 when I was working with Control Data, I was assigned to SRI and installed the first version of the Augment system based on a CDC 3100. So the system has been around for a long time and it has been developing during that period of time.

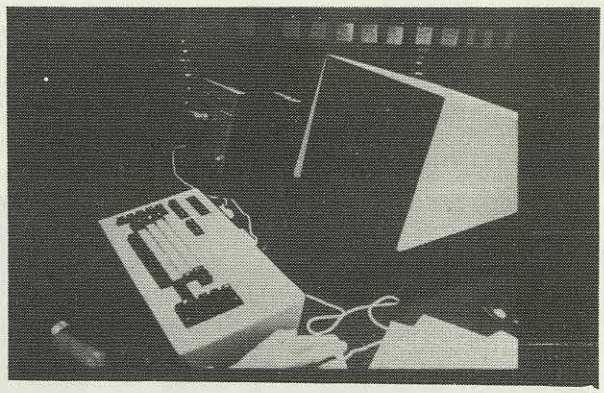

The terminals are of the type that are shown here. It is a 1200 baud terminal. It has a keypad, and a thing called a mouse. On this side you can see a 5-keypad, one key for each of the five fingers. This system extends beyond the concept of just being able to prepare documents easily and distribute them electronically. It allows the more difficult

process of information sharing to be executed. It is a dynamic process that allows a collaboration of people within an environment.

The system differs from most administrative support tools in that the operators are not secretaries or support people, but are the administrators, engineers and scientists that actually create the documents. They do not use the terminal just as a method of preventing duplication of effort, but as an augmentation to the creative process.

The 5-key keyset that is shown there has the effect that you can generate letters. If I press down one key, that's an A; if I press key 5 and key 1 at the same time, that would be a Q. So I can generate any of the letters of the alphabet by pressing down various combinations of these keys. This may seem like a hard way to generate letters rather than just typing them, but the people who become very versed in doing this can work extremely fast that way. It is like striking chords on a piano; it becomes very natural.

At the same time this little mouse unit is used. There is the mouse unit down at the bottom and this causes a cursor to move on the screen. That is how you move round the cursor on the screen.

Using these two in combination, you can edit text extremely fast. The little mouse also has three keys on it that can be used in conjunction with the five to give you sets and supersets, so that 93 different characters can be generated.

The file structure that is used with this system is a hierarchical file structure. It is like a tree structure. So in addition to the normal file structure in which stuff is just linear, you can go through this process looking through your file and just look at headings, or just look at level 1, just look at level 2, or just handle it like a normal text processor. So there is a lot of different flexibility involved.

Since it uses a very large central database, you can have as elements of your file a 10,000 page report. You can extract things from a variety of reports, merging them all together. The screen itself can be split into eight separate pieces and you can work on various sections of different reports, or various sections of the same report, bouncing from one section to another.

It eliminates the idea of doing things sequentially. It allows the user to create text material as he really thinks. You think of parts of structures and you think of other parts of the same subject, and this allows you to bounce back and forth between them.

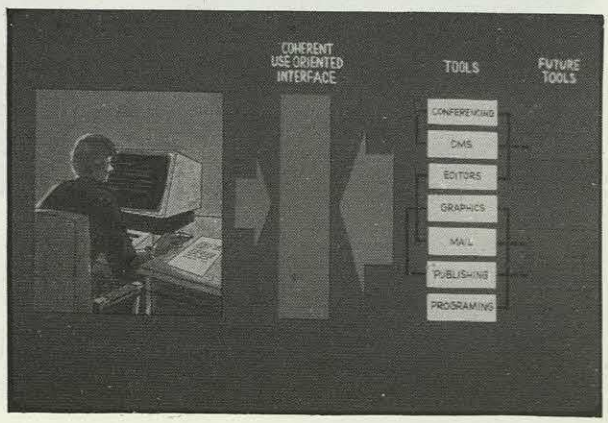

So in addition to creating this text, we can now go through a bunch of different processes using it. It can be used as part of a mail system which is within this system. It uses graphics. There is a really neat calculator mode that I saw demonstrated, in which you can display a table that you have just created. Maybe in the course of writing a report you have created a table, and you can now take your little cursor and aim it at a column. You can type on a thing "Multiply this column by 15%", and the column just changes in front of you. All of a sudden it is increased by 15%. Or you could ask it to total a column. Or you could give it a whole series of instructions on how to work on this report, and it will just happen.

The recent report claims that this system and the Xerox system, both of which are bidders in the executive offices to the White House, bear a great deal of operational resemblance to each other. This is probably very true since several of the same people worked on both systems.

The Xerox system is based on a minicomputer, with video terminals, which should enhance the speed over this system. This system does have a tendency to be kind of slow, although the centralised database may have more features than the minicomputer versions. I wish that I could be with you later this week when you go to the White House and will be able to see the system operate.

David Butler pointed out in his presentation that networks such as ACS, X10 and TYMNET provide only a very basic form of connection. This is very true. Systems like this are really neat, except that I as a user of these kind of systems would like them to be able to communicate with the other systems that I work on. I have databases that operate on other computers and it would be neat if I could use this kind of a system to grab information, pull it back, and then work with it on a system like this.

Those types of processors do not yet exist; and those are the kinds of things that we should be working on to make systems and automated offices totally complete in their scope.

RAY: If there are no questions, I will conclude by thanking Neil for his presentation, and particularly for looking at those three case studies of how TYMNET has been used by organisations.

DEVELOPMENTS IN EQUIPMENT AND FACILITIES CONCERNED WITH THE AUTOMATED OFFICE -2

Tony Mallia
Wang Laboratories Inc.

Tony Mallia is the Product Manager for Office Automation at Wang Laboratories Inc. of Lowell, Massachussets. He was previously with Sanders Data Systems of Nashua, New Hampshire in product planning and development, and prior to that was in sales in IBM (DPD) UK.

He has a BSc in Physics from Imperial College of Science and Technology, London University.

RAY: I would now like to introduce our fourth speaker this morning, Tony Mallia. He is British and studied in the UK. He worked for IBM in the UK before coming to the States. He is going to talk about specific aspects of the office and office automation, particularly looking at the creation, filing and distribution of documents.

MALLIA: I should like to cover three aspects in this presentation. The first is to look at the office itself. Secondly, I should like to cover a little bit of the technology and what we are supplying from Wang; and thirdly, to look at some of the management and development functions in management that may have to occur before one would truly get to an automated office.

I should like to discuss people. My belief is that the approach to the automated office is by understanding the office. If we understand what goes on in the office, then we have a good chance of applying the right solutions to it.

People are in the office. The office is a place where people work. There are some projections in the future that maybe we will not have the office, that people will work at home, and that the office will be an information processing plant. I think that is a long way off, although we are seeing signs now in the United States that the cost of the word processor has dropped to such a point that the monthly rental of the unit is less than the difference in salaries and wages for paying somebody to work at home, as opposed to paying somebody to work in the office. There is a considerable interest now in using the vast resource of people who have to stay at home, housewives who have to look after children, handicapped people, for doing some of the typing work that has to be done.

The office is the environment for information transfer. It is a place where people communicate. They communicate about lots of different things. Generally people like to be in the office because they like to communicate. They like to have arguments; they like politics; they like having discussions. It is a person to person environment. It is a place where people can fulfil themselves in terms of their own personality, in terms of their own presentation styles, and so on.

Let us look a little bit more at what happens there. This idea of the structured and non-structured was discussed earlier. I should like to go into that a little more. The first is

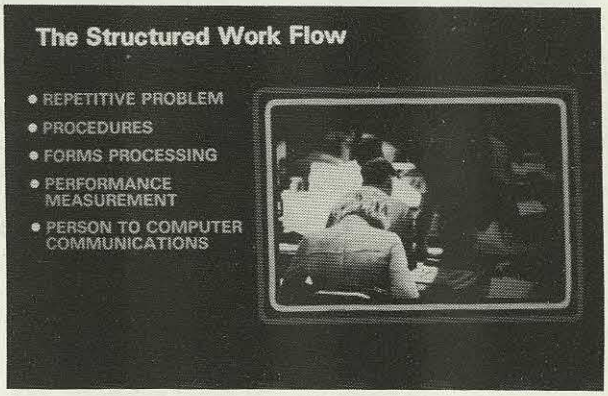

really the structured work flow within the office. This is where the office is part of the business environment; that it is a repetitive problem; procedures have been defined by management of how that problem is to be dealt with. The information itself tends to be structured into elements, and therefore we get forms used in this type of environment.

The form is the control in which the information can be structured so that you know you will have a name, that there will be an address, an age, date of birth and so on.

It tends to be reasonably easy to measure the performance within this environment, because you can have a repetitive unit going through and how many forms you process in the month, or how many orders you take, or whatever it is, is a reasonably easy thing to measure.

The automation aspect of this type of work flow tends to become a person to computer communication. It tends to lead to the interactive work station — whatever sort of

terminal that may be — in the office, with people making enquiries or doing data entry functions into some computer. That computer may in fact be a local one to the office or a departmental computer, or it may be a centralised, large mainframe; but it tends to lead us towards that type of communication.

The sorts of industries and business for which this is particularly suitable are the businesses which themselves are structured and part of a machinery — the banks to a great extent. Therefore you will find, when you look at office automation, that a lot of the banking functions readily fit into the person to computer type of communication.

However, as we move up within the organisation or we move into different types of organisation, we tend to have a looser work flow. The looser work flow is because generally

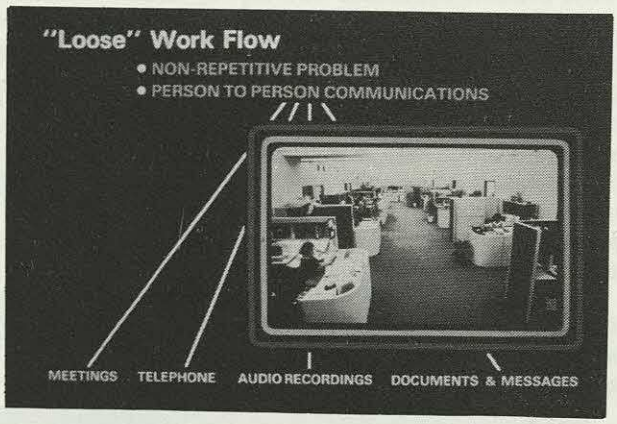

the problems are not repetitive, they have come up for the first time; and in order to solve these problems one has to go to somebody else to say, "Do we know about this problem? How do we sort it out?" So the emphasis of the office activity tends to person to person communications.

This breaks down into four major communication modes:

- meetings,
- telephone,
- audio recordings,
- documents and messages.

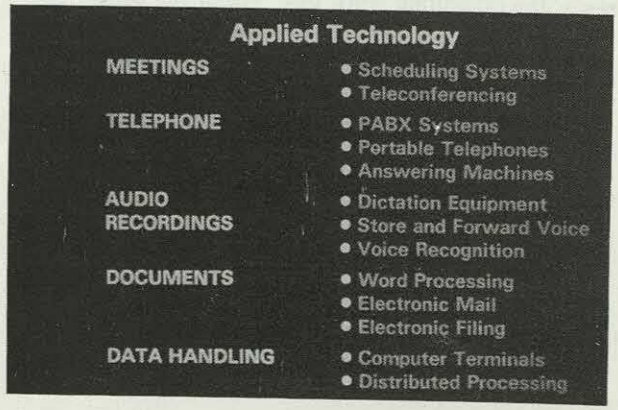

Let us look at those four areas a little more closely. The meetings and telephones are what we call synchronous communication; in other words, both parties have to be present at the same time. We are now in a meeting. I had to travel here; you had to travel here. The communication is a verbal one. The expense is quite high. And we have some problems, as mentioned earlier, scheduling meetings and getting people into the right place at the right time.

The telephone is the same. Both parties need to be talking at the same time. There is the situation of the connect problem on the telephone which is running in the United States now at about 70%; in other words, there is a 30% chance when you call somebody that they will not be there. There is also a 30% chance that when you have left a message and they call you back, you will not be there either.

What we see very much here is a movement out of synchronous communications because of the scheduling problem into maybe documents and messages. This is where maybe the computer based message system, as was described earlier, fits in with the frustration of getting through and scheduling and finding somebody at the other end of the telephone.

Both audio recordings and documents and messages are non-synchronous communications. They allow the work to be queued up and for the person who is receiving the communication to take it at the time that he wants.

Let us look a little bit at the office activities now and what is going on. The studies that are currently around right now — and it will vary depending on who you read and how you want to take it — you do not need to take these figures as gospel, although they do seem to indicate what is happening. That about 75% of the secretarial and clerical

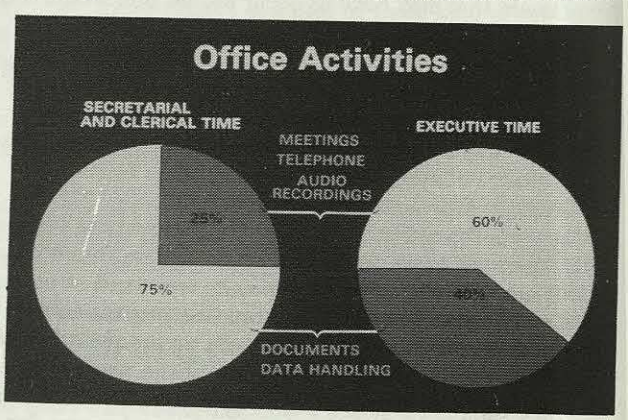

time seems to be spent on documents and data handling; 25% on meetings, telephones, audio recordings. With the executive time, however, as you go up the organisation, 60% is spent in the personal type of communications and audio recordings and 40% on documents and data handling.

So if we were to approach some of the automation aspects of the office, we might want to see a different approach based on the type of activity — a different approach for the executive or a different approach for the secretary.

We had the five areas, if you remember: four in the loose work flow, and really the data handling which is the structured work flow solution. The sorts of technologies that we have around in the market now are indicated on the right. For meetings we are talking about calendar management systems, for scheduling meetings, for getting people to find out what time to call the meeting, to find out whether there is a conference room available. There are some interesting developments in teleconferencing, which is basically aimed at competing against the travel industry.

With the telephone system we are seeing intelligent exchanges on the market, particularly in the US; obviously with the interconnect situation and the regulation of the European PTTs the situation is somewhat behind in some cases or is more of a monopoly situation. In this country we are starting to see a big growth in portable telephones. And we have always got the answering machine. This is one of the solutions to the 70% connect problem.

In audio recordings we have dictation equipment. We are starting to see some store and forward voice networks being built up. And maybe down the road — although the technology is here and now — voice recognition which can, at a reasonable cost, transfer voice into data or character text.

In the document area we have word processing, electronic mail and electronic filing. I will be dealing with that in a little more detail. In data handling, this moves into the person to computer communication. We have computer terminals moving into distributed processing for access to local data in the department, and for response time considerations when you are looking at the office.

Let me just go back to that. Just to explain a little bit about where we are in Wang, we have about half our products split on data handling in terms of distributed processing, and about half within the word processing area. We also have a subsidiary company, which is called Wang Voice Communications, which we recently acquired. They produce products for radio broadcast stations right now.

They have a particularly interesting device which is called a torque tunnel. The torque tunnel is a semiconductor device that gives you a six-second delay for real time broadcasts. So it is just enough time for somebody to push the button to kill the swear word. This is selling particularly well right now. They also have some radio paging units.

I am going to concentrate, though, on the document side.

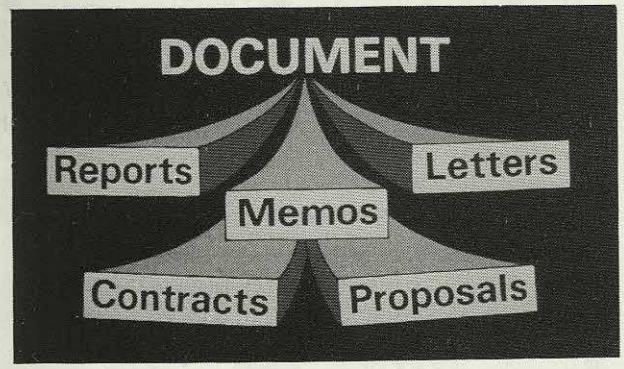

Let us look at what documents are. They could be anything. They could be a letter, a memo, a report, contracts, proposals. Let us see what happens in this area.

The way we look at the functions in the office is to translate this into what we call the document life cycle, which is all the stages through which a document moves from its original creation to its destruction. This is a particularly appropriate way of looking at the functions in the office because, given a single document, you can apply costs to that document for each of these stages. It is analogous to some of the standard costing functions found in manufacturing.

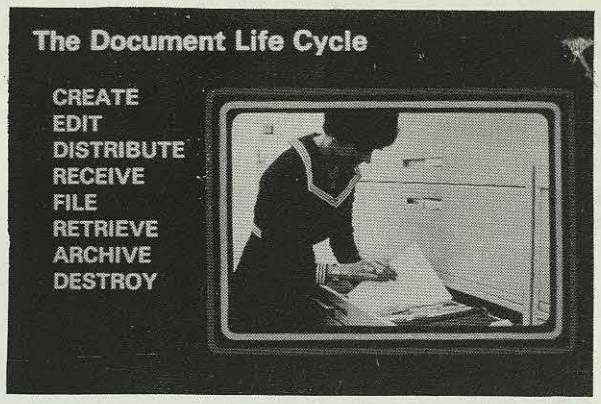

Let us quickly look at what these costs might be. To create — we are talking about author time primarily — you could work out, based on technical writers, some idea of what it costs per page. My figure was that it could cost something like $25 per page to create a document. The editing function is really a combination of the transcription on the part of the typist, and the review cycle through which that document goes before being in final proof form for distribution. That, depending on the type of document, could be of the order of $4 to $5.

The distribution function or the combined distribute and receive is probably another $4 to $5 on top of that. There are some interesting studies now on internal mail distribution which show that the handling cost is particularly high. Stanford University has done some studies recently, showing that an internal mail memo through the campus costs them $5, so that postage costs are obviously incidental to the cost of the actual mail and physical handling.

Filing and retrieval. A document will go through multiple filings mainly because, at the distribution point, it has been replicated. The studies from the copier industry right now indicate that the mean copying run has gone up to something like 10 times. So that for each unit that we did a creation for and an edit, we are now talking about 10 filing functions. We are also talking about the multiple retrieval and the cost that it takes to retrieve a document from a file, let alone the cost of failing to retrieve the document from the file, which could be very significant by the time you have put in the time taken to find out whether it existed, who has it, and where it is now kept.

Archiving and destroying. Depending on the security or the classification of the document, the destruction can in fact be quite expensive. We are even looking at the document cost of recycling paper or whatever costs to do with the disposable material.

Let us look at where word processing fits in. The very

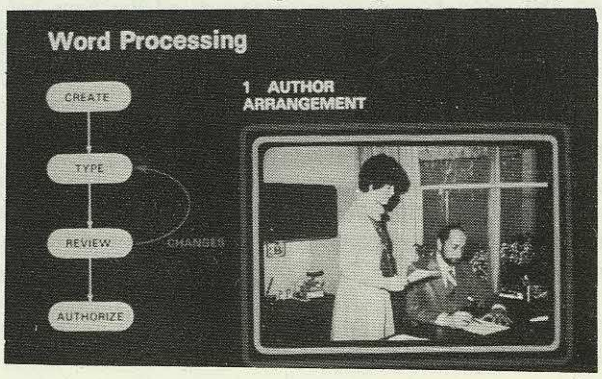

first stage of word processing was really aimed around the letter. This was particularly a one author/typist arrangement whereby the letter was created. Now you could argue that if the author got it right first time and the typist got it right first time, then the number of changes is minimal. That does not happen to be the case. But more significantly, we find that word processing fits into the organisation more at the multiple author arrangement.

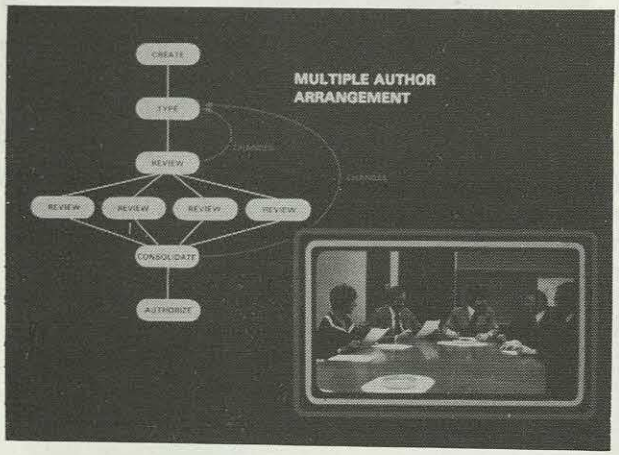

The multiple author arrangement comes in at a more significant document, when it reflects a corporate or company position; and that the review cycle of that document might include legal, technical, and other department management. You can quite easily see in this cycle four or five cycles of revision in that document. Proposals, contracts, that type of document, really do cycle through there; and not until those cycles are complete does it get authorised.

The justification for word processing is based on the unnecessary page retypes. Those pages are retyped either

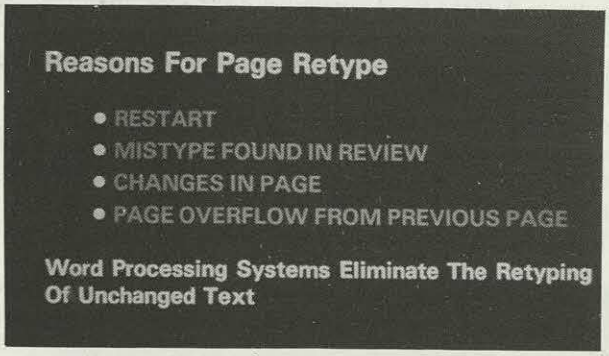

because the operator made an irrecoverable error; there is a miss-type found in the review process; there are changes in the page, which can be very, very significant, particularly with the multiple author arrangement; or the page has overflowed from a previous page. Somebody has inserted a paragraph on one page; it could go through to the next 10 pages.

You can work out some factors in determining the cost justification of word processing based on the type of the document and the number of cycles that it goes through in the review and the percentage of pages that need to be retyped.

But let us move on a little bit. I just want to show you what products we have in this area.

A quick sales pitch: that is the only sales slide. We have a range of word processing units, which range from a single station, a minimal system selling at about $5000 in the United States, to significant clustered systems which, with our office information system, will now go up to 24 work stations or 32 devices.

The system has three basic components:

— first, the work station, which is a simple keyboard and a CRT, 24 lines, some status on the top two

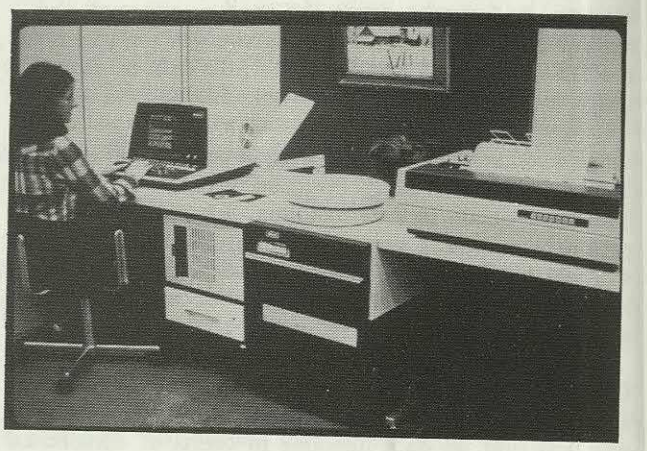

lines to tell you where your positions are;

— second is the central storage which is shared amongst the multiple work stations, which can either be diskette based or can be disc based. We now have products which range from the 5, 10 megabytes of storage up to 160 million characters of storage;

— the third element is the printer. I will be showing more of these products further on.

Let us look at the work station first. There is an interesting problem that happens in the office that is not common with data processing. It is the human factors requirements for a secretary or an operator of a typewriter moving from a typewriter to a word processing unit. The significant things are, firstly, that the computer languages and the commands are not familiar to that operator; and secondly, and more importantly, which will affect networks and the communications area, the response time that is expected from the operator is far faster than the normal data processing user. When you hit a new page key looking through a document, you are expecting to get the new page in about half a second or three-quarters of a second. The interesting thing about the human factor engineering in the office is that it creates totally new architectures. It creates requirements for local processors and local disc storage in the office.

We have various functions that can be performed. It has all the basic editing features. There are global searches where you can search through the whole of the document for particular words or character strings. There are global replaces which allow you, if you have miss-spelt a name through the whole document, to find it and replace it in every instance. We have various human factors to help you as you go through some critical process like deleting text, to make sure that you have got it right first before you actually commit it for deletion.

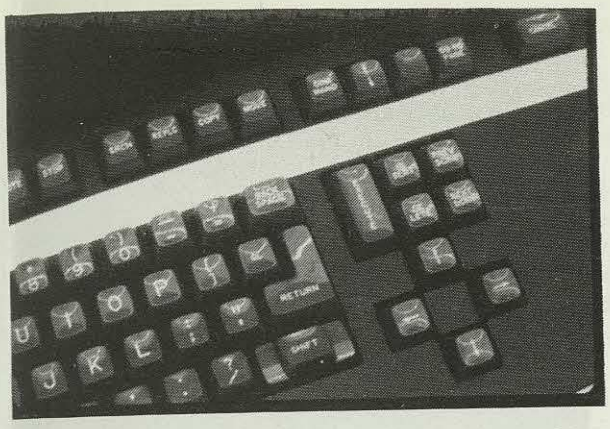

It is a simple keyboard. It is very similar to a typewriter. It has a few extra keys down in the main body, but you will notice on the top layer that we have the keys specially oriented towards the functions associated with the office: to search and replace; copy text from a particular point; move text. There are some extra commands. GO TO PAGE as a single key operation. This is the sort of thing that an operator expects.

We also have the ability to merge text from two documents, particularly for creating form letters where you have a name and address list on one document, and the standard letter goes out on the other.

I talked about repagination earlier, about the overflow. We have automatic repagination, paragraphing, page numbering, titling, and footnoting. These are the things that are generally required for such a system.

This is a very interesting feature. This is the closest to programming that we allow a user to get. The glossary is an ability to store key-strokes which would normally be keyed on the keyboard, as a procedure, and to call them with two key-strokes. One key hits the glossary key, and the second one is the index key to that stream of characters. You can, for example, put in your own name and address under a single key. You can put in a whole document or a whole form under a single key. This can become very extensive, because you can automate a whole lot of the procedures of converting, moving, and re-editing documents under a single key-stroke; so that the operator hits that key and watches the system do all these functions.

The central storage unit can vary in size. We also have diskettes which allow you to archive documents and hold

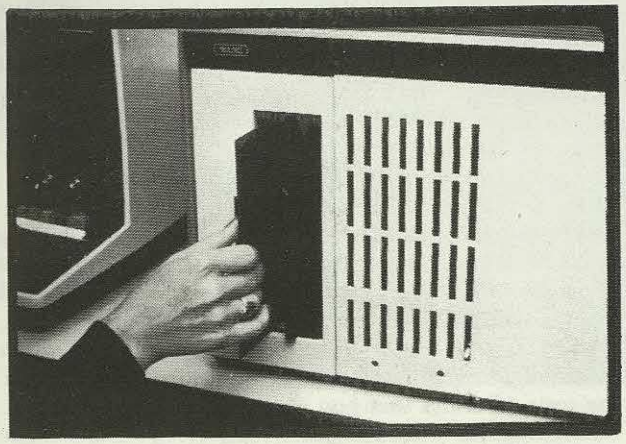

them off line. This allows you to move documents between different systems so that two operators working on two different systems could in fact share the same document, maybe working on different areas of it.

The printers start to become interesting — the whole range of printers here.

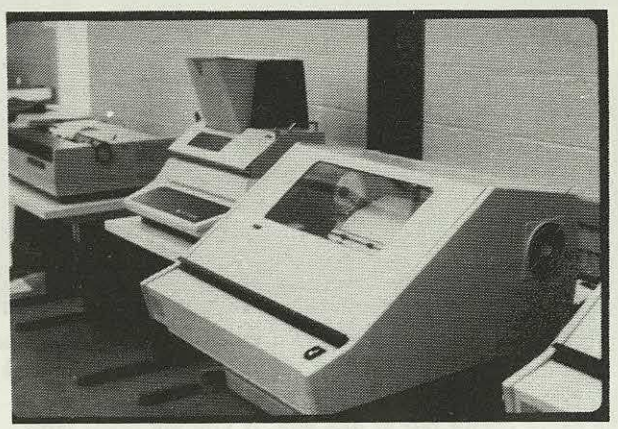

The most recent one that we have announced is called the 'image' printer. The image printer is a fibre optic device.

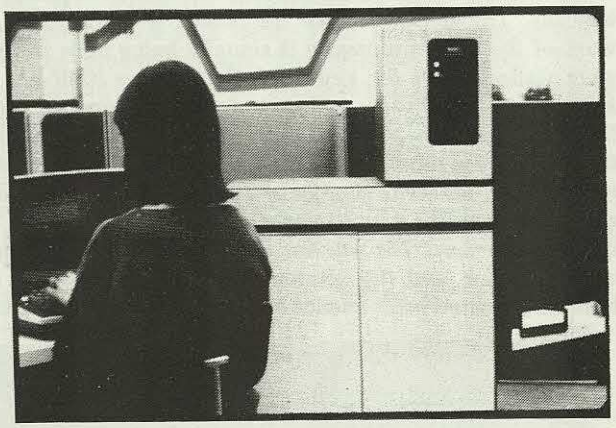

The fibre optic bundle sits on the face of a CRT; in fact the CRT is in the column on the right. Basically the fibre optic bundle is able to convey the resolution of dots that would be on the phosphor directly down to a copying device, a plain paper copier immediately under it. This device forms characters as a dot matrix of a density of 300 dots per inch. This type of quality makes it almost exactly the same as the high quality daisy wheel printers or other types of impact printers that are on the market.

The other advantage of this is that it is able to do this at a pretty fast rate. We are able to generate about 18 pages per minute off this type of device. It just feeds the paper through, images it on to the paper, it goes through the zerographic process and comes out in the stacker at the end; and there is no impact printing at all in that. This is particularly good for multiple document origination where you are going to generate the multiple copies immediately from the system, because of course it sorts them already, because you have copied one document through from page 1 to page 13 into there, then you are doing the second copy page 1 through page 13, also into there, so that it is already collated by the time that it comes out into the stacker.

This is the other area of printing. This is a typesetter which, just like another printer on the system, allows you to output

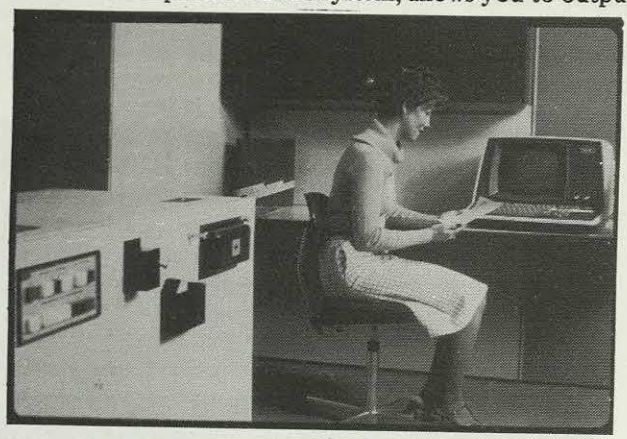

a document straight across into the typesetter. The interface between the two allows you to translate commands within the document into photo typesetting commands, and the paper comes out at one end. If you are doing large production work which requires that sort of quality of printing, this allows you an extremely versatile approach to editing documents and creating them to deadlines.

The printers and all the functions work at the same time basically. The architecture of the system is that the real work of the word processing is actually being done in the work station under the keyboard. Each device itself has a microprocessor inside, and each device performs its own function. Therefore the more work stations you add on to the system, the more capacity you have for meeting that requirement. The central units really support only the disc access. So it is like a multiplexed disc going on to the individual devices. The printers themselves also have their own processor, and the printer goes and gets all its information from print queues on the actual disc or diskette.

This is another feature which we have on the printer, which is a dual forms feeder. It allows you to put in two forms at the back, one form which would be your heading page of your document and others being subsequent pages. This is all cut paper. It feeds it automatically through the printer and comes up into the stacker which is right at the front. It looks rather like a mincing machine to me.

All these systems get combined and figured together into a total system. This is in fact one of our later office information systems. We have two product ranges: the word processing system and the office information system. The office information system is leading us more towards what

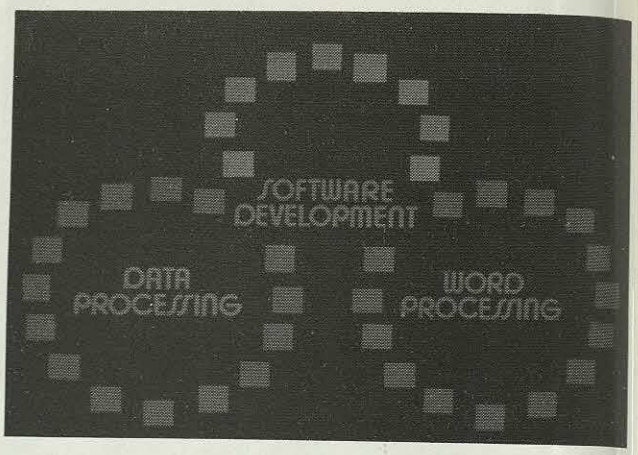

we call an integrated approach, where you can mix data processing functions together with word processing from the same work station. We are approaching this from two directions: from the word processing angle of adding computer languages on to the word processor; and from our computer line in the distributed processing area of coming into the word processing.

The interesting thing at this point in time within the product line is that the work station and the work station technology is identical for both our major computer line and the word processing. So in the near future you will see some major integration of those two functions together.

Now on top of that, we also have telecommunications devices. Currently, 50% of the orders that are taken for

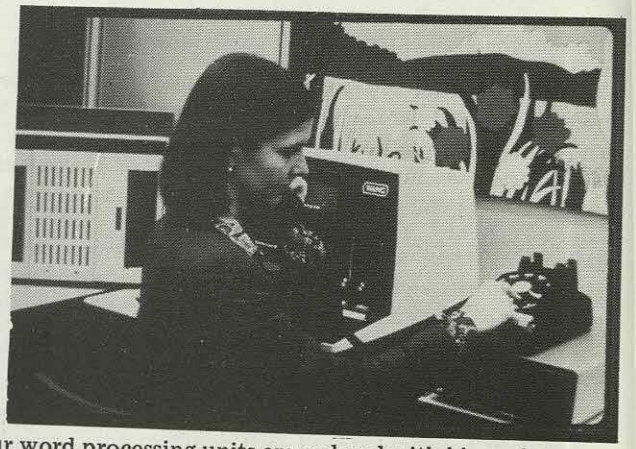

our word processing units are ordered with bi-synchronous communication. We have seen about a 900% growth in the last eight or nine months of telecommunications within word processing. We have found on studies that about 60% of those people who have telecommunications are using it for data entry, remote job entry into central systems. There are some interesting applications there.

There is one that I should like to mention, where a programming department is using one of our word processors to create and edit source Cobol programs. The cost per hour of using a word processor work station like that is something like $3. The cost of using an interactive terminal to, say, a TSO environment is about $20. The editing functions on the word processing are extremely good and give you all the sorts of things that you would need for generating a Cobol program — like the global replace; if you have a label incorrect you can go through and re-name it all the way through the program.

The other advantage is that the response time is very fast. The whole system is designed for the half second/three-quarters of a second response time to the operator, so the human factors are very much improved.

We see some interesting applications where we did not expect them. There is a lot of straight data entry being done on word processors, where the secretary or the clerical worker would have filled in a punching document and is now translating that directly on to the screen. It is interesting because the screen is in the same location that a lot of the typing has to be done as well; it is in the same office environment.

We are moving closer to some degree of integration between these two functions. In some cases people feel that this is office automation. We really see this as a triangle. There is a word processing function at one point, and the data processing tends to divide itself into both central and distributed data processing. We feel that you will find a mixture of all three. With the telecommunications options that we have, we have batch communications into host systems. We also have on one of our product lines 3270 compatibility right now, and we are developing that across all the product lines, to give interactive access to central computing. So that mixture can be provided from the single work station within the office.

That is the sort of technology — that is just about where we are today. Down the road we see a whole lot of developments in terms of better human factors, more graphic capability and so on. However, as was mentioned earlier, we seem to have all the technology we can handle right now. We have arithmetic functions on the system. We can add; we can do all the major multiplication on the word processor. We can perform invoices, do order entry. We can sort columns within the word processor. We have multi-language capabilities that allow you to have the character sets that are appropriate to the different languages. We also have the necessary security on the system to prevent somebody else getting in to the documents.

We have a lot of technology, so what is the problem? I think that this sums it up.

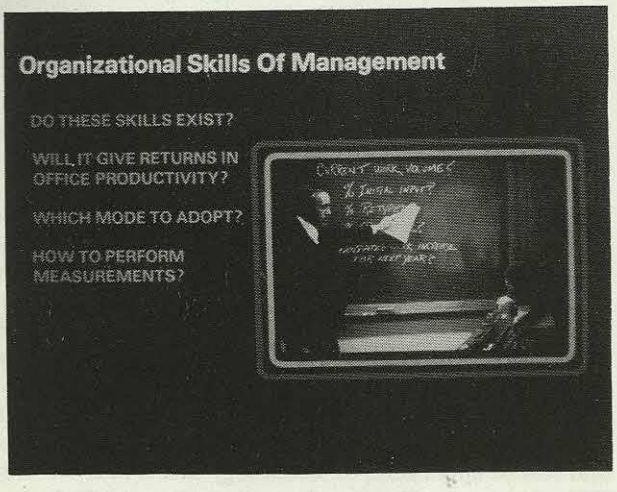

Organisation skills and management. Where is that management? The management is primarily in the office. It is the department manager. They are running the show right now. The question is: what skills do they need and what help do they need in technology? How does the management services help in this area? Do the skills exist? Probably not. We are finding that even to install word processing and make that productive, a lot of the skills are still missing and people's understanding at ground roots of what word processing is.

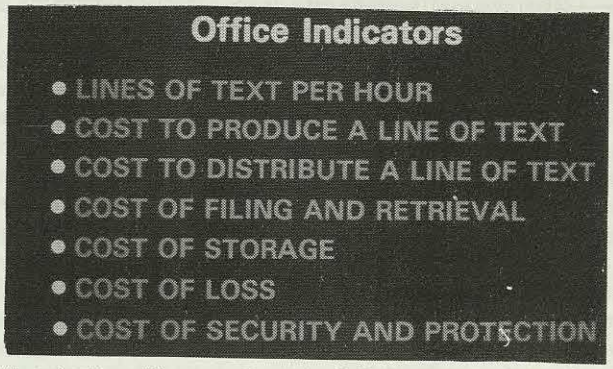

What is the office productivity? Is it even measured? Do we even know what it is right now? Which mode to adopt? Are we going into a structured type of work flow or are we keeping a loose work flow? There are two analogies elsewhere in the organisation. The structured work flow was very much the manufacturing area. We saw the production line with Henry Ford, and the consequences of the production line. The division of labour. The specialised use of people within that system. We are seeing some movement away from that in some of the large organisations over here, towards project or total vertical functions performed within an office — the team basis. Just like some of the developments in Swedish car manufacturing, we are seeing the same concept starting in the office.

What about the loose work flow? How do we perform the measurements in the office? Manufacturing was very measured to start with. The key breakthrough in manufacturing was perhaps some of the standard costing methods, which allowed us to track back and find out what a particular function cost, what a particular part cost to manufacture.

There is another area in research and development which is really the one-off problem. We see project control and time analysis being used in those areas. But these concepts have not yet arrived in the office. So we are not yet ready to perform this measurement and get the direct trade-off of whether or not it is an advantage, whether or not it saves anything. We are still based on two equivalent models, one without the technology and one with, to try to find out whether we would save anything.

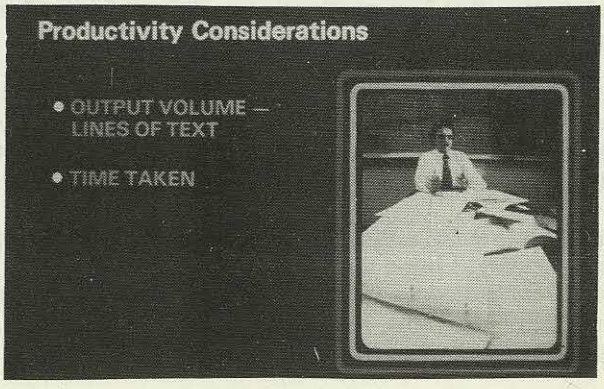

What is this productivity consideration? Let us look at pure output. Let us take it as lines of text. Let us look at the time taken. The time taken by whom? The author? The typist? Who is it? It is a very important aspect, that if you take the lines of output the major considerations for productivity are, firstly, how is the work submitted? Is it submitted as a handwritten document? Or is it submitted as a cross reference to other documents that already exist? Or some standard structure, boiler plate type of operation? Or standard paragraph type of work?

What is the extent of the changes? We went through the revision cycles in the multiple author arrangement. Each time that goes through revision, it is actually decreasing the productivity of the unit. So how can we minimise that revision cycle? Only last do we really come to the operator skill in terms of the number of key-strokes. We find that if you can submit work by using reference in existing material, the word processing systems as I have described are a very sophisticated cut and paced machine. If you can get to that area and that level of use, then you are really increasing the productivity.

Further Considerations

- HOW IS WORK SUBMITTED?
  - REFERENCED DOCUMENTS
  - BOILER PLATE
  - STANDARD PARAGRAPHS
- EXTENT OF CHANGES
- OPERATOR SKILL

Some Productivity Factors Are Under The Control Of The Author

Should The Author Be Measured In Document Productivity?

That takes us to the question of where should the word processing unit reside? I think that the movement in the future will be back towards the author. At the moment it tends to be put into the word processing centre. We are seeing also the movement into the administrative secretary. I think that what we will see after that point is first runs being done in the word processing centre, with the work being passed back to the administrative secretary electronically for the fast review cycles. Finally, after that point, since the author is used in the review cycle, we will see the beginning of the executive work station in terms of reviewing and correcting documents. I actually do all my own typing on a work station. I find that I can create and think directly on that work station, type the information in, go back and correct it immediately, and it is ready for final printing. So I can generate something of the order — depending on how much reference documents I can get at at the time that I am creating that — of about four or five pages an hour, created in this manner.

The interesting question is: what is the psychology in the office and when will people recognise that they can in fact type? And when will executives learn to touch type? That is going to be a very interesting time.

I believe that the author has to be measured in the document productivity. I think that we are kidding ourselves with the word processing systems now to totally measure key-strokes of the typist. In fact we do not want the typist to do too many key-strokes.

With the glossary functions and these sorts of automation, the operator can hit two key-strokes and wait for five minutes while the word processing system actually does the work. So some of the measurements of key-stroke productivity are actually contrary to the direction that we are looking for, which is output.

We can use a set of indicators in the office. These may be things that we start to measure and start to get some indication of. An interesting one here is the ratio of the number of key-strokes keyed into a document against the lines of text produced in the final document. This gives you an idea about the efficiency of the way that the author has given the information to the typist; also how many review cycles it went through.

Time Analysis

- ROUTINE MEASUREMENTS NOT SAMPLE
- ALLOCATE TIME BY PROJECT NOT ACTIVITY
- FEEDBACK OF REAL COSTS TO MANAGEMENT
- FUTURE PLANNING ESTIMATES
- AWARENESS OF HIDDEN COSTS

I mentioned time analysis. I believe that this is going to be the most significant in getting the office to some stage of automation. I believe that it is the equivalent of the R&D group where people in fact account for their time within the environment. They account for their time by project, not by the fact that they have answered the telephone but by which project they have been working on. This gives a feedback of real costs to management. This is what it costs in the office to produce a particular report of this sort. It therefore gives future planning estimates back to management in terms of what it really costs them to do that, and do they want to do it this time? It brings out the hidden costs into some real time.

Some cases of application of technology result in the movement of costs from visible costs into hidden costs. I would like to give you just one example. Some of the advanced telephone systems. We have just installed one, and it does optimisation of the outgoing calls over the long-distance trunk lines. The way it does this is to queue the outgoing calls with the telephone number, and when you have got your outside line it calls you back with two rings. It has stored the number that you wanted, but it calls you back with two rings. You pick up the telephone and you hear it dialling away outside.

The interesting thing about this is that, firstly, it takes about 21 digits to dial the system, first of all to get out, followed by an average of something like a 20 or 25 seconds wait, when you do not really want to do anything because you know that telephone is going to ring. On the other hand, you are seeing the visible cost reduction of the Bell charges at the end of the month. What you are doing is maybe transforming those costs into a non-productive time that is spent by all the people who make the telephone calls. It is that sort of application of technology that is likely to move. You can cost justify it on the saving of your telephone charges, but you do not know what it is doing to your internal efficiency.

I should like to finish with this slide, which is that the real productivity will arrive when we know what we are dealing with in the office; when we know what it costs us to produce something; and when those costs actually become real in terms of management's awareness and is not the general overhead of office administration. I think that this will particularly happen in the loose structured area. I think that the structured work flow is probably fair game for data processing and we will continue to see interactive work stations and that type of terminal within the office.

QUESTION: Do you recommend stand-alone or shared logic systems?

MALLIA: That is one of the big questions. We have both stand-alone and shared logic. The criterion between those is really one of the fear of failure within the system as against the cost effectiveness of sharing peripherals. For example, if you have individual units you probably want to share some printing device or you are going to have multiple printers on each system. As you start to replicate those peripherals across the stand-alone units, you are putting more cost than if you were able to share and queue up documents for a printer.

On the other hand, the argument is, "Well, if one of the systems goes down, then I've only lost one operator." So there is some balance here in what size of shared logic you are prepared to accept and what the effect will be if, for instance, that central unit goes down. We are particularly sensitive to that, but that is the general response that we have. If you want to go single stations then we have them; what we are looking at right now are more sophisticated methods of back up, which allow the work station to continue operation when one part of that shared logic is down. The most likely things to go down are probably the physical movement devices, such as discs and diskettes. Since we have distributed architecture within our own product line, and since the word processor is working in the work station not on the central unit, we can route those messages across to a secondary master which might have another disc based on it.

So we are looking right now at providing for two clustered systems — a buddy arrangement between the two where one could back up the other in the case of failure. It is really a trade-off. There are a lot of emotional issues to do with that failure and a lot of productivity considerations about what does it really cost you. I see those as the significant balances between the two.

RAY: Before we break for lunch, I should like you to join me in thanking not just Tony but all of our speakers this morning for their presentations.

MANAGING WORD PROCESSING

Randy Goldfield
Booz, Allen & Hamilton Inc.

Randy Goldfield is a senior associate in the Administrative Management Services group of Booz, Allen & Hamilton.

Since joining the firm, Ms Goldfield has headed the word processing practice nationwide and specialised in office automation, including feasibility and diagnostic studies and implementation planning. Her areas of expertise include electronic mail, communications systems, office of the future planning, word processing, secretarial career path planning, job enrichment, work measurement, job simplification, and professional productivity improvement. She has participated in assignments for many banks, insurance firms, pharmaceutical firms and utilities.

Ms Goldfield is also well known for the articles which she contributes to journals such as Datamation, Management World and Time Magazine.

COX: Good afternoon, and welcome to this afternoon's session. Our first speaker this afternoon is Randy Goldfield. I think that it is fair to say that over the last two or three years, as interest in word processing has grown, a number of people have jumped on the bandwagon of running conferences and seminars on word processing. I can assure you that it is the aim of everyone in Europe organising a conference on the subject to get Miss Goldfield as the keynote speaker. She is a very well-known speaker on the subject, and a very well-known author on the subject, but more important than that, to my knowledge, she is a very sound practitioner in this area. So without further ado, let me introduce the first speaker this afternoon, Miss Randy Goldfield.

GOLDFIELD: The subject of my talk today is managing the word processing environment or facility. I think that before I begin I ought to start with a disclaimer, which is that although I have worked for vendors in the field, and I have worked for manufacturers of equipment, and I have worked for several different consulting firms and have had my own consulting firm, I have never actually been a manager of a word processing facility, except as a consultant, in which case we have managed many different ones that we have implemented in turnkey operations. However, I think that it is not a disadvantage in that it gives me a broader overview. I have run into most of the problems without having to be quite as limited in scope as a typical user who is managing the facility for only one company.

I would like to say, however, that nobody will ever understand your company and your own situation as well as you within it do. Obviously, if you make clever use of consultants as you all are, you learn to take that which is applicable and will work within your organisation and leave behind that which is not.

Recently, I came back from England, where I was in Rye, on the Sussex coast, and previously to that I had been doing quite a bit of work in London in word processing and office automation. The situation there is analogous to that in the States about four or five years ago. I think I ought to take a step aside and say that I have been through the development cycle several times. It is very nice right now being in demand as a keynote speaker, but 12 years ago when I started out in this business and called up virtually every senator on Capitol Hill to tell them that they could come, free, to a word processing seminar, everybody responded with the same thing, which was not "Terrific!" but, "What? What's word processing?". It has taken about a dozen years for word processing to be a reasonably recognised term within business in the States. So for me to say that I see a lot of analogies in Britain, and in Germany as well, with the situation here about five years ago, is not necessarily undiplomatic — necessarily.

However, you have some major advantages. Many of the mistakes and, quite frankly, the stupidities that the States went through in terms of trends of popularity, are things that seem to be avoided right now in your country. We have skipped over a lot of the brouhaha that initially got a lot of press coverage for word processing and office automation, but did not do a heck of a lot to create effective, well-run, smoothly operating facilities. So there are major advantages in not being a pioneer. You know the old saying that "Pioneers are the people that wind up with arrows in their back and their faces in the dust". It is not necessarily a good position to be in.

I do not know who or if anybody has defined word processing, and I am certainly not going to get into that because, quite honestly, there are as many definitions as there are vested interests in this business. A vendor will tell you that primarily it is the movement and processing of information on paper, and then try to sell you an automatic typewriter, or a copier, or an intelligent facsimile transmitter or some such thing. The users will tell you that it is the management of a very complicated and sophisticated method to get work done more efficiently and effectively and increase productivity; and then tell you that they need more

funds to do that, and more personnel working for them to get it done. Of course, the consultants will tell you that this is a highly sophisticated technology that very few people understand, which has tremendous benefits but also tremendous potential risks and "You really need the guidance of somebody like myself to assist you". Therefore, when it comes to definitions, I will leave that to you.

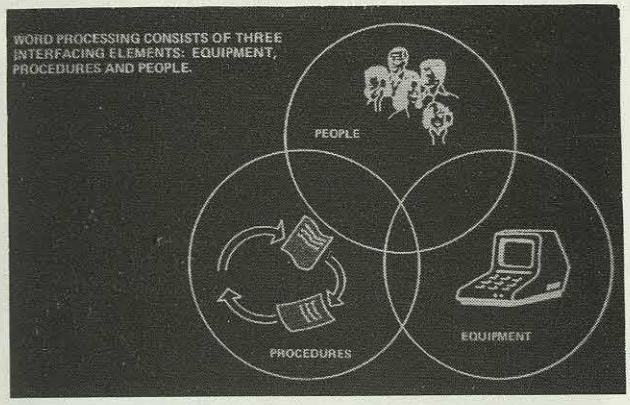

But we have to consider three important factors. If you have not run into these already, then whoever has been speaking earlier has not run through the whole gamut because people, procedures and equipment are the basis. Now I know that you are going to cover a lot on the equipment side of things and see quite a few facilities, but I think that we have to concentrate, when it comes to management of a facility, on people and procedures; which again is different from the trend initially seen in the United States.

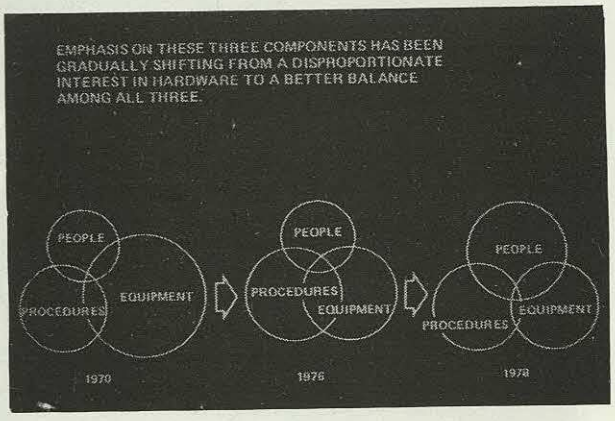

Emphasis on these three components has been gradually shifting from a disproportionate amount of interest in hardware to a better balance between all three. In 1968, nobody knew anything about it or what was going on. In 1970, when people who knew anything at all about it talked about word processing, they were talking about automatic typewriters. At that time that meant the MTST, the IBM Magnetic Tape Selectric Typewriter, for the most part. Equipment was what drew people; equipment was what was interesting. But unfortunately, when I see facilities that fail, it is not because they have poor equipment.

You can take regular electric typewriters and run a pretty good word processing establishment, if you have the proper procedures, controls and personnel; whereas the best equipment in the world will not help you if you are not managing the other two effectively.

I should like to continue with this trend for a moment because, in 1976, we found that these things started to equalise to a certain extent, although 'people' was the least impressive side or the least focussed side. This was mainly because of the vendors, the manufacturers of word processing and other office equipment, who had a vested interest in selling high technology, not necessarily in implementing it effectively. Obviously it is very expensive to provide the kind of system support that allows you to change an organisation effectively if you are not getting paid for it, as consultants are. So vendors who had suggested, for instance, that a company centralise all of its word processing facilities, did not tell how or what was the best way, or how many, or how often, or where, or when this centralisation should take place. All they did was realise that for one sales call you could sell a lot more machines if you sold a centre.

This in fact was the impetus for the early concept of centralisation, not because it was necessarily more efficient — although in many cases it will be — but because there was a very high cost of making a sale.

To sell a $17,000 typewriter took an awful lot of selling, and if you could sell 10 to the same guy instead of one, you had a bigger return on the investment of your salesman's time. This thrust towards equipment came from the vendors.

There is something that you should all remember. When we talk about keyboarding, copy reading and editing, that in terms of secretarial time accounts for only 20% of the secretary's day. That is fine and dandy overall, but that means that all of this interest in word processing, if you take a look at it in its narrowest sense of automated typing and editing, is looking at only 20% of the secretary's potential.

What about that other 80%? That is part of the trend that we see, that there is much greater interest right now in the States in all of the other activities that the secretary is doing that do not necessarily involve typing, and how to control those. Obviously there is no vendor out there pushing this, because there is no way very easily to sell equipment in this area. One of the things that we find that is most difficult is the quantification of this kind of productivity, because although you can say very easily, "Well, a secretary typed 300 lines today on a non-automatic machine and we've brought in this automatic one and she types 600 lines, so she's doubled her productivity," you cannot very easily say, "Well, the secretary took 20 phone calls today, and she took 27 yesterday, which means her productivity has gone down." Obviously it is not quite the same.

However, there are ways to measure even that sort of productivity. Then you are looking at the entire secretarial pie, and that is nice because that is the single largest information handling cost that companies have to deal with — the secretarial salary. That includes EDP, telecommunications, printing and photocopying, and other areas; so it is quite a large piece of the information handling segment.

However, when you take a look at salaries overall, in England we are only talking about 6% of the overall salary pie of any

company spent on secretarial salaries. What about the 48% that is the professional's time — the professional's salary? That obviously is the end of this trend, and that is something that you should all be most aware of, because saving secretarial time in improving office systems is good, but saving professional time is the key to advanced economic advantages.

The situation is, however, that if you think it is hard to quantify a secretary's time, try quantifying a professional's time. I put together two proposals this month versus four proposals last month. Since basically that is my job, not giving these lovely talks, I am half as productive this month as I was last month. But that is not true, because both of my proposals this month were accepted, whereas last month of the four only one was accepted. So I am twice as productive this month as last month. But that is not true, because the two this month are smaller in terms of total dollars than the one last month, which was a very large government account. So that means that I was more productive last month. But that is not true, because the governmental account has very little profit margin and the profitability on that job will be virtually nil, and the profitability on the two private sector jobs is pretty large. Therefore, I have made most money for my company, ergo being most productive, this month.

Obviously that can go on forever. How are you going to take the activity that the professional does during the day and translate that into productivity? How will you then be able to assess what impact automation will have on that person's productivity, so that you can determine the cost effectiveness of all of this wonderful equipment?

An interactive work station is just a wonderful thing. That sounds great. I can call up information from the database. I can send messages. I can receive messages. I can keep my calendar. I can access my files. I can schedule meetings. Will it save anything in time? How much will it cost? How much time of mine will it save? How much is my time worth in dollars? Will it make me more productive?

Now if you find anybody in your sojourns over here who has found a way to quantify that productivity, he is either a liar or he is crazy — or I'd like to meet him. So keep your eye out, because right now there is nobody who has done this very effectively. There are loads of pie charts that say that executives spend 14% of their time on the telephone and 22% of their time in meetings. That is not the answer. Obviously we have to find some way to correlate the activity the professional does with the kinds of tasks that he does, to be able to define how word processing, telecommunications and other new technologies will impact that productivity, because that is the 48% of the pie.

Now you may say that this will never be done and professionals simply cannot be quantified. That is what they were saying about secretaries 10 years ago here in the States, and we can quantify them right now down to the eyeball blink. So I think that it is possible and I know that it is coming. The question is when and how effectively, and how effectively will organisations use this to make themselves most competitive. Because as we see it here in the States, and you have seen it even longer in Europe, the vast, open horizons in business are closing. There are not too many pioneering areas left. You cannot simply say, "Well, let's go and open up an office in Liverpool. No one's ever thought of selling refrigerators up there." Obviously there are no new virgin territories, and therefore the tremendous opportunity in profitability and growth will not be in new horizons per se, but in more effective internal procedures so that you can cut operating costs which are now astronomical. When you get your internal costs down, your profitability will be just as highly impacted as if you had more salesmen out there, selling whatever it is that you make. That is the direction that we see things taking.

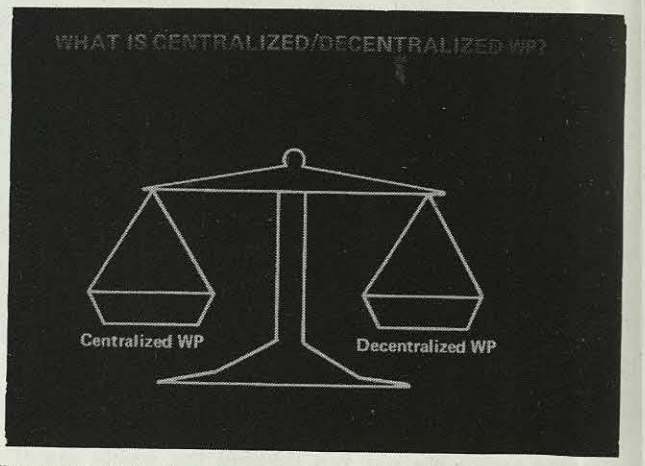

When you talk about word processing, it is important to understand that we are not simply talking about equipment

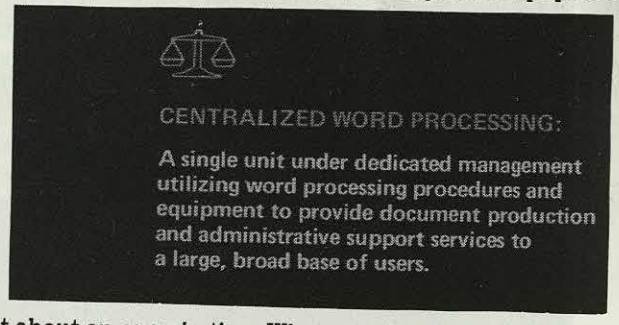

but about an organisation. When you decide on whether or not you want to centralise and have many word processors

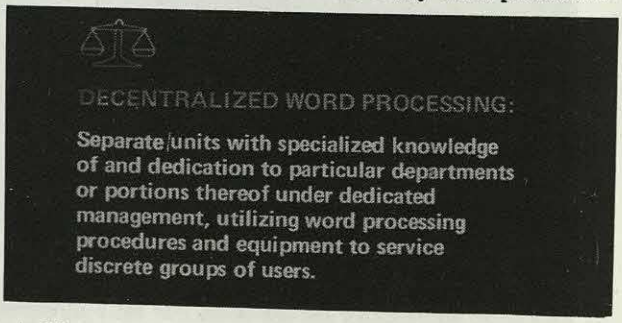

located in a single area, or decentralise and have them spread about, you have to base your decision on many different things. Let us go through them quickly.

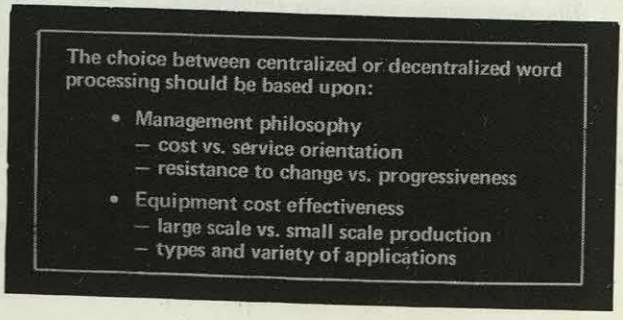

There is the management philosophy of your company. There are some groups who say, "We don't believe in things like typing pools. We want our people very closely involved in the work, therefore we want a secretary to every professional," or that sort of thing. So there is a cost versus service orientation of any organisation, and a resistance to change versus progressiveness. We are working now for one of the few private banks left in the States and obviously they are not the type of group known for their progressiveness. To walk in there and suggest some very avant garde kind of office environment will not be a very successful ploy. Again, you know your organisation better than anybody else, to determine what would be.

Equipment cost-effectiveness. Large scale versus small scale production, and types and varieties of application have to be considered.

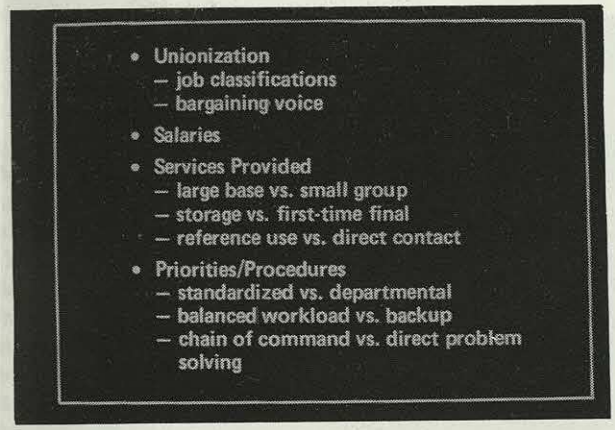

Unionisation. This is certainly a problem for many groups, particularly in the UK right now. What will this do to job classifications and what will it do to the bargaining voice of people involved? Will there be resentment and concern about redundancy and that sort of thing, or is this something that can be accepted?

Salaries. Are there any changes that will take place in salary structure because of word processing and, if so, what?

Services provided. Will you have a large base of users or small groups, and define them that way?

What about storage of information? Will you store all documents that are typed? You can create a tremendously huge storage problem, physically as well as in terms of dollars by beginning with word processing, because all of those letters that you had no record of at all except for the paper copy, you now have some kind of magnetic medium that has to be dealt with.

Reference use versus direct contact. We talk about priorities and procedure setting. Will you have standardised procedures, or departmental procedures, or perhaps none at all? Whatever the user wants, the user gets. Balanced work loads versus back up.

You always have a trade-off problem here in priorities and procedures, and with a large centre you have a pretty good opportunity of balancing your work load; with a decentralised environment you do not. A chain of command versus direct problem solving. In a typical centralised environment you have a chain of command. In a decentralised environment you have direct problem solving by the operators of equipment. I am not saying that one is particularly preferable to the other, you have to make the decision for your own group.

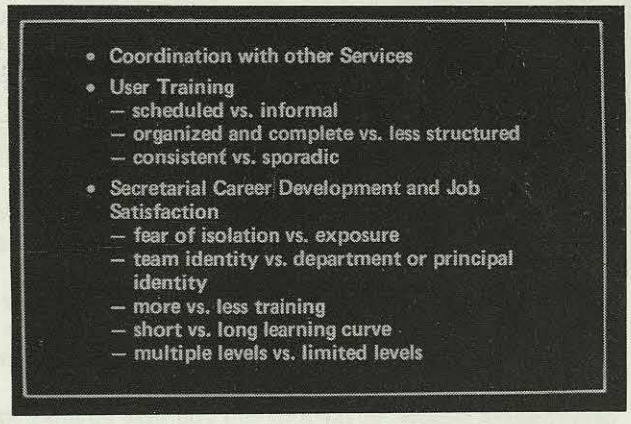

You have to coordinate with other services. I have seen situations where one group does the typing and editing and has no interaction with the printing and photocopying group, which is a very silly system where you find a lot of duplication of effort.

User training. Scheduled versus informal. You find that you can train people who are supposed to be users of the system when you have a centralised facility with people that have the time to train, versus a decentralised environment where people may not understand and there may not be anyone dedicated to training them to understand about the advantages of use of the system.

Organised and complete versus less structured and consistent; usually in a centre you find it is consistent versus sporadic. Secretarial career development and job satisfaction either is a problem or need not be a problem in either a centralised or decentralised facility. But you have to deal with the fear of isolation that many operators will have. "I don't want to be in a centralised environment, I'm going to lose touch with what's going on in the company," versus exposure when you put an operator here and there, they are more a part of the department within which they work than a part of the word processing group. The team identity versus departmental or principal identity.

More versus less training. You always have less training in a decentralised environment.

Short versus long learning curve. In a centralised environment where you can ensure the use of the equipment, you find that the learning curve is shorter. People learn how to operate the equipment more competently and more completely, in a shorter period of time.

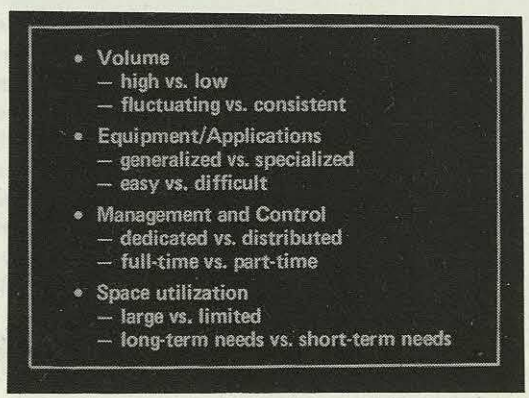

Multiple levels versus limited levels of volume. If you have high volume, centralisation might be a better alternative; low volume perhaps decentralisation. Fluctuating volumes, again centralisation because there the fluctuations can be evened out. In a decentralised environment, if the legal department is overworked every Friday and under-utilised Monday and Tuesday, you will have operators who are not kept busy.

Equipment applications. Generalised versus specialised. If you have some special applications going, you have to consider decentralised word processing to meet those. And easy versus difficult. Management and control. Dedicated versus distributed. Full time versus part time. Space utilisation. Sometimes the decisions to decentralise or centralise is based simply on whether or not you have the room for a centralised facility, and when you do not you have to decentralise. But you ought to know that in a decentralised facility, in the long run you take up more space because each operator needs approximately 95 square feet versus about 70 square feet in a centralised environment.

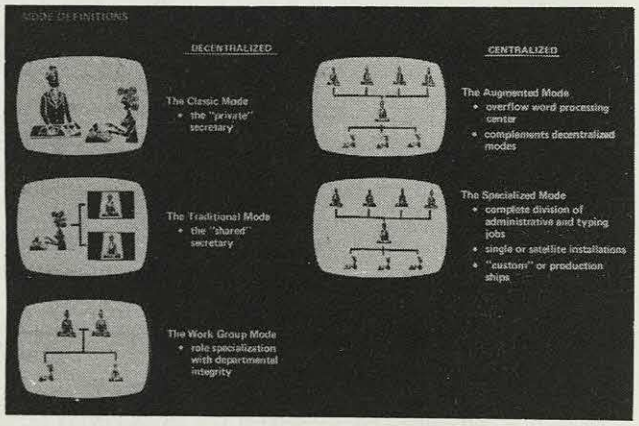

Now that does not mean that your only options in managing a word processing facility are simply a centre, or a bunch of operators spread around indiscriminately. In fact, there are basically five modes that we can go through that are either one or the other, or a combination of them. If you take a look at these five modes, in the next few days as you tour facilities, you will see that most of them are a breakdown, or a combination, or exactly one of these described.

There is the classic mode: the private secretary. Now this can be a word processing mode if the secretary has automatic equipment, or has access to automatic word processing equipment. The traditional mode. That is the shared secretary, where two professionals typically, sometimes more, share a single support person. The workroom mode. Now this is very important. It is the first mode here where we find a specialisation of work taking place, and that is critical in word processing. We are not talking simply about getting in automated equipment, we are talking about the need to specialise two tasks: the document production task and the administrative support task. What you find is that when you have someone who specialises in typing and proof reading, they become much more efficient at both of those tasks and their productivity increases tremendously. With or without equipment, that takes place.

The only reason why the secretarial job is so unspecialised is because of the historical situation that has been created over the years. Initially, when women were hired as typewriters — which is what the people who operated typing machines were called, which is why there is a song that was popular back in the 1880s call "Sitting with a Typewriter on my Knee", which does not refer to a Remington Rand — when the typewriters first came into the office, they were only used for typing and they were very efficient. Now when you typically hire a typist you have her take a typing test and she types in a very specialised environment. Then you put her behind a desk where the phones ring and there are distractions, people walking by and interruptions, and the productivity of the average typist goes from 60 words per minute down to 13 words per minute. You are paying her for her 60-words-a-minute skills, and you are getting 13 words per minute in terms of final output. So in a workroom you have a dedicated typist, and an administrative secretary who does the other tasks, and they work for two or more professionals.

The centralised alternatives are basically two: the augmented mode, which is an overflow word processing centre that can take work from the classic, traditional or workroom mode; so that if the private secretary has work that has to be typed, she can send it down to an augmenting word processing centre. The specialised mode is really the totally centralised division of labour, where you have administrative secretaries and typing secretaries who are pooled. You can have single or satellite installations, and basically you have a choice between a custom or production shop, which means that you can have a word processing centre that produces work as needed, in the manner specified by the individuals, or that has standards and will only do the work in the manner that they outline, so that a user can have his choice of single-spaced or double-spaced, but they cannot have one-and-a-half spaced pages.

To make your life easier, I have put together a listing of each of these modes, with the 13 criteria that I have quickly discussed. Management philosophy and support; management control and supervision; space utilisation; workload balancing, etc; and a check list as to whether or not these modes are effective in those areas. You find that typically some are effective and others ineffective, and you can select the one that seems most appropriate for you.

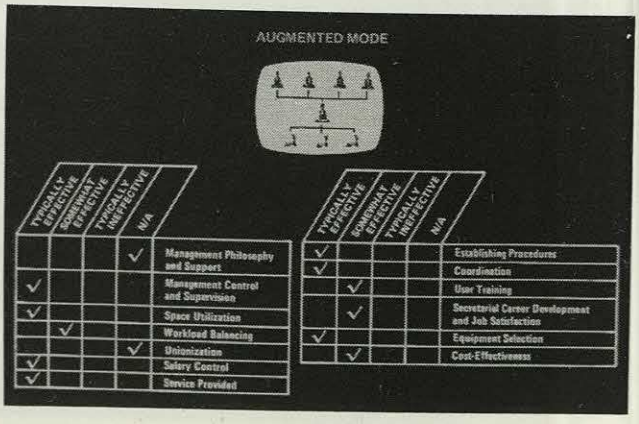

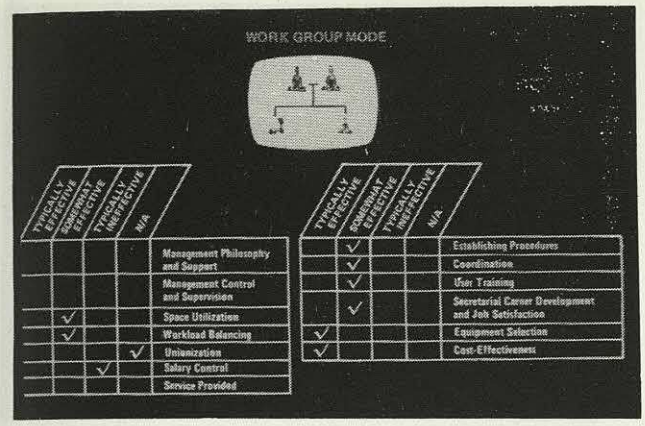

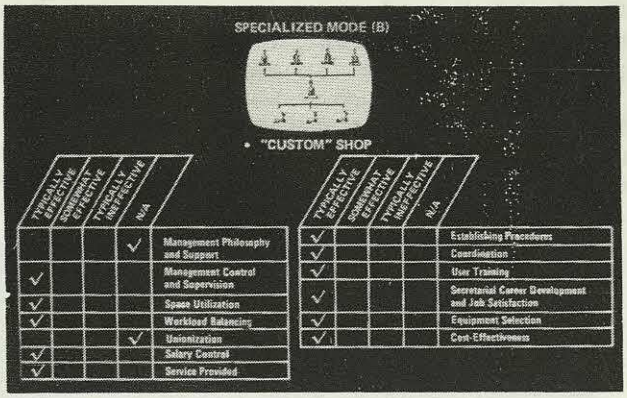

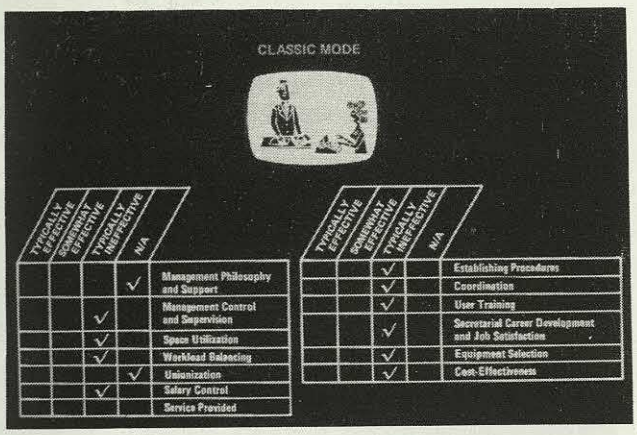

One ought to note that although centralisation was in vogue for a while and is now out of vogue, and everybody is talking about decentralised mini centres and satellite centres, to take the broader overview is most important. What we have found is that in most organisations you will need a combination of many of those different modes that we discussed, to fit the different needs of different groups. Obviously, the top executives very often can generate enough work to keep a secretary busy full time, except when they are on business trips, which might be a large proportion of the time, leaving a lot of idleness. On the other hand, there are engineers, for instance, who might be able to work very effectively in an augmented mode situation. So that decision is for you to make, based not only on the needs of your company but the different needs of different departments.

Now how about the people in word processing? We have talked a little bit about the organisation procedures; let us talk for a few minutes about managing the people.

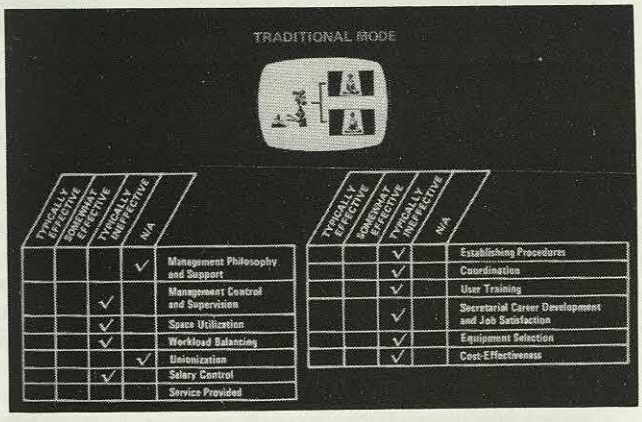

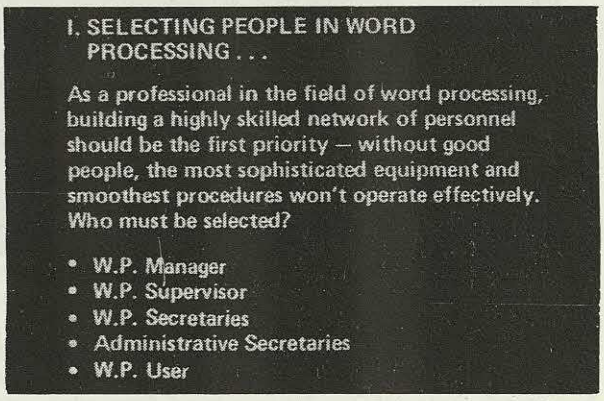

Who are they? Typically, in the States we have four different groups: the administrative and word processing secretaries;

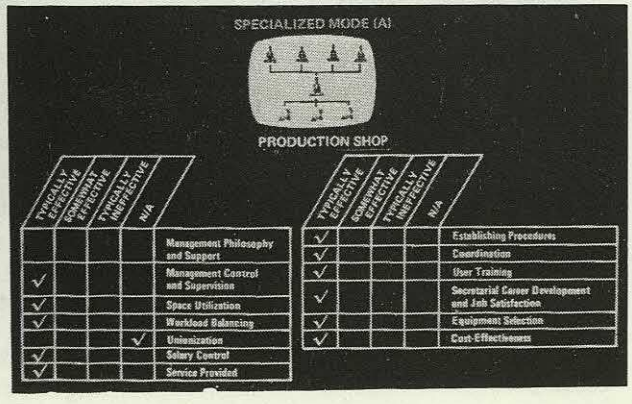

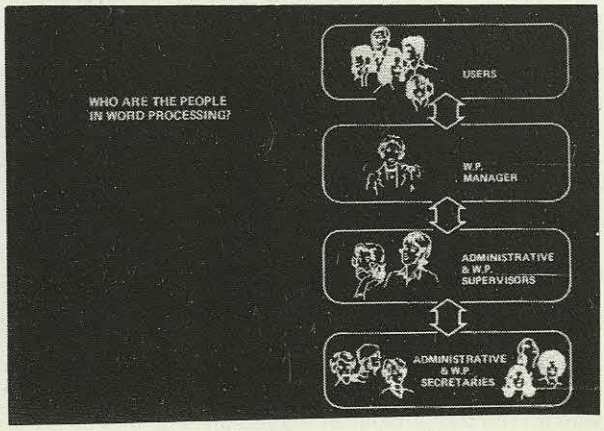

the administrative and word processing supervisors; the managers; and the users. Those are the people that you are going to be involved with at your home companies as well.

Selecting people in word processing is the key to effective management of any word processing facility, and how you do that is still a mystery to many.

To select word processing personnel effectively, what data is needed and should be prepared?

- Functional role definition
- Effective, realistic selection criteria
- Selection tools, if possible

What I should like to go through is in which ways you can best come up with a good personnel group for any word processing operation, because that is the basis for anything you will build on. First, you need a functional role definition. Whatever jobs you are going to design, you have to functionally describe before you can fill them. Effective, realistic selection criteria and selection tools, if possible. I can tell you about those that some of the more successful companies are using here in the States.

SELECTING THE W.P. MANAGER

- Role Definition:
  The management and organization of the word processing system including:
  - Management of the secretarial and supervisory staffs
  - Interaction with users and management
  - Development of systems and procedures
  - Long-term planning
  - Equipment and personnel selection

First of all, the manager. What is the definition of the manager's role? Basically we are talking about someone who is the strategic planner as well as the implementer, and has overall responsibility for any word processing, and sometimes all office automation technologies. The management should include management of the secretarial and supervisory staffs, interaction with users and management, development of systems and procedures, long term planning, and equipment and personnel selection. Here we are not talking about the old typing pool supervisor, we are talking about somebody with a sophisticated understanding not only of management techniques and cost justification, but also the technology involved and some of the long term planning potential that will be needed to allow you to keep up with office automation.

What are some of the selection criteria? Excellent people handling skills, previous supervisory management experience, good personal packaging. You have got to take somebody who is effective in front of management as well as the users,

SELECTING THE W.P. MANAGER . . .

- Selection Criteria:
  - Excellent people handling skills
  - Previous supervisory/management experience
  - Good "personal packaging"
  - Methods and procedures background
  - Budgeting and cost justification skills
  - Previous experience in word processing

as well as the secretaries.

We have a lot of jargon. I like the company for which I work very much. Booz Allen is a good group, but we have an awful lot of jargon. One of the things that we use internally, which took me aback the first time I heard it, is 'personal packaging'. When you talk about somebody who looks good, who looks right, you say they "have a good personal package". When you talk about somebody who does not, you say, "Well, his personal packaging isn't what it ought to be." I know about good usage and crazy words, and I was taken aback the first time I heard it as well, but it has been so inculcated in me that I do not even think about it any more.

I was having lunch with my sister about three weeks ago, and some businesswoman she knew walked up to us and said, "Hallo," and introduced herself and walked away. My sister said, "What do you think of her?" and I said, "Well, she's got great personal packaging." My sister said, "The job's finally gotten to you! What does that mean?" Anyway, what I mean is that they have to look right.

Methods and procedures background is very helpful. Budgeting and cost justification skills and previous experience in word processing. If you can get somebody who has all these things, you are almost totally assured of a successful word processing facility — or anything else for that matter. If you do find one of these people, give them my card because I will pay them more to work for me here in the States as a consultant. It is impossible to find someone with all of these skills right now.

SELECTING THE CORRESPONDENCE SECRETARY

- Selection Tools
  - Interview Guide
  - Typing test
  - "In basket" test

What selection tools? An interview guide; a resume review and a reference check. That is pretty standard. Role definition. Supervision and organisation of a word processing centre.

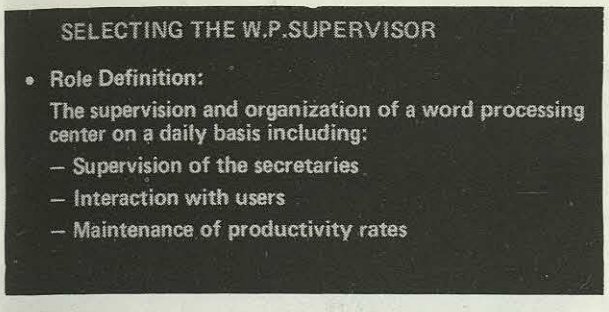

This is the supervisor on a daily basis. If you have a centre, or even a decentralised facility, you must have somebody there to manage the staff; to keep historical data so that you know how productive they are today versus yesterday, and how the work loads are changing; to balance the work loads and move them from place to place. That is one of the main ways in which you can improve efficiency — by balancing secretarial work loads. What are their daily tasks? The supervisor is needed to do all these things: supervise the secretaries; interact with users; and maintain productivity rates.

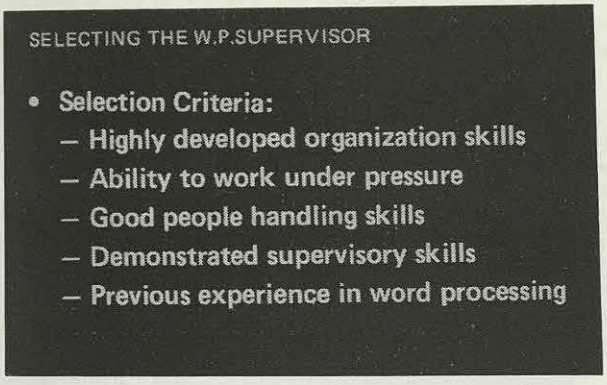

Selection criteria. Highly developed organisational skills; ability to work under pressure; good people handling skills; demonstrated supervisory skills; and previous experience in word processing. Again, if you can find somebody with all of these skills they ought not to take the job; they should know better.

There is a story about a consultant who, instead of being given a bonus for a job well done, his boss says to him, "Time is money. I don't have any money to give you, so I'm going to give you time. You can have a 100-day vacation." The guy says, "Well, that's great," and he goes home and he thinks, "If I have to go home for 100 days, I think I'll go out of my mind." He sees in the newspaper that night an advertisement that says, "100-day cruise — $100." So he thinks, "This must be some kind of publicity gimmick. You send it in and they write back saying, 'Well, we're all out of $100 cruises but we do have $2,000 ones'."

But he tries. He sends in a cheque and next day he gets a telex confirmation. It says, "Yes, you're confirmed". He packs his bag, shows up at the wharf. It is a sleek white ship. It looks terrific. They come down in their white uniforms and they put him on the cruise. They unpack him in his state room and he is all excited. It pulls out past the three-mile limit and they grab him and drag him down to the hold There are hundreds of people sitting at benches, chained up, and they are all pulling at big oars. In the prow of the ship there is a drummer, with his shirt off, banging the drum, and every time he bangs the drum one of the guys pulls on the oars. They all pull, pull, pull for 50 days, then they turn the ship around and start to bring it back. The guy turns to the man on his right and says, "Excuse me, but are you going to give a gratuity to the drummer?" and the man sitting next to him says, "Well, we did last year." And that is the situation here. If you can find somebody who has done all of this and is willing to do it again, they are obviously masochistic and I would not hire them if I were you.

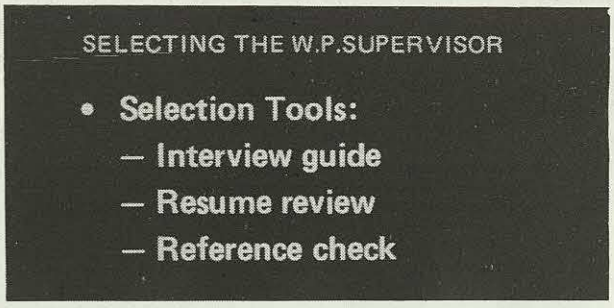

Selection tools are basically the same.

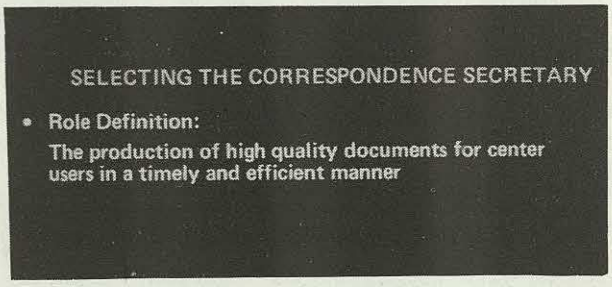

What about the correspondence secretary? That is a very fancy name that IBM has basically landed us with for the operator of the word processing machines. There is a reason for this, however, and that is because they do not want to be called operators. I think the situation is very similar to yours in the UK or in Europe overall, in that blue collar workers basically had a choice of becoming white collar secretaries or continuing to work in factories; and they did not want to work in factories. They sometimes take lower pay rates in not working in factories, and they do not want to be operators of machines. Therefore they came up with this term 'correspondence secretary' which you will find used all the time. They do not type correspondence most of the time, and they are not secretaries any more, but that is what they are called. The production of high quality documents for centre users in a timely and efficient manner is the role of the correspondence secretary.

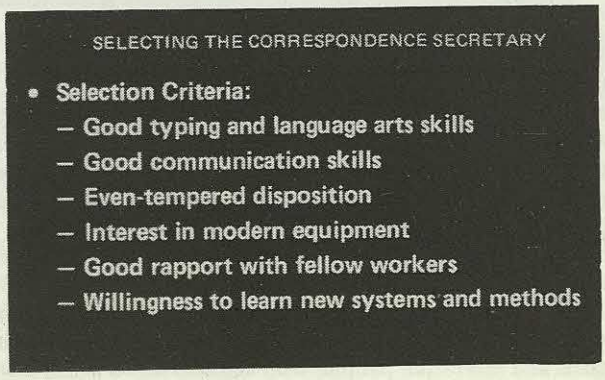

Selection criteria. Good typing and language art skills — and the language art skills are far more important in managing a good centre or any other word processing facility than is the typing. You can improve someone's typing speed, but it is very difficult to improve their language skills. Good communication skills. Even tempered disposition. You do not want a prima donna in the midst of these secretaries because you will very quickly have problems. Interest in modern equipment. A good rapport with fellow workers and a willingness to learn new systems and methods. The type who says, "Wow, that looks fascinating," is worth her weight in gold versus the one who says, "I don't know anything about computers and I'm afraid — I don't want . . . ". Those sorts are not going to work out very well.

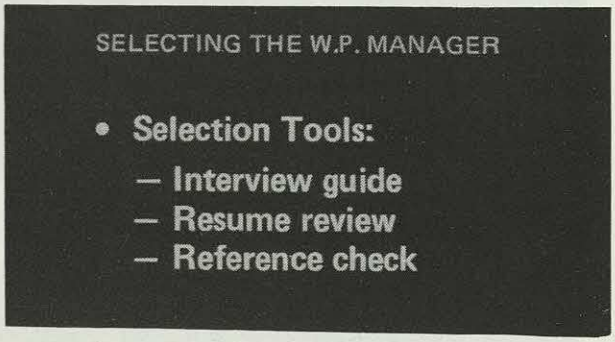

One of the things that we use as a selection criterion which has not yet been made illegal, although many of our testing has been, is 'in basket' tests, which are still acceptable. That is simply, when you have a potential candidate, to ask them to do some of the kind of work that you typically would be doing in the environment in which they would be working. This works pretty well — be it asking them to type something, or proof reading it, or whatever, making corrections as they go along. We cannot give grammar or spelling tests any more because they are now illegal, but you can ask somebody to correct the grammar and spelling of something that they are typing, because that is not. I do not know why that is, but it is. As George says, I just have a practical knowledge of this business, I do not know why things are the way they are.

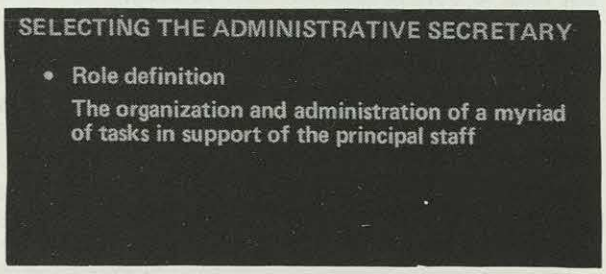

Role definition of the administrative secretary: the organisation and administration of a myriad of tasks in support of the principal staff. Now the only reason that the administrative secretary — the secretary that does not do the typing in a specialised environment — exists is to make the professional more productive. If she is a great typist, if she is a wonderful file clerk, if she answers the phones very nicely, if she makes good coffee — those things do not matter unless that secretary is very effective at improving the efficiency of the boss. When you specialise tasks in this way you allow a situation to exist where you can put college grads and interested people, who want to move ahead and out of the secretarial area, in this job because they no longer have to have typing or stenographic skills, and they can in fact become part of the mainstream of corporate tasks. I know that you will be visiting IBM. You ought to ask them about all the success stories they have of people who have moved up through the ranks from this position. I think they are true, so far as I know.

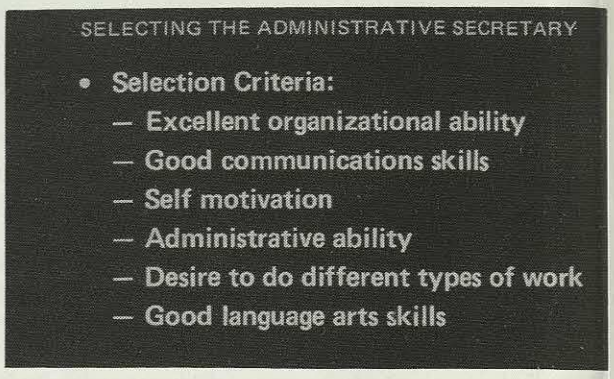

Selection criteria. Excellent organisational ability; good communication skills; self-motivation; administrative ability; the desire to do different types of work, because the administrative secretarial role is not easy and is very much driven by the professional. And good language art skills.

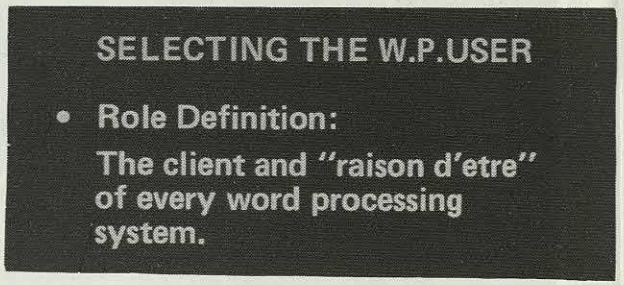

Finally, the user. This is the key to good management and, quite frankly, is terrifically important. If you pick the wrong user as a pilot or as overall user of your word processing facility, you will have a failure on your hands. If you pick the right group, you are 50% there towards success. They are the entire reason why the system is being built. If you put together a beautiful centre that is very difficult to use, because you have to fill out forms and queue up to get there, and wait a long time for the stuff to come back, it does not matter how cost effective it is — in the long run it will not be a productive centre because you are undermining your whole reason for being, which is to make the professional more productive.

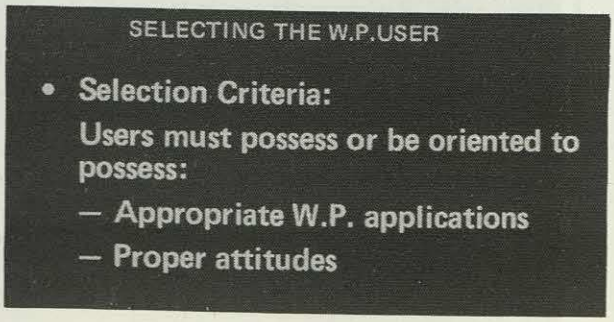

How do you select the word processing user? How many of you are already using word processing in your organisations currently? Of that group, how many have only pilot groups? You all extend it to all users basically. How many of you are considering beginning pilot groups? One. Good luck. Listen to this very carefully. Users must possess or be oriented to possess the appropriate applications and the proper attitudes. You will find that if you create a word processing facility there are those that will immediately begin using it. I can tell you right away who they will be: they are the people on the bottom of the ladder who do an awful lot of work and get very little support. So you find that the junior guys who do not have secretaries are the ones who will jump on to the bandwagon because at last they are going to have some support. Now if you can take a lesson from that, what you find is that there are those who can use word processing as a service; and if they get better service than they have currently available to them, they will be pleased.

That means that someone with a private secretary may feel that he can never get better service than he already has. But that is not necessarily the case because in a typical secretarial environment we find that a private secretary is away from her desk 34% of the time. If you want the phone answered, or you want a copy gotten for you, or if you want something out of the files, you have a problem because one third of the time your secretary is away from her desk. She may be making a copy for you, but she cannot also answer the phone at the same time.

A word processing system can, in fact, change that because you create clusters of secretaries who continue to have primary relationships with the professionals that they work for, but are also responsible for the service of an entire group. That way you can balance the work load and ensure better service for everybody.

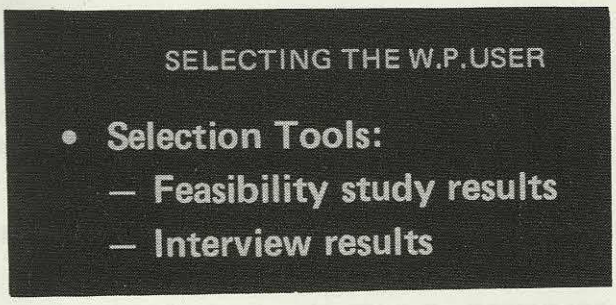

Let us take a look at the users and how you sell them on this. The best way to select a word processing user is to do a feasibility study, to determine exactly what kind of applications they have, what kind of needs they have, how legitimate the needs that they have are; to interview them and to see what they feel about the current level of service and what they need as a final, or best, level of service.

The thing that you have to worry about here is actual and perceived needs. People very often think that they need typed things back in an hour because they get them back in an hour. But in fact, since it is not mailed till the end of the day, they do not need it back in an hour at all. So that is something that must be considered.

A feasibility study typically will include work sampling of all the secretaries; a pretty good sample of all of the typing that is done; observations of the secretaries to make sure that what they record on daily logs is accurate; an analysis of all of those things and cross-referencing to check them, as well as estimates on the part of the secretaries and professionals as to how secretarial time is spent. Typically we find tremendous discrepancies between what the professional thinks a secretary does, what the secretary thinks he or she does, and what the secretary actually does. We do not know why some of these discrepancies exist, but we do know in some cases that they are very, very consistent.

For instance, we find that typing work is typically over-estimated by 100%, almost right on the nose, in virtually every job we have ever done. We do not know why that is, but we do know that the manufacturers of word processing equipment took advantage of that for years and did free feasibility studies. Now all they ever did was take estimates and based everything on the estimates. What people would say was, "Well, that's what your staff told us, and that's what we did." So you typically got 100% more equipment than you needed because everyone had over-estimated that; and, at the same time, had underestimated the administrative support tasks that have to be done.

What does all of this add up to? If you decide, based on those 13 issues, what kind of organisation you need — centralised or decentralised; the classic traditional or workroom mode; augmentation system or perhaps a centralised facility, or a combination of them — if you go from there to looking at personnel and putting together a pretty effective group of manager, supervisors and secretaries, you have 90% of your problem licked. Quite honestly, it really then becomes icing on the cake as to whether you should select a shared logic system, or a stand-alone system, or a computer text editing system, mainframe based, because that is not going to be the decisive factor in your success; nor will that make the tremendous cost difference.

The key to effective word processing management is knowing how secretarial time is spent and how effectively it is spent, and then applying automation and, more importantly, reorganisation and management to an area that has in the past predominantly gone under-managed. It is a resource that is sometimes completely unmanaged. When we ask the professionals to estimate how the secretaries spend their time, and they are sometimes off by as much as 700% in some categories, particularly idleness for instance, and you realise that the professionals are the ones who are charged with managing the secretaries — who is watching the store?

The situation is really very simple. Why should we ask a very expensive professional that you have hired to be an engineer, or an attorney, or a physician, or salesperson within your organisation, to also be a secretarial supervisor? And how effective is that as a use of their time?

Managing word processing basically needs three key things. If you have those, you will have a successful system.

Number 1, you have to specialise, because specialisation is the key. You must divide the document production task from the administrative support task, or else you will not really achieve high equipment utilisation or effective unit times in activities that are done.

The next thing you have to do is to reduce duplication of effort. Obviously, when a document is typed and then edited and retyped, you have a tremendous duplication of work. When the telephone is answered by a receptionist, who then has to transfer the call, where it is answered again by a secretary, who then puts you on 'hold' so that she can send it to the professional, and all three of those people have spent time answering the same phone call, you have duplication of effort.

Finally, you have to reduce the consequence of error. You must look at a word processing system, always having in mind that error is what costs the most money in any of these systems. It is the revision and editing that takes up 64% of the cost of producing documents, not the original typing, although that certainly seems to be the most substantial part initially.

How do you reduce the consequence of making an error? Mainly through technology, because when somebody is typing, for instance, and makes a mistake, they have to stop and backspace and white over, and retype versus, for instance, typing on a CRT where they catch the error and simply re-keyboard and the error is never even made on a piece of paper. The consequence of making an error in dictation, for instance, to a secretary is less, very often, than the consequence of making an error in dictating to a piece of dictation equipment; which is why sometimes dictation equipment is still not an effective time-saver, although it very often can be. So consequence of error is the third thing.

If you consider specialisation and duplication of effort, your consequence of error will mostly be handled by automation and you can manage a word processing facility pretty effectively. The main thing is to keep your objectives very clear in mind at all times. You may think that you are there to reduce costs, when in fact your company is interested more in improved service, or vice versa. You may find that, rather than reducing costs and reducing heads, you are interested in a longer term cost avoidance scenario where you simply do not have to hire any additional help for X amount of time.

I had one client recently who called and said, "We have consistently maintained for the last three years a 120% growth rate in our company." They are located in Colorado and in building, and obviously that is an area that has been booming lately. He said, "I'm not interested in saving the number of secretaries in this organisation. We can hire three for every professional. My objective is to make sure that we don't have to hire as many professionals as we did last year because that's a tremendously expensive operation. The finding of them, the relocating, the losing of several that constantly goes on. The best thing that could happen to us is being able to make those professionals more productive. And if it takes word processing equipment or more secretaries or whatever to do so, that's just fine."

In managing a word processing facility, one way or another you will not please everybody, but you have to make the decision that the objectives that you follow are those that are most appropriate for your company. Do be sure that those that you choose to follow are those that have top management's full support because, as a final thing, in managing word processing facilities you can be very sure that without top management support, no matter how good a manager you are, no matter how good a manager you select, you will not be very effective; and the dollar and cents savings and whatever will not gain you support if you do not have it to begin with and there is no real commitment on the part of management to automation.

COX: We will quickly take one or two questions.

QUESTION: I am intrigued that you make the distinction between a professional and a secretary. I am also intrigued that the UK are five years behind the US. The UK is equal or slightly more advanced than the US.

GOLDFIELD: As a matter of fact, SRI, just outside London, has 50 interactive work stations that the professionals and secretaries use indiscriminately as needed and which are tied into a satellite communications facility. It is the state of the art. You cannot see a better system any place, for what they are doing. So I do not mean to say that. It was terribly rude, and I am sorry.

What I meant to say was in terms of content, the kinds of things that I am hearing from management in the UK right now and the kinds of problems that people are concerned about are those that we heard in 1972 and 1973 here, when one would begin to talk about word processing; such as, "Well, the secretaries would never stand for that," or, "Our company doesn't have a policy that would allow for that sort of thing," or, "We don't believe it can really save us a great deal of money," or, "We don't find that it's necessary right now." Those are the kinds of things that we do not hear very much in the States any more. Virtually every company has accepted word processing and begun to implement it, and realises that not only is it cost effective today, but that that actual keyboard may very well be the keystone to the Office of the Future; because you can cost justify it right now and, with the technology being what it is, it can very easily communicate with and be the access point to any far more complicated office automation system.

By the way, about the secretaries not being professionals, I do not mean that at all, but I need some way to categorise the other people. I cannot just call them the 'other people'. You cannot call them executives because they are not all executives. And though — God knows — they are not all professional, that is the best I can do.

QUESTION: What is the rule for determining the desired productivity in feasibility studies?

GOLDFIELD: Quite frankly, when I started we called up 300 of the Fortune 500 firms and offered to do a free study so that we could get that information, so that we could start a database; and only one company said "Yes". So we had a very small database. But since then I have probably done about 120 or so of these studies, and we keep the information by category, by the kind of department that we are looking at, by the kind of company and whatever, and then have some guidelines for productivity.

But we find, surprisingly enough, that most secretarial activities, regardless of the company they are working for or even the department that they are in, are very similar in make up. In other words, the legal department secretaries' work is not terribly different in terms of breakdown of activities, although obviously the actual work is different.

But the breakdown is very similar to that of the secretary in the accounting department, for instance.

We have pie charts and graphs and whatever that not only give the percentage of time spent on each task — we break them down into about 26 different activities — but also the unit time for each of these activities. For instance, if it takes 1.6 minutes on the average for someone to access a file, we then find that in a particular group the average for that company in a study comes out to be 1.9. We break groups down into zones and we find that within a zone, within a department for instance, that average is a little higher. When we find that we figure out why, and it is sometimes because the files are located in a bad spot, or because they do not have a good system, or a lot of other things. But we look for irregularities, obviously, from those.

COX: Unfortunately, at that point we must close the session. Randy, I should like to thank you very much for a most informative and thoroughly engaging presentation. Thank you for coming.

EXPERIENCES IN PLANNING AND DEVELOPING AN AUTOMATED OFFICE -1

John Gosden
The Equitable Life Assurance Society of the United States

John Gosden is the Vice President for Telecommunications for the Equitable Life Assurance Society. He is responsible for planning, providing and recommending management policy for all electronic communications both voice and data. He joined Equitable in 1970 as a second Vice President in charge of the Technical Support Group and was responsible for EDP technical planning, development, maintenance and training services for computer operations and programming. From 1975 to 1978 he was in charge of Corporate Computer Services providing computing and systems development services to Equitable, was responsible for effective use of computers, chaired the EDP Policy Committee, and was responsible for institutional policy, standards, and system assurance.

Mr Gosden has participated in several special study groups, notably the National Library of Medicine MEDLARS system; the Manned Orbiting Laboratory Support System; and the Air Traffic Control En Route System. In 1977 he was chairman of a Federal Advisory Group on White House Information Systems advising Frank Press, Science Advisor, and Richard Harden, Special Assistant to the President, on ways to use information systems to improve the decision-making processes of the President.

COX: Gentlemen, having discussed so far some of the general moves in the area of office technology, what we are going to do now, for the final two sessions of the day, is to look at the experience of others in the area. We will look at two organisations to hear how they have been tackling the technical systems and the organisational problems presented.

The first of our speakers is John Gosden, from The Equitable Life Assurance Society of the United States — one of the major institutions in this area in the United States — who to my knowledge are very forward system users, and have been for some time. John's background, as with all our speakers, is summarised in the notes you have been given. But one thing I learned about him over tea, which I didn't suspect, was that he was actually engaged at Lyons, which most people here will recall was one of the first commercial computer installations in the United Kingdom, way back in the 1950s.

So he has been in the business a very long time indeed, and it is a pleasure to have him here this afternoon.

GOSDEN: I need to preface my remarks with several caveats. I found preparing this talk very, very difficult and I have a fair amount of material, some of which is really only background. I will put that up and you can read it about ten times as fast as I can read it to you. And I don't intend to work my way through all those slides — they are really just background.

Secondly, I usually tell a number of humorous stories that I wouldn't like to see in print. I therefore intend to embellish them with scatterlogical swear words. In order to see that they get properly deleted. I apologise; this is not my usual business style, but I have been told that I have no editorial control over the results!

Thirdly, I will point out to you that the major part of my talk comes towards the end when I tell you why we are doing all this and how we expect to do it, so if you do want to go to sleep the early part is the best time.

THE EQUITABLE LIFE ASSURANCE SOCIETY OF THE U.S.

SALES	$6 Billion+
PERSONNEL	21,000+
SUBSIDIARIES	Variable Life Casualty/Property Leasing Real Estate Investment Time Sharing/Program Products Environmental Health

EQUITABLE
A LARGE INSURANCE COMPANY

$100,000,000,000	LIFE INSURANCE
$ 20,000,000,000	ASSETS
3,500,000	INDIVIDUALS
17,000,000	GROUP

You have already seen how large we are and how important we are, and 20 million people, one way and another in this country of about 200 million, have substantial insurance coverage with us.

We have for many years had a concept that our head office is nothing more than a paper factory, and when you talk

A PAPER FACTORY

400,000,000 TRANSACTIONS/YEAR

7% GROWTH/YEAR

about office automation we have been essentially in that business from the early 1950s. I have been using this slide now for about six years steadily, so you can do the compound gross at 7%, you now know that it is more like 550 million transactions a year. Don't ask me what a transaction is. I got somebody to sit down and try to estimate this some time ago, because we were trying to figure out what we were really doing. Sometimes this is something like a real estate deal that takes eighteen lawyers and twenty two real estate men six months, and sometimes it's just paying a claim.

But it's a phenomenal amount of work, if you look at it in this way. And to give you some idea, 14% of our operating budget other than the sales commissions and costs goes into data processing and communications. About 7% is data processing, and about 7% is communications. That is one of the reasons why I'm going to start with some funny slides.

VOICE VOLUMES

17,000 Stations Nation Wide

1,000,000 + Calls Month

$2,000,000 Private Line

$1,500,000 WATS

$5,000,000 M.T.S.

4.2 Avg.call length

CHURN 1.3

Let me tell you about our telephone system and how large it is, and it's the thing on which we want to build our networks. It's a fairly large system. It uses up about 30 million dollars. The telephone itself is about 25 million dollars a year, and about 5 million is data at the moment, expanding rapidly. And the average call length is about 4.2 minutes. Churn means how often do we move an extension or do we change it. Every telephone moves or changes approximately 1.3 times a year. And when you realise that the Bell telephone company in recent years has changed the charge for moving a telephone from about 20 dollars to 120 dollars, you can understand why we think that's a very interesting number.

PROBLEMS

* Call Capability
* Optimum Use
* Dialing Procedures
* Tariff Increases

The other part of all of this, if we want to keep on moving, is to worry about the problems in all that. We were worried about our call capability — our telephone system was antique; we didn't see how it was going to take us down to the future, let alone be a base for what we wanted to do. We had a very old 701 switch, which we couldn't get optimum use out of this kind of thing for dealing with long distance traffic and so on. We thought we were paying too much for telephone calls. You have heard about dialling procedure from somebody else. You get your own network up, you dial busily through twenty one odd digits, and if you can get that right you are in good shape. Also tariff increases were going out of sight.

	DATA Characters (Billions)	VOLUME Cost ¢/1000	GROWTH Terminals
1977	14.	13.0	500
1980	89.	4.5	2,400
1984	132	3.8	5,000

Meanwhile, we got into the data business. We started about 1970 with on-line systems and, as you can see, it's been sort of climbing along fairly fast and we are starting to go through a rather explosive growth rate. And to date as we look at ourselves going on down the road, the costs of communication have sort of dropped off into this area, and we are not sure they are going to drop off very much more rapidly.

MAJOR TECHNICAL TRENDS

* ON–LINE APPLICATIONS EASIER
* USEFUL MINI-COMPUTERS
* USER-ORIENTED LANGUAGES
* CHEAP MASS STORAGE
* GOOD COMMUNICATIONS NETWORKS
* MORE USEFUL SERVICES VIA NETWORKS

It still pays to try and cut your communication costs. But look at the overall way data processing has been going. One of the things that has been happening is that on-line applications are getting easier and easier. So communications

are becoming more and more important and on-line work of some kind or another is becoming pervasive. Minicomputers have at last become useful. I could spend a whole session on that — maybe a whole conference — and if you want to talk to me about some of these throwaway lines, do it afterwards.

Cheap mass storage is really important. There are a whole lot of things which are just beginning to happen now as the cost of keeping masses of data on-line in some reasonable way is becoming comparable with keeping it in all those vaults and cabinets and things around the place.

Good communication networks are now possible and we are building them, and we are starting to be able to buy useful services via networks. One of the things we hate as a sort of a semi-pioneering firm (that means we like to be about third or fourth, not first or second) is to buy services somebody else has paid all the development for. OK, another part behind all of this — push harder, push harder again. I need to change the next slide. The next one is the mass

MASS STORAGE

	Cost per 1,000 char.
1961	$140.00
1973	4.50
1975	2.29
1976	$ 0.50

storage — a nice little table, and it really shows you how even in the last few years since '73 to '76, the prices dropped off very dramatically.

I prepared this slide in '76 and I haven't bothered to prepare another one since. Once we got down to the 50 cents in there, that was good enough. It's still going down, and that has been a very dramatic kind of number as far as we were concerned.

TERMINAL TYPES

IBM 3767 DATA 100	TRENDATA
IBM 3774 DEX	ATLANTIS
IBM 3270 EDT 300	WANG
IBM 3790 RCA TTY	PDP-11
IBM 3777 GTE 7800	TEXAS INSTRUMENTS
IBM 1050 GTE 1514B	DATAPOINT
PIX SYSTEM	CDC

One of the things I've done is to look and see how fast I thought growth was going to occur in practise in our organisation — for some of those different technologies — as opposed to how much more rapidly it was going to go from the technologists point of view.

- Terminals, very fast. Once you have got a few in you can start to expand, and very rapidly, in growth.
- Minicomputers, moderate. You still need software, you still need lots of other stuff that isn't around, and you need to tailor applications.
- Communications, extremely fast. Once you have got terminals out there you can really start to build up traffic. Communication growth can go really fast indeed on you.
- On-line storage growth, moderate. Not because the technology isn't there, but it takes a while to build useful databases that you want to pay for.

CURRENT PROBLEMS

MINI-MAXI SYSTEMS COMPATIBILITY
HIERARCHY LIMITATIONS
INTERFACE PROTOCOLS INCOMPLETE
PORTABLE SOFTWARE
SYSTEM IMPLEMENTATION IN PHASES
SPECIFIC HOST SYSTEM LIMITATIONS

- Large computers, slow. We have got most of the things up. There will still be some small growth in that for very large applications.
- Printing, again slow growth. We don't need that much more printing, in fact we need less. And we are trying hard to make that happen.
- Microform, moderate. We are still putting more stuff away on microfiche and microfilm and so on.
- Applications packages, moderate. We would like to be doing a lot better. The things we see moving very fast are in the communications area and in the terminal area. And the other ones that we are interested in are the minicomputers and the on-line storage.

What are the current problems with all of this, from us as the users point of view? First of all there is one that I don't have on here. The applications, where we are trying to expand were built on database systems. Unfortunately our mini and and our maxi systems are basically incompatible. Things like that are just starting to change. We're starting to get some interfaces in between them like the kind of things you have heard from Wang this morning. The way systems are built in major hierarchies is also a big problem. The interface protocol and portable software we don't really have; system implementation in phases is a pain — it takes too long, it is very difficult to do, and nearly every host system has some kind of limitation.

REGULATORY MESS

* WATS ?
* EMFIA FX'S ?
* MTS HI/LO ?
* NETWORK SERVICES !
* BULK RATES ??

If you turn to the communications side there is a regulatory mess, and I won't go into that unless you want me to take another week on all of that.

What are our major users actually doing? First of all there are two kinds of users. There are sophisticated users, and what I call the uncontrolled and naive users.

MAJOR USER TRENDS

SOPHISTICATED USERS
* MAJOR RELIANCE ON ON–LINE SYSTEMS
* DISTRIBUTED EDP WITH NO LOCAL EXPERTS
* PHASED DEVELOPMENT-CONTINUAL CHANGE
* PACKAGE USE TO CUT LEAD TIME
* USER MANAGEMENT OF EDP SYSTEMS

UN-CONTROLLED (NAIVE) USERS
* PROLIFERATION
* NO RECOVERY, CONTROL, COMPATIBILITY
* INADEQUATE CONTRACT COVERAGE
* RISK EXPOSURE UNKNOWN

And some of those are doing the best things around, of course — much to our annoyance. First of all, the sophisticated users have got themselves into a trap. By and large they have gone into on-line systems. They are very good and they are tied into them like crazy — there is nowhere else to go. We've gone for distributed data processing, often without the necessary local experts to do the things we'd like to do. And we are pushing even more out there.

Phase development and continual change are going on all the time. It just doesn't stop. There isn't a beginning or an end to a system any more, just continual change. Packages are used to cut lead times. I had a classic example last year — a beautiful example. We had a big new system we wanted to put in and it was important to us from a competitive point of view. And we went to a lot of trouble evaluating alternatives to decide which would be technically best. We did a good cost benefits analysis. The decision was between going the package route in which we would have to modify our unique needs, and building our own system, both of which would have cost about the same amount of money. However, we thought we could get in about six months earlier with the package system, and I thought that it would be less hassle overall.

It also turned out that our user estimated that the competitive advantage to him when he was in would be about a million dollars a month, and the package got us in six months early. Lead time for very big systems really is the most crucial thing in this. And this is what a lot of technical people really don't understand. Not only do you have to cut that lead time down, but you have to cut all those slippages down that double the lead time. And a good way to do that is start from a package.

First of all, if you offer a user a clean sheet you're going to spend a year negotiating requirements. If you say, "I've got a package and we can just do these few extra things or not", it doesn't take him so very long to make up his mind, and you can save a year negotiating. Now with uncontrolled naive users we have proliferation. Each one knows exactly which minicomputer will do best for him — the salesmen told him so. They put in systems without recovery, control and compatibility, there is inadequate coverage, and I'm expected to bail him out.

So we have problems with all of this and we are now facing office automation and even more things going on.

SUMMARY

* Trend to "Work Stations"
* Changes in office Management
  - less paperwork
  - higher service expectations
  - better control of processes
* Systems Development

Meantime, I can give you a whole talk that ends up with a summary, and I've only brought the last slide — that says what's happening is, we are moving more and more to work stations, and there is a whole change in office management coming about. There is in fact less paper work coming around, as we are moving through all of this. There are much higher service expectations — people don't write letters of complaint any more, they pick up that phone and they expect somebody to answer it and tell them what is wrong. You can't go pull the file and say, "I'll call you back later." They call up somebody else and complain about your non-responsiveness. People have grown up using the telephone. The telephone company has been marvellous in this, and has generated huge business.

These are the things that are driving the way offices are managed. Meanwhile, systems development in the classical sense hasn't really done very much about this. It has sort of sat there and cried about it. Meanwhile, the naive users have come trotting along, minicomputer people have come trotting along, and the office automation people have come trotting along, and they are going to solve the whole

STRONG NEED TO
PROTECT THE DOWN-SIDE

DISASTER–RECOVERY–BACKUP
BECOMES A CRITICAL USER
(NOT DP) PROBLEM

problem. Meanwhile, I'm sitting there, worrying about protecting the down side. I can't afford systems that are built today, on which we build our business, to go bust.

I can very easily lose the company a lot of money by having a system that fails. It is very hard for me, from a technical point of view, to save the company a lot of money except by cutting that lead time so that they can get their competitive edge in earlier. That's where the saving is, the real importance is in that competitive edge in the business, not the technology. In fact, disaster recovery and backup is becoming a critical user, not a data processing problem. They know what backup they need. We can do our side of it: I can supply replacement lines; I can supply replacement processing. Can they find a building to put it in? Can they find people and get them there? There are a whole series of other problems in their side.

EVOLUTIONARY DEVELOPMENT

* Business Needs
* Technology

On top of this we have got this continual change, we have evolutionary development that's driven by the business need, and technology is merely a tool to help us get there.

BUSINESS NEEDS

* COMPETITIVE NEEDS
* EFFICIENCY
* MANAGEMENT CONTROL
* MANAGEMENT DECISION INPUT

What are our business needs? Essentially, first of all there are the competitive needs. We want to be right up there with our share of the market, or whatever it is. Efficiency is secondary in all of this, particularly within expense control itself. But it is very important. Management control turns out to be the third most important, and management decision input turns out to be the fourth, when you really get down to gut issues on all this.

Now I am suddenly going to get to office automation. Why are we looking at office automation? One of our problems is, having automated a great deal, having got the guts of our business up there, we now have a lot of data about what is going on. We can go look at that data, and in fact we use it to manage better. And the data we can put out about what we're doing and how we're doing it is starting to grow.

We now have tools go to into our IMS databases and pull out almost anything you want to know with a GIS run or something, overnight, very easily. What kind of managers are going to be able to use this data? Some are really going to start to use it, and that is going to be where the good managers of the future are going to pull ahead of the lousy managers. That's where the edge is going to be. We're going to need managers who can handle information and who have good tools to handle it with. This means that we have to bring together the raw operating data and the tools for these guys to do it. And that, I think, is where office automation is really going to help us. It's when we bring the data, the word processing and tools on the desk together, not necessarily for that executive sitting at the terminal, but his staff people and others. We don't want to have to go into different systems and manually move it across. That is what we think is important. We think that's going to happen over the next five years.

So we are in office automation because we believe it's going to be an important management tool for running our business down the road. We expect to see the good managers move in to grab the tools, as they see it's necessary from peer pressure, and other managers will start to use it.

MARKET ENVIRONMENT

MANY ALTERNATIVES
- DATA PROCESSING
- VOICE
- DATA COMMUNICATIONS

PROBLEMS
- RESTRICTIVE SOLUTIONS
- VENDOR SURVIVAL

Let's look at the market environment behind all of this. There are too many alternatives in the data processing area, the voice area and data communications area. One of our problems is that there are too many restrictive solutions — none of them do what we want. And another is, when you go out there for a lot of these nice alternatives you've got the whole problem of vendor survival.

GENERAL OPINION TODAY

* WORD PROCESSING — MORE OR LESS PAPER
* SOCIAL CHANGE SLOW —
  * SECRETARIES OBSOLETE ?
  * EXECUTIVES AT TERMINALS ?
* OPPORTUNITIES
  * FASTER, CHEAPER MAIL
  * PAPER REDUCTION ON INFORMATION MEMOS
  * BROWSING
  * FILING
* PIGGY-BACK: Calendars, messages, follow-ups

What's our sort of general opinion today? The classic way of looking at word processing that we see in many places is still more paper and not less. I would agree with the speaker somewhat earlier today, who said that the major use of the word processing stuff today is to perform letters and major reports that need to go through lots of revisions.

The other intermediate stuff; we don't even know why it needs to be on a word processor anyway, or why it needs to be polished up by anybody with literary arts, because 99% of the people who are going to read it wouldn't appreciate the literary arts anyway.

I have even seen some funny studies that carefully show how many pages are done by a word processor each day. My secretary is much smarter — she types a memo and if there is one thing wrong in it she snopakes it out and types the correction in. She only produces one piece of paper. If that was on the word processor system I would have produced two, and I would have doubled productivity. We then xerox it and send it out, and nobody knows who snopaked it. And if she isn't in that day, I cross it out, write it over by hand, and send it over for xeroxing anyway. And my boss has never objected to that.

Social change is slow. We have been told that secretaries are going to become obsolete first of all. Now people are beginning to realise that that isn't true. Executives at terminals? They have been trying that for years and that isn't going to happen either.

Executives worry about the problems they have today. As soon as they have figured out what data they need, they have solved the problem before we can write the system to give them the answers, and by then they have figured out what the standard answer is and they have delegated it. I think we are going to be a long way from executives at terminals for a long time to come. Except those who are really running things in the standard way, instead of worrying about the future.

Now in the opportunities area we see faster, cheaper mail. Our mail costs are rising so fast and the service is getting so slow that that's just got to be worth it for all kinds of psychological and other reasons, even if it seemed like it was going to be a good buy. The trouble with an electronic mail system is it's like starting up a telephone system — you have to have enough people on it to win; the first person on has nobody to talk to. So it's getting in there.

It is a problem. We see a tremendous opportunity on paper reduction on information memos. We have lots of memos. If there are fifty people, sometimes thirty page memos; you read the three pages you are worried about, and throw the rest away, or you file it which is even worse. What we would like to do is have all that stuck away and you can come in and browse through it on your terminal, or whatever, and just save all of that; save all the filing costs. That's what we would like, a good browsing facility and a centralised filing system. We can see a lot of paper work savings just in there. But again, that won't work unless there are a lot of people on it together. The whole of that distribution list.

After that we would see piggy-backing — you know, the calendars, the messages, the follow ups — none of those will pay for themselves, they are just piggy back frills that have come on afterwards. A speaker in just twelve years from now will tell you that these were the main justifications for the whole thing, but I am certainly not going to be able to tell it today.

So we have a long term strategy. First of all, the one in white is the obvious one, that we have been doing forever. And that's why it is very hard to do long range planning, because you have to throw away your long range plan next year when IBM does something different, or somebody else comes out with some new thing, or the Federal Communications Commission issues a different ruling. And so we stay loose. At the moment we have stayed for about three years now IBM/Bell compatible, so there is always Ma somewhere or other to run home to.

LONG – TERM GENERAL STRATEGY

EXPLOIT OPPORTUNITIES BUT STAY LOOSE

* Stay IBM – Bell Compatible
* Shun Vendor Development
* Combine Voice – Data
* Beware Special Technology
* Protect Downside
* Favour Early Payback

Shun vendor development; we're not in the development business — they are. That doesn't mean to say that we don't get stuck in there sometimes, for competitive reasons, but we do it with our eyes wide open. We've combined our voice and data. In fact, we have a whole talk on this one too. It's turned out to be very useful and attractive to us, and it will enable us to bundle. Where you want one line into some remote place, if you have your voice and your data bundled you're providing all these services down it, then you can really start to make it marginally possible to do things which enable you to get the big network going.

We protect our down sides, which I talked about earlier. For example, we have stayed away from package systems. First of all, they require the special protocols. We need interfaces into them and out of them, and we don't want to get locked into something we are not sure we are going to stay with on down the road. And we favour early pay backs. We don't look ten years down the road — we want our money soon while we think we can still get it before conditions change and then move on from there.

MERGED VOICE/DATA PLAN

NETWORK	Point-to-Point Switched
COSTS	Reduced 12%

We merged voice and data, and we've actually gone into a point-to-point switching network. We have not gone the package route. We run and manage our own, and where we need it we go out and buy auxilliary services. We reduce our costs, just by doing that, by 12%, and as we start to go through data explosion and to put in more automation at the end of those lines, that's going to move up about 20%.

Computerised Telephone System

* Universal Numbering Plan
* Optimised Routing
* Predictive Call
* Call Queuing
* Computerised PBX
* Auxiliary Data Exchange

We have put in a computerised telephone system. This is the fun one we give another one hour talk on. We put in a very modern system, we leapt the whole technology. It's a system by Danray, now Northern Telcom. First of all we put in the network control, then we put in our home office internal system. They build a system for big people. They don't build small ones. And it can handle digital data. It's a six wire system, and we don't have to dial twenty digits — the system itself will do all the optimising and the routing, and it does predictive calling. The telephone company, if you dial across the country, will hunt its way through, trying to find its way to the other end. Our system keeps a table of status of all lines to the computers in New York, and we don't even start to put you over that net until we have got the line all the way through.

This, we figure, in peak hours saves us about 10% of our capacity. Not only that, we optimise and the optimising routines are very sophisticated. We have twenty different patterns in there and they change at different times of day, depending on the way we are trying to do things, and which services are useful. We save about another 10% on that. We have call queueing, and we run that so that the average peak is about three minutes. What we do is, we offer a service and say, "If you stick with us and go queueing, we will give you rates which are 40% off Bell companies rates." If you request a priority, when we want to put you in the queue, which only happens at peak hours, we say, "OK," and we put you straight through and we charge you Bell rates, even if we don't use Bell facilities. That's the kind of thing you can do with this. The user is making a choice, we make the price for him depending on his choice. Give him both ways.

We save about a million dollars a year, essentially by being able to do the queueing and balancing in there. Our PABX is computerised and we are putting in an auxilliary data exchange. Our home office has about 10,000 people in it — 8,000 phones. We don't have to re-wire for different kinds of sets, you can put in an adaptor between the jack and the telephone to plug in your terminal. The adaptor costs us $250, we don't have to rent it for $250 a month, and we can put it on any of the telephones in our building and it does not take away the telephone capability. Our building is now wired via the telephone, automatically. We don't have to go wire it separately for data — we can put it anywhere.

That will save us immediately. We don't do a lot of outside stuff in the building at the moment. That will save us, immediately it goes in, about one hundred thousand dollars a year, right off the top. Then we are going to extend this out through our network.

The telephone company actually provide a four wire line. Just before it comes into your switchboard it drops down to two wires. We jump around that and we have four wire capability all the way across the country, just by tapping in the switchboard at each end, and we put our own monitors on there as well. So we are really building a telephone network out there, which is going to be our data network, for everything except the heavy traffic. This is where we can put all our marginal traffic without extra cost. And technology has just made this possible for us.

What are our short term tactics in all of this? Well, we are building a good network on the applications we have got today and on our telephone system, so that we can go do the things, so that we can put a lot of people in at once and get the traffic in there.

SHORT – TERM – TACTICS

* Build Good Network On Applications
* Incorporate Easy Terminal Access
* Provide Cheap (Mass) Storage
* Use Integrated – Easy – Capable System
* Look For $100 Terminal
* Begin With Good Pay Off Pilots
* Then Piggy-Back

We are trying to incorporate easy terminal access anywhere we have telephones, to make it cheap and easy so that people or the head office don'thave to pay for wiring in a new system. We are providing cheap mass storage on our major system so that people can put files up. If they have very big files we can go into the browsing or the report distribution, or into the filing thing.

What we also want is an integrated, easy and capable system that people can use who don't have to worry about jargon and various other things, don't have to worry about different kinds of terminals around them all over the place. We have just done a major study looking to see who has the best system for us to run our pilots on, as opposed to the networks we already have up, that's integrated and easy and has capability. We looked at about eight systems and we are now making a decision about which way to go on that. We are not looking for the cheapest, we are not looking for the technically fanciest, we are looking for the one which has the best capabilities from our users' point of view. We think the dominant thing there is what it can do and how well it can do it, rather than the costs.

And we would really like to have a $100 dollar terminal. We think that's possible in the next five years. You can buy a black and white television set today for $70 retail. We will be able to have a keyboard and auto mechanical keyboards for probably about $5 or $10. We believe that they will be offering us $100 simple CRT and keyboard terminals within the next few years. Then we will really take off, and that's when electronic mail and other things will really start to fly. Our agents will buy five out of their own pockets — I'd be tempted to buy one myself if my boss didn't elect to give me one. We are starting to begin in this area, even though we've already been into this whole thing. It's hard to draw the line where our kind of automation and office automation cross over.

Talking about supporting management directly, we are beginning by looking for good pay off pilots. That is, users out there who really want to do something, who think it's a good idea and are prepared to pay for it out of their budget. In fact, the proposal we put up for management was that if we got about $500,000 this year in our budget we would use it as a revolving fund, if for every project we put up we can have the first two years' savings. We believe that we could build up our business over five years into $10,000,000 a year quite easily. We believe there are really good savings in there, where the line management really

want to go in and do things on this. Then we will piggy-back other things on that; all the goodies and the marginal things are on the outside of it.

PILOTS

Scheduling
Publications – Photo Composition
Legal Documentation
Regional Mail
Operating Data Browsing (DP & C)

I want to tell you about the kinds of pilots we are actually looking at right now, though they change every month as budgets get altered and different managers get moved around.

First of all, we have quite a number of different problems in our building just related to scheduling — simple ones, such as scheduling all the classrooms and meeting rooms that we have, and our dining facilities and our training organisation. We believe that we have a high pay off thing in there. It's all done by hand today, and there are frequent mess-ups and slips. The irritation caused when a senior executive vice president and six senior VPs have a meeting and then confuse five other meetings as everybody bumps somebody else's conference room down, is enough to pay for this one in just trauma. Apart from the direct savings we can see in that, our publications area could use a lot of the word processing and the message distribution out to the field and photo composition to their stuff. They are very keen on this, and we are starting to tie this into other things. It's a very good application for us. Our whole legal area is already full of IBM magnetic card things, and we are planning to cut it over into this, and to start to track their activities. They are going to use the electronic mail feature among themselves to track what they are doing about things in their file — it's their chrono file about what they have actually done, and people they have talked to and so on — and they are going to use it to be able to index it with professionals, and then be able to go back and find things. Have they dealt with the customer before? Have they dealt with this corporation before? All kinds of things in there. Our legal department is about ninety people.

We can immediately start to do something good with our regional mail. We have six regional centres. Even if we only did the mail to them we would be way ahead — way ahead of the Post Office service, and way under the Post Office price. Then we can start to expand out from that. In our own areas — data processing and communications areas — we have a tremendous amount of management information and resource scheduling, scheduling problems on the computers, the charging problems, what happens to things and so on. The amount of paper work we produce to run our own business is enormous, and there is a huge pay off in there that we expect to be picked off by lots of other people when they start to see how we do it. If you look at internal operations in our company today you will find that this area is one of the best managed from knowing what data is available and going after it, and then producing reports. We can see tremendous savings and improvements to be made in this general area.

So what we are saying is, we think that a lot of this office automation is an extension of the way we want to manage the business and we want our operating managers to drag it out of us. All these people on pilots, essentially, either came to us because of what they read in the press, or listened to briefings that we had given and got excited about it.

They were knowledgeable enough to understand what the potential was for them, and then we went to talk to them and we essentially told them it was different to the old days. We don't have corporate money to do this; they have got to fund their own way to go. But what we have tried to do, as a background, is to provide a very healthy base on which to do it. We are making it easy for them to have a network by building our network so that it can take the traffic. We are going to make it easy to have the interconnections into the big systems, because that's where we have to pull the data out. And we are going to make it easy for them by picking facilities that are easy for users to use and that do fit together. We are expecting that the cost isn't going to be the major problem in that, and what we are finding is that there are a goodly number of people out there who really want to do this. We believe that it's going to start rolling.

The big hurdles we had were concerned with being able to produce what we thought was a cheap network where people are only going to do a few things initially. We couldn't afford to put in a line for that, but if we can just add it on over the top of the telephone line as a sort of zero marginal cost, it solves a huge problem. The second one was those facilities, and we believe that today those kinds of facilities are here and are possible to use. I will tell you a little bit more about that a year from now. But we think it's just about here, and that it will fit in with our major systems, where they have to get the data from it.

That's our experience in general. If you want to ask me particular questions now or later, I'll be very happy to answer them.

QUESTION: I was impressed with not only the smooth way you put it across, but with the impression of speed that obviously goes on in your firm. In terms of a system development, do you normally think that if you can't get it off running in six months, forget it? Do you (the company) have a criteria like that behind you?

GOSDEN: Well that's a question that is very hard to answer. I do, I believe that any time I look at a project — and we are getting more and more people who think this way — as it is slowly spreading, it's taking about two years. Essentially, I say that if there is any big project that takes more than a year then we want it in phases, and this is the early pay off thing. What can I have in six months that will do something for me, that I will be happy with if we couldn't go on and do the rest?

Now the big systems we are putting in now; it had to be a year for other reasons. But that was the choice between one year and three years. But we do hassle for that, and we hassle for it corporately, and we now have enough line managers who know that that's the right question to ask when they are on a steering committee for something. Let's have some pay offs, and then others later. Because by the time you have done six months the world may have moved. A different manager may be in there who doesn't want this whole thing at all.

I have a very good example. We are trying to automate our financial system. The general plan was to take five years and take $8,000,000, and that had to come out of the hides of the operating line divisions. Finance doesn't do anything for the company as a whole, but we have to produce this legal statement every year, and there are certain schedules in it that have to be done by law, and some of them are very tricky. We do them several times over with different parameters to see which one looks the right one.

They were going to re-do this as one part of a big system. I told them they could only have $100,000 for the first year, because that's all I thought we could get, unnoticed, in our budget. They then came up with a way of doing just that piece of the whole financial thing. Not only that, because I had harassed them so much they decided to get even with me and put up a trivial version of it on time-sharing in one month for $10,000. That almost caused me to cancel the $100,000 version. But that would have been silly. What they did was, they omitted the automated version of getting the data out of the other files into this file. They did it by hand quickly, just for that year. But to go on and do the other things we needed those bridges, so we built them.

Next year we gave them $200,000, and that's a typical kind of way to go after getting that pay off, and getting their eyes focused on something. So we are moving to that; not everybody does it yet, but they are beginning to catch on.

QUESTION: At one stage in your presentation you made reference to the use of package switching, and then later you stated that you stayed away from packet switching as a communications technology at the moment.

GOSDEN: So far it hasn't been useful to me in my business. And it also requires a different technology from what I have today and I don't want the pain of selling it.

We are unpopular for that, but I have users who tell me what to do too, and I have to unsell them, so I try to restrict their mail and the advertisements they get and things like that.

I was trying to explain partly why we are not into packet switching. It doesn't do anything for us at the moment. It may do, and when it does we will go into it. But we tend to be very electric and use the things which are around, trying to keep them all compatible.

QUESTION: Please could you tell us more about the implementation of your telephone system.

GOSDEN: We went after that in phases too. The first thing we did was just automate the long distance traffic out of the home office. That's about $5,000,000 in tolls a year and WAT services. So we had an old 701 switch and instead of a 701 switch going straight to the central office, it went straight into this other device. It was what we call an RSS, and we have a whole series of WHY switches. There are about 106 trunks going up there and they are all on WHY switches, so we can throw it backwards and forwards.

We switched to it over a period of about three weekends of trials, and then we switched it over in about a week. Until then we had 701 or Dimension type switches. We did the tandem switches. When we put those in, we also put queueing in at the same time from the home office. We did that to the field, and that's where we ran into our biggest problems, because we were interfacing with 105 different telephone companies. We started to find out that the telephone company didn't even live up to its own specs, because its old electro-mechanical equipment was extremely tolerant of things. We found out that there were Dimension systems in the field that gave out signals that weren't in the specifications, and so did the telephone company, and we had a real hassle with that. Half of the problems were really telephone companies' problems and half of them were ours. It led to an extreme lack of user confidence in the field on that system, so we didn't dare put the queueing in, and queueing is going to go in later.

That took about five months to settle down. We also had a software problem in there, and it took forever to find. It took about a month to find. One card has about ten circuits on it and if circuit zero disconnected, it disconnected all the other circuits as well, so we had random disconnects going on. But it was even funnier than that, because the way the network was set up, only one side of the conversation got disconnected. So somebody would be busily talking to somebody else for five minutes before they discovered they hadn't been listening. And you can imagine what happened to credibility. That was a disaster — we have just about pulled out of that.

The next thing we did was to put our home office system in. Everybody has two telephones at the moment. We are going through the building and training people. Forty per cent of the people don't turn up to the training classes because they know how to use the system. The way you use the functions on these things is entirely different, and we don't use a switch hook the way the Bell telephone company does. You can put people on hold, you can queue them up and leave them there and go home, and you can call forward. You can programme call forwards and automatic call forwards and conditional call forwards and temporary call forwards. And you can set up daisy chains!

That's going to take us about five months because we are phasing the whole thing. We are doing one thing at a time and we looked at everybody else who changed over a system smaller than ours and we didn't like any of them. So we changed our plan about five times, which made management very confident about us!

First we tell them we are going to do this, that and the other, and then come back and say, "That won't work, we are actually going to do it this way, and that's very much better." That's very good. Then we come back and say, "Well, we thought about that but it's taken us a long time to find a way to do that. It's going to take about five months to do all of that, then we are going to put in that data exchange which will go in over about one month, we think. Then we are going to put in at the system at the end that ties all the things together and makes it easier for us, and that's going to be transparent to everybody." The big day comes when we send out the disconnect orders to the Bell telephone company for the 8,000 telephones in our building. After five weeks of quiet operation, guess when we are going to get our first big crash? And if you know how to avoid that, I don't. So that's taking us about two years, overall. And we are about three quarters of the

way through.

QUESTION: You mentioned GIS earlier, and your up and coming managers making better use of it. What's the performance so far? Do you expect the number of people to grow fast over the next two or three years, or is it a slow process?

GOSDEN: It varies. Let me tell you how it varies. First of all the GIS users cannot change the databases, and we didn't put that service up until we could make sure that happened. All our major databases are controlled. There is a users' group for each database, and no changes are made until agreements are reached about what the definitions are and clear responsibility is set up about who maintains them. GIS users cannot screw up the system. The use of GIS just suddenly started to grow in different communities very fast. Our personnel system was one of the first ones that just had explosive growth and the reason being, all those important reports for affirmative action and the other things we have to keep filling out — without GIS it would have been a real pain.

So that system gets extremely heavy use, just for that reason. It met a particular need at a certain time. We had another explosion when we suddenly put our on-line system up for individual policies. For years we hadn't really known what was happening to our market, and then we started to do pattern searches on the data and so on. We suddenly brought 52 files together and you saw a huge spurge of stuff. Now it's slowed down and we have about five runs a month, something like that.

We are now getting much more bottom-line orientated. If we had our financial database up we could see the huge spurge in the financial areas. Our financial database is useless to them, so we don't see it. What we see is a spurge in their timesharing activity, building their own little accounting systems to do what they really need as opposed to the corporate one, which doesn't. And that's another candidate for our pilot.

So it varies all over the shop, and as managers migrate around from area to area and see what other people are doing you get little explosions here and there.

People learn. My wife couldn't use tools until she became an artist and was doing wood blocks. I was away for months and she needed some wood desperately — she suddenly learned how to use an electric saw. She went away, and I suddenly learned how to cook. It works both ways. It also works with this kind of stuff. You would be surprised how quickly people will learn when there is enough incentive and motivation there.

COX: At that stage we must draw the session to a close, and ask you to retain any further questions to put to John later in the day.

I would like to thank you very much indeed for that very crisp presentation — very clear insight into what you do and why you do things that way. I am very much indebted to you John. Thank you.

EXPERIENCES IN PLANNING AND DEVELOPING AN AUTOMATED OFFICE -2

Richard McClelland
Exxon Corporation

Rick McClelland is project manager in the Office Systems Technology group of Exxon Corporation. In this role he has been concerned with the development of methodologies for analysing office systems and defining the requirements of these systems.

Recently he has been concerned with the evaluation of a multifunction office system with a range of capabilities including electronic mail, filing, photocomposition facilities and user programming.

Mr McClelland is an MBA graduate of Brigham Young University.

COX: We now come to our second look at the experience of a company in this area. The Exxon Corporation, like many others, has been facing the problem of how to deal with office automation, an area which in the past has been a low level administrative function; and how you start bringing into it the kind of skills and expertise with which we have been familiar in computing for some years, and similarly in the field of telecommunications. To unify the control of these information handling technologies, Richard McClelland has been moved into the Office Systems Technology Group with Exxon Corporation. He will talk to us for the next hour about what they have been doing and the lessons they have been learning.

McCLELLAND: I am glad to be here. I spent two very memorable years in England, most of that time up on the Scottish border in Carlisle, and I enjoyed that very much. This is my first opportunity to speak before a major group, and I am enjoying the opportunity to make this presentation in front of men from the country of Great Britain. I was told when I came back that I had an English accent, by my brother who was in Texas. I think that my accent and my English has probably been corrupted, and I'm back to American again.

EXPERIENCES IN PLANNING & DEVELOPING
AN AUTOMATED OFFICE

INTRODUCTION

- HISTORY
- BUSINESS COMMUNICATIONS SYSTEM (BCS) CONCEPT
- ACTION PLAN
- COSTS OF BCS
- BENEFITS OF BCS
- STUDY METHODOLOGIES
- PROBLEMS & ADVICE

I would like to go through a little bit of the history and development, starting with word processing in Exxon and going into the advanced office technology function, which is now called the office systems technology function that I belong to in Exxon. Then I want to introduce a business communications system concept, where we look at the document flow and the stages that a document goes through during its life cycle.

I would like to talk about two different approaches that we have used within Exxon to quantify the costs of the business communications system; then talk about ways which we have used to show benefits of the business communications system; talk about some methodologies; and then also a few of the problems that we have had and we have seen within our organisation.

EXXON ORGANIZATION

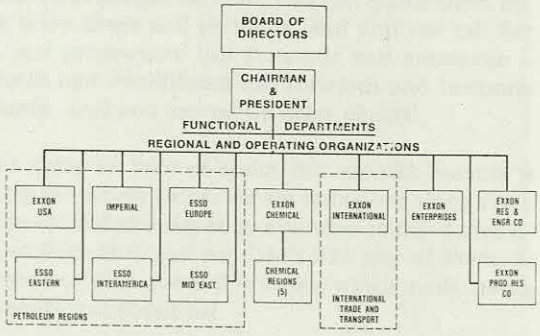

Before we get into the presentation, I would like to put this organisational chart on the board. Exxon is a large company. It is a very diverse company. The corporate headquarters — which consist of the board of directors, chairman, and the functional departments — are headquartered in New York City, and those functional departments coordinate the activities of many regional and operating organisations all over the world. There are the petroleum regions throughout the world, and Exxon Chemical Company has five other regions throughout the world. All of these organisations are almost run like separate companies, and some are more separate than others.

I want also to point out that Exxon Enterprises, that many of you might have heard of in the word processing field, is a separate division of Exxon. The group that I belong

to is within one of the functional departments of Exxon Corporation. That means that as we act as consultants within Exxon and coordinate the activities of advanced office systems, we hold at arm's length Exxon Enterprises. That has worked to their chagrin many times. We have many other competing word processing systems — Wang, Xerox, IBM, and just about anything you can name we probably have somewhere within our organisation.

HISTORY IN EXXON

- CORPORATE
  - INITIAL STUDY (1972-73)
  - ASD OVERFLOW CENTER + MINI CENTERS (WORD PROCESSING STUDIES)
  - FORMATION OF AOT
- EXXON CO. U.S.A.
  - OVERFLOW CENTER 1974
  - INITIAL STUDIES WITH CONSULTANT 1975
  - DEVELOPMENT OF "HOW TO" HANDBOOK
- OTHER AFFILIATES
  - IMPERIAL OIL
  - EPRC

The initial study of word processing was made in the 1972/73 time frame. This was the time of the IBM push, and what Randy Goldfield referred to earlier as the vendor orientation, trying to sell as many systems in one place as possible. The concept there was to take typewriters away from secretaries; and that was part of the original recommendations in the initial study that was performed within Exxon. That was not very popular. They did modify the concept and did not centralise all the typing for the headquarters building in one location, as the study did not make that recommendation. But it did make the recommendation that the typing be done on each of the separate floors by a small group, and that the typewriters be taken away from the secretaries. That met with a lot of resistance in Exxon. People did not want to lose their secretaries and, as a fallback position, an overflow typing centre was organised, staffed with Vydecs in the administrative services department. After that administrative services typing centre was mature and established, after about a year, they started branching out and establishing mini centres.

These mini centres were one Vydec and one operator. It was located close to the professionals, but it was still on the administrative services staff budget. They charged a flat rate charge to the using department. This allowed the administrative services group to train all of the operators. The only way that an operator would be put in one of these decentralised mini centres was to have had at least six months in the overflow centre. Therefore you knew that you had a good, qualified person coming to serve you in this particular centre.

Before a mini centre was established, a word processing study of the type but probably not the same depth that Randy was talking about earlier, was conducted. The application for that centre was judged and then it was sold to management.

In 1976, because of many reasons which I will discuss later, the advanced office technology team was formed. As I indicated before, Exxon is a very large organisation and it is run almost like separate companies. Activities were going on in many of these other domestic affiliates and all over the world in the word processing function. I came out of Exxon Company USA, which is the domestic affiliate, headquartered in Houston and in 1974, soon after the corporate headquarters established an overflow centre, we established an overflow centre also. We did not go into the mini centre concept, however. We decided that we wanted to form a team and do word processing studies for using departments ourselves. We figured that we did not know very much about it and so we contracted with a consultant to come in and perform the study, and we would ride his coat tails and learn how to do word processing office system type studies.

In November 1974 we completed and made recommendations to two large departments, and those recommendations were basically accepted. We then embarked on forming our own team to do these types of studies. As part of that development, we developed a "How to Conduct an Office System" study handbook. It was our feeling that we would not be able to go all over the Exxon circuit and conduct these office system studies ourselves, and we did not want to staff up for that large team. Therefore, we developed this "How to . . . " handbook and we used it with one organisation before I left, where we only went in and made a presentation to them on how to use the guidelines; and then later, after they had collected all their data, we came back in and discussed the kinds of recommendations that they would be coming up with.

That today is a very successful operation. It is a word processing operation today. At the same time Imperial Oil, in Canada, was developing its own methodology on doing office system studies. They patterned their methodology after the Xerox methodology. Esso Production Research Company — the last item on the slide — was also a very sophisticated user of word processing technology and photocomposition technology and was working with interfaces between the two and with other advanced technologies. So we had a large number of sophisticated users.

AOT ORGANIZATION

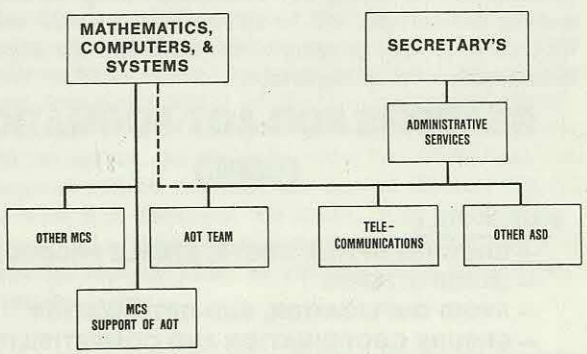

A survey was conducted which surveyed the movement in the change of pace in the office technology market. As a result of that survey our team, or the Advanced Office Technology organisation, was established. It was established within the administrative services department, which is underneath the secretary's department within Exxon. Housed also in administrative services is the telecommunications function.

My boss also had another reporting relationship; it was to the head of the mathematics, computers and systems group, which I understand most of you represent within your own companies. So he had a dual functional reporting relationship, and only a stewardship relationship to the administrative services function. On the 1st of this year that was changed, and the telecommunications function and the advanced office technology team were reorganised and cleaved off of administrative services and put into a new department called the communications and computer sciences department.

REASONS FOR AOT FORMATION

- TECHNOLOGY PUSH
  - — RAPIDLY CHANGING TECHNOLOGY
  - — RAPIDLY SPREADING TECHNOLOGY
  - — MERGING/INTEGRATION OF TECHNOLOGIES
    - + WP/DP/TELECOMMUNICATIONS
    - + MICROGRAPHICS/FACSIMILE/ELECTRONIC INDEXING & RETRIEVAL
    - + VOICE/DISPLAY
    - + ETC.

There were many reasons for the formation of the advanced office technology team. As you know today, and I am sure that many of these have been discussed, there is rapidly changing technology. There is rapidly spreading technology. In Exxon we stopped and did a survey after our team was formed, and we had a hundred word processors of some type within our office building in down town Houston. It was a "me too" type technology like when computers first came out and everybody wanted one. It did not take much justification to get one; all it took was the secretary to convince her boss that she had to have one, and it came in the door. Sometimes he said, "You go ahead and sign it, I don't care what your reasons are. Whatever you want you can have."

There were merging technologies and integration of many new technologies. Word processing, data processing and telecommunications were all merging. Micrographics, facsimile and electronic indexing and retrieval technologies were all merging. We see this continuing and we felt that we needed to have some way of keeping a gauge on those merging technologies and be able to keep a handle on it. We look to see the merging of the voice and the display technologies more in the future.

REASONS FOR AOT FORMATION (cont.)

- USER PULL
  - — GROWING OFFICE COSTS, STABLE PRODUCTIVITY
  - — LABOR INTENSE
  - — AVOID DUPLICATION, SUB-OPTIMIZATION
  - — ENSURE COORDINATION AND COMPATIBILITY
  - — INCREASE EFFICIENCY AND EFFECTIVENESS
  - — MANAGERIAL/PROFESSIONAL VS. CLERICAL/SECRETARIAL

There are also many reasons for the formation of our advanced office technology group because of user demands or user pull. There are growing office costs. I am sure that you have heard before that while there was stable productivity in the United States, the average costs are going up 8% to 12% for office workers, while productivity is staying fairly stable at around 4%; while in the manufacturing industries in the United States there are 80% to 85% gains over the past few years; and in agriculture 150% gains in productivity. These gains are directly attributable to capital investment in the workers in these different industries. In the United States the average capital investment per office employee is around $2,000, while on manufacturing it is more like $25,000.

Half of the employees within Exxon are office workers. The office, of course, is a very labour intensive place. With this many people within Exxon employed in the office small percentage gains in their jobs could result in large gains because of the large multipliers that we have of employees working in the offices. We wanted to avoid duplication and sub-optimisation of system solutions. We want to ensure coordination and compatibility and increase efficiency and effectiveness.

Randy talked a little bit about the managerial/professional versus the clerical/secretarial problem. Advanced office technology, and now our office systems technology, is very interested in concentrating on managerial and professional productivity increases as opposed to clerical and secretarial productivity increases, the reason being that that is where the largest part of the bill is.

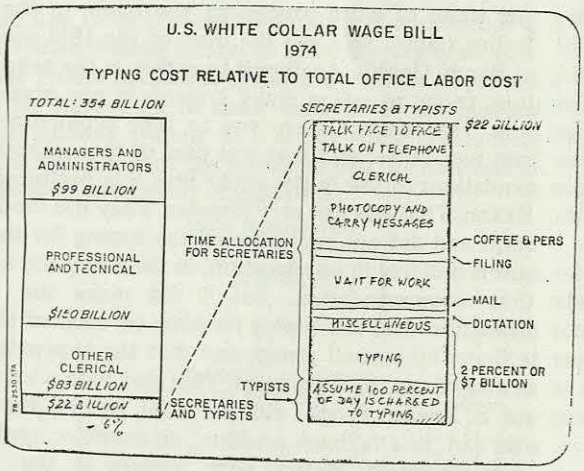

In 1974 the white collar wage bill in the United States totalled $354 billion. Randy quoted the $22 billion figure for secretaries and typists right down here, which is only 6% of the total white collar wage bill in the United States. She also mentioned that out of the secretarys' time only approximately a third of their time, if you throw in all of the typists who type 100% of their day, is spent in the typing function, which of the overall white collar wage bill is only 2% or $7 billion. This is where most of the word processing technology is all aimed at — the typing function. We in advanced office technology and office systems technology are trying to get away from just word processing studies by themselves.

Our activities are broken down into six major areas. Because we are such a wide and diverse organisation we have a planning activity. We not only have the responsibility of setting up a plan for our own group and the New York headquarters building, but we are trying to take on the monumental task of organising a plan for Exxon worldwide; which means that we are sending out and asking the regional

REORGANIZATION AOT TO OST

OST ACTIVITIES

- PLANNING
- COORDINATION
- CONSULTING
- TECHNOLOGY EVALUATION
- METHODOLOGY - TOOLS AND TECHNIQUES
- INFORMATION DISSEMINATION

contacts that we have all over the world to put together their plans on what they are planning to do with office automation, and then we are summarising those, consolidating them, bringing them up and presenting those along with our own plan for office automation for the corporate building in New York.

We have a big coordination function where, because we have so many diverse organisations, each one of us in our office is assigned a different organisation, and we are to act as the liaison between those organisation, to help and provide them with anything that they may need. It may be a visit to Citycorp, which is a very popular visit because of all the publicity that they have had on the system that they have installed; or it may be a visit to one of the Exxon enterprises enterprises — we do arrange those kinds of visits — or other vendors within the New York area.

We do limited consulting. We have the responsibility for the New York headquarters building, and we do some limited consulting in the outlying regions. Next week I will be going to Houston to consult with Exxon Chemical USA. We try to limit our consulting activities and doing these office systems studies to areas where we feel there will be either new knowledge that we will gain about the office of the future and about the integrated electronic office and user needs, or areas where there may be real high pay-offs because of the implementation of a system.

We are involved in technology evaluation. My next slide will talk a little bit more in depth about a number of prototype projects that we have going in our office. We are involved in trying to come up with a methodology. Randy indicated earlier that she would like to see the people who will be able to measure the professional's time. That is a task that we will be trying to undertake. We have seen some tools that have been developed in the area, mainly random sampling type tools, some logging tools, but this summer we will be trying to put together the methodologies that have been developed in the different organisations within Exxon, and trying to extend those methodologies to cover the professionals because, as Randy said, there has been very little work done in trying to estimate what has been done or what is done by professionals. One of the biggest problems is that it is a moving target. Six months ago I am sure that many of you did not do the same work that you are doing today, or in the same way that you are doing it.

The last one is information dissemination. We develop materials; we gain knowledge; and then, as we gain it, we try to disseminate it to the regional contacts that we have, to aid them so that they do not have to reinvent the wheel and make all the evaluations of the technology that is on the market today.

OST PROTO-TYPE PROJECTS

- CORRESPONDENCE CONTROL SYSTEM
- HERMES ELECTRONIC MESSAGE SYSTEM
- CWP INTERFACE/TRANSLATOR SYSTEM
- SLOW SCAN TELECONFERENCING
- INTERACTIVE SYSTEMS ONE (IS/1)
  - -- TEXT EDITING
  - -- TEXT PROCESSING
  - -- ELECTRONIC MESSAGE SYSTEM
  - -- PROGRAMMABLE
  - -- PROGRAMMER'S WORK BENCH
  - -- PHOTOCOMPOSITION INTERFACE

I would like to look at the prototype projects that we are involved in today. The first one is a correspondence control system. We have installed this in our office and we are using only a small piece of this correspondence control system. We are using the piece that allows us to index documents, and then later be able to retrieve them on line in a limited number of categories.

One of the most popular areas of office automation turned out to be, in Houston, the automation of files — the file room. There were a number of large, central file rooms in Houston where I was working that were repositories of records, and if anybody wanted to find something they always went to their desk, or their secretary, or their chrono-file. If it went to the central file room, it was pretty well forgotten, but they still kept a copy in central files. This was a real concern to those people who were the file room clerks and the people who were managing that, and they wanted to be able to improve their service and find documents without having to file multiple copies. This correspondence control system allowed that type of thing where documents would be filed in a sequential manner, and then an abstract would be developed of that document with the author, the recipient, some key words that were associated with it, the date, and a small abstract that told you what that document was about. Then you had the opportunity of retrieving that document on line or in 370 batch-like reports. Many of you may have those types of systems today.

We are also involved with a Hermes electronic message system. Equitable and Exxon and a number of other large, leading edge organisations in office automation have banded together and formed something called the Office Automation Round Table. Through that Office Automation Round Table a proposal was made to that group of people to adopt this electronic message system for their internal communications amongst that particular group of people. Since we were on the Round Table we got a Silent 700 Texas Instruments

terminal and installed it in our headquarters building. I think that we have received maybe five messages in three months from the Round Table group; and, of course, this is one of the problems with electronic message systems; just like the telephone, if you do not have somebody on the other end there is not much sense in sending messages.

But with the reorganisation that took place on the 1st of the year, my boss now has two offices. He has an office in New York City and one in New Jersey; and he has used the Hermes electronic message system to keep himself in touch between those two offices. He sits down at his terminal and keys in messages. After reading some of the literature and being convinced of some of the things that have been said in the literature, that when you pick up the phone and make the telephone call, and the fellow at the other end is not there because he is out of the office, because he is temporarily away from his desk, or because he has a busy line, he has found that he has had to go back and forth and keep remembering that he has got to call him back. It is a whole lot easier for him to sit down and type out a few quick lines on his terminal and send that over to us, and he can forget about it then. He has found it very effective and he uses it — probably more than I would like to see him use it.

We have been involved, ever since I joined the group over a year and a half ago, in trying to find a universal black box that will make two word processing machines talk to each other. I would imagine that subject has been covered to some extent. We have found that users have naively bought the communication options on a Vydec machine, tried to dial up a Wang machine and make them talk to each other, and found out that all they get is garbage on each end. It has caused them some real headaches and frustrations thinking, of course, that if you have a communicating word processor you can talk to everybody in the world. I think most everybody here knows that that is not true. There are many problems attendant with that, as John Gosden was telling us. There are many presentations that can be made within here, and we have people in our group who are working closely on this particular problem.

At one time we had a person who was forming a company to make this black box. He did not make it! I understand that we now have contacts with another company which has a lot of experience in translating and making conversions between different word processing codes. They are supposed to be producing a black box in the near future that will enable us to do this interface between dissimilar word processors. We are also involved in a slow scan teleconferencing business test.

When you hear the scenario of the office of the future, you hear about full motion video teleconferencing and the fact that that is around the corner and it is going to be here. The cost of that today, just in telephone line charges, is around $4,000 an hour. With the slow scan teleconferencing, we use normal telephone lines to do the slow scan teleconferencing, and so the costs come down significantly to more like the area of $25 an hour. This slow scan teleconferencing is mainly used for meetings where people know each other. They are meetings where people come together on a fairly periodic basis, and they know each other, so there is not that "I need to get to know you on the other end of the line" type problem.

The slow scan sends a picture of a viewgraph such as this down the line every 30 seconds; a new picture can be generated every 30 seconds. Another area in which we are involved is interactive Systems One or IS/1, which is an advanced office technology or integrated office system. It has text editing capabilities, text processing capabilities, and a very fine electronic message system. It is programmable. It has a programmer's work bench. One of the main purposes of the development of this system was to aid programmers as they develop programs. This system is based on the Unix system which was developed by Bell Labs.

I saw some heads nod when I mentioned Unix. This company bought a marketing licence from Bell Labs and is now marketing the system for them, selling object code for it.

One of the real benefits that we have found in these prototypes is the ability to get a better handle on what we feel the office of the future is for Exxon. The White House did a great thing for the advanced office technology market. About the middle of last year they published an RFP where they asked for the ultimate system. Most of the time we are a vendor driven market. We live in a vendor driven market. The word processing people put out packages and they want to sell systems. The same with the data processing people; they sell more of a general purpose type of system where you make your own. It has been a vendor driven type of thing. The White House said, "This is the kind of system we want." It was a document about that thick. It was very detailed and a very well-written document. One of the things that we have done is to look at that document, look at all of these prototype projects, and look at the needs that we have seen within Exxon. We have developed our own start on some functional specifications for an integrated office system. I do not have time to go into those right now, but I have extra copies of those specifications. One is an overall list of all of the functions that we feel ought to be in the integrated office system. The next one is the text editing comparison criteria list. The third one is a description of a good electronic message system. The fourth is an interim electronic filing system. The last is a calendar management system.

We have found that if you wait for the vendors to put out a product it will not be what you want. They will usually put out a product that is not close to your specifications and it will take a year or two of their getting feedback on their product before they will really come up with what will be usable in the market place.

ADVANCED OFFICE TECHNOLOGY PROJECT
SCOPE

- BUSINESS COMMUNICATIONS SYSTEM
  - -- INCLUDES ALL BUSINESS COMMUNICATIONS WHICH ARE DOCUMENTED OR RECORDED.
  - -- INCLUDES INTERNAL BUSINESS COMMUNICATIONS AND BUSINESS COMMUNICATIONS FROM/TO THE OUTSIDE.
- CONSISTS OF SEVEN COMPONENTS
  - -- CREATION - ACT OF THINKING AND FORMULATING A COMMUNICATION
  - -- CAPTURE - PLACING THE COMMUNICATION ONTO A MEDIUM

Let me quickly go into the business communications system as an introduction to the development of costs within an organisation of the documented communications. The business communications system includes, as far as we are concerned, all documented business communications. It includes all internal and external communications. We look at the seven components of the business communications system as being, first of all, creation; that time that you take as you are thinking about a document and creating it in your mind before you actually take the next step, which is to capture that document in some type of medium, whether it be handwritten, dictation, or actual keyboarding.

In Exxon 90% of the documents that are created are captured through handwriting; 5% are dictated to secretaries; 3% are dictated to machines; and 2% are keyboarded directly. Many of those are composed by secretaries themselves, some by professionals.

ADVANCED OFFICE TECHNOLOGY PROJECT

SCOPE

- BUSINESS COMMUNICATIONS SYSTEM COMPONENTS (CONTINUED)
  - -- "KEYBOARDING" - ENTRY TO/PROCESSING BY/OUTPUT FROM A KEYBOARD
  - -- DISTRIBUTION - MESSAGE CARRYING, MAIL HANDLING, ELECTRONIC TRANSMISSION
  - -- EXPANSION - COPYING, PRINTING, MICROFILMING, DUPLICATION OF MAGNETIC RECORDS
  - -- STORAGE AND RETRIEVAL - INDEXING, STORING, SEARCHING FOR AND FINDING INFORMATION
  - -- DISPOSAL

The keyboarding of the document is the next stage in the life of a business communication — these are documented business communications now. Distribution is the next step. It may be distributed through the mails, by carrying a document. We have found around 10% of documents produced are hand-carried to their ultimate destination; or, in a very few cases at this time, electronic transmission. Expansion of document is another stage in the life of a document. In Exxon we have about 19 copies of every original that we produce made in our quick copy facilities and in our central reproduction rooms. So there is quite an explosion factor of paper. I am sure that in your company it is somewhat similar.

Out of those 19 copies that are made on the original, eight to nine of those are stored somewhere, in somebody's files. So a lot of storage goes on. The last step of the business communications system is the disposal of those documents. Of course, in here all of the documents that you want are disposed of right away, and the ones that you do not need are kept forever. We have a lot of 'kept forever' documents in Exxon; because of some of the suits that are being brought against us, many of our documents are kept forever.

With that type of background we have used those stages of a document life to try to develop the costs of the business communications system within Exxon. The reason that we have done this is because we feel that once we have the costs of the documented business communications system, we have a basis for measuring our incentives that we might be able to claim or project, or measure, if we have been able

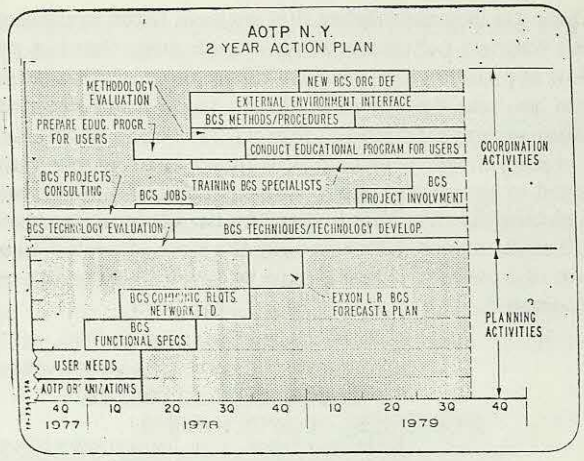

to be that sophisticated.

It helps to define the scope and size of the problem. If most of your cost is in the storage and retrieval of documents, then that is where you need to spend the money and the dollars in order to get the pay-offs. However, if most of your information is distributed and sent out, then maybe the distribution channels are the ones that need to get your main attention in your company.

So by developing the costs of the business communications system we were able to home in more closely on the problem areas and the areas where we can get the biggest pay-offs. This also enhances our ability to communicate with managers, because we can tell them and we can talk to them about how much it costs in each of these areas of our business communications system.

COST MEASUREMENTS

- REASONS FOR MEASURING COSTS
  - -- BASIS FOR MEASURING INCENTIVES
  - -- HELPS DEFINE SCOPE OF PROBLEM, SIZE, EFFORTS
  - -- ENHANCES MANAGEMENT COMMUNICATIONS
- APPROACHES TO MEASURING COSTS
  - -- PAPER FLOW
  - -- PEOPLE COSTS

ESTIMATING BENEFITS OF OFFICE AUTOMATION

SUMMARY OF TIME SPENT BY OFFICE WORKERS (%)*

BCS COMPONENT	MPT	OCS SECY	OCS CLERICAL
CREATION	45	13	26
CAPTURE	9	6	19
KEY BOARD	—	38	8
EXPANSION	1	10	8
DISTRIBUTION	5	14	2
STORAGE/RETRIEVAL	18	11	27
DISPOSAL	—	2	2
OTHER (NON-WRITTEN COMMUNICATIONS, ETC.)	22	6	8
*NOT ACTUAL FIGURES	100	100	100

There are two approaches that we have taken in developing this business communications system cost. One is a paper flow approach where we have taken many of the statistics that we have developed in the 40 to 50 word processing office system type studies that we have done within Exxon, and used those to project how much time people would spend in generating paper within Exxon. The other way is to take the time that is spent by people in various tasks and make projections based on the time allocated to each task of how much time would be saved if the office was automated.

PAPER FLOW APPROACH

ASSUMPTIONS AND COMPUTATIONS

- 1,000 PROFESSIONALS IN ORGANIZATION
- 300 SECRETARIES
- 100 CLERICAL AND OTHER OFFICE WORKERS
- 50 LINES OF TYPING/DAY/PROFESSIONAL
- 25 LINES OF TYPING ON AVERAGE PAGE
- 10 WORDS PER AVERAGE LINE OF TYPE
- 225 WORKING DAYS PER YEAR

I should like to briefly go into a mythical organisation. This is not Exxon. One thousand professionals in an organisation, with 300 secretaries, which is about a 3 to 1 relationship; with 100 clerical people and other office workers. We found in the studies that we performed that professionals, managers, all the people in that category, produce around 50 lines of text a day. There are 26 lines of typing on the average page, we have found within our company, and 10 words on the average line, and professionals and office workers work an average of 225 days out of the year.

PAPER FLOW APPROACH

ASSUMPTIONS AND COMPUTATIONS

- AVERAGE SALARY, BENEFIT AND BUILDING OVERHEAD COSTS:

  -- PROFESSIONALS

  $30,000/YEAR; $18/HR; 30¢/MIN

  -- SECRETARIES AND OTHER EMPLOYEES

  $15,000/YEAR; $9/HR; 15¢/MIN

1000 PROFESSIONALS X 50 LINES/DAY ÷ 25 LINES/PAGE X 225 DAYS/YR = 450,000 PAGES TYPED PER YEAR

The salary is shown here. These again are not necessarily Exxon figures, but are used in this example to help to develop the costs of the business communications system. The professionals would earn around $30,000 a year or $18 an hour. That includes all of the building overhead costs and all of the costs for professionals. Secretaries earn around half that amount of money. If you take all of those assumptions that we have just made and multiply them together, we find that this particular organisation would produce 450,000 pages of paper each year.

Taking those 450,000 pieces of paper, and these assumptions that we are making on those 450,000 pieces of paper, this last assumption is fairly well backed up by the statistics that we have found in the word processing studies and office system studies that we have done, where 65% of the documents that we produce are original documents, 31% are revised documents, and 6% are repetitive letters or boiler plate type material.

In looking at the creation, the time it takes to think and formulate a document in your mind, if we make the assumption that one half hour of professional time is taken for each page that is produced, and one-sixth hour of creation time is taken for each revised page, then we can compute the costs of creation by applying those factors and assumptions that we have made, and come up with about $2½ million for original pages, about $½ million for revised pages, and repetitive letters taking a small amount, coming up with a total cost of about $3 million for this type of mythical organisation to just create a document.

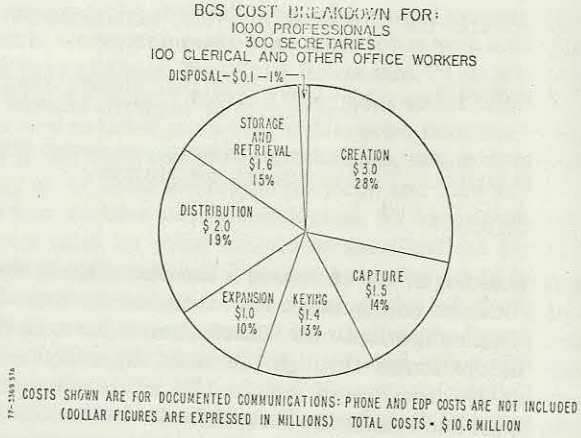

I am not going to go through the whole business communications system and show you all of the assumptions that we have made in coming up with this particular pie chart. I have included those assumptions in the viewgraphs that I have left with Butler Cox and they will be distributing those in the conference materials. (Editor's note: this material is included in the transcript at the end of the session.) But this mythical organisation would come up with a total of about $10 million for their business communications costs, making the assumptions that we have made here. We feel that these are fairly conservative assumptions: 28% creation; 14% capture and so forth. This would be the way that the business communications costs would break down on a paper flow approach.

Now taking the other route and using the people approach, we have verified some studies that have been done in other places. The studies that were used divided the tasks being performed by the MPTs (managerial, professional and technical people) and the secretarial and the clerical people and broke the work that these people did down into 40 tasks. We took those tasks and recombined them into the seven components of the business communications system. We came up with these figures.

COST MEASUREMENT BASED ON PAPER FLOW

- NEED TO KNOW
  - PERSONNEL NUMBERS: PROFESSIONAL, SECRETARIAL, OTHER
  - COMPENSATION
  - LINES OF TEXT PRODUCED/TYPED BY PROFESSIONALS/SECRETARIES
  - ESTIMATED PROFESSIONAL/SECRETARIAL TIME ALLOCATION

The same type of thing can be done in your companies by doing some time studies on the secretaries, which is a fairly well developed science; and also doing some studies and using some knowledge of your own company operations and developing the professional time allocation that they spend in the different tasks that they perform.

Now taking the latter approach and using the costs that were developed, using these percentage figures and applying those against the costs of the salaries of each of those groups of people, you can come up very easily with the costs of the business communications system within your company.

Once that is done, it is a fairly easy process — well, it is not a fairly easy process, but there is a process whereby you can go through and look at each area of the business communications system and project the benefits that will accrue to your office through the implementation of an automated office system.

We have done that in Exxon and we have come up with the following figures. The figures represent the time spent in each of the categories, MPTs, secretarial and clerical people, and are the categories that came from the previous slide. The time savings are shown in percentage figures and are the projections that have been made by people who are very close to the situation, have done a number of studies and have used studies done by other people.

ESTIMATING BENEFITS OF OFFICE AUTOMATION

SUMMARY OF TIME SAVINGS
BY SHIFTING TO AN ELECTRONIC ENVIRONMENT*

	MPT		SECY		CLERICAL	
BCS COMPONENT	TIME SPENT	TIME SAVED	TIME SPENT	TIME SAVED	TIME SPENT	TIME SAVED
CREATION	45	7	13	6	26	8
CAPTURE	9	2	6	4	19	15
.	.	.	.	.	.	.
.	.	.	.	.	.	.
.	.	.	.	.	.	.
DISPOSAL	—	—	2	2	2	2
OTHER	22	—	6	—	8	—
	100	22	100	39	100	41
IMPLEMENTATION FACTOR		.7		.7		.7
% TIME SAVED		15		27		29

*NOT ACTUAL FIGURES

IBM 10-20% 30-40% 25-40%

We feel that not all of the benefits of an integrated electronic office can be captured, and so therefore we are applying a 70% fudge factor against those in coming up with the percentage that we feel could be captured of an integrated electronic office, of 15% for MPTs, 27% for secretaries, and 29% for the clerical staff.

IBM made a presentation to us about three months ago, where they said that their estimation was that the implementation of an electronic office would save MPT people about 10% to 20% of their time, and secretarial and clerical would be in the range of 30% to 40%, and the clerical people would be in the range of 25% to 40%. Other studies that we have seen in the same area are right in that same ballpark. So we felt fairly verified in the figures that we came up with when IBM came up to us and we were right around their mid point ranges.

ESTIMATING BENEFITS OF OFFICE AUTOMATION

	MPT	SECY	CLERICAL	TOTAL
TOTAL BCS COSTS	7,600	100	900	10,600
% TIME SAVINGS	15	27	29	19
COST REDUCTION (DISPLACEMENT) $M	1,100	600	300	2,000

Once those figures have been developed, those percentage projections have been developed, they can be applied against the total business communication costs. Then totals can be developed where we can see that out of the total business communication costs within this mythical organisation of $10 million, 19% of that could be saved through the implementation of an integrated electronic office, or about $2 million.

Now those are all sort of 'iffy'. Presentations like this have been made within Exxon to different management groups to try to create an awareness of the potential of this field, not projecting that these are actual hard dollar savings that can accrue to Exxon.

The next area that I would like to approach is that of the study tools and the methodologies that we have employed within Exxon to perform these office system studies.

STUDY TOOLS

- PROFESSIONAL QUESTIONNAIRES
- SECRETARIAL QUESTIONNAIRE
- TIME ANALYSIS TOOLS
  - -- TIME LADDER
  - -- TIME LINE
  - -- RANDOM SAMPLING
  - -- SUPPORT ACTIVITIES LOG
- DOCUMENT PRODUCTION (TYPING) LOG
- MAIL LOG
- COPIER LOG
- MUSTS/WANTS LIST

We have developed professional and secretarial questionnaires. For each one of the tools that I show on this slide we have developed computer programs so that

we can have the results that come back from the people who are using these tools. They will be key punched and then fed into some programs which will do some analysis work on them. The professional and secretarial questionnaires are mainly used in our studies as pre-interview questionnaires, to try to bring out of the management groups and the people that we will be surveying the preferences and needs that they have. We have a number of time analysis tools that have been developed within Exxon: time ladders, time lines, random sampling, and support activity logs. We also have a document typing survey, which asks that secretaries take a carbon copy of everything they type for two years, and then take one of those pages and code on the page what they did with that document, how long it took them to do it, what type of document it was, and so forth.

We also have a mail log that we have used quite successfully. I will talk a bit more about the importance of that later. We also have a copier log. I did not want to take time to go through each one of these study tools and so I have prepared a packet of all of the study tools that we use in these office system studies. I have left them with the receptionist, along with the other package that I talked about earlier on the functions within an integrated electronic office, and any of you are free to pick that up if you would like.

One of the last things that we do is to develop a musts/wants list, which I am sure is no new thing to data processing people. It was a new thing to me up to about a year ago, and it was a very valuable tool that we used. It is something that was used effectively in one particular organisation, where we developed a musts/wants list of what people wanted in the office. One of the nice things about it is that since the office of the future is so far away, most everything that is on their musts/wants list is also in the office of the future. You can talk about it, but you cannot give it to them.

We are involved in a three person project to develop a study methodology. We are in the middle of trying to put that together. We have gone about a third of the way. We are right at the point now where we have defined all of the data that we would like to gather, and now we want to take the data that we want to gather and go back to the tools that we have developed, and try to make sure that the tools gather the data that we have, and then look to see what other kinds of tools we need to develop. The main area where we will need to be developing tools will be in the analysis of professional time. We already know that. I would like to show you quickly the study methodology that was developed and is put forward in the "How to . . . " book that was developed in Houston.

Study Activities

PREPARING FOR THE STUDY	**1. Department Management Presentation**
	2. Developing the Study Team
	3. Announcement of Study
	4. Orientation Meetings
CONDUCTING THE STUDY	**5. Time Ladders (optional)**
	6. Secretarial Support Survey (optional)
	7. Secretarial Interviews
	8. Typing Survey
	9. Random Sampling
	10. Professional Interviews

STUDY ACTIVITIES (CONT'D)

SYSTEMS DESIGN & PRESENTATION	**11. Data Analysis**
	12. System Design
	13. Prepare Report
	14. Management Presentation & Approval
IMPLEMENTATION	**15. Presentation of System to Professionals & Secretaries**
	16. Staff New System
	17. Coordinate Equipment & Space Changes
	18. Train Staff
	19. Start-Up
	20. Evaluate for Phase 1 Changes **Evaluate for Phase 2 Changes** **Evaluate for Phase 3 Changes**

This is basically word processing study methodology. I am sure that it is very similar to data processing system methodology in many cases. However, I have been asked that question many times: what is the difference between a data processing study and an office systems study? I am trying to answer that question. One of the answers that I have is that a data processing system study very often is a top down type of approach, whereas an office systems study combines not only the top down type of approach but also the bottom up type of approach, to look at the basic detailed tasks that are performed and try to find out what is actually happening in the office. Preparing for the study, developing the study team, announcing the study, orientation meetings.

Then we go into the data gathering phase where we have different tools that were used within Exxon. We always did our professional interviews last because we felt they were the most crucial. We wanted to be prepared for those, and we usually had a lot of our data messaged by the time we got to those, so that we could then talk a little more intelligently to the professionals. After the data was gathered we went into a systems design and presentation of the recommendations that were developed, and then into an implementation phase. We tried to do it in an evolutionary rather than a revolutionary way.

ORGANIZATIONAL ALTERNATIVES

- CENTRAL
- CLUSTERS
- DEDICATED OPERATORS
- SHARED/CASUAL USERS
- DEDICATED EQUIPMENT

I have seen a number of outcomes of word processing studies and I have seen all of them used effectively. A lot of it has to do with the management backing and how much they want the system to work.

You have heard about centralised word processing systems where all typing is done in one place. You have heard of clustered centres where two, or three, or four word processors are dedicated in one particular place and perform the typing function. I have talked about the mini centre concept where there is only one word processor and an operator that is dedicated in a particular location, close to a department. Often I have seen it developed within that department, when they feel that that is the way for them to do it.

Another one that I have seen that has increased in popularity is the shared or casual users, where you get a word processor and it almost always has to be a CRT device that is very easy to learn. CRT devices are very easy to learn. An operator can come up in a very short amount of time, can learn how to use that system, and can use it effectively. I have seen it where two secretaries have shared one word processor, and I have also seen it where five secretaries shared two word processors, very effectively and very happy about what they had.

There is also the case where the word processor is bought and installed for a secretary who may be using it only 25% of the time, but when she has to get something out for the controller of the company, it has to go out, and it has to go out now, and it is going to be revised many times. So you justify it on the fact that his job is important and therefore you give him the service that he needs.

PROBLEMS

- MANAGEMENT BACKING
- CLIENT INVOLVEMENT
- BUZ WORD SYNDROME
- CRITICAL MASS
- LACK OF STANDARDIZATION
- TECHNOLOGY GAPS

There are many problems. One is management backing. It is important and crucial in all areas of the implementation of an automated office. It is important in the beginning. It is important in the middle. It is important in the end.

Just a quick description of the two word processing studies that we undertook in Exxon Company USA with the consultant. One of the two has captured the savings in people that was projected because the management took an active interest in the implementation of that particular system. The other department did not capture the projected people savings. It was to the tune of five secretaries that were no longer needed. Fortunately, in Exxon we have 25% attrition and therefore we did not have to fire anybody; they were absorbed.

Client involvement. When we do an office systems study we are extremely interested in having a management contact and having his backing all the way throughout it. If we can get him on the study team we feel much better about it. If we cannot get him on the study team we ask him for somebody else to be involved on the study team with us, so that they can be a part of the recommendation that is made.

We feel that it is important to avoid the buzz word syndrome. It is easy to get caught up in talking a new computer-ese language which is called office systems technology. It turns users off very quickly and it is something that needs to be watched out for. The critical mass concept has been alluded to before. I heard about an installation where the people went out and bought an electronic mail system on a "Me too" basis — I want one too — and in one case it flopped and in the other it did not. In one case they had 70% of their communications that stayed internal within their group, and in the other case they had 10% of the communications that stayed within the group. It was not too hard, after they went back and took a strong look at what they were doing, to find out why their electronic message system was never used.

Lack of standardisation within the industry, not only in communications but also in training. Once you train an operator on one system they cannot use another system because they are so dissimilar. There are some real problems in the lack of compatibility. There are technology gaps within the industry. The office of the future is somewhere around the corner. I have heard my boss say a couple of times that he thinks it is further away the more he gets into it.

TECHNOLOGY GAPS

- INFORMATION CAPTURE
  - — SPOKEN WORD RECOGNITION
  - — WRITTEN/PRINTED WORD RECOGNITION (OCR)
- INDEXING AND RETRIEVAL
  - — AUTOMATIC INDEXING
  - — "FULL TEXT" INDEXING
  - — "PLAIN ENGLISH" RETRIEVAL
- INTERFACES/COMPATIBILITY

There is the spoken and written word. One of these days we will be able to speak into a microphone and see that projected on a screen. That is coming. It is here today. It is in the labs. They say that they have them with about 80% recognition. That is not good enough to put on the market, so they are not on the market today. So that is somewhere around the corner, in the five to ten years time frame.

OCRs. All those documents that are generated extenally would be nice to have in your system to be able to retrieve, but you have to re-key them in order to get them into your system. You can buy an OCR that will get them into your system, but it only costs $2 million for each of our departments to buy to put documents into a system. So that technology leaves a real gap for us. We anticipate that will be solved somewhere in the future.

The storage and retrieval of documents is also an area where there are some real technology gaps. Once you have a

document that is captured in a word processing system or an integrated office system, how are you able to retrieve that document? You have a lot of key words that you would like to use, but being able to take 500 pages of text and be able to find the right key words and the right handles that you want to retrieve that document are not easy. We have not seen any software on the market today that really does a good job of automatic indexing of documents so that they can be searched for on some key words without human intervention. Compatibility is an important issue and a real area of technology gap.

Don't believe anything until you see it; and then don't believe it — until you have had in installed for at least six months. We have had some grief with vendors who have made all kinds of wonderful claims about the way that their products were going to perform. We start banking on some of the things that they are saying. We go to the demonstration and we ask about three pertinent questions, and we walk away, very disappointed. Even though they may be telling a good story, it is probably not there, mainly because it is some engineer in the background, in the labs, who has been creating a product in a vacuum. They have finally got it to the market and it is being pushed by the vendor.

We are interested in Exxon in trying to influence vendors. One of the reasons that we developed this integrated office functions list was because we wanted to influence vendors. We have already shared it with a number of vendors and plan to share that document with still more vendors in the near future.

COX: Rick, we are very grateful to you for your contribution, particularly for all the work you have done in preparing the sample figures that we will be distributing as a guide to analysing the costs and the benefits, and for the package of analytical tools, which I am looking forward to exploring with great interest. Many thanks for sharing your experience with us.

APPENDIX

INCENTIVES ASSOCIATED WITH:

CREATION

1. REDUCE CREATION TIME
2. MINIMIZE RE-CREATION
3. PROMOTE JOB ENRICHMENT

CAPTURE AND KEYBOARDING

1. LESS TIME REQUIRED TO GENERATE FINAL DOCUMENT.
2. SECRETARIAL TIME SAVINGS WILL RESULT IN:
   - INCREASED ABILITY TO ASSUME ADMIN DUTIES
   - REDUCED PERSONNEL REQUIREMENTS
3. SAVINGS IN PROFESSIONAL TIME FROM DICTATION EQUIPMENT.
4. POTENTIAL REDUCTION WITH ADVENT OF VOICE INPUT.
5. REDUCED RE-KEYING OF DOCUMENTS NEEDED TO BE COMMUNICATED, PHOTOCOMPOSED, OR STORED ELECTRONICALLY.
6. REDUCED COSTS IN GRAPHIC ARTS PREPARATION.
7. OCR FACILITATES HANDLING OF EXTERNALLY GENERATED INFORMATION.

INCENTIVES ASSOCIATED WITH:

STORAGE AND RETRIEVAL

1. SHARPLY DECREASED COSTS.
2. "FRIENDLY" PLAIN ENGLISH RETRIEVAL.
3. OCR INPUT OF EXTERNALLY GENERATED INFORMATION.
4. ENHANCED RETRIEVAL.

INCENTIVES ASSOCIATED WITH:

DISTRIBUTION

1. INCREASED TIMELINESS
2. IMPROVEMENT IN CREATION
3. EASIER/MORE TIMELY RETRIEVAL
4. REDUCED MAILING COSTS, EXPANSION, STORAGE

INCENTIVES ASSOCIATED WITH:

EXPANSION

1. FEWER HARD COPIES FOR OFFICE USE OF DOCUMENTS ELECTRONICALLY STORED.
2. REDUCED STORAGE AND DISTRIBUTION COSTS.
3. FEWER HARD COPIES FOR USE OUTSIDE OF OFFICE WHEN ELECTRONICALLY ACCESSIBLE THROUGH PORTABLE TERMINALS.

INCENTIVES ASSOCIATED WITH:

DISPOSAL

1. EFFICIENT RETENTION PROCEDURES

2. LOWER DIRECT DISPOSAL COSTS

CAPTURE

A. COST ESTIMATES

- HANDWRITTEN ORIGINALS
  450,000 PAGES X 56% X 250 WORDS/PAGE ÷ 15 WPM X 30¢/MIN = $1,260,000
- HANDWRITTEN REVISIONS
  450,000 PAGES X 31% X 30 WORDS/PAGE/REVISION ÷ 15 WPM X 30¢/MIN = 80,000
- STENO DICTATION
  450,000 PAGES X 5% X 250 WORDS/PAGE ÷ 30 WPM X (30¢/MIN + 15¢/MIN BOTH PROFESSIONAL & SECRETARY ARE PRESENT) = 80,000
- MACHINE DICTATION
  - ROUGH NOTES TO DICTATE FROM
    450,000 PAGES X 2% X 100 WORDS/PAGE FOR ROUGH NOTES ÷ 15 WPM X 30¢/MIN = 20,000

CAPTURE

A. COST ESTIMATES (CONTINUED)

- MACHINE DICTATION (CONT'D)
  - ACTUAL DICTATION
    450,000 PAGES X 2% X 250 WORDS/PAGE ÷ 60 WPM X 30¢/MIN = 10,000
  - DICTATION EQUIPMENT COSTS
    1000 PROFESSIONALS X 10% X $140/YR = 10,000
- HANDWRITTEN REPETITIVE
  450,000 PAGES X 6% ÷ 10 X 250 WORDS/PAGE ÷ 15 WPM X 30¢/MIN = 10,000

$1,500,000

CAPTURE

B. ASSUMPTIONS

- 56%, 31%, 5%, 2%, 6% = ORIGINAL, REVISION STENO, MACHINE DICTATION, REPETITIVE TYPING, RESPECTIVELY
- 15 WPM, 30 WPM, 60 WPM = HANDWRITING, STENO, MACHINE DICTATION SPEEDS, RESPECTIVELY
- 10% OF PROFESSIONALS HAVE DICTATING MACHINES

KEYING

A. COST ESTIMATES

- STANDARD TYPEWRITERS USED AT SECRETARIAL STATIONS
  - SECRETARIAL TIME
    450,000 PAGES X 50% X 25 LINES/PAGE ÷ 80 LINES/HR X $9/HR = $ 630,000
  - EQUIPMENT COSTS
    450,000 PAGES X 50% ÷ 1,800 PAGES/YR X $200/YR = 30,000
- NON-CRT WP USED AT SECRETARIAL STATIONS
  - SECRETARIAL TIME
    450,000 PAGES X 30% X 25 LINES/PAGE ÷ 95 LINES/HR X $9/HR = 320,000
  - EQUIPMENT COSTS
    450,000 PAGES X 30% ÷ 2,600 PAGES/YR X $3,000/YR = 160,000

KEYING

A. COST ESTIMATES (CONTINUED)

- GRAPHICS DESIGN COSTS (INCLUDES FORMS DESIGN PHOTOCOMPOSITION AND VISUALS)
  - 1000 PROFESSIONALS X $80 PER PROFESSIONAL PER YEAR = $ 80,000

TOTAL COST ESTIMATES $1,400,000

KEYING

A. COST ESTIMATES (CONTINUED)

- NON-CRT WP USED IN CENTRAL CENTER OR MINI CENTER
  - OPERATOR TIME
    450,000 PAGES X 10% X 25 LINES/PAGE ÷ 120 LINES/HR X $9/HR = $ 80,000
  - EQUIPMENT COSTS
    450,000 PAGES X 10% ÷ 8,600 PAGES/YR X $3,000/YR = 20,000
- CRT WP USED IN CENTRAL CENTER OF MINI CENTER
  - OPERATOR TIME
    450,000 PAGES X 10% X 25 LINES/PAGE ÷ 200 LINES/HR X $9/HR = 50,000
  - EQUIPMENT COSTS
    450,000 PAGES X 10% ÷ 14,400 PAGES/YR X $6,600/YR = 20,000

B. ASSUMPTIONS

	STANDARD TYPEWRITER AT SEC'Y. STATION	NON-CRT WP AT SEC'Y. STATION	NON-CRT WP IN CENTRAL OR MINI CENTER	CRT WP IN CENTRAL OR MINI CENTER
VOLUME OF TYPING ACCOUNTED FOR	50%	30%	10%	10%
EFFECTING TYPING RATE (ALLOWING FOR INTERRUPTIONS/PERSONAL TIME/SUPERVISION)	80 LINES/ HOUR	95 LINES/ HOUR	120 LINES/ HOUR	200 LINES/ HOUR
TIME SPENT ENGAGED IN TYPING FUNCTION EACH DAY	2.5 HR./DAY	3 HR./DAY	8 HR./DAY	8 HR./DAY
ANNUAL PRODUCTION	1,800 PAGES	2,600 PAGES	8,600 PAGES	14,400 PAGES
EQUIPMENT COST	$200/YR. (PURCHASED)	$3000/YR. (RENTED)	$3000/YR. (RENTED)	$6,600/YR. (RENTED)

- GRAPHICS DESIGN COSTS EQUAL ABOUT $80 PER YEAR PER PROFESSIONAL.

EXPANSION

A. COST ESTIMATES

- PROFESSIONAL TIME
  1,000 PEOPLE X 1% X $30,000/YR = $ 300,000
- OTHER EMPLOYEES
  400 PEOPLE X 6% X $15,000/YR = 360,000
- CENTRAL PRINTING FACILITY COSTS
  1,000 PEOPLE X $150/PROFESSIONAL = 150,000
- CONVENIENCE QUICK COPIER COSTS
  1,400 PEOPLE X $150/EMPLOYEE = 210,000

$1,000,000

DISTRIBUTION

A. COST ESTIMATES

- PROFESSIONAL TIME
  1,000 PEOPLE X 2% X $30,000/YR = $ 600,000
- OTHER EMPLOYEES
  400 PEOPLE X 10% X $15,000/YR = 600,000
- COST OF DISTRIBUTION CLERKS
  1,000 PROFESSIONALS X $100/PROFES. = 100,000
- MAIL ROOM COSTS
  1,000 PROFESSIONALS X $65/YR = 70,000
- POSTAGE/DELIVERY COSTS
  450,000 PAGES X 25¢/PAGE = 110,000
- ELECTRONIC COMMUNICATIONS INCLUDING CABLES AND FACSIMILE TRANSMISSIONS
  1,000 PEOPLE X $500/PROFESSIONAL = 500,000

$2,000,000

INFORMATION STORAGE AND RETRIEVAL

A. COST ESTIMATES

- PROFESSIONAL TIME
  1,000 PEOPLE X 3% X $30,000 = $ 900,000
- OTHER EMPLOYEES
  400 PEOPLE X 7% X $15,000 = 420,000

CENTRALIZED REFERENCE INFORMATION COSTS
1,000 PROFESSIONALS X $250/YR = 250,000

$1,600,000

EXPANSION

B. ASSUMPTIONS

- 1% OF PROFESSIONAL TIME SPENT IN EXPANSION ACTIVITY.
- 6% OF OTHER EMPLOYEE'S TIME SPENT IN EXPANSION.
- MANPOWER/EQUIPMENT/SUPPLIES/SPACE COSTS FOR OPERATING LITHOGRAPHY/DUPLICATING/QUICK COPY SERVICES IN A CENTRAL PRINTING FACILITY EQUALS ABOUT $150 PER YEAR PER PROFESSIONAL.
- EQUIPMENT AND SUPPLIES COST FOR CONVENIENCE QUICK COPIER EQUIPMENT EQUALS ABOUT $150 PER YEAR PER EMPLOYEE.

DISTRIBUTION

B. ASSUMPTIONS

- 2% OF PROFESSIONAL TIME SPENT IN DISTRIBUTION ACTIVITY.
- 10% OF OTHER EMPLOYEE'S TIME SPENT IN DISTRIBUTION.
- COST OF DISTRIBUTION CLERKS EQUALS ABOUT $100 PER YEAR PER PROFESSIONAL.
- MAIL ROOM COSTS EQUALS ABOUT $65 PER YEAR PER PROFESSIONAL.
- POSTAGE/DELIVERY COSTS EQUAL ABOUT 25¢ PER PAGE OF TYPING.
- ELECTRONIC COMMUNICATIONS COSTS (CABLES AND FACSIMILE TRANSMISSIONS) EQUALS ABOUT $500 PER YEAR PER PROFESSIONAL. THESE COSTS EXCLUDE PHONE AND DATA TRANSMISSION COSTS.

INFORMATION STORAGE AND RETRIEVAL

B. ASSUMPTIONS

- 3% OF PROFESSIONAL TIME SPENT IN INFORMATION STORAGE AND RETRIEVAL ACTIVITY.
- 7% OF OTHER EMPLOYEE'S TIME SPENT IN INFORMATION STORAGE AND RETRIEVAL ACTIVITY.
- CENTRALIZED REFERENCE INFORMATION COSTS EQUAL ABOUT $250 PER YEAR PER PROFESSIONAL. THESE COSTS INCLUDE THE LIBRARY, REFERENCE SERVICES AVAILABLE FROM ON-LINE DATA BASES, CENTRAL FILE ROOM OPERATIONS AND ARCHIVAL STORAGE.

DISPOSAL

A. COST ESTIMATES

400 OTHER EMPLOYEES X
2% X $15,000/YR = $120,000

B. ASSUMPTION

- 2% OF OTHER EMPLOYEE'S TIME IS ASSOCIATED WITH THE DISPOSAL FUNCTION

SUMMARY

David Butler
Chairman, Butler Cox & Partners Limited

COX: My colleague, David Butler, will give a brief summary of some of the points that have been covered during the day.

BUTLER: Thank you, George. What I would like to do is not in any sense to try to recapitulate or compress into ten minutes the messages which have been delivered to us by the various speakers today. I think the effort that they have put in to prepare their presentations and to supply a wealth of detail to support their different views is such that it would be quite impossible and inappropriate for me to attempt to do that.

Rather I want to try to spin out of today one or two lessons, to try to put the messages of the day into some kind of framework that we can think about and see whether that helps to relate what's been said to our own particular situation in each of the member companies. If I fail in my attempt, it won't be the end of the world because I promise to take only ten minutes of your time.

Sources of Guidance

- A. 4th. Century BC
  The Peloponnesian War — Πρόφασις / Αἴτια
- B. The 1960's in DP
  - The respectable illusion
    'clerical labour savings' (Ha ha!)
  - The real-world conspiracy
    The salesman and the DP manager

There are certain sources of guidance, which I think are available to us in seeking to understand the message of today. The first, obviously, is the Peleponnesian war between Athens and Sparta. There may be one or two of you who haven't recently re-read Thucydides. The rest of you will remember that Thucycdides distinguished between two things, both of which in a sense could be said to be the cause of the war.

One was what he called the prophasis — the pretext, the argument which was put forward by the Spartans for going to war with the Athenians.

The other was what he called aitia, the root cause, the real unspoken reason that lay beneath the outbreak of the Peloponnesian war. I suppose if one looked at any great movement or event in history, it is always possible to distinguish between the rhetoric and the reality. That's what I'd like to try to do right now.

The second source of guidance, I think, is the 1960s in data processing. There was a respectable illusion, wasn't there, that in the 1960s the advent of data processing was going to be the source of clerical labour savings. Payroll systems, inventory control systems, production control systems and so forth implemented in batch mode on large, relatively expensive computers, were going to produce clerical savings. And I think, with hindsight, the comment that I have written next to that is not wholly inappropriate because, in general, those clerical savings did not materialise. This is the world of prophasis.

In the real world — the world of aitia — there was a conspiracy between the salesman and the data processing manager to bring into companies technology which they couldn't really justify on the basis of any reasonable forecast of saving. I think if one were to look at the history of those companies since then, one's verdict on that conspiracy would have to be, thank God that the majority of the salesmen and the managers concerned had the courage to lie to their bosses. To introduce technology which, if one had known what was to happen with inflation and labour rates in the succeeding fifteen or twenty years, turned out to be the smartest thing they could do.

But it really does bear thinking upon that in many cases there was an element of at least unconscious deception of senior management that paved the way for that highly desirable development.

The need for O.A. (Πρόφασις)

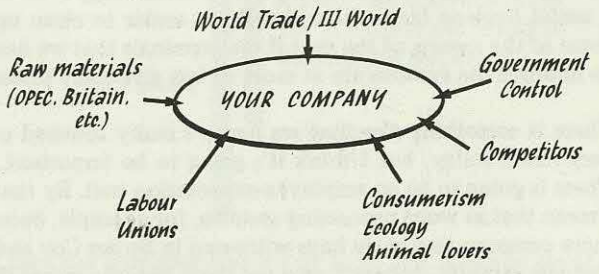

Now what do we see today? We have another respectable illusion, I suggest — another prophasis or pretext — which lies conveniently to hand. We can depict the world without too much straining of the imagination to top management in this way.

Here is our company at the centre. It's afflicted by all the things I've shown on the slide there. Shortages of raw materials, particularly energy materials, due to the aggressive unprincipled behaviour of certain countries.

There is the enormous pressure on our companies coming from the relative decline in world trade since 1973. There is the determination (which I suppose we must all applaud on humanitarian grounds but which nevertheless creates real problems for us) of the third world (not to beat about the bush) not to be screwed by the rest of us for the rest of their short, nasty and brutish lives.

There is government control. Do you know the most shocking thing that I've heard today? I always thought that we in Europe understood every way to interfere with a company's business, to stop it from making a profit, but when Randy said that here in the United States you are not even allowed to find out if people can read or write, I really did think that "the torch has passed to a new generation of Americans."

There are unions and there are all the forces of rampant consumerism, dog lovers and the rest of it. And then, if there are any of them left, there are your company's competitors. Now I think there is enough of a seed of truth in that picture to make it a rather useful tool in our conspiracy. I think what it shows is probably that there are no easy triumphs to be won over the next twenty years, and that the battle to defend margins which are being eroded in markets which are under attack is a real battle, and one that needs to be fought with every tool which is at our disposal.

The need for office automation (αἰτια)

- *Low cost a hygiene factor*
- *Employee expectation pull*
- *Technical zeal*
- *Sales pressure*
- *Qualitative change*
- *Picking winners - then exploting marginal cost*

In the 1980's the main virtue is agility

Let me argue that at a more banal level the need for automation in the office really comes down to these facts among others; that the low cost of this technology is indeed a useful hygiene factor. It will make it easier to clean up some of the messes of the past if the terminals that we need to improve the systems are at more or less give-away prices.

There is something else that we haven't really touched on very much today, but I think it's going to be important. There is going to be an employee-expectation pull. By that I mean that as word processing systems, for example, become more common and as we have witnessed in Butler Cox and Partners recently, the staff who use them become rather firm enthusiasts of word processing.

Imagine how difficult it would be to get girls accustomed to systems like that to go back to ordinary typewriters. Almost as difficult, I think, to get them to go back from electric typewriters to manual. You will all know that in the middle of cities at least it is virtually impossible to recruit girls nowadays to work manual typewriters.

So I do think, contrary to our expectation, that at least part of what we are talking about will be people-led, rather than technology-led or price-led. It will become the norm to have many of these systems and it will be extremely difficult to recruit people without them.

The technical zeal of the enthusiast and the sales pressure from the vendor are going to be, once again, as they were in the case of data processing in the 1960s, powerful weapons. All we must hope is that the deceptions which are practised this time will turn out to be as creative as the ones which were practised last time.

I believe that we are going to be looking not only for quantitative change, not only for ways of doing things cheaper, faster or more often, but also for qualitative change. One of the factors which is often left out of account when we discuss the microprocessor and employment is the vast hinterland of untapped potential that lies behind almost every job, just for doing it better, just for improved quality in jobs. If people had more time to think about what they are doing I believe that will be an attractive bonus of these systems.

From a practical point of view, every presentation that we have had today really spelt out one message. The approach which is likely to succeed is really sheer opportunism — looking for a winning application, an application which is so promising that even if you make rather a technical mess-up, it's almost certain to be well received by the users. When you have done that, you can build on the fact that you have such a system to further exploit at marginal cost.

If I may quote my colleague, George Cox, he told me once an excellent piece of advice that his father gave him when he was a child — never get into a fight that you are not certain you can win. I think the same tactics will pay off here. If we don't get into fights that we're not pretty certain we are going to win, if we look out for applications which really are going to pay off (and most of those seem to be arising now in the area of word processing) then we can also exploit the systems which we shall have at marginal cost.

It's not perhaps a grand design. It's not perhaps a terribly romantic or flashy view of how the future might arrive. But if I could go back to my opening point, what you have to remember is that the Athenians, for all their cleverness in philosophy and art and music and all the rest of it, lost the war. It was the Spartans, with their quiet, steady persistence over a long period of time, who won it.

I hope that some of the lessons that have come out of this day will help you, ladies and gentlemen, to win your war with steady, quiet persistence.

That is the end of our proceedings today. Tomorrow we turn our attention to the subject of distributed processing — another great theme for the future.

Let me once again, on your bahalf, thank all the speakers today and I look forward to seeing you again first thing in the morning.

A PLAIN MAN'S GUIDE TO DISTRIBUTED PROCESSING

Hal Becker
Advanced Computer Techniques Corporation

Hal Becker is Manager — Communications and Distributed Systems for the Technical Analysis Group (TAG) of Advanced Computer Techniques (ACT), located in their Phoenix, Arizona office. He joined ACT in 1979, following nineteen years with General Electric and Honeywell.

In 1963 he transferred to General Electric's Computer Department in Phoenix and joined the team installing the first data communications processor shipped.

His activities in Phoenix have been devoted exclusively to the development of Data Communications and Information Network Technologies.

Mr Becker has lectured extensively in major United States, European, United Kingdom and Australian cities on the Functional Approach to Data Communications and Information Network philosophies. He has also written a book, "Functional Analysis of Information Networks" and is currently writing a second one, "The Distributed Environment".

ZEDLITZ: My name is Chris Zedlitz. May I introduce you to today's theme which is distributed processing. Hal Becker will be the first speaker today. Hal is an expert in data communications and information network technologies. He will structure the problem of the distributed environment for us.

I asked him what "A Plain Man's Guide to Distributed Processing" meant. He told me that it should really be called "A guide to distributed processing for someone who lives in the plains of the States."

BECKER: Good morning. It is a pleasure to be with you this morning. We made the arrangements for this last year, and I was at that time employed by Honeywell Information Systems in Phoenix, Arizona. I left Honeywell in January and joined Advanced Computer Techniques, which is a consulting firm and has been in the industry for 17 years. Charles Lecht is the founder. Some of you may have heard of him, listened to some of his talks, or read some of his material.

I have been in the computer industry for 20 years. I started in 1959 with General Electric in Detroit, Michigan; moved to Phoenix in 1963 with General Electric and joined the computer department at that time. I joined the group in Phoenix that was installing the first data communications processor that they installed. It was a Datanet 30. I have been in Phoenix since 1963, and have spent all my time in the area of data communications and network technologies.

I am considered an expert in the field. The definition of the word expert that I like to apply there is: an expert is a person who has made every conceivable mistake in an extremely narrow field. I think that, with that definition, I qualify as an expert.

Our topic here for the next hour is the distributed environment. It is a very popular topic in the industry. There are a lot of people writing about it, a lot of people using the terminology, a lot of people selling and installing equipment that is called distributed processing equipment.

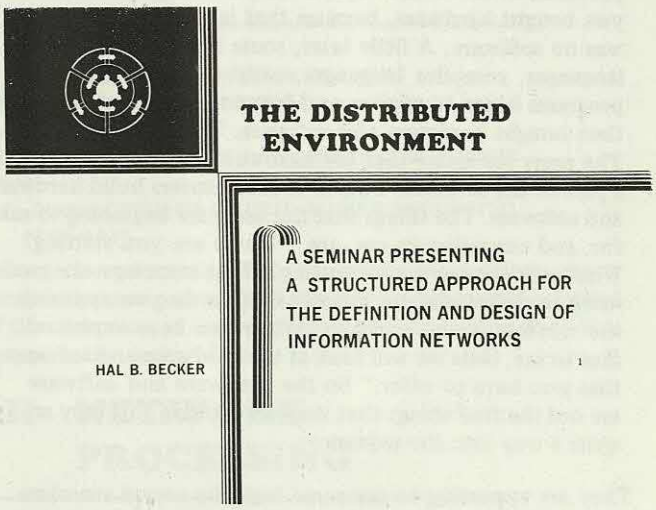

I used this title slide and, a year or so ago, a person about ten rows down into the audience looked up and said, "Oh yes, the disturbed environment." It occurred to me that if you rearrange the letters just a little bit, you get the word "disturbed". That is part of the problem with this topic; it is disturbing because everyone defines the term "distributed processing" a little differently. It does not mean the same thing to everybody who uses the term.

So what I will do in the next hour is to present a structure that allows us to define the distributed environment in terms of its basic functions or basic building blocks. I believe that the advantage of a structured approach to any problem is that it allows us to take the problem itself, and break

it down into a series of smaller, more manageable pieces. That is what we will be doing here. We will provide you with a structure that will allow you to take the distributed environment, break it down into its smaller pieces, each of which can be managed a little more effectively than they can if you try to work with them as a whole.

We will do this in four sections. The first section is a brief introduction: some of the common user complaints, criticisms, problem areas of concern, the questions that they are asking. In the second section we will briefly define the environment itself. We will identify the basic building blocks that go into a network of any kind. We will see throughout the session here that there is a common approach. The building blocks that we use for a distributed network are no different than the building blocks that we use for the classical or traditional centralised networks.

In the third section we will look at the user community and explore briefly the requirements statement which is the definition of the problem that the user expects this network to solve. We will provide here a look at the elements of information that are necessary for those who are going to design and install the network. Finally, in the fourth section we will look very briefly at the analysis and design sequence which is the way in which we design the network, given the requirements statement from the users, and we will design the network with some combination of the building blocks that we have discussed in section 2. So it will all come together in section 4.

The user community recognises now that the hardware and software that they purchase is a part of the solution to their problem. In the earlier days, when you bought a computer you bought hardware, because that is all there was; there was no software. A little later, some primitive assembly languages, compiler languages, early canned application programs began to appear; and for some period of time you then bought hardware and software. Today it is different. The users recognise that the hardware and software is just a part of the solution. Lots of manufacturers build hardware and software. The things that the users are beginning to ask for, and expecting to see, are, "Where are you starting? What architecture are you using? What structure are you using to define the environment? How do you approach the solution of my problem? When you have explained that to me, then we will look at the hardware and software that you have to offer." So the hardware and software are not the first things that they are looking for, they are quite a way into the sequence.

They are expecting to see some logically sound structure for defining the problem, defining the environment, and the approach that will be used in solving the problem. This generally takes the form of looking at the functions that are required. I will use this term "function" several times and we will explore these functions a little later. There are three basic functions that are used in constructing a network of any size.

The users recognise that the design of a network is a very complex thing and that they usually, as users, do not have the kind of talent on their staffs necessary to design and install a network. They are also recognising that the design is very quickly leaving the realm where it can be done successfully using strictly manually oriented design practices.

THE DISTRIBUTION OF THE THREE BASIC FUNCTION SETS:

1) INFORMATION PROCESSING

2) NETWORK PROCESSING

2) DATA BASE PROCESSING

- CONSIDERATIONS/OBJECTIVES
- ADVANTAGES
- DISADVANTAGES

We recognise that we must learn to apply the computer itself to the design problem of information networks. If we are looking at a network that contains several thousands of terminals — and there are quite a number of those in place and growing already — the number of different design options that you have is staggering; and to explore them fully, using manual techniques, is frequently impossible. So we as an industry are learning to apply the computer itself to the design of computer networks.

Another question that the user community has is that of predicting the performance of the installation. Back in the days when you were buying a batch processing installation you could expect the vendors to benchmark the configuration; put a copy of it together, run it, and show you how it performed. Benchmarking a network is impractical because of the time, the distances, the resources required and the complexities involved. So the users are left with the question still in their minds: how do I know that the configuration that you are proposing will work?

The industry is beginning to provide some higher level modelling and design capabilities that will allow them to explore the behaviour of a proposed network and come up with some prediction of its performance. The user community is beginning to ask for this in lieu of the ability to benchmark the network and see how it works in that sense.

A sixth point is that the user community is asking for assurance that the installation that they buy and install will be capable of adapting to the changing technology. Hardware and software technology continues to change. The user community is very much aware of this and is asking, "Put an installation in that will grow and adapt to the changing environment. Don't sell me an installation that I can put in and use for some two or three years or so, and then I find I'm faced with another massive conversion because it can grow no more. Put something in that's capable of growth in a series of logical, planned steps for quite a way out into the future."

Finally, the users are beginning to ask for higher and higher levels of security in these installations. Computer fraud is big business already. The FBI in this country has recently released information statistics where they find that the average computer fraud nets the perpetrator in excess of $3½ million. The average armed bank robber, on the other

hand, gets $2,300. The armed bank robber goes to jail for 15 or 20 years if apprehended; the computer fraud, if he is apprehended, goes to jail for 18 to 20 months. They believe further that the amount of computer fraud that is detected represents just 1% of that which exists. In the earlier days the computer fraud took the form of manipulation of existing programs for some financial gain. They are finding now that whole series of programs are being written with fraud in mind.

Computer fraud is big business. It is becoming a bigger business and, as corporate auditing staff become aware of the extent of computer fraud in the industry, some of them are getting absolutely paranoid when they see what kinds of things their people are putting on the machines. Auditing in the computer environment is becoming a whole new speciality industry of its own in this country, because they are recognising that traditional corporate auditing techniques and procedures are frequently incapable of detecting computer fraud. So that is an area that is getting a considerable amount of attention.

Now any structure that we use for approaching the network environment must do a number of things for us. First, it must recognise that we will almost always have existing installed equipment. It is a very rare network that is put in from the ground up where nothing existed before. Once in my career did a prospect come in and say, "We don't have any computers installed at all, but we're going to put 25 great big ones in. What do you have?" That happened once. Most of the time they come in and say, "We have 25 computers," or 60, or 80, or whatever the number is — "we'd like to connect them together now in some kind of a network." So our structure must recognise that we will almost always have existing installed equipment. It should provide for the integration of that equipment with the new equipment that will make the network a reality; and it must be capable of assimilating new hardware and software as it becomes available. Again, the ability to adapt to change. Heraclitus said it 2500 years ago: the only thing that is constant is change itself. That could not be truer in the computer industry today.

Let us look at the environment itself. Here we are going to define the basic building blocks that are used to configure any kind of an information network, whether it is the classic centralised variety or one of the newer distributed types.

COMPUTER USER REQUIREMENTS CAN BE SATISFIED WITH VARIOUS COMBINATIONS OF THREE BASIC CAPABILITIES:

1. INFORMATION PROCESSING
   THE MANIPULATION OF INFORMATION TO PRODUCE THE DESIRED RESULTS

2. NETWORK PROCESSING
   THE MOVEMENT OF INFORMATION BETWEEN THE VARIOUS SITES

3. DATA BASE PROCESSING
   THE STORAGE OF INFORMATION IN FORMS APPROPRIATE TO THE USER.

Any computer installation consists of some combination of just three basic building blocks. Look at your own as an example. There are three things that we are doing:

- One, we are processing information; information processing is the manipulation of information to produce the desired results.

- Secondly, we are moving information — the network processing function — between the various locations of the network. That is network processing, the term that I will use, sometimes called data communications or just communications.

- Thirdly, database processing. The database function is the ability to store potentially large amounts of traffic or large amounts of information in a form appropriate for use by the network and its users.

These are the three functions that we are concerned with. The network processing or data communications function has become a utility to the other two; and as such the amount of application dependent logic that exists in the network processing function is becoming almost non-existent, and that is the way it should be if it is to be a utility to the other two. If it is to allow the user community to access the information processing resources and access the database resources, no matter where they are, it must be a very efficient, adaptable capability, which it is becoming.

-I- INFORMATION PROCESSING

- COMPILATION, ASSEMBLY OF USER APPLICATION PROGRAMS
- EXECUTION OF USER APPLICATION PROGRAMS
- PRODUCTION OF OUTPUT IN USER REQUESTED FORMATS

Information processing includes, among other things, the compilation and assembly of the user's programs; the control of the execution of the programs; and the production of the output in the desired form.

-II- NETWORK PROCESSING

- CONTROL OF TERMINAL DEVICE INTERFACE
- CONTROL OF HOST INFORMATION PROCESSOR INTERFACE
- CONTROL OF INFORMATION FLOW BETWEEN TERMINALS AND/OR HOST INFORMATION PROCESSORS

Network processing is the control of the interface between terminal devices and the network; on the one side the control of the interface between the host processors and the network. On the other side, and perhaps most visibly, control of the movement of information between the various locations; terminals to the host computers, the host

computers back to terminals, one terminal to another terminal, one computer to another computer; these are all the forms of traffic that can occur in a network.

-III- DATA BASE PROCESSING

- GENERATION OF DATABASE(S) IN APPROPRIATE FORM
- PROVISION FOR ACCESS TO DATABASE(S)
- MAINTENANCE OF DATABASE(S)

The database function is responsible for generating the database in the appropriate form, using an architecture that suits the installation. Secondly, it must provide for an efficient level of access to the database if it is to be useful to the network and its users; and thirdly, it must maintain the database from an accuracy, integrity, security point of view. These are the three functions that any computer installation today will consist of; some combination of these three.

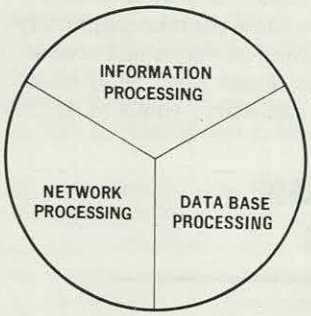

THE FUNCTION SETS

We can define the three functions using this kind of diagram. This does not mean to imply that we always have equal amounts of the three functions; that just is not true. I speak at the National Conferences every year, and someone will make a statement like, "The average computer user will spend 46% of their money on the information processing function." Everybody writes down "46%" and then rushes home to see "How much are we spending on information processing function? Are we spending too much money? Not enough money? Or just right?" I do not believe in that. I do not think that the average network exists. Yet you see people trying to design to this average all the time. We may see one network where the information processing function represents 60% of the expenditure, because that is the kind of installation it is. In another installation the database function may be the biggest one, it may be getting 60% of the activity, because that is the kind of installation it is. It is a database centred network, whereas the other one is an information processing centred network. Or maybe we have some other mixture. There are no useful rules of thumb that I am aware of that will tell us how much we should be spending in each of these three areas, so I quit looking for them a long time ago.

It is the interaction between these basic building blocks that allows the network as such to function. As such, these interactions or intersections must be an integral part of the approach to the network that you use. If it is not, recognise that you are assuming a significant risk. Frequently, these intersections are left until last in the design of a network, only to discover that because they have been left until last they become the biggest bottlenecks of the installation.

THE THREE NETWORK FUNCTIONS

Prior Technology

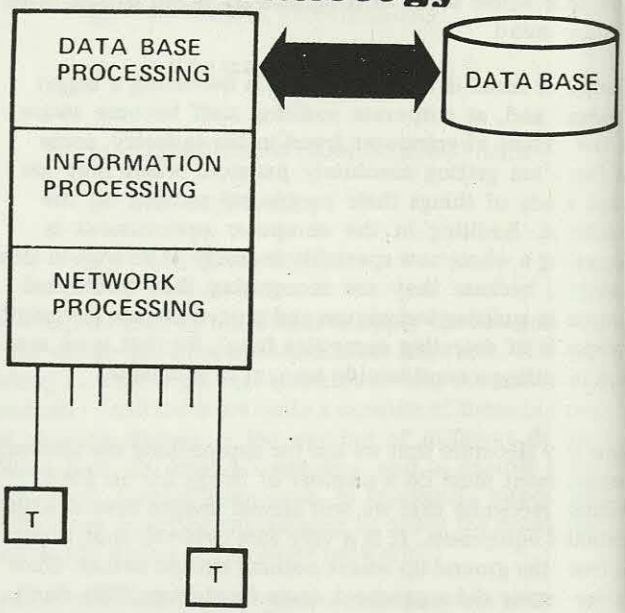

Let us look at these three functions a little differently. This is a somewhat historical approach here. Back in the early 1960s when data communications first began, we accomplished it by adding to our existing host processor some primitive line interfaces that allowed it to talk to distant terminals. So at this point we then had the one computer, the host processor, responsible now for all three functions. We did the information processing in the machine; we did the database processing; we did the data communications. What very quickly happened was that we ran out of capacity. Sometimes we found ourselves spending so much of our time and energies doing the communications functions that we did not have enough time in the processor left to manage the information itself once we got it in. So this conflict was resolved in the industry with a rather pioneering step.

THE THREE NETWORK FUNCTIONS

Current Technology

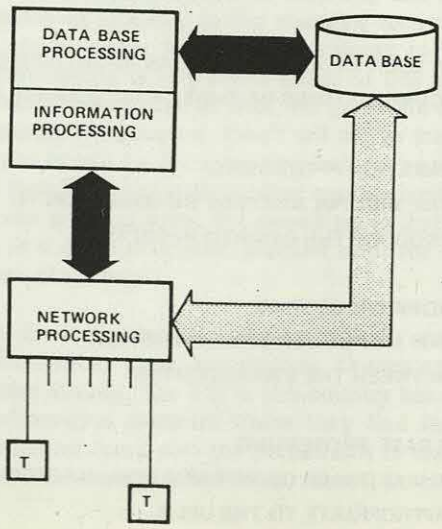

This happened about 1964. General Electric, and Honeywell, and others made what was called a 'separation of function' where the communications activity was pulled out of the host processor and put into a second machine that was very quickly labelled the front end. So the front end network processor then took over the entire communication load, leaving the host processor up there free to do what it was designed for — that is, processing of information.

So at this point we had the communication function executing in the front end; we had the database and information processing remaining back in the host. That is how it existed up until quite recently. This was the state of the art up until a couple of years ago.

What has happened is that in many installations the fact that the user is moving more and more of the applications from the batch world into the on-line world means that a database gets dragged on-line with it. So in many installations that we look at we see the database function growing more rapidly than the other two.

What has happened is that we now have another conflict. The host processor is handling the database function and the information processing function, and it frequently cannot do an acceptable job of both. So how will this be resolved?

THE THREE NETWORK FUNCTIONS

Coming Technology

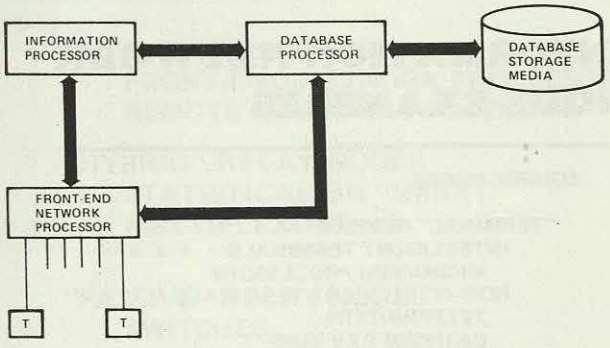

We have seen it already. There is occurring now a second separation of function, wherein the database function will be split off from the host processor and put into a separate machine of its own that is called a database processor. That database processor has already been labelled the back end.

There are installations in place today exploring this approach. You cannot buy these as standard, off-the-shelf products yet, but there are a number of them in place that are being explored to see just what additional efficiencies and throughputs we can achieve by separating these functions. These are the three basic functions that any kind of a network today will consist of; whether it is the traditional, centralised network or one of the newer, distributed configurations, we are still talking about combinations of these same three basic building blocks.

The small arrows between the three sectors, between the three functions, and those arrows represent the need for efficient levels of intersection or interaction between these functions themselves. Again, these are too frequently left until last with the result that the intersections are very

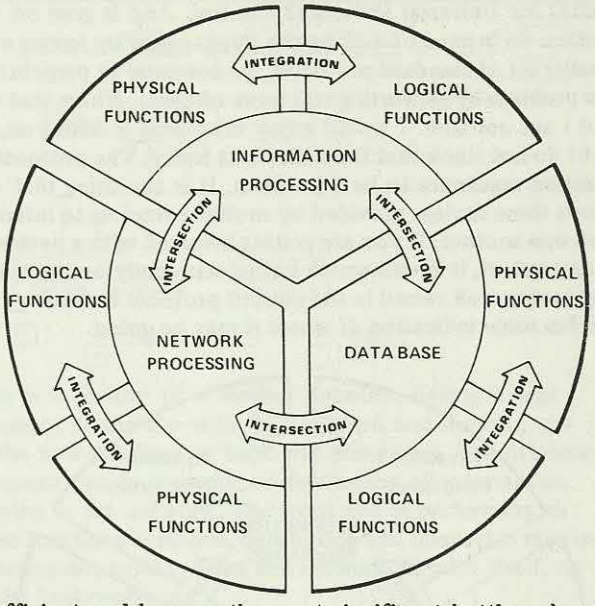

inefficient and become the most significant bottlenecks of the network.

THREE PROTOCOL LEVELS

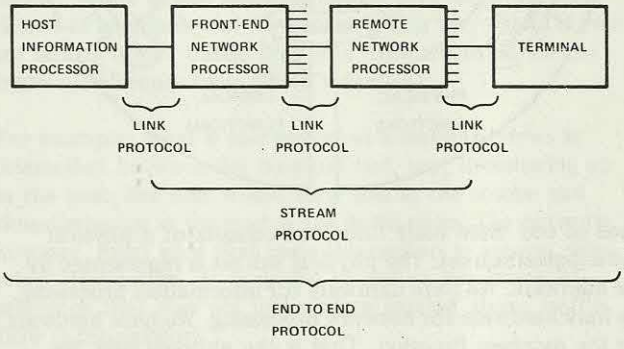

The protocol question continues to be a problem. A protocol in a computer network is a language that allows dissimilar devices to communicate with one another. It is just like walking around in various countries in Europe; I cannot talk to the people or communicate with somebody until we find a language that we both agree to use and can use. It is no different in a computer network. The protocol is the language that these devices use to intersect with one another. That terminal device on the far right cannot communicate with the network processor over there until they have a mutually agreed upon language which is called a protocol. That remote network processor cannot communicate with its front end until they have a mutually ageed upon language, which may be a different link protocol. Similarly, the front end cannot communicate with the host until they have a language that they both agree to use.

Now the problem is that every time we turn around there is another "standard" link protocol appearing. Everybody thought that BSC would be it for a long time. Well, the next thing we know there is SDLC (Synchronous Data Link Control), IBM's linked protocol. The International Standards Organisation, Honeywell and others, said, "No, we don't like SDLC, we're going to go with a thing called HDLC (High Level Data Link Control) which is different". This went on for a couple of years. Univac looked around and said, "Well, we don't like any of those that we've seen

so far. We're going to design our own." So Univac announced UDLC for Universal Data Link Control. And it goes on and on. So instead of making the things easier by having a smaller set of standard protocols, we continue to perpetuate the problem by generating still more of them. Where that will end I am not sure. I would enjoy discussing it with you, but I do not think that time allows us today. The protocol question continues to be a problem. It is the thing that allows these devices provided by multiple vendors to interact with one another. If you are getting involved with a network to any extent, it is well worth having somebody on the team who is very well versed in the current protocol technology and has some indication of where it may be going.

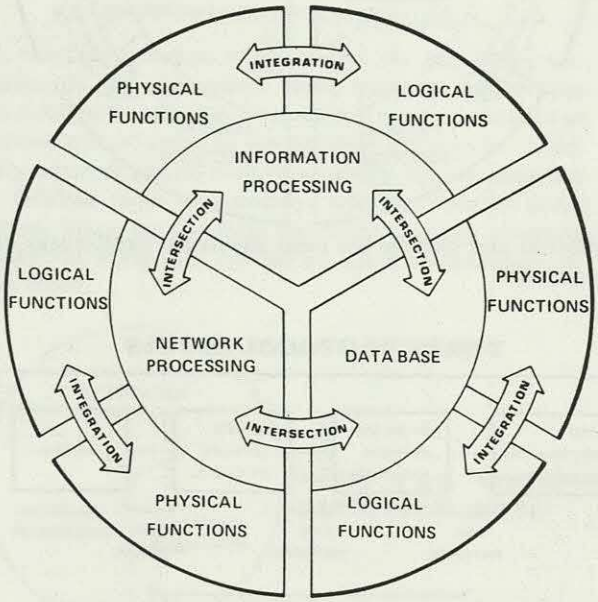

Each of our three basic functions consists of a physical and a logical sub set. The physical sub set is represented by the hardware. We have hardware for information processing. We have hardware for network processing. We have hardware for the database function. That is the physical side. We are coming down the next level in our structure here. These are the devices and facilities that process information, move information and store information for us, again corresponding to our three basic functions.

The logical functions are represented by software that executes in the various processors of the configuration. We have software executing in the host processor; software executing in the communication processors; and software executing in the database processors. This software always does two things for us. First, it controls the physical devices themselves; and, secondly, it controls the flow of information within, through and between the physical devices themselves.

We can add another piece to our picture. This time we have added the outer ring, showing each of the three functions having its physical and logical sub set; the physical becomes the hardware for the function, and the logical becomes the software for that function.

We can add one more piece to it. Here we have added the arrows that are labelled 'integration'. The design of a network involves the selection and integration of appropriate sub sets of the hardware and of the software. We will explore a little more later on this selection and integration of an appropriate set of hardware and software. Notice that I have said nothing yet about a centralised network or a distributed network. What I have said so far applies equally well to all of them. We do not have to treat the new, distributed configurations with a completely different set of rules and guidelines than we used in the classical, centralised approach. They are not that different.

DEFINITION OF INFORMATION NETWORK "NODES"

- A "NODE" WITHIN A NETWORK IS:
  1) A POINT AT WHICH INFORMATION MAY ENTER THE NETWORK (SOURCE NODE)
  2) A POINT AT WHICH INFORMATION MAY LEAVE THE NETWORK (DESTINATION NODE)
  3) A POINT THROUGH WHICH INFORMATION MAY PASS IN TRANSIT BETWEEN THE SOURCE AND DESTINATION NODES (RELAY NODES)
- ANY INFORMATION NETWORK CAN BE DEFINED IN TERMS OF SOURCE, DESTINATION AND RELAY NODES

Now we will define another term — 'node'. A node in a network is a point at which information may enter the network.

INFORMATION NETWORK NODE EXAMPLES

- SOURCE NODES
  - "TERMINAL" DEVICES
    - INTELLIGENT TERMINALS
      - MICRO/MINI PROCESSORS
    - NON-INTELLIGENT TERMINALS
      - TELEPRINTERS
      - CATHODE RAY TUBE
      - FACTORY FEEDBACK
      - INDUSTRY/BUSINESS SPECIFIC
      - * * *
  - INFORMATION PROCESSORS
    - "HOST" INFORMATION PROCESSOR
    - "SATELLITE" INFORMATION PROCESSOR
  - NETWORK PROCESSORS
    - "FRONT-END" NETWORK PROCESSOR
    - "REMOTE" NETWORK PROCESSOR
  - DATA BASE PROCESSORS

That is a source node — a source of information to the network. Source nodes can be terminal devices; they can be information processors; they can be database processors; they can be intelligent terminals; satellite processors. All of those things can be sources of information to the network.

At the other end is the destination node, the point to which we deliver information; and it may or may not leave the network. Anything that can be a source can usually be a destination also.

NODE EXAMPLES - CONTINUED

- DESTINATION NODES

  "TERMINAL" DEVICES
  - INTELLIGENT TERMINALS
  - * * *
  - NON-INTELLIGENT TERMINALS
  - * * *

  INFORMATION PROCESSORS
  * * *

  NETWORK PROCESSORS
  * * *

  DATA BASE PROCESSORS
  * * *

A third type of node is necessary in designing a network today, and that is a relay node. It is a point through which

NODE EXAMPLES - CONTINUED

- RELAY NODES

  "ACTIVE" RELAY NODES
  - FRONT-END NETWORK PROCESSOR
  - REMOTE NETWORK PROCESSOR

  "HYBRID" RELAY NODES
  - STATISTICAL OR "SMART" MULTIPLEXORS
  - * * *

  "PASSIVE" RELAY NODES
  - SWITCHES
    - RECONFIGURATION
    - LINE/TRUNK
    - * * *
  - MULTIPLEXORS
    - TIME DIVISION
    - FREQUENCY DIVISION
    - * * *

information travels on its way from the source to the destination. A smaller network may traverse just a single link in getting from the source to the destination. A larger network with thousands of terminals in it may have to pass through many links before it gets to the destination. Every time it goes from one link to another it passes through a relay node. The relay node technology is fairly complex all by itself, but it allows us the degrees of freedom that are necessary in designing an efficient, workable installation. Any network that we want to view, small through very large, centralised as well as distributed, will be some combination of these three basic types of nodes — source, destination and relay. So again there are not that many functions or building blocks or definitions that we have to deal with.

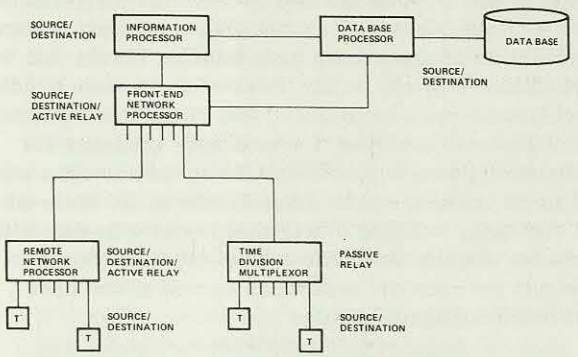

Here is a picture of a modest network, having a host processor at the top with its front end and showing one of the new database or back end processors. The database processor can be a source or destination of information flowing in the network. The front end is performing all three functions — source, destination and relay. It is relaying information, coming from the network through itself, up to the host processor.

Down here we have a couple of relay nodes; on the left a remote network processor which is handling the traffic flow from this cluster of terminals here into the network. There we have a simpler relay node, a time division multiplexor, which is also handling traffic flow between that cluster of terminals and the network. It is very useful if you are talking about traffic flow in a network to define the frame of reference from which you speak.

For example, what a user seated at a terminal views as interaction between the terminal and, say, timesharing up in the host, the user would view this as the source and timesharing up in the host as the destination. The network, however, must view that same interaction as source, relay, relay, destination. It is useful to define the frame of reference that you are using when talking about traffic flow in the network.

I use an analogy in a longer version of this talk that I give. I have a couple of slides to cement this idea. The first one has a picture on the left of a big, fluffy chicken; in the centre is an egg; and on the right side is another chicken. It says across the bottom, "An egg is a chicken's way of making another chicken."

The second slide has a big egg on the left, a single chicken in the centre, and another big egg over on the right. It says at the bottom, "A chicken is an egg's way of making another egg." It is just as logical. The same event but viewed from two quite different frames of reference comes up with a little different interpretation.

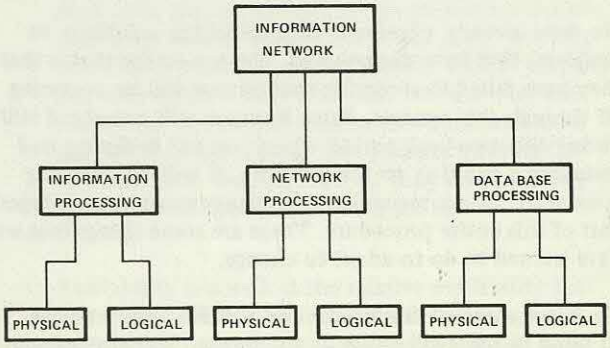

We can summarise the structure for defining these building blocks in this way. At the top we have the information network itself, which will consist at the next level of some combination of these three basic building blocks that we have discussed briefly so far. Each of these three building blocks, coming one more level down, consists of a physical and logical sub set. Here I would have hardware for information processing, software for information processing, and so on across the three basic functions. So there are not that many building blocks that we have to deal with. Again we can use these three basic building blocks to construct the classical, centralised as well as the newer, distributed configurations.

Now let us look at the requirement statement. The requirement statement frequently takes the form of a document called the request for proposal, request for information, request for quotation, request for tender — all of those terms that are used. This is an area that has some very significant pitfalls in it that we are learning to recognise.

The requirement statement should be a comprehensive definition of the user's problem. Too often, the requirement statement is a very careful explanation of the currently installed and running applications. That is fine, we need that information; but the risk is that you will get somebody who will design a network just to handle those existing installed applications and the coming future requirements are not that well handled.

It should be clear, without ambiguity. When I was on a design team we would occasionally have to send everybody home with a copy of a certain section of the document, saying, "Read this, and when we get in tomorrow morning we're going to vote and see what we think it means." It is not that the user did not do a good job of writing it and defining it, it is just that we could not understand it.

Do not consider a requirement statement or a request for tender as a static definition of a static problem. Again, one of the big problems that we have come to understand is that change is occurring constantly throughout this cycle. There is a thing that occurs that I call the 'vanishing problem syndrome'. It works something like this. The user writes a document called request for tender. It takes six months to do it. He issues it to all the interested vendors and they take six months to respond. Then you get all of the responses back, 15 or 20 of them, and you spend six more months evaluating them to decide which one you want to select and install. Then you enter a final six month period of final negotiation with the selected vendor, getting the contract worked out and the legal requirements and resources and so on. Finally, two years after you sat down and defined the problem, you tell the vendor to go ahead and install it.

We have already observed users installing solutions to problems that have disappeared. The reason for that is that they have failed to recognise that change will be occurring all through this process. Your business will not stand still during this two-year period when you are designing and installing a solution to the problem, it will be changing constantly. So we recognise that change must be an integral part of this entire procedure. There are some things that we have learned to do to adapt to change.

We know some of the hooks and handles that we have to build in that will allow us to change. The requirement statement is then the definition of the problem that is to be solved. It should also define the anticipated usage of this installation. Different users will develop quite different usage patterns and profiles on a given network. If this usage is not anticipated and defined fairly clearly, you force the designers to guess or interpret it as best they can.

The most significant point of all is that the document should define how this network's performance is going to be measured. One user may say to me, "Security is the most important thing. That's what I'm looking for in this installation. You provide me with the level of security I want and I'm happy." Another user says, "No, to me response time is the key. We have all those terminals out there on line, with people sitting there in front of screens, and telephone headsets on, answering telephone enquiries. Response time is our thing." Somebody else says, "No, to us it's availability. The system has to be up and running 24 hours a day, seven days a week," and so on and so on.

Every user will have a different interpretation of what is performance. They will have a different perception of the network. If that is not explained and defined how you are going to measure this performance in the network, you again force the vendors to interpret as best they can.

The user community is represented by two broad classes. We have the human users, which is where we tend to concentrate all our attentions. But I believe that we are ignoring a class of users that are just as important if the network is to operate successfully — and that is the non-human users.

NON-HUMAN USER CATEGORIES

- INFORMATION PROCESSORS
  - HOST
  - SATELLITE
- NETWORK PROCESSORS
  - FRONT-END
  - REMOTE
- DATA BASE PROCESSORS
- INTELLIGENT TERMINALS

What do we mean by that? The non-human users include the host processors, the satellite processors, the front end and remote network processors, the database processors, as well as the intelligent terminals. These are users of the network and its resources and, as such, they have a set of requirements that must also be met if the network is to be successful. We assume that somehow these non-human users' needs will be met. That is not always the case, as we are learning. So recognise this as a class of user as well.

The human users — I have selected half a dozen or so here — range from high level corporate management. They are users of the network. They are using the resources of the network, getting useful work out of it. They will, over time, develop a fairly pointed idea about what is acceptable performance.

HUMAN "USER" CATEGORIES

LEVEL	TYPE OF USAGE
• CORPORATE/ORGANIZATIONAL MANAGEMENT	EVALUATION/UTILIZATION OF NETWORK OUTPUT
• DATA PROCESSING MANAGEMENT	AVAILABILITY, RELIABILITY, COSTS . . .
• SYSTEM OPERATORS/SUPERVISORS	CONTROL/SUPERVISION OF OPERATING CONFIGURATION
• SYSTEM DESIGNERS	INFORMATION PROCESSING, NETWORK PROCESSING, DATABASE PROCESSING RESOURCE UTILIZATION
• SYSTEM PROGRAMMERS	PROGRAM DEVELOPMENT
• END USER – FULL TIME	TERMINAL IS INTEGRAL PART OF JOB ACTIVITY
• END USER – PART TIME	OCCASIONAL TERMINAL USAGE IN SUPPORT OF JOB ACTIVITY

Data processing management, using the same installation, will have a different set of ideas about performance. They are concerned with availability, reliability, schedules, costs, response time and so on. The system operators and supervisors are users also, and they will have a still different idea about performance.

So will the systems designers, people designing new applications, growing and expanding the system. They will have a different idea about performance from their point of view. Programmers are also users, doing assemblies, compilations, debugging runs, parallel runs, cut-overs and so on. They will develop their own idea about performance, as will the end users down here — the people who are not computer experts but are nevertheless involved with the system in that they may be using a terminal device, either on a part time or full time basis.

The point is that all of these people and more — if we wanted to put other levels in here we could — will have their own idea about what is acceptable and what is non-acceptable performance. What we are asking the designers to do is to come up with one set of hardware and one set of software that keeps all of these people happy all of the time — a seemingly impossible task.

It gets back to the idea that the better we define the problem of these users, the more able the designers are to come up with an installation that meets their needs and requirements.

- THE USER REQUIREMENT STATEMENT IS CLASSIFIED INTO SEVEN BASIC AREAS:

1. TOPOLOGY
2. VOLUME
3. INFORMATION PROCESSING
4. DATA BASE PROCESSING
5. RESPONSE
6. AVAILABILITY
7. SECURITY

The requirement statement covers seven basic areas. There is not time to go into them in any depth, but I will say just a few words about each of them.

— Topology is a definition of the boundaries within which the network will be contained. We will see one network that exists in a single building. There are some. Very large corporate research and develop development centres will have a complete network contained in one building; a fairly sizeable computer with quite a number of terminal devices connected to it. So topology is the boundaries within which we are contained. Other networks will be covering a state, a group of states, a country, a continent, and on through those that are truly global in nature.

— Volume is a look at the relative amounts of traffic that will be moving through the network. We need to know a little bit more than the number of characters a day you are going to ship. I have seen requests for proposals where they will have one paragraph in there that says, "The network is expected to handle 5 billion characters a day, period." That is a fine number but it really does not tell a designer much about how to design to handle that 5 billion characters. Are we to assume that the traffic is distributed uniformly across the hours of the day? Probably not. That is a bad assumption to make. Are we to assume that the traffic is distributed evenly across the processors and terminal devices? Equally faulty. I suspect that the average hour in a network never happens, but you see designers designing to averages all the time. It is risky, and getting more so.

— Information processing is a look at the kinds of information processing that you are doing. Heavy scientific work; heavy financial, commercial, industrial kinds of applications. What kinds of information processing are you doing?

— Database the same way. Are you very heavily dependent on a large database that you have to access very efficiently and effectively? Or are you not so dependent on a database? That gives us a look at the database requirements of the user.

— The response time is a look at the relative levels of urgency that the users associate with the different kinds of traffic flowing in the network. This will range from applications where they have a three to four second response time as the expected response time, other applications will have a 60-second response time, still others will be 60 minutes — an hour or more, and on through those that might have 12 or 24 hour response times associated with the traffic. A very frequent mistake that is made by network designers is to say, "I'm going to take the peak hour volume and design the network to handle peak hour volume every hour of the day," thinking that that is probably a safe decision. Well, it may be, but then again it may not. Just because you can handle the peak hour volume in 60 minutes does not, by definition, mean that you are meeting the response time requirements of that volume. They are two somewhat independent problems.

— Availability is a look at the relative availability that you expect. Is this an installation that is up and running 24 hours a day, seven days a week? OK,

fine. Be advised that that level of availability will entail a higher level of duplicate and redundant hardware and software to make that availability happen. The higher the level of availability you request, the higher the level of duplicate hardware and software and the higher the cost. Nobody yet builds equipment that does not fail. A close look at the aviailability requirements allows the designers an opportunity to optimise the cost involved in providing that level of availability.

- The document should also define the security requirements. There are three basic levels of security that we can identify rather quickly that are useful here. The highest level of security is that typically associated with government and military installations. A very high level of security. The kind of thing where every piece of paper in the organisation has across the top of it big red letters that say, "Burn before reading". That level of security is typically not available with the standard, off-the-shelf hardware and software. You expect to pay a premium price for special hardware and special software to achieve that kind of security.
  The second level of security is that associated with banks and financial institutions. They have an above-average concern for security, but they are not quite as paranoid as the government and military people. With that level of security much of it is available with the off-the-shelf products, but they may have a small feature or two that they are willing to pay a modest premium to get.
  The third level of security is that associated with the standard, off-the-shelf hardware and software. That level of security is increasing. There are more and more devices available that will allow a relatively high level of security using standard, off-the-shelf products.

This is a very quick look at the requirement statement. These are the elements of the requirement statement that the network designers need the most. This is the information that they need to do an adequate job of designing this installation.

ANALYSIS/DESIGN SEQUENCE OBJECTIVES

- IDENTIFICATION OF WHICH OF THE THREE BASIC FUNCTION SETS (OR COMBINATIONS THEREOF)
  1) INFORMATION PROCESSING
  2) NETWORK PROCESSING
  3) DATA BASE PROCESSING

  WILL BE CONFIGURED AT EACH OF THE SELECTED SITES IN THE NETWORK = FUNCTIONAL DISTRIBUTION

Let us look at the analysis and design sequence. In here we are going to explore some various levels of distribution. We have not yet talked about centralised or distributed configurations. Let us look at them briefly.

The requirement statement that we just looked at is the definition of the problem that the network is to solve. As such, it becomes input to the design sequence. The goal of the design sequence is to select and integrate a set of hardware and software across each of the three functions that provides us with a workable, cost effective solution to the problem.

OBJECTIVES — CONTINUED

- IDENTIFICATION OF THE RELATIVE AMOUNT OF EACH OF THE THREE BASIC FUNCTION SETS THAT WILL BE CONFIGURED AT EACH SELECTED SITE IN THE NETWORK = FUNCTIONAL DENSITY
- IDENTIFICATION AND INTEGRATION OF APPROPRIATE SUBSETS OF THE PHYSICAL AND LOGICAL ELEMENTS OF THE THREE BASIC FUNCTION SETS

The first objective of the design sequence is: given these three functions, how many locations should each of them be configured at? That is, how many locations should have information processing functions executing? How many locations should have communications? How many locations should have database? This objective, boiled down to two words or less, becomes one of functional distribution. How should I distribute the three functions?

OBJECTIVES – CONTINUED

- IDENTIFICATION OF THE RELATIVE LEVELS OF INTERSECTION (INTERACTION) BETWEEN THE BASIC FUNCTIONS AS CONFIGURED AND AN INDICATION OF "STANDARD" PROTOCOLS THAT MAY BE DESIRABLE AND/OR NECESSARY
- DERIVATION OF A PERFORMANCE PREDICTION FOR THE PROPOSED CONFIGURATION

A second objective is: given this distribution of the functions, how much of each function? Each database site: How much database? Great big database? Smaller partitional database? Each information processing site: Big host installations? Small or satellite installations? And so on. This objective in two words or less becomes functional density. How much of each of the functions? Then and only then am I ready to explore the selection of hardware and software. Now I can ask the question: what is an appropriate set of hardware and software that will provide me with this functional density in this functional distribution?

So you see that the selection of hardware and software is not the first thing we do. It is at least the third or fourth thing that we do. Many people designing networks, their first impulse is to get to the chalk board and start drawing boxes and connecting them with lines. That is not the place to start. We are learning that. I learned that the hard way also.

We will also, in the design sequence, identify the levels of interactions between these functions which are necessary. Also we will be deriving a prediction of the performance of the configuration. This must be in many cases a part of your response to the prospect as a vendor.

"Here's how we modelled it, simulated it. Here's how we interpret the results. Here's how we think the installation will perform." So this performance prediction becomes another one of the things that comes out of the analysis and design sequence.

THE "DISTRIBUTED" ENVIRONMENT AND ITS RELATION TO THE THREE BASIC FUNCTION SETS:

1) INFORMATION PROCESSING
2) NETWORK PROCESSING
3) DATA BASE PROCESSING:

Let us look at the three functions and discuss some levels of distribution here. There is no single answer to the user community's problems. Just because everybody is talking about distributed processing does not mean that the vendor community is going to abandon the traditional, centralised approach. Quite the reverse. There is a significant class of users who, for years to come, will be satisfied with the classical, centralised approach. Other users are forced into exploring the distributed approach for a variety of reasons, a few of which we will explore here to the extent that time allows.

TWO EXTREMES EXIST RELATIVE TO THE DISTRIBUTION OF THE THREE BASIC FUNCTION SETS

1. ALL THREE FUNCTIONS – INFORMATION PROCESSING, NETWORK PROCESSING, AND DATA BASE PROCESSING – ARE CONFIGURED AT A SINGLE, CENTRALLY LOCATED SITE. SUCH CONFIGURATIONS ARE TOTALLY CENTRALIZED.

2. ALL THREE FUNCTIONS ARE CONFIGURED AT TWO OR MORE GEOGRAPHICALLY SEPARATED SITES. SUCH CONFIGURATIONS ARE TOTALLY DISTRIBUTED.

SEVERAL OTHER COMBINATIONS EXIST BETWEEN THESE TWO EXTREMES – ONE FUNCTION CENTRALIZED AND THE OTHER TWO DISTRIBUTED, TWO FUNCTIONS CENTRALIZED AND THE THIRD DISTRIBUTED . . .

Let us look at the three function sets from the point of view of distribution of the three functions. There are two extremes. In the one extreme all three functions will be configured at a single, centralised site. This is again the classic, centralised network. We have been building these for 15 or 20 years. The other extreme is one where each of the three functions is configured at more than one site. I have multiple sites executing information processing; multiple sites executing database logic; multiple sites handling data communications logic and so on. The same three basic functions, but now each existing at multiple sites, this is the new, totally distributed configuration.

If these are the two extremes, there must be some shades of grey in between, and indeed there are. Let us explore them a little further.

In this installation we have all three functions configured at the centralised site A. These again are the ones that we have been doing for 15 years or so. You can put these in with off-the-shelf components with relative ease from a variety of manufacturers. These pictures reflect the current technology; that is the database function is still embedded in and controlled by the same processor that handles the

TOTALLY CENTRALIZED CONFIGURATION

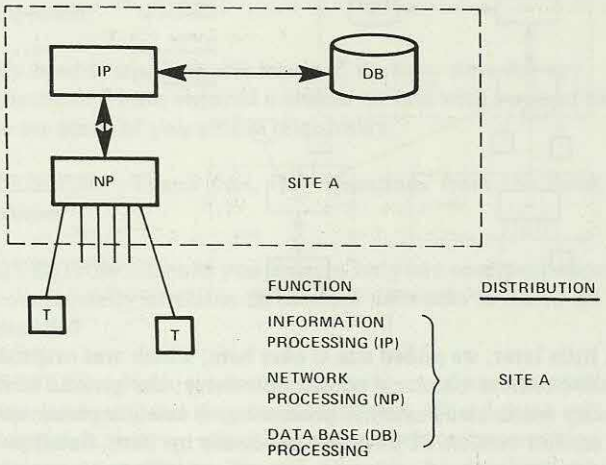

information processing function. When we get our database processor or our back end processor, these pictures will change. So this is our totally centralised network because all three functions exist at the same centralised location.

PARTIALLY DISTRIBUTED CONFIGURATION

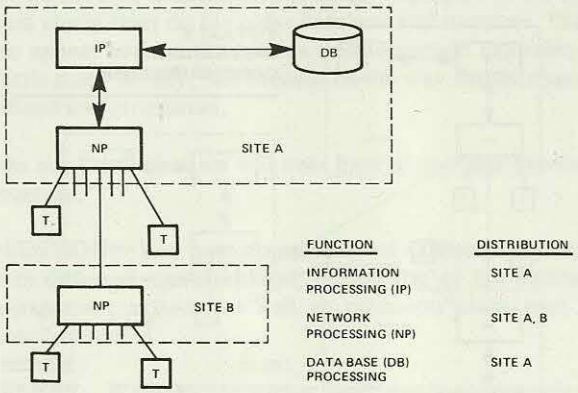

Not too long after that, about 1964 or 1965, the users began to ask us for the ability to allow this cluster of remote terminals to access the distant processing resources. Giving each terminal out here its own link all the way into the host site was prohibitively expensive, so the remote network processor or remote concentrator appeared about 1964. This in my mind marked the beginning of the distributed environment, because in this installation here the information processing function is centralised at the host; it is all done there. The database function is also centralised at the host. But the communication function is split between the front end, at site A, and the remote network processor here, at site B. So that function is distributed across two sites. This is what I would call a partially distributed configuration. We have been building these for 15 years. Again, most of the vendors can provide you with this capability relatively easily with off-the-shelf products.

PARTIALLY DISTRIBUTED CONFIGURATION

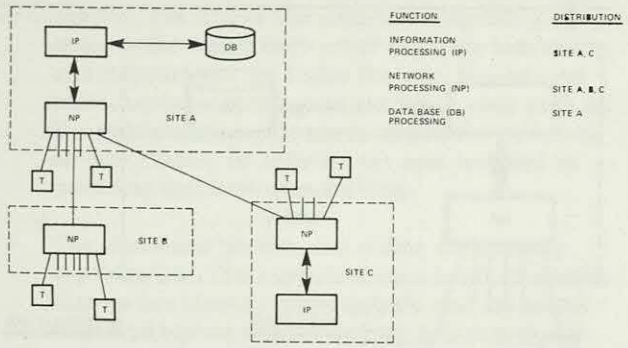

A little later, we added site C over here, which was originally called remote batch or remote job entry; the generic term today seems to be satellite processing. A satellite processor is a smaller version of our host processor up here. Satellite processors generally adopt a subservient role to the host. We have a master/slave or a host/satellite relationship. They do not view each other as co-equals.

This installation is another version of a partially distributed configuration. The database function remains centralised at site A; the information processing function is split between the host at site A and the satellite at site C; the communication function covers all three of the locations. So this is another version of a partially distributed configuration. These have been around for a long time also.

TOTALLY DISTRIBUTED CONFIGURATION

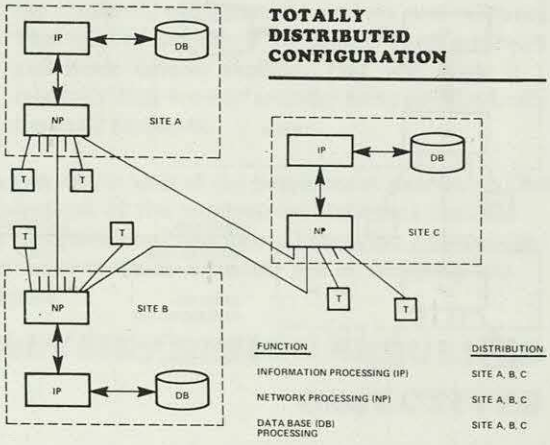

Now the one that people like to talk about — users and vendors alike — may look like this. I have chosen to limit it to just three sites here. The user says, "If I could have three processing locations, each one identical to the other, and have them interconnected with some suitable communications capability, I could distribute my data processing load uniformly across those resources, and get much better throughput and performance and response times than I can with independent locations." We are talking today about multiple dimension networks. A multi-dimension network is a set of hardware and software, perhaps like this, that provides the user community with timesharing as a dimension, transaction processing as a dimension, satellite processing as a dimension, and word processing. You have heard a little bit about office automation and word processing; that may be the next dimension that we will see on these kinds of networks.

This configuration is totally distributed. Each of the three functions exists at multiple locations. The users again like to think in terms of distributing the processing load very uniformly around these resources. It is very difficult to do yet, for two fundamental reasons. The operating systems that we use today in our host processors. All of the major manufacturers' operating systems are built around centralised philosophies, not distributed. We do not yet have mature distributed operating systems. They are coming, but they are not here yet. You do not take a mature centralised operating system and overnight convert it into a mature distributed operating system. By "mature distributed operating system" I mean an operating system that recognises the existence of co-equal resources. Sure, we can have host processors talking to each other today through a network, but they typically do it by treating each other as terminal devices. In the earlier days that was an acceptable level of interaction between them, but in today's networks that is unacceptable. It is not efficient or effective.

A second obstacle to this kind of installation is the database problem. Our current database architectures, the architectures that we use for designing, building, accessing and maintaining databases, are also built around the traditional centralised philosophy. You do not take a database architecture, a mature centralised database architecture, and overnight convert it into a mature distributed database architecture. All of the problems that we have today with centralised databases, as far as multiple update, concurrent update, conflicting update, reconciliation, deadly embrace and so on — all of the problems that we have with databases today in a centralised mode take on orders of magnitude more complex when we start distributing the database.

There are two approaches to distributing a database. In the first approach I will partition the data: that is the portion of the data at site A is different from the portion of the data stored at site B, which is different from the portion at site C. There are certain kinds of applications that lend themselves nicely to a partitioned database.

The other form of distributed database is replicated; which means that I have multiple copies of the data. I have a complete copy of the data at site A, and another complete copy at site B, and another complete copy at site C. Other types of applications lend themselves more effectively to a replicated database. Naturally, many installations in coming years will do both; will have portions of the data that are partitioned to serve those applications, and other portions of the data will be replicated to serve other sets of applications. So you begin to see the complexities involved in deriving mature distributed database architectures.

There is a lot of work going on in this area. Where it will end and how long it will be before we see these things as off-the-shelf products I do not know. I am guessing that distributed database architectures will probably be available in 12 to 24 months. The evolution of the back end processor, the free-standing database processor, is one of the key elements in deriving distributed database architectures.

One of the things that you should look for and start thinking about is: if we have a configuration like this, how do we determine how to partition the traffic? How do we determine how to partition the processing load uniformly

across the resources that are available? A given application program may one day execute at site A; the next day it may execute at site B. The next day it may execute at site C. Due to the changing load and usage profiles, the best place to perform a given function may vary.

We are starting to see the directory function occur. The directory function will be logic that may be distributed itself, that is we may have a sub set of the directory at sites A, B and C. The directory logic is that which determines where to send this job to get it executed most efficiently at this point. We may at one point move the process to the data. The next time we may move the data to the process; and the next time we may move both the data and the process to a third location. The directory function will be the one that does that.

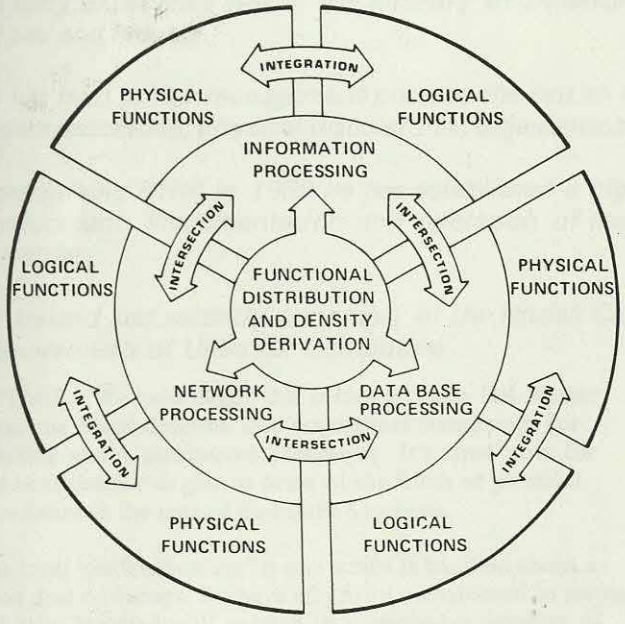

Let me add the last piece to the picture here, which is this bullet in the centre. We have identified the three basic building blocks, discussed the level of intersection between them, discussed the hardware and software sub sets of each of them — the physical and the logical; and the need for a selection and integration of the hardware and the software during the design phase. The bullet in the centre now adds the two highest level objectives of the analysis and design sequence; that is, how do I distribute the three functions; totally centralised, partially distributed, totally distributed? What is an appropriate functional density? That is, how much information processing at each of the sites? How much database processing at each of the sites?

So if we can, we can boil the structure here down to one diagram and summarise it here with this chart. Again, the distributed processing approach is not for everybody. There is a substantial class of users who will for years to come be satisfied with the traditional, centralised approach. Some users are forced into exploring the distributed configuration, for the problem that they are facing is called complexity of scale. For 20 years this industry has been saying, "We're going to build bigger, better, faster, large scale mainframes to make it easier for you to do all of your processing in one big super centre." Well, there are a number of users today who, in pursuing the economy of scale they have been promised in that approach, are running headlong into the complexity of scale inherent in that approach.

For many users today, the technology just does not build a single processor big enough and fast enough and inexpensive enough to do all of their processing in one machine. So they are forced into exploring the distributed approach.

My hour is up. I do not know if we have time for any questions. I have enjoyed a chance to talk with you and hope to see some of you a little later today.

ZEDLITZ: Thank you, Hal. Questions from the floor, please.

QUESTION: Could you enlarge on your comment about commercially available distributed networks in 12 to 24 months?

BECKER: Each vendor has its own set of architectures for databases, and they are all different. What I think will happen is that we will learn how to achieve intersecting database capabilities. We will, through the development of additional higher level languages — vendor independent in a sense — be able to have intersecting database capabilities. We cannot expect one vendor to modify or translate their entire database into the structure of another. The user community just will not stand still for that, and I do not think they should. What I think the user community should do, and is doing, is to ask for the ability to have this database intersect to some extent with this database.

The new relational databases are getting a fair amount of play over here. The relational database concept appears to lend itself to a distributed database organisation a little more nicely than do the older database architectures. They also appear to lend themselves to intersection capability a little more readily, just because of the way the database is defined and structured.

I do not think that we will ever have a standard database structure.

QUESTION: You have complexity of different protocols from different manufacturers, complexity of the database management system, as well as protocols from each manufacturer.

BECKER: We have had higher level, vendor-independent languages for information processing for years — Cobol, Fortran and so on. We have had higher level languages for databases. Not vendor-independent, but there are higher level languages. IBM has higher level database languages; Honeywell has them, and so on. What we do not have yet is a higher level language for communications, but there are some glimmers of it beginning here. There is an ad hoc committee being formed by one of the publishers here. I am involved with the project. We are going to sit down soon and start defining what we think is a first step towards a higher level language for communications, which may or may not be vendor-independent. But if we can define the structure for a higher level language, then maybe the various vendor implementations of that standard language will ease at least the communication problem that we face — the problem today is that in a given network it is not uncommon to have four or five different code sets rumbling around. Five level Baudot, 6-bit BCD, EBCDIC, ASCII, and so on. The different link protocols. It is not uncommon today in a network to find yourself having to deal with

two, three or more different link protocols. The single vendor network does not exist yet, in spite of IBM's advertisements. Any network you look at is a combination of a number of different vendors' equipment. It is not uncommon to find yourself as a user dealing with half a dozen different vendors. Of course, we are all familiar with the finger-pointing that goes on when there is a problem; everybody says that it is somebody else's problem, and you go round and round.

This is a long answer to your question, but I think that the higher level language for communications will help the process. The relational database approach will help us to see our way into intersecting database capabilities. Now it is obviously a long time out before that capability will exist with standard, off-the-shelf products. That is not something that I would say would be here within 12 months, clearly.

ZEDLITZ: Thank you, Hal.

POLICY ISSUES ARISING FROM DISTRIBUTED PROCESSING

Jim Ireland
RHM Management Services Limited

Jim Ireland is a founder Fellow of the British Computer Society and has been actively involved with large scale data processing systems using computers for the last twenty-six years.

His early experience was in the Ministry of Defence of the development of the major computer installations for pay and records.

He has held senior management posts at the British Rail Board, with the Burmah Castrol Group of Companies in data processing, physical distribution, organisation and systems planning.

Since joining RHM in 1969 he has established a highly professional multi-disciplined team specialising in the development, implementation and operation of management science techniques across a broad spectrum of industries.

Mr Ireland was recently a member of the British Computer Society committee which reported on the 'Future Requirements of Users for Computing'.

BUTLER: We have heard this morning, from Hal Becker, what one could describe as a conceptual framework for thinking about distributed processing. It's time now for our next speaker to give us some of the fruits of practical experience in the area of distributed systems.

The term 'professionalism' is one which is bandied about a good deal nowadays. We have all grown accustomed to seeing the term 'professional' applied to a particular category of foul-football, where it seems to mean the blatant, deliberate and totally unfair, so it's a term which is abused a good deal now. But I think we need to remind ourselves from time to time what we mean by professionalism and I think that our next speaker, Jim Ireland, has devoted a good deal of his effort over the last 25 years to giving meaning to the term 'professional' in relation to computer systems — both as a founder member of the British Computer Society and as the Managing Director of RHM Management Services which, as most of you already know, is RHM's vehicle for providing truly professional services not only within its group but also outside, a model which is also followed by one or two other companies in this room.

To hear what the fruits of the work at RHM on distributed processing have been and what problems have been encountered and how they have been solved, let me now hand you over to my friend, Jim Ireland.

IRELAND: Good morning lady and gentlemen. Before such a distinguished and experienced audience, and following a day and Hal this morning of competent and polished speakers, I have some sympathy with how Elizabeth Taylor's eighth husband will feel on his wedding night. I know what to do, but I don't know how to make it more interesting.

Rank Hovis McDougall is a British company. It was founded about a hundred years ago on the base of the old Joseph Rank flour milling business.

Just to give you some idea of the scale, it has a turnover of £1,500 million per year, it employs 55,000 people, operates from 200 locations in the UK and is a largely UK based company but does have international operations in 35 countries.

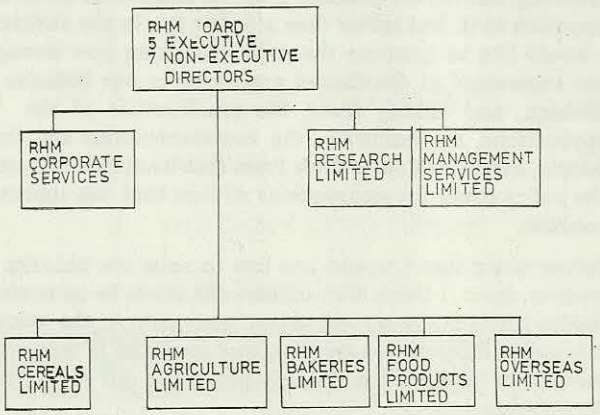

It is not a Freudian slip, the top of the slide is the RHM board. There are five executive and seven non-executive directors. The main trading divisions are along the bottom and they are supported by two service divisions, RHM Research and RHM Management Services, which is my company, and a very small corporate department. In fact the Corporate Services Department and the board are only 110 people including secretaries, so it's a very small unit for a fairly large company.

Very briefly, talking about the company — because I think to understand distributed processing in RHM you need to get a little bit inside the culture of the company — the Cereals Division is concerned with milling flours from wheat,

production of pastas and the supply of gluton and starch to produce syrups to other primary manufacturers and producers. The Cereal Division operates as 18 profit-orientated companies.

The Agricultural Division is concerned with a total supply and service to the farming community, including cereal and herbal seed, spray chemicals and fertilizers, re-purchasing of farmers' crops, animal feeds compound, life stock breeding and marketing, and sales of farm machinery. It operates as 20 companies.

RHM Bakeries is concerned with the manufacture of bread, breakfast rolls and other morning goods, and confectionery and cakes, and the distribution of its production to super-markets, smaller shops and stores and to individual housewives. It operates as 80 companies.

RHM Food Products is concerned with the manufacture and distribution of food products to supermarkets and shops, and the products include such things as salt, household flours, breakfast cereals, soups, stuffings, soft drinks, cheese, butter, conserves, chutneys and sauces. It operates as nine companies.

RHM Overseas operates in 35 countries around the world. The Overseas Division is only concerned with the operations in non-EEC countries. Where there are operations in EEC countries they are aligned to the divisions.

Yesterday, David Butler talked about the strategic issues in office automation. If I was to change that to the policy issues in distributed processing I would probably talk about the same six things that he talked about — project evaluation, the human dimension, planning the foundations from which you build, the timing and the technology — and I will seek to try and develop some of the technical and people issues affecting distributed processing and management services' approach to it. But rather than attempt this in the abstract, I would like to illustrate the point by taking you through our experience of distributed processing in our Bakeries Division, and talking about the practicalities of the applications, the hardware, the communications and the people, and leave you to draw from that what you feel are the policy issues for management services that this approach contains.

Before doing that I would just like to scale the bakeries because, again, I think it's a culture one needs to understand. Baking bread is the second oldest profession in the world. It is highly labour-intensive. It is very unskilled in terms of the type of labour it has. Twenty percent of our employees are illiterate.

RHM BAKERIES LIMITED

OPERATING UNITS	-	100
EMPLOYEES	-	38,000
VEHICLES	-	9,000
PRODUCT TYPES	-	2,000
SHOPS	-	2,500
CUSTOMERS-CREDIT	-	10,000
-CASH	-	7 MILLION

In many bakeries the quality control is dealt with by colour charts, rather than chemical specifications — if it's black you throw it away, if it's white you haven't cooked it enough. In something like 12% of our plants in the UK, all the notices in the bakery are in Urdu. We are the most government investigated industry in the UK. We have had 12 major government investigations in as many years. The management style of the bakeries is called fear.

RHM Bakeries has 100 operating units. It employs 38,000 people, it operates 9,000 road vehicles, it has 2,000 products and it operates 2,500 of our own shops. It has something like 10,000 credit customers and it has something like seven million cash customers.

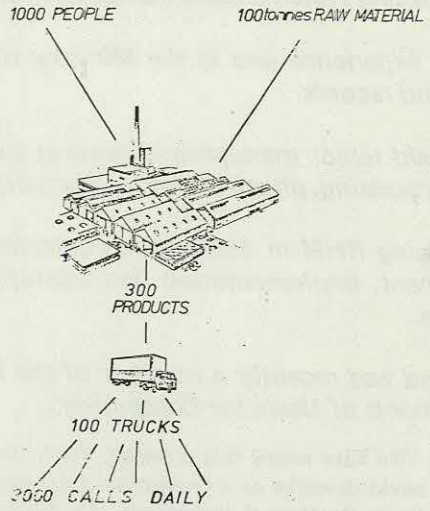

A typical production bakery will employ about 1,000 people. This is slightly above average — it's the size we are working towards rather than the average size. It puts through about 100 tons of raw materials in a week. It has 300 products, of which about 150 are national products and 150 are regional products. It may have about 100 vehicles and those 100 vehicles will make about 3,000 calls daily.

The time scale of operation is short cycle. The vehicles start returning at about noon each day. We have the business under tension.

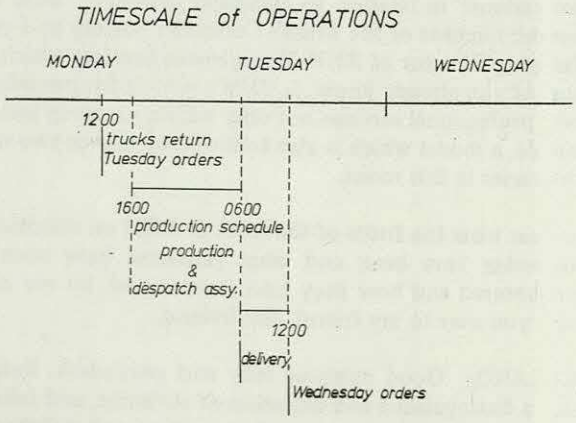

Production sells to the vehicle drivers and the vehicle drivers are in fact the primary outlet of the bakery, so they are the first level customer, if you like. The driver in turn has his customers, who are either shops, supermarkets or house-wives. So when the drivers come back the first thing you have to do is to make sure that you receive a reconciliation

of the product they took out from the bakery that morning, and the second thing they do is to firm up their orders for the following day on which the afternoon and the overnight production will be based.

So between noon and 4 p.m. drivers are coming in and firming up their following day requirements. Between 4 p.m. and 6.00 a.m. on the following day production and despatch assembly and vehicle loading takes place, and the cycle starts repeating itself again. And it is a six day cycle.

During the 18 hours that we have available to us, the whole of the production and the distribution preparation process has to take place. Many feasibility studies were conducted into computing in bakeries during the '60s and early '70s, but none of them could meet the design criteria laid down for the successful integration of computers into bakery operations.

This situation changed in 1972 and 1973 with the emergence of intelligent terminals, minicomputers and improved communication capability to mainframe machines. In other words, we had the capability to consider distributed processing.

Much of my talk this morning is going to be from the bakery management point of view rather than from the data processing professional point of view. Distributed processing brought a change in the approach to designing computer systems; unlike the old central batch-oriented systems the user actually became important, and meeting his requirements became mandatory rather than a hopeful by-product of the technical elegant computer solution.

COMPUTING COMPATIBLE WITH

THE ORGANISATION, STRUCTURE,

AND OBJECTIVES OF THE

BAKERIES AS A DIVISION

So if we perhaps look at the management objectives that were set, the first was that computing should be compatible with the organisation structure and objectives of the bakeries as a division — meeting regional and local needs as well as national in a way convenient to the bakeries and not to the computer systems designer — and should be in line with, and indeed embedded in, the daily production and despatch cycle.

IMPROVED ADMINISTRATION
- containing costs

-increasing demands & complexity

-increased information & quality

-reduced dependance on clerical staff and turnover

-reduced accounting equipment

A secondary objective was to improve the bakery unit administrative capability while containing costs and reducing the number employed, and at the same time meeting increased demands and complexity of information requirement, mainly from Government (which was estimated to be growing at the rate of three to five percent per annum compound and was leading to greater requirements of clerical labour).

Secondly, to meet and improve upon the timescale and quality of information required by national management for discussions with national customers, Government, trade unions, shareholders and employees.

Thirdly, to reduce the dependence on low level clerical staff, and in particular the high rate of turnover and rapidly inflating costs, without any attendant increase in the added value that those clerical staff brought to the work.

And finally, to reduce the multiplicity of accounting equipment of different ages, types and systems in use throughout the Bakery Division, and to achieve standard but flexible systems.

At that point I resigned!

On the basis of those management objectives, I think the first thing to look at is the applications, and then to follow that with a look at the hardware and communications, and then multiple suppliers, and finish with a bit of the people side.

Firstly, the applications selected had to primae facie meet the management objectives and provide a base and a framework upon which future applications could be bolted. So the applications were considered in two groups.

COMPUTER ORDERING SYSTEMS
IN OPERATION

1 ORDERING

2 PRODUCTION SUMMARISATION

3 DESPATCH

4 CHARGEOUT (& SETTLEMENT)

5 CREDIT INVOICING

The first group were those concerned with the process of ordering, the preparation of production and despatch summaries, the calculation of chargeout of the value of product put on a van, and to settle the van man on the basis of what he has sold, either for cash or credit, or what was spoilt, or what he brought back in the way of returns, and finally, to present information to the central computer systems for credit invoicing of the national customers — particularly the supermarket groups such as Co-ops, Tesco and Sainsburys.

A sub group of this first set of applications was that concerned with our own 2,500 shops and was concerned with monitoring and controlling their trading patterns. So in effect what we were doing was treating our shop as though it was a van, and making sure that it was screwed down to the ground every day in terms of what it had received, what it had sold, what it had returned and what it had spoilt, etc.

PRODUCTION SYSTEMS

Main Processes

-costed recipe

-stock reconciliation

-resource usage

-production/despatch reconciliation

The second group of applications was concerned with the production process itself, which is batch-oriented — and with its base dough mixes and secondary mixes, wrapping, raw materials and labour, it is not dissimilar to a conventional assembly process in light industrial manufacturing. In fact, if you ask a baker what baking is all about, it is how to make water stand up, sliced, in a wrapped envelope.

The component parts of our production sub systems — in our jargon — are costed recipes, which are bill of material processing, stock shrinkage reconciliation, resource usage of both materials, labour and energy, production and despatch.

It was felt that these two main applications and sub-systems — ordering and production, and operating at local bakery level — would provide all data capture and the framework onto which other systems could be bolted, such as local and national management information systems, payroll and accounting systems, trading and profitability reporting, which is done on a weekly basis in the Bakery Division.

So compared with the classical approach to computer applications, we were starting at the wrong end. We were doing the in-line things first and putting the payroll and accounting things on at a later date.

It might be useful if I were to go through the 12 applications, mentioning some of their features and purposes — and I say this because I am using the words of a bakery management that set the terms of reference — and I will flip through very quickly, but I think you might find it worthwhile.

ORDERING

Purpose

-maximise sales,controlled return.

-responsibility at point of sale

-demand on production/despatch

If we look at the ordering systems, the purpose was to maximise sales while controlling returns. It must be designed to leave the responsibility at the point of sale — that is, the vanman or the shop manageress — and it must result in a demand on production and subsequently despatch.

Some of the features were:

- that the starting point ought to be the actual sales of the same day and the previous week, to minimise the amount of data that the salesman would have to put down on the terminal and the amount of processing that would have to be done by the distributed computer system.
- There would have to be freedom in ordering, in that the salesman or the shop manageress started the ordering cycle but that the shop manager or the retail or wholesale sales manager could, in fact, impose his will on it — either in the form of upping everything 10% or downing it, or pushing particular lines on which promotional activity was planned to coincide with it.
- And finally, it triggers production.

PRODUCTION SUMMARISATION

Purpose

-demand on production

-production orientated

If we look at production summarisation, the purpose was to place demand on production. The sales product, prior to creating a production demand, would be converted by the computer process into the units that production like to work with compared with the units that salesmen like to work with.

It would have to itemise all finished goods and bought in items. The production manager ought to be able to walk along and get an interim production command off the computer, even though all salesmen hadn't finished their following day's orders, as well as final demands on production. It should be by product and product group; it should produce finished goods summaries as well — in production-oriented terms — not only for products and product groups but also by our outlet types, whether they were bands or shops and, where a bakery had a particularly large wholesale business, perhaps splitting it into individual customer requirements for wrapping, because Sainsbury's have their own particular wrappers compared with Tesco's, compared with British Home Stores, and so on.

Very much production-oriented, pack extensions, units, sacks of flour, loaves of bread and the coding structure would be decided and set up by production management according to the production profile of that individual bakery. And some bakeries were bread bakeries, some were confectionery bakeries, some would be mixed, some would have tweedy mixers, which are semi-automated mixers, some would have manual mixers, some would have turbo ovens, some would have gas fired ovens.

So we were talking really of writing software packages, and an individual bakery would pull out those that it needed for its own particular profile.

DESPATCH

Purpose

-controlled allocation of products

-despatch adjustments

-method of vehicle loading

-entry of late orders

Despatch — the purpose was the controlled allocation of products to points of despatch to be able to report on adjustments carried out at despatch, to provide a method of loading vehicles, and to provide a facility for bringing late orders into the system or, if there had been problems in production, to provide a facility to treat equally across all the vans, shortages of production.

CHARGEOUT

Purpose

-charging of goods issued

-provision of management information

-production/gross despatches reconciliation

Charge out is the fear bit of bakery management — that when a vanman loads his van he needs to know the exact value of all the goods issued and loaded onto that van prior to the time he is allowed out of the bakery. He is also the starting point for management information and it should facilitate a reconciliation between gross despatches and the actual production from the previous night.

CREDIT INVOICING

Purpose

-invoicing of credit customers

-input to settlement

-local control

-links to national systems

Credit invoicing — the purpose of the local operation was to prepare data, the accuracy for which the local management could be responsible, ready for subsequent transmission and incorporation into the national credit invoicing system. It should also provide the input of credit items to the daily settlement system, provide local control and, lastly, the links to the national system.

SETTLEMENT

Purpose

-daily round accounting

-report on claims & allowances

-basis of sales analysis

Daily settlement is, I think, something peculiar to the bread and milk industries in the UK. Its purpose was to achieve daily round accounting by striking a balance each day for each vanman. It reports on all claims and allowances and promotional incentive schemes, and it provides a basis at local level of sales analysis by product groups and individual products.

SHOPS SYSTEMS

Purpose

-shops settlement trading accounts

-performance reporting
+shop catering controls
+trading trends

-basic information to
+national shops systems
+mgt./financial accounts

-'ad hoc' management information

For those bakeries which are concerned with operating some of the 2,500 shops we run ourselves, the system should provide shop settlement trading accounts that include all basic financial trading requirements. It should provide performance reporting, shop and catering controls, and trading trends. It should provide basic information for transmission and incorporation into national shop systems, and it should provide bakery unit management with their financial accounts. And finally, it should have a facility for ad hoc management information things like trading trends, sell out times of various lines, etc.

PRODUCTION SYSTEMS

Purpose

-reporting on entire production process

-promotion of management action on abnormal production situations

-standardisation of report content

-applicable over varied bakeries

Turning to the requirements of some of the production systems, the purposes of the production sub-systems are that they should report on the entire production process, from raw materials to finished goods; that they should promote management action on abnormal situations, i.e. those that

deviate from an acceptable normal standard; they should promote standardisation of report content and format as an aid to production. Management training and mobility from bakery to bakery and its production sub-systems should be applicable over a variety of bakeries; it should be flexible and be capable of encompassing different physical profiles and management structures and different bakery production units.

COSTED RECIPE

Features

-structured recipes

-ease of change

-aid to management

-basis of production systems

Looking at the detail of some of the four component parts: costed recipe — which is our bill of material processor — the ability to structure recipes with a production bias into processes and resources, quantity and financial terms, and make provision for sub-assemblies and assemblies or batch mixes and various dough mixes. They should be easy to change in updating recipes for changes in price, changes in resource levels, resource substitutions in situations of shortage. They should provide an aid to management in hypothesis testing — the 'what if'?' question. What if apples are cheaper than dates this week? It should quantify the implications of recipe changes for substitutions, permit the separation of variances into those attributable to price and those attributable to other factors, and finally to provide a basis for production systems in that all production sub-systems draw upon the costed recipe information and to achieve, promote and maintain standardisation.

STOCK RECONCILIATION
(Raw Material Movement)

Features

-traces stock movement

-identifies variances

-automatic valuation

Stock reconciliation. The objective was to trace the stock movement throughout the production process, achieving reconciliation of actual stock shrinkage against standard usage, to identify variances and stock position to aid management control, and to identify those variances at many points in the production process and therefore — independent of any one bakery's physical layout and management structure — to provide automatic valuation of stocks and stock shrinkage in both quality and price terms by production cost centre at actual versus standard, in pricing terms, in units for production purposes, and in value for accounting purposes.

RESOURCE USAGE

Features

-actual v standard comparison

-identification of variances

Resource usage. To provide actual to standard comparison throughout the entire production process of all resources, ingredients, packaging, labour at many points in the production process, the bulk and breakdown stores, mixes and part and wholly finished products. Also to identify variances on an ingredient-by-ingredient basis to promote management action, separate price and quantity effects by cost centre and at three levels in the production process.

PRODUCTION/DESPATCH RECONCILIATION

Features

-daily highlighting of variances

-streamlined standard documentation

Production despatch reconciliation — daily, to highlight variances to lead to faster and more effective management action; to streamline documentation, making the task of recording easier and therefore achieving greater accuracy, automatically transferring volume and value information to weekly management accounting and profit reporting systems on the central computer.

When we were told all that by the bakery management, the rest of my staff resigned!

While the establishment of these various sub-systems' needs was going on and the purposes, features and architectures were being progressed, at the same time we were evaluating the available hardware and communication facilities.

Now please bear in mind that the time was July 1974, and the management criteria for the hardware — and this is bakery management criteria — was as follows:

- It should be powerful enough to support all time critical applications without recourse to the central computer.
- It must be proven equipment — the 'if I can touch it it works' syndrome.
- It must come from a financially stable company.
- It must have an ability to carry out installation and maintenance of the equipment on a nationwide basis, i.e. from Inverness to Truro.

A number of manufacturers were considered, each of whom from a technical viewpoint had equipment that could meet the basic applications specification, but the list dwindled to two or three when the management requirements were applied and to one on cost evaluation. And the one selected was the ICL 2903 minicomputer.

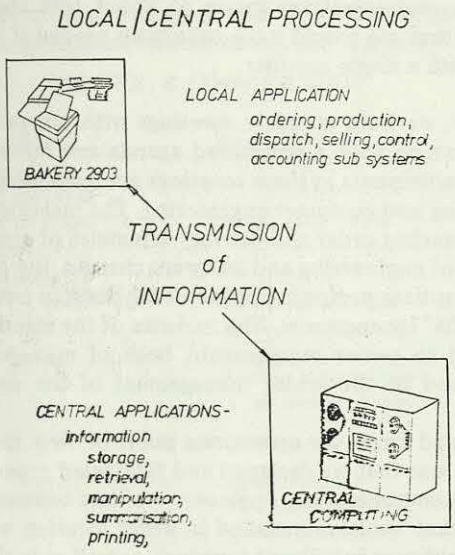

The 2903 minicomputer specified consisted of a 24K word processor, 300 line per minute printer, card reader, operator's console and 20 megabytes of disc (ten of which were fixed and ten demountable). Data entry was to be done by between five and eight direct entry keyboard and screen devices according to the bakery size.

All time-critical applications would be processed locally on the 2903 and any information required for central applications such as credit invoicing, shops and rounds, management information, weekly profit reporting, would be transmitted to the central installation. The method of transmission would be the public switched network or, if that particular bakery was on the group's private voice network, then the data would be moved over that voice network in switched mode.

Similarly, as all program development was centrally controlled all program changes would be down loaded from the central installation to the system's resident disc pack, as would all changes to national product prices and codes. And this was done once a week — Monday morning first thing.

The initial capital authorisation for 2903s was for 14 sites at a capital cost of £50,000 each, to have a viable payback through a net reduction of 25% in the clerical and administrative cost of the bakery unit. It was assessed that 40 out of the 80 bakery units would meet this criteria and it was hoped that a combination of people/cost inflation, bakery rationalisation into fewer but larger units and falling hardware costs over time would bring all units into a situation where they could meet the payback criteria.

The decision to order was made in February 1975, and the programme called for a new installation every six weeks commencing in August 1975, and each installation required 12 application sub-systems to be implemented and post operational savings to be achieved.

The following areas were success stories:

- Physical planning and equipment installation.
- The application development time and cost (in fact the order processing application took four months to write and test and implement in the first site. In fact, although it was done in four months it was two man years of effort and when we compared it with what it would cost to develop for our 370 mainframes we worked out that it would have been 80 man years and 18 months. So the effect was dramatic in terms of system development time).
- Application implementation was a success story.
- The performance of the 2903 in terms of hardware uptime was excellent. In fact now, from hindsight after about 2½ years, we are getting better than about 98.5 uptime from our 2903s.
- The Facility Management concept which we implemented at the same time, which I will refer to later in my talk, was also successful.

However, we had two disaster stories.

Firstly, the hardware performance, in terms of throughput, was nowhere near the manufacturer's specification and the result of that meant that the achievement of the planned savings was impossible. The principle problem was that they said we could do all this in single shift, and we found we were running 18 hours a day to do it, and so we were not only double-shifted but any of you who are ICL users will know that ICL maintenance outside normal working hours is prohibitive in cost.

The cause and effects are typical and I don't need to elaborate, but the net effect on the performance of the first six bakeries implemented was that only 20 bakery units instead of the expected 40 would meet the financial payback criteria.

Fence-sitting management got ready to fall off on the side of the "I told you so"s, and the committed managers got ready to wipe the egg off their faces, or worse.

The solution lay in either getting the hardware to perform to its specification or to use cheaper hardware that would eliminate some of the scheduling bottlenecks and manning requirements without negating the design criteria and the cost that had already been incurred in program development — all 12 application sub-systems were virtually complete.

The solution lay in upgrading the 2903 and going to a secondary level of distribution using a smaller ICL machine which by then had been announced (this is now early 1976) which was the 7502 minicomputer.

The 7502 consisted of a 20K word processor, 150 line per minute printer, 5 kilobytes of floppy disc and three to four VDU screen and keyboards (which replaced the five to eight direct entry devices on the 2903).

The 2903 was upgraded from 24 to 48K words, from 20 to 40 kilobytes of disc and from the model 30 to model 40

LOCAL/LOCAL PROCESSING

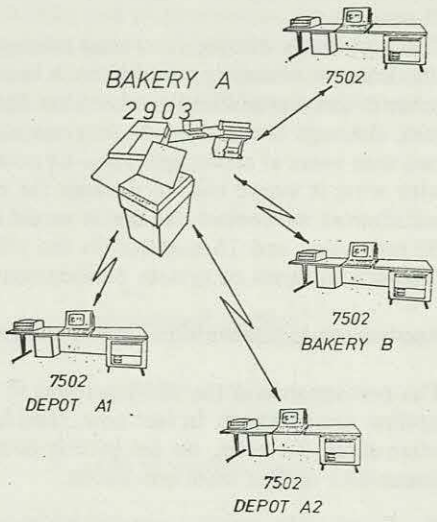

to give a faster cycle time.

The second level distribution clustered either a couple of bakeries or a couple of depots, or perhaps three bakeries to a host 2903. The most time-critical applications were performed on the 7502s without recourse to the 2903 — for example ordering and production despatch summarisation — while the 2903 dealt with all other systems for the bakeries in the cluster and communicated to and from the 7502s by means of either, in the case of the host bakery — hard wired — in the case of remote bakeries to the host bakeries' 2903 — private leased line, and of course the 2903 itself communicating to the mainframe computer either by PSN or by group private voice network.

Again, without going into all the arithmetic, the net effect of the hardware change was that to achieve the financial payback target, we now required to reduce the bakery administrative cost by a net 20% instead of the original 25% when we were considering the 2903s on their own. The lower overall hardware cost, the ability to share a lot of our facility management resources between a number of sites in a group rather than dedicate them to an individual site, in fact brought the number of sites in which the system would be viable up to 50 from the original 40.

I think this is what David referred to by the unholy alliance of computer salesmen, and indeed data processing management — because we were able to present it as a better solution than the one we originally thought of.

A word on the multiple vendor situation. For many years in RHM we have had a number of vendors in our central computer installation to achieve what we believe is the most economic equipment profile — IBM, Itel, Memorex, DEC, Racal-Milgo and the PTT. With the advent of distributed processing and telecommunications we added, over a period of 18 months, ICL, Olivetti, Data 100, Interscan, Motorola and GEC.

The problems of many people contributing to a single end result is no different in data processing than in any other field, but the effect of lack of cooperation is perhaps more dramatic in data processing than in some other fields and certainly more immediate and more costly.

We began to suffer from the "No it's not my problem, it's his," syndrome.

We did three distinct things to minimise the adverse effects of multiple vendor situations.

Firstly, by a conscious decision we manned our software and telecommunications group to about 50% above the strength that we would have otherwise needed if we had stayed with a single supplier.

Secondly, we hold monthly meetings with our principle vendors which have a structured agenda and minutes. The vendor participants at these meetings are from sales, systems engineering and customer engineering. The meetings review the outstanding order and delivery schedules of equipment, the state of engineering and software changes, the previous month's uptime performance including 'time to arrive' and 'time to fix' by engineers. The minutes of the meetings are circulated to senior management both of management services and to the senior management of the vendors.

Thirdly, and from the operations point of view the most effective, was that we designed and fabricated — using our own telecommunications engineers — a data communications monitor that we have installed in every location with a minicomputer or intelligent terminal, as well as in the main computer centre. The data monitor sits between the modem and the terminal or the minicomputer, but it is transparent to it. It monitors line speeds, error rates, re-tries and transmission, it pinpoints malfunction to device, network or main centre, and it acts as an independent alarm device that can be triggered from the main centre to all outstations simultaneously.

It was enormous help in getting the right manufacturer to the right piece of kit when it broke down.

Finally, as my time is running out, I would like to say some words without pre-empting the next speaker's presentation on the crucial people aspects associated with distributed computing.

Centrally processed batch-orientated systems have required a low level of user involvement in either the design implementation and operations phases. In fact the only reason most of the users were there was to pay for the system. Often it was limited to training user clerical staff to fill in data entry forms and training user management to use the output.

One of the objectives of distributing computer power to the users was to make the users feel more committed and responsible for their systems and the equipment because it is in their budget, their location, and under their control.

Control and communication of a data processing system that is distributed across the RHM bakeries was one of the most

CONTROL & COMMUNICATION

'WHO IS RESPONSIBLE
FOR WHAT ?'

significant areas for our consideration. Well done it would assure a successful outcome but, badly done, disaster was inevitable. Some people thought disaster was inevitable anyway.

So who is responsible for what in a project that affects 100 sites, ten regional managers, the national bakery management, three hardware vendors, one software house, the PTT, 30 local operational managers and over 100 systems designers, programmers, implementors and trainers?

CONTROL & COMMUNICATION

PROJECT	'Bakeries Project'
STEERING COMMITTEE	– overall policy control
USER COMMITTEE	– bakery representation – systems/implementation policy – function representation via Working Parties
PROJECT REVIEW	– technical support & control – system development & implementation – manufacturer liaison

Firstly, the project must be seen to be user driven — a bakeries project in this case, not a computer project.

The Steering Committee was concerned with overall policy control. It met monthly, it was chaired by a director of the bakeries, and it had a formalised agenda and minutes. We compared the actual against the budgeted performance in terms of time, resources and cost, for both development implementation operations and engineering economics. And finally, their main role — in addition to reviewing that — was problem areas such as unplugging bottlenecks.

The user committee was formed from representatives of Bakery Division Management. It was responsible for representing the bakery view and for system implementation policy and timescale. System functional requirements were initially prepared by a number of working parties whose members were drawn from the Bakery Division, such as sales and accounting people for the order processing sub-systems, production and accounting people for the production sub-system.

The user committee had the final sign off responsibility for system specifications that were produced by the working parties that were reporting to them.

The Project Review Committee was formed from the Management Services and data processing professionals, and was responsible for technical support and control of systems development, programming and implementation and vendor liaison.

CONTROL & COMMUNICATION

REGIONAL SYSTEMS ACCOUNTANTS	– regionally responsible for implementation – post implementation audit – interfacing clerical systems

The fourth group of people were the regional systems accountants, who were bakery people, responsible for any preparatory work required in the bakery of their region — which was having a minicomputer installed — other than physical planning. They were not responsible for any physical planning.

They were also responsible for post implementation audit and the dismantling of superceded systems and equipment and for identifying planned savings at unit level with the site managing directors or general managers, and they were also responsible for the interfacing of the computer-aided systems with the other non-computer clerical based systems at unit level.

FACILITIES MANAGEMENT ROLE

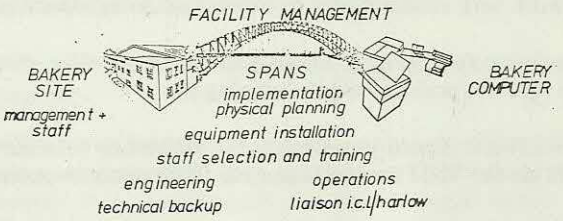

The final people in the piece were the facility managers. Now you may have picked up that though the bakeries have gone for computing in a big way, they haven't taken on one technical data processing man, either in development operations or systems, and that was part of their policy — they didn't want to know the problems of having data processing professionals inside their division.

Well that was fine in terms of design and programming and all those other aspects of the technical side, and the system implementation side. The problem lay in finding how we could bridge the gap between putting a computer in the bakery and the site management.

As a result of seeing this as a problem, it was decided that Management Services would recruit and train some of its central unit shift managers and shift leaders to act as on-site operations managers to effectively bridge this gap between the computer in the bakery and the bakery management.

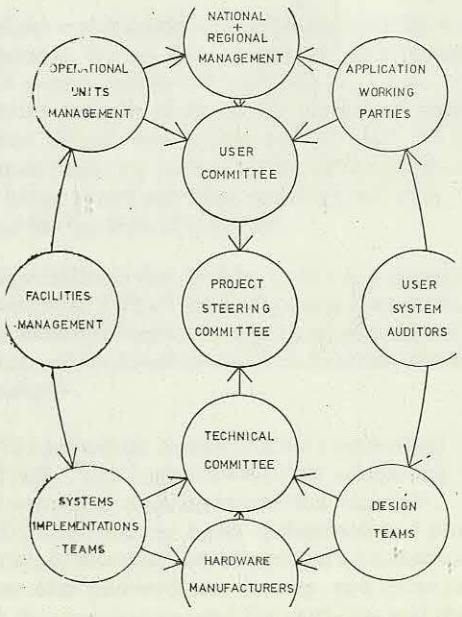

The facility manager was initially responsible for one site; they are now responsible for two or three, and ultimately — when we have completed the programme — about four or five sites per manager. So it's a hand-holding job, not an on-site all the time management job.

The role spans implementation and physical planning, equipment installation, staff selection and training, scheduling of operations, liaison with vendors and central management services staff, and on-site technical back-up to the bakery staff.

To sum up, a schematic of the control on the people side, for control and communications. It looks like this.

It was set up for the development stage but it has been found to work very satisfactorily in normal on-going operations.

Thank you for listening. I hope you found the nitty gritty of a user's experience worth listening to.

QUESTION: Could you explain the difference between the work on the 7502 terminals and the 2903 minicomputers?

IRELAND: Those jobs which were embedded in the 18 hour ordering production summarisation, despatch summarisation had to be done without recourse to a different site, so they had to be the things that were on an individual bakery 7502, without reference to the 2903.

So the 2903 is really more concerned with the production systems (that is, costed recipe, production reconciliation, stock shrinkage) and also controls the transmission to the group computer centre of all credit invoicing and management information for further processing.

QUESTION: How many systems are now installed?

IRELAND: Our intention is that we will go to 19 2903s, of which we have 17 installed, and we will go to 46 7502s, of which we have 34 installed. So we are about half way through the programme in site terms. We are probably about 80% through the programme in terms of volume of business.

We are installing either a host or a satellite once every four weeks, and that will go on till the end of 1980 or early 1981.

BUTLER: Thank you very much Jim. Ladies and gentlemen, I am afraid we will have to stop there.

I am sure that the members would like to join me in thanking you for a truly professional presentation and a very balanced picture of the difficult problem areas. Many thanks.

DISTRIBUTED DATA PROCESSING - SOME OF THE PEOPLE ISSUES

Carl Reynolds
Hughes Aircraft Corporation

Carl Reynolds is a graduate of Harvard, where he read physics.

Mr. Reynolds worked for IBM in sales, engineering and system development roles during his twelve years with them. In particular he was Manager of Systems Programming between 1962 and 1966.

Mr Reynolds is currently Director of Communications and Data Processing for the Hughes Aircraft Company. In addition he is a member of the Advisory Board of Datamation.

BUTLER: One of the criticisms often levelled at the data processing industry is its preoccupation with its own problems and the tendency to completely ignore the question of how ordinary human beings working in the real world are going to live with the systems which it produces. These problems — the human aspects of systems — are becoming more and more worth our attention, particularly as in some countries in Europe already there are legislative rules which deliver a good deal of control into the hands of people and their organisations, who can inhibit the introduction of systems unless it can be proved that these human aspects have been properly considered. Therefore I think that our agenda today, looking at distributed processing, would have been deficient if we had not had at least one session on some of the people issues that arise from distributed processing.

We have invited to give this session a speaker from a company, the Hughes Aircraft Corporation, that has always had a very high and deserved reputation for making the most creative and imaginative use of manpower. Therefore I should like to introduce to you, with great pleasure, Carl Reynolds, from Hughes Aircraft Corporation, to speak about distributed data processing and some of the people issues.

REYNOLDS: From what I have been able to gather, I assume that most of you here have jobs like mine, interfacing between management which has important things to do, and data processing people. Our job is to somehow make them come together. So you are all pretty much the same sort of person that I am.

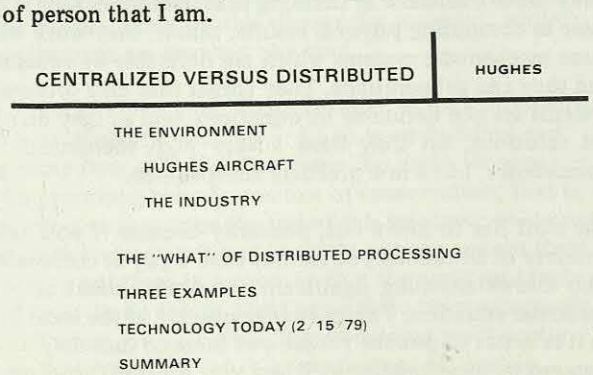

I should like to talk a little bit about the environment that I face at Hughes. I think that it is unique, but I find at coffee break that we all have some of the same problems. I have a somewhat different approach to distributed data processing than the technical one that we heard earlier. I will talk about three examples and highlight some of the people problems. I began to think that I should have been a psychiatrist when David was introducing me. I am not. I am just a worker.

Hughes has a turnover of something like $2,000 million (we call it $2 billion) in sales. We are a highly diverse product company. We make everything from little amphenol connectors to connect up cables, under water or out in the field, up to satellites, of which we make one or one and a half a year. Our major sales are to the US Government, mostly the Department of Defense. However, we have a variety of customers within the Government — NASA, all the different branches of the service, which all have different requirements.

We have about 50,000 people today. However, they are mostly concentrated in that geography that you see on the slide; 50 to 60 miles of Los Angeles contains 80% to 90% of all our people.

One of our major environmental features is the decentralised management. We are divided into six main groups which contain approximately 40 operating divisions, and we are sometimes thought of as 50,000 people in business for ourselves. We are really a very large job shop. We operate on a contract basis; we have a couple of thousand contracts in the house at any one time, although just a few of those account for the bulk of the sales.

One big advantage that we have is that our management is very technical; 75% of the top management of the company has an advanced degree in some technical speciality. Our Chairman of the Board is a PhD in Aeronautical Engineering, for example.

In 1970 I joined the company after a consultant had come in and said, "You ought to centralise computing, Hughes"; and it was pretty clear that his advice was right. Grosch's Law was just being understood and proven, day after day. It was clear that if you had a production control program that worked in one factory, you could just plonk it into the next factory, and the next one, and the next one,

DECENTRALIZATION/ CENTRALIZATION

HUGHES

1970—CENTRALIZATION

"OBVIOUS" COST SAVINGS

GROSCH'S LAW

MULTIPLE USE OF APPLICATIONS PROGRAMS

MANAGEMENT ATTENTION

DP IMAGE ENHANCEMENT

and you could save lots of money with programmers. We did not have many of them. When you talk about savings that seem rational in large numbers, like $2 million a year, it is easy to get management attention; and furthermore, the DP people love it. That is really important, we can pile ourselves up in one big heap and be four times as important.

One thing that I did not realise at the time was that all the staff in the company get their own importance amplified also. The central staffs then have the tool of the central data processing to wield their power indirectly over the people who would not pay any attention to them in the first place.

We centralised everything from 1962 to 1972. In fact in 1973 we had two 165s. This is as of next month because our 3033 is not in. But we have not quite started to decentralise yet. In fact we are still running half of the company's work on those two Amdahls in Fullerton, and twice as much work

1980—DECENTRALIZATION

HUGHES

NO OBVIOUS COST SAVINGS

SOFTWARE COST

DUPLICATION/TRIPLICATION

INEFFICIENCY

MANAGEMENT WARY

DP IMAGE DEVALUATION

FUNCTIONAL STAFF DEVALUATION

as we were running in 1971 for the whole company. So it is hard to say what we are talking about in decentralisation.

What I mean is getting the computing power out to the people, to the bakeries, out where the operating people are. It is hard to do. One reason is that the cost savings are not all that obvious. To the people who want to promote this activity it is obvious, but top management is a little nervous about the savings you will achieve with hardware and doing things twice as well as a central staff — because this is what it amounts to; top management does not buy that there are any significant cost savings yet. They know for a fact that there will be lots of duplication and that leads to inefficiency, so the top management is a little wary of this mini revolution.

Furthermore, it is not good for the DP image. I mean, if everybody can do it, it cannot be as hard as we have been telling them! The staff has had lots of tasks over the years that they do not want to give up.

I should like to mention what we were talking about — distributing. I feel that we distribute three things: the hardware, the technical capability, and the responsibility.

DISTRIBUTED PROCESSING

HUGHES

HARDWARE

TECHNICAL STAFF

RESPONSIBILITY

The thing that I really want to achieve is the distribution of the responsibility for what is done. How it is done, the hardware selection, is a technical problem; but the important problem from a company's point of view is to get the management involved, to get the job done that they need doing.

HARDWARE DECENTRALIZATION

HUGHES

MAJOR THRUST: LOCAL CONTROL
HARDWARE PRICE/PERFORMANCE

RESULT: MULTIPLICITY OF ARCHITECTURES
"STATE-OF-THE-ART" SOLUTIONS
DUPLICATION—SOFTWARE
—TRAINING
COMPETENCE LEVEL—THEORY
—PRACTICE

Let me briefly say what I feel some of the motivations are of the distribution of these three things. All our line managers would like to get their control on their own data processing. In a decentralised company where you have performance evaluation rather than staff kinds of control, the guy with the job wants the resources; he will use them and decide on the allocations and just measure him on performance. When the hardware started coming down in price, they started in a major way to ask for hardware.

When that happens, the debates are, "Well, this machine is the best possible machine for my particular application, which is different from every other application in the business. I'm all by myself on this one contract, and I can do it better and cheaper this way." Since you are measuring him on his cost/performance the management sits there and swallows, and says, "OK," and you end up with a multiplicity of architectures.

In our company we are full of technicians. They are all geniuses in their field. As a result, they know all the things that you can do with a computer. They know nothing about doing it in practice, but they know theoretically. They make a mistake in thinking that data processing is the same as computing physical results; that is, they work with these mechanistic systems which are definable by equations and they can get solutions. They forget that data processing systems are not definable by equations, and so they do not get solutions. So they have a very high theoretical competence, but a low practical competence.

The staff has to move out, primarily because if you have a variety of situations you cannot have a guy at corporate who knows anything significant, in depth, about any particular situation. Things change quickly at the local site, so it is better to put the people out there so that they can respond to local problems. When you do that, however, lots of things fall apart. John yesterday mentioned the lack of back up as a simple example of the kinds of things that are

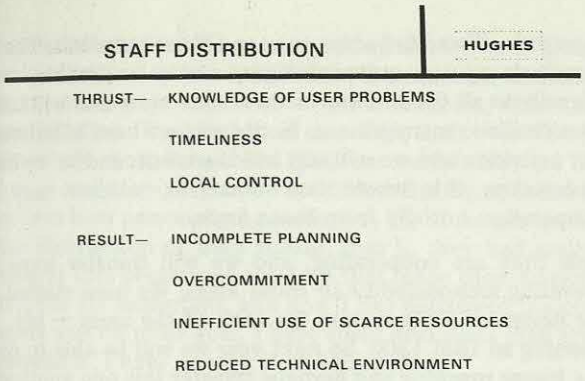

normally seen to at a high level of centralised DP management which are lost initially when you go out to a local site.

You spread your staff very thin and you have to take a fellow who maybe did accounts payable for five divisions or groups, and I have to stick him in one group, so you have four divisions that have no capability. Finally, if you are one man out of four in a DP organisation, presumably you do not have as much peer support and growth potential as you would in a central location.

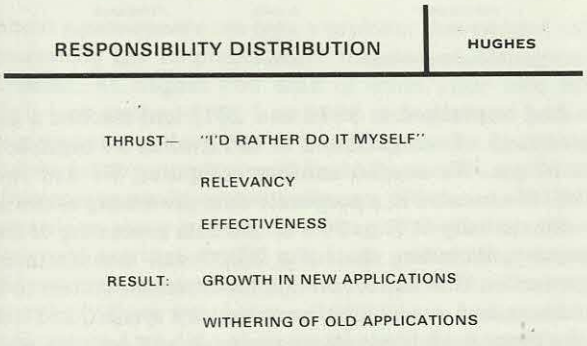

I anticipated all those problems. We have not solved them but we are evolving around them. But I had not anticipated what would happen to the responsibility. I assumed that when a man said, "Give me my own staff and give me my own hardware, I'll solve my own problems," he would in fact solve the whole problem. But it turns out that this is not really what he means. He says, "Give me my own staff and my own equipment, and you handle all the stuff that I don't know about and that has been going on, and let me do all the new things." So every time we put in a mini we get more work at the centre, because we have to do what we used to do plus analyse everything that he is doing automatically now that he used to do by hand. So the old applications are dying on the vine at the moment, while new applications are growing like mad.

I have a theory about that. When I started in data processing, which is only eight years ago — I had been building and selling computers, which is a lot more fun and a lot easier — the data processing charter was one of conservatism; that is, we are trying to maximise the use of this resource; we do not have enough trained staff and the only way we can get them trained is to put them in a room so that they will get taught the right way to do things, and they will do them effectively in the DP sense. We have a lot of systems that work pretty well on that basis but what is happening, at least in California, is that we are running out of people. We are running out of trained people.

My favourite example of that is that the telephone company hires Spanish speaking people to give out telephone information, which makes it difficult for about half the people in LA; I suppose the other half are all right. Today, the job we have to do is to get some things done now. At least in our company we are faced with tremendous growth. We grew at 4% to 5% a year for some years, and in the past couple of years we have been up in the 15% range and that is projected to go on. As a result, simple things like processing resumes to hire 10,000 people — that is 50,000 to 60,000 resumes that you must keep track of.

These people are trying to do it with file cards. Now they have word processors but they have queues behind the word processors keeping track of what goes in. So we have a very high pressure to get new things done. The classic approach up until a few years ago was, "Well, we'll re-design the system and we'll add all these new features to the system." It is two years before you ever get to the new features because you are still doing a better job of re-writing the old stuff. So again, as John pointed out, we are in the situation now where we have to find ways to add to the capability of the people. I believe that the naive users now have demands which are overwhelming and you have to do that as well as plan long term for bigger and better systems in place of the ones you now have.

Where we ended up with Burroughs, Data General, DEC, IBM and Honeywell, is that we have a lot of stuff that is unique. Almost all of those installations today, with the exception of the Honeywell, is overloaded and over-committed; and we are not quite sure about who is supposed to do what. But we are going to keep going. The reason that you have to keep going is that the pressure is worse every year to go smaller. Every time IBM cuts price, at least up until lately, the price performance of the minis has come out faster, sooner, and more often than in the maxis.

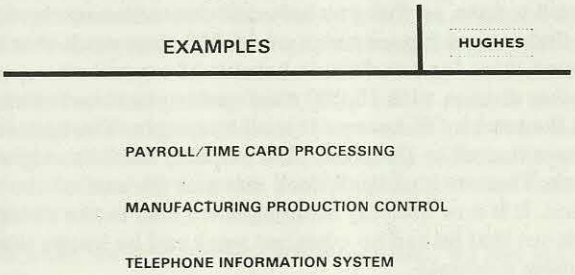

I would like to talk about three things we have done. There is payroll/time card processing which involves a micro level machine, and now a mini; a Datapoint 5500 is where we started on that. There is manufacturing production control where we distributed an IBM 158; and a telephone information system which is on a Data General Nova.

The first one illustrates to me a couple of the problems of working with a staff and with users who do not want the whole job, and also trying to get something done in a hurry. The second one was our first effort at distributing and illustrates some of the difficulty with the DP staff. The last one is my own pet attempt to find a distributed database approach that will work, and it has raised several issues.

The company grew from 10,000 or 12,000 people back around 1955 to 1960, to its present state of 45,000 to 50,000 projecting to go to 55,000 by 1982. This system

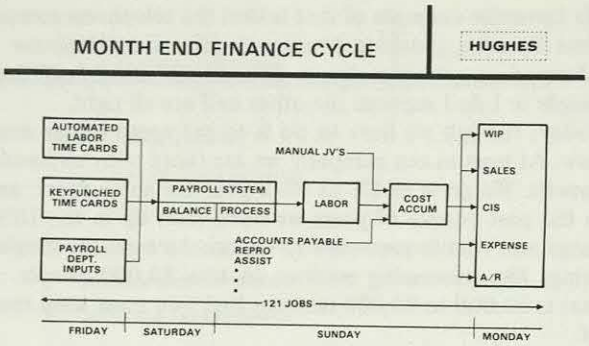

of payroll processing has persisted since the very beginning; in fact I believe that most of the systems were initially designed around 1965. The reason that it is important in this work in process is that most of our work in process inventory, over half of it is labour. We have projects, a thousand or so individual contracts; and we have to cost out those projects every week. So starting Friday afternoon, we have 140,000 transactions of time cards, labour information, to key punch, verify and balance; and then, if we are lucky, we get all the way through that and get all that labour priced by Sunday morning. If we work really hard we can get all of that cost accumulation done so that by eight o'clock Monday morning at least the basic information is known to every project manager about what was spent on his project.

It turns out that one of the things that happens to you is that you do not know everything that is going on. In the growth of the company, to save money we stopped time-keepers; so this cut off period was delayed until Saturday night. The reason that this all came about was that we missed the year end processing which we were supposed to have done on 15th January, and we did not get it done until 15th March this year. So now we have 140,000 transactions down here, dated in order by person, because that is the way payroll is done, so that you have no reasonableness check. By that we can have a batch of 15,000 time cards that is entered twice for one division because of control errors, and another division with 15,000 time cards not entered at all, and the total is OK because it is all by people. You cannot balance that all to the penny, but certainly there is no gross check. Then we sort it all back down at the end of the period. It is now Monday morning and a man in the division finds out that he had no costs last week and he knows that we made a mistake.

It is never that easy and we have had a lot of troubles. So we developed a method to distribute this input to nine minis in nine locations. We put it on the first one, and it took us longer to do than Jim said it took to install a whole bakery system — which says something about the UK versus the US, or at least California. When we got it done we thought that we would go round to these other eight controllers and say, "Here's a nifty thing to improve the accuracy of your system and the timeliness of the reporting we can give. Wouldn't you like to take on the task of doing your own payroll processing?" and the answer was, "Yes, but I've got these other things to do so we need a study." We studied five different places for nine months and they came up with nine different configurations of hardware and software that was the only thing they could possibly use in their organisation.

So we stopped and now we are putting this on to Harris RJE terminals. We added some core and some terminals. The RJE terminals are now quite intelligent and quite flexible. We will do all the running in these nine areas and will take on a facilities management. In the past we have billed out all our activities and we will still bill them out; and so we have undertaken this whole task ourselves, without any cooperation initially from those areas.

Now they are cooperating, and we will transfer some operating responsibility to those areas. We have started on the design of a big system for each of the areas — big meaning an IBM 4300. So next year we will be able to put in a bigger machine and perhaps transfer this one application to that new machine if they want to do that. But the point is that they have to be weaned into being responsible for the whole job; they will not take it gladly in most cases.

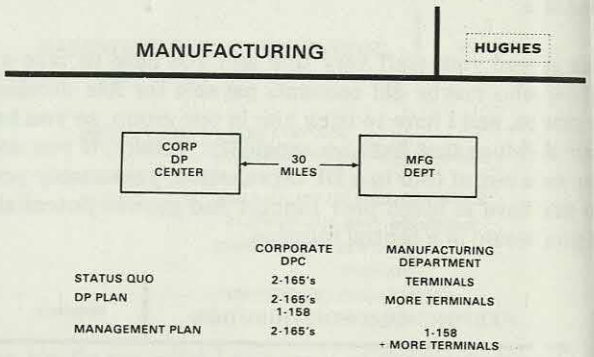

We had centralised in 1970 and 1971 and reached a high-water mark of centralisation in 1972, when we began to run out of gas. We needed another computer. We had two IBM/165s located in a corporate data processing centre and we did virtually 80% to 90% of the data processing of the company, including that of a 3,000-man manufacturing organisation that had an on-line procurement system to track purchases and availability, an inventory system, and were in the process of developing a materials and net requirements production control system. So the obvious solution from the DP department was to get a /158 and stick it into our ASP configuration. We used ASP to control these two /165s. The other big guys in Los Angeles like Rockwell had three-headed ASPs and we only had a two-headed ASP, so everybody thought that we really ought to try for a three.

The rationale behind that is that the manufacturing department would use most of the /158 during the day and half of it over the weekends to get their net requirements through, but they would use hardly any of it at night during the week, and they might even have 10% or 15% left over during the day, so we could really optimise that. So I said, "Well, that's reasonable, but gee, the manufacturing guy says the last time he had a problem he had to wait for somebody to put out a proposal and get the payroll run before he could get his production schedule. He doesn't want to do that any more, he wants to run his own show." That seemed perfectly rational to me, so I said, "That's what we're going to do. We're going to keep the two /165s here and we'll offload the manufacturing division's work on to their own machine and we'll waste 25% or 30% of a /158." That seemed a reasonable management decision to me. But I lost the head systems programmer over that. Of course, this was back in 1974, and it was an emotional blow that he would not live with; he just quit.

So we set up a group within our centre at Fullerton, and separated the work physically as well as logically inside the

one set of systems. We got something like 30 or 40 people trained on that whole application, then sent all of them over to the manufacturing division over a weekend, when the /158 was installed. They had their own IMS, their own TSO, and the whole shop. It has been quite successful; they are doing just as well or poorly as we are today. But I just could not believe how personally he took it. A lot of people took it as the destruction of their future; that is, they had really bought that centralised stuff, so that made them important, it made DP important, and now I had destroyed the goal and there was no more goal. I have spent ever since then trying to rebuild it. I have no idea how I am doing.

I have an intellectual problem with databases. Everybody tells me that they are it, that you have just got to have database management systems, and that is the future, and there is a nagging undercurrent that it is one big thing. I know that they will learn how to distribute them some day, with those back end processors, relational overlays and other things, but nobody can tell me exactly when. I also have not seen all that much program and data independence in IMS. I am sure that part of our problem is ourselves, but I am also convinced that in our shop at least there is no single database management system that will do everything that needs to be done today.

About a year ago we ran into a problem that needed solving, concerning the corporate staff Director of Management Systems. At Hughes you kind of write your own title. He is very similar to me. I am more technical and he is more management, and we report to different vice presidents; but we get along pretty well. He had a need to write letters. His manager would occasionally say, "Write a memo to all supervisors above a certain level to give them this message." It turned out that we did not know where they were; we did not even know who they were for sure. We knew who everybody there was, but we did not know today what his rank was and we did not know what his work station was; and so we had to send these letters home. That cost us $50,000 a year in postage, for just two mailings.

So we started out to build a locator system. We struggled with the problem of 'how do you get the input in?'. If you do not have an automatic input, the data is never any good; automatic in the sense that it is either to everybody's advantage or they cannot live without putting it in. So it looked as though the telephone was the solution to this problem.

The telephone operators already got little pieces of paper every time a guy moved his telephone, so that is a pretty good source of input. But we wanted to add a few things. Name; initial; telephone number is what they had.

But if we could get that form changed a little we could add his mail station, the organisation code that he worked for, and his employee number. That would be useful information.

As long as we had that we might as well put his job title and his job code. We can get a lot of this out of the existing batch personnel system.

It would also be nice to know where he goes on alternative days, what we charge for him, his pay rate, his minority code and so on . . . and pretty soon this fellow had visions of an on-line personnel database just to write letters to the supervisors.

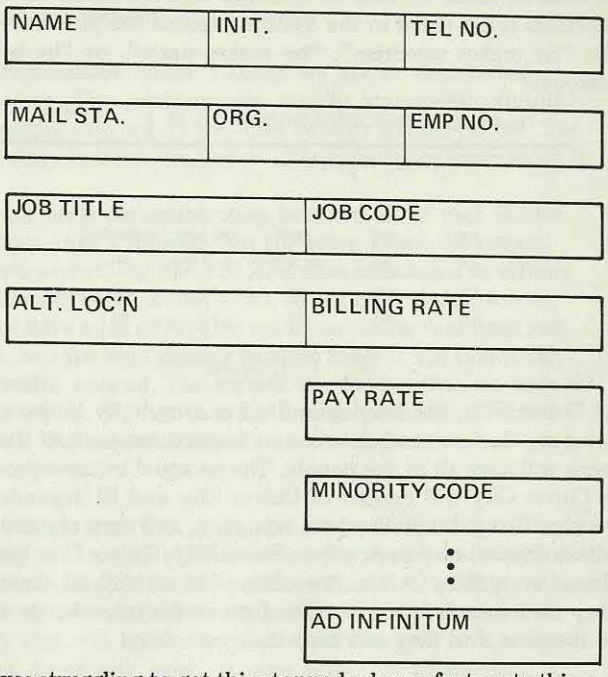

I was struggling to get this stopped when a fortunate thing happened. What we did was to put it on IMS. That is all right, but it turns out that IMS in our organisation is not reliable enough to be a truly on-line system; that is, you would not want your telephone operator not to be able to give out numbers when IMS was not available because that would be a significant fraction of every day. Not a lot, but enough to cause trouble. So you have to have a back-up system. If you have a back-up system, you might as well not have on-line inquiry because that is expensive.

So we went to IMS input and microfiche delivery. I hate to tell you, but we used to print every quarter a listing of a bunch of cards. We would hand the listings out to the telephone operators. When these little slips of paper came in, the first operator would enter the change on her book and pass it to the girl next to her. I did not find out about this until late: it is not my fault. This would cascade down through three operators, and they would give it to the telephone book operator who would see to it that it got key-punched and filed in this box of cards; and then once a quarter we would run that.

So when we first started gathering in IMS we had better data, so we started running a book once a week. Then somebody said, "Gee, we can save a lot of money. Put that on microfiche and run it every day and it's much less expensive." So we put in a cost improvement program-saving slip for the difference between a weekly batch of four copies of these for a daily copy of microfiche.

I was not getting anywhere with confining this problem; they were just going to make it bigger. They were going to add to it when they transferred communications to me and the telephone book to me away from my friend.

What we are trying to do is this. He has his big database on people. It is a batch IMS database. It is fed by a weekly cycle on a processing personnel data. We also feed it from a Data General Nova. We have five files on the Nova and we just have the first site installed. We have one file which is name and site. We have a lot of detailed data on each person. We have location data on each person. We have a lot of data on each of the organisations within the company, such

as what its name is, who its manager is, what kinds of functions it has listed in the Yellow Pages of the phone book, like "he makes satellites", "he makes parts", or "he is personnel".

TELEPHONE INFORMATION SYSTEM | HUGHES

SITE	EMP.	PERS	EMP. LOC.	CLASSIFIED
CC	ALL	CC & ES	CC & ES	ALL
F.U.	ALL	F.U.	F.U.	ALL
N.B.	ALL	N.B.	N.B.	ALL

CC CULVER CITY
ES EL SEGUNDO
FU FULLERTON

At Culver City, the employee file has everybody in the company. In fact at all the sites we expect that each of these minis will have all of the people. The personal information in Culver City will consist of Culver City and El Segundo people. That is the bulk of the company, and they are also within five miles of each other. Essentially, Culver City has almost everything; but at the other sites we will let them keep their own detailed file. The formats for interchange will be the same, but they will keep their own data.

The thing that got to me here is that everyone I put on this kept expanding the problem to fit the concept that if you have data in more than one place that is bad; and that the only way to do a good job is to put it centrally.

To illustrate, if you did want to send a letter to all supervisors, then they wanted to know in Fullerton who all the supervisors were and what their current address was inside the company as well as out, so that they could send from Fullerton a letter to every supervisor.

Quite a different approach is to say that each of these sites will keep track of their own supervisors, and in this case you write three or four letters to the site people, who are all line management, and tell the line management to write a letter to their supervisors.

Now of all the data in the personnel directory, less than 1% of it ever comes from corporate. I am excluding now balances in savings plans and that sort of thing, but as to the employee status, all of that data is entered by his department manager. The department manager is the first and only one who knows accurately what the status of that guy is. The physical effort was great. I had to get physical control over that project myself in order to keep people from designing a large database problem. I call that the large database in the sky. I do not think that will happen this year.

PEOPLE PROBLEMS | HUGHES

N. I. H.

FEAR OF CHANGE

IMPATIENCE

DATA PROCESSING

LINE

FUNCTIONAL STAFF

People are people and they are no different in this sort of operation than in any other. The things that I have run into the most . . . N.I.H. is not invented here, I am sure everybody knows that, but it happens now on all sides of us. The line manager does not believe that we know anything; we do not believe that he knows anything; and it is a barrier to progress. I have a young man who is my Number 2 guy and in a meeting we had the other day he said, "It's so frustrating to try to keep these people from making the same damned mistakes that I've taken ten years to figure out how not to make, and they just won't listen." I said, "Jim, the difference between you and me is that my youngest kid is 19 and the oldest is 30, and yours are four and 12. I know that you can't teach 'em and I also know it doesn't matter, they'll survive anyway." So we have some of those problems.

In the staff areas, at least in our company, we have had a long-established staff. We are a very dynamic business. We work on the edge of product technology. Up until TI cut the price, we made most of the world's watches. We are in the forefront of integrated technology for watches. We make parts for Amdahl. We do a lot of very high technology work. In that environment you want the administrative support not to be too innovative. There is enough trouble hanging on to the technology without having the world's first of anything in the administrative area. That has bred a staff that does not want to do anything at all — myself excepted, of course!

So that is a big problem. A guy has been filing time cards for 25 years and we said, "You're not going to have a time card any more, you're going to have a time sheet in this new system." That new payroll system works because 75% to 80% of the data this week is identical to last week's. So we have a local database saying what each man charged. We print out each Monday morning what he charged last week, and all he has to do is check each day. If he works on the same think he does not have to do anything, he only has to write in the changes and then we only have to key the changes. Those do not fit in a card filing drawer, and I thought the guy was going to come unglued about that problem. He wanted to stop the whole process because he could not file cards any more.

We have a very elaborate procedure now for filing time sheets. We have to use heavy stock, so you have to get around those things. Some of us are impatient. Mostly the line management is impatient because they are under the gun and they cannot wait forever. It is hard for me to distinguish when I am making them wait for good reasons and when I am making them wait for a bureaucratic reason. If I cannot figure it out, I know they cannot. So that is a big problem.

All of these problems affect all of the people involved. The only trouble is that they are all in different ways and you get around them in different ways. They are different ways because everybody's goals are different and everybody's perception of their problem, the company's problem, and what the other people ought to do to make their job easier is different. There is a whole bunch of chickens and eggs of different sizes.

I think it is kind of fun. It is different than programming used to be, but I have not done that in a long time so I guess it is not there. This chart indicates to me why all these problems will get worse before they get better. The horizontal axis is our estimate of the performance of the various machines shown there in MIPS, and the vertical axis is our estimate of that performance divided by their price.

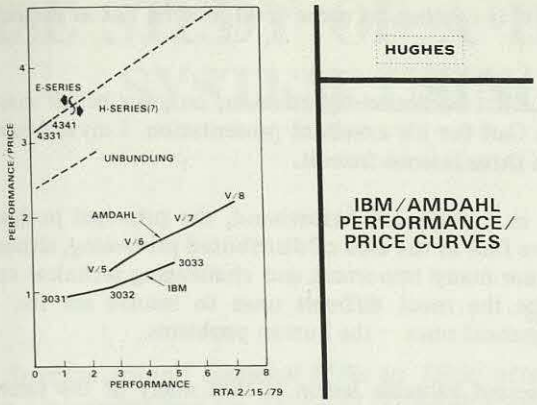

It is a relative performance and a relative price. There is some debate about how close that line is or what will happen up here, but the startling thing was the 4300 which is a little machine. Well, we older folk do not think that it is all that little, but today everybody says that it is little. That is the first time that IBM has ever destroyed Grosch's Law in public.

If you put in the price of software, we estimate that brings it down some, but it is still twice as good as the high end of the line. That says two things to me. The justification for equipment is changing. Clearly the 3000 series will get banged up there. Either the H series will be on that same price/performance curve or they will cut price again on the 3031, 3032 and 3033, and bring it back in. But it is beginning not to matter that it is cheaper to do it in a big pile, because it is now down at a cost where anybody can afford it.

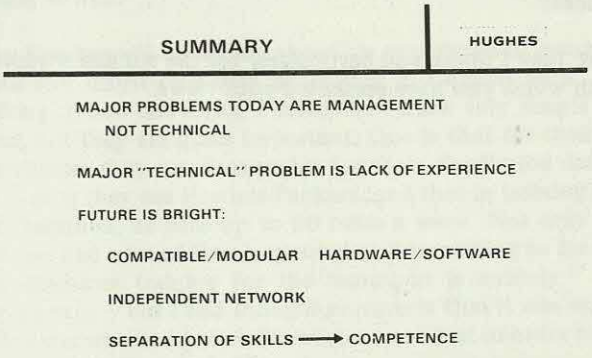

I had an interesting experience in that my wife has just gone back to work. She works for a classical computer user, a vice president of Human Resources, Industrial Relations, Personnel and whatever. It has been a revelation to me how little users know about data processing unless they happen to be somebody in Finance that has had 20 years of frustration working with data processing. They just do not know what can be done, and they do not know how to go about finding out. So I told my wife, "Call the DP department," and she did. They had a big meeting with a director of DP, her boss, and two or three other people. It was two or three hours and she came home that night and I asked, "How did it go?" and she said, "Oh, so-so." I said, "What's the matter?" and she said, "All they told me about was their problems."

We went through their budget and all they had dedicated to Human Resources was two people — what we call 'below the line'. They are not actually in the budget, but if somebody gives us some more money that is the next thing we will go to work on. If you now get the price of equipment down so that is not an obstacle and anybody's budget can accommodate it, then I think that it will happen there, and I think that it should happen there.

This summarises where I think we are in distributed processing. The problems are mostly managerial, except for the fact that we do not have enough trained people. The technology is there, we just need somebody to learn about it.

We have been communicating between all of that diverse hardware that I showed you for some years. At present we have something like 120 RJE sites connected in various versions of HASP, 2780, 3780, to the central site, and we do not have a lot of trouble communicating data back and forth. Not for load sharing because there is not that much bandwidth around. The highest bandwidth that we have is 56 kilobytes. We use that for remote printing of volume reports. We have not found out how to print microfiche remotely yet.

But I think that IBM really did an amazing thing in the pricing of their 4300. You now have an architecture that goes from, say, $100,000 purchase to several million dollars, in which the instruction set is completely compatible. I think that will have a tremendous impact on how we go about doing our work. It now offers us tremendous opportunities to move off selected things that fit and to retain things, either because they are too big, too hard, or they are too intertwined with some existing things, and it really gives us the opportunity to evolve.

We have taken the approach for some years now that the networks should be totally independent of the host computer. We were forced into that by the proliferation of protocols that everybody is using. DEC has DECNET; Burroughs has Burroughs NET. We do not have any Univac, but we have CDC. We have IBM with SNA, and without SNA. The only way out of that that I could see was to treat the network as a facility quite independent. We could not standardise on computer hardware, but I felt that if we provided a stable data network interface that had a chance of surviving two or three generations of computer software and a chance of evolving itself, that would provide one floor of compatibility that would be useful to the company.

So we are in the process of building internally — or rather TELENET is building for us — a packet data network, private, to be run by them for us to connect all those computers together.

What I see happening in the business is a further level of skill differentiation. We have more and more specialists within my organisation, and there are more and more people who are operating at a shallower and shallower depth of speciality out in the user organisations. I think that will continue. We are trying to give the tools to the naive users to build their own systems and reserve the high-class, sophisticated people for high-class, sophisticated problems.

In our particular case, I think that we will evolve a lot of naive inputs/outputs alternatives to the large, sophisticated problems that we currently process; and we will be going through a situation in which, in the next five years — or ten years it will probably take — we will not be doing any more of the classical applications that we were set up to do. I do not think that we will do payroll centrally, I think we will do payroll in the department with one of the micros that you will hear about this afternoon. I think that we will

do big database searches, unstructured searches centrally. I think that we will do hardly any of the current work centrally.

I think that it is an exciting time. There are a lot of people problems, but there always have been; and if it were not for those, I personally think that it would be very dull.

BUTLER: Thank you very much, Carl. Ladies and gentlemen, we have time for one or two questions before lunch.

QUESTION: Thank you for a most interesting and honest talk. Would you say that your experience is typical of other large American organisations like Boeing or McDonnell Douglas?

REYNOLDS: The nature of a Douglas or a Boeing is a few, very large products. Their situation then is that they are naturally concerned about a bigger thing than we ever are. So the approach to solving the problem of producing a 747 is a lot like mine, except that it is on a big scale and it dominates the hardware, software and staffing. What little information I have indicates that these fringe problems are ignored. Their style is much control over acquisition of resources and much less concern or awareness of the fringe problems, so the fringe problems are getting less treatment than they get at Hughes, but it is only a matter of time.

In one of those companies which is very tightly controlled — in fact it is my wife's — her boss went to his guy and said, "I've got to have some data processing support," and the fellow said, "Why don't you buy a mini?" which is heresy in that particular company. But the lack of responsiveness of data processing to anything but their central thing has penetrated that organisation. So I would say that if it is not there it is coming, to some level, maybe not as radical as mine.

BUTLER: Ladies and gentlemen, on your behalf may I thank Carl for his excellent presentation. I myself have drawn three lessons from it.

First, as I suspected beforehand, the principal problems that we face in the area of distributed processing, although there are many important and challenging technical ones, perhaps the most difficult ones to resolve are the management ones — the human problems.

The second valuable lesson is that many of the same problems that Carl mentioned — the N.I.H. problem and fear of change — exist no less in Europe than they do here in the USA.

The third valuable lesson that I learned is never, ever to employ Carl's wife!

I like to have an unfair advantage over our speakers. As it turned out, Lionel Green and I had lunch recently with one of Carl's customers, in a pretty visible application. We had lunch with an astronaut, Dave Scott, who drove the moon buggy on the moon. I just thought that I should like to tell Carl that one of the questions that we asked him was, "There must be many leery moments in a mission: what's the leeriest?" He said, "Well, it comes about ten seconds before blast off. At that point, you hear the countdown and you think, 'I'm sitting on top of 120,000 individually-made components, and each one of them was made by the lowest bidder!' "

Carl, may I thank you particularly for the wit and wisdom with which you have presented your views.

CASE STUDY OF A DISTRIBUTED PROCESSING SYSTEM USING MINI COMPUTERS

Gary Specker
General Mills Inc.

Gary Specker joined General Mills in 1966 after receiving his MS in Industrial Administration from Carnegie Mellen University and an AB from Grinnell College.

Mr Specker's current position is Director of Systems and Data Processing for General Mills Consumer Food activities. For several years he was active in the standards arena as a member of the ANSI committee that reviewed proposals for data processing standards. He is currently a member of the executive board of the Grocery Manufacturers of America Administrative Systems Committee.

He is actively involved in a study of the impact of 'mini' and 'micro' computer technology on large corporate data processing activities and has spoken on this subject at several national conferences during the past year.

MIKE COLIN: This morning we heard a lot about policy and people issues in the use of minis and distributed networks. There is no logical reason for me to be sitting up here, chairing this first session this afternoon, other than the fact that the organisation from which I come uses quite a lot of minis; and if we had known that it was called distributed data processing we would have been a user of that as well.

Gary Specker will be talking about his experience of using minis and distributed data processing in General Mills. Talking to him last night, I discovered three very simple facts, but they are quite important. One is that for their mainframes they use Burroughs; for their distributed data processing they use Hewlett-Packard; and that in training for the marathon, he runs up to 60 miles a week. Not only have we had a lot of Greek spouted at this meeting so far — and therefore training for the marathon is entirely appropriate — but I had thought previously that it was only IBM users who had to run 60 miles a week just in order to keep up with the latest diktacs!

It gives me great pleasure to introduce Gary Specker to talk about minis and distributed data processing.

SPECKER: Thank you for inviting me to speak to you today. It is an unusual situation. The last time I spoke on this subject to a group similar to this, I also followed Carl Reynolds. I find myself in that same unfortunate situation again. I hope that I can bring at least a little new information to this meeting.

I think that many of the experiences that we have had at General Mills are in some ways very similar to those that some of the other speakers have talked about, but perhaps we have had a little different perspective on the distributed issue than some of the other users.

Just to bring you into focus, in particular as to what type of organisation we are and our size, my activity services the consumer food group of General Mills, which is about a $1½ billion operation. General Mills itself is about a $4 billion company. Much of the rest of it is made up of a very wide-ranging assortment of consumer products — companies ranging from toys, fashion, restaurants, speciality retailing type outlets, most of which were purchased some time during the past ten years.

The rest of the company is highly decentralised in its operation. Looking to the future, the Group appears to be headed on a course that, from a management style and philosophy point of view, will lead it to be a more and more decentralised type of operation. The food group is the core of the company. It was the traditional business that was there when this diversification programme started. The operation that I had was at one time the corporate data processing organisation; but as we began to diversify it was moved in to be part of the consumer food group. It serves that group as well as continuing to serve some corporate headquarters types of functions.

Before I start on our plan and how we are doing on it, I should like to offer my definition of distributed processing. I think that everybody has tried to define that term. This is really a definition that I came to after we had already embarked on this programme, and one that grew on me as it really became clear what we were doing.

In my mind, distributed processing is going back to doing things the way we would have done them had we not gone through this somewhat temporary phenomenon of centralising data processing activities, as a result of some unfortunate economies of scale that existed in the very early stages of data processing. I think that as people look back on the history of this particular technology, they will look at

this last 15 or 20 years that we have been through in business as exactly that kind of phenomenon: that it was unfortunate and we are now going to pay a heavy price for moving the data processing resource back close to the basic business functions that it has to service, and that much of the work that we did over the past 20 years we will have to undo and will want to undo, during the next ten years.

GENERAL MILLS IS 2 YEARS INTO IMPLEMENTATION OF A LONG RANGE DISTRIBUTED PROCESSING PLAN DESIGNED TO:

1. BETTER SERVE THE CURRENT NEEDS OF OUR USERS.

2. PROVIDE A INFORMATION SYSTEM CAPABILITY WHICH ADAPTS MORE READILY TO OUR CONSTANTLY CHANGING BUSINESS ENVIRONEMENT.

3. FLATTEN DATA PROCESSING COST CURVE BY PROVIDING INCREASED FUNCTIONALITY AT LITTLE OR NO ADDITIONAL COST.

Basically, we are two years into our long-range distributed processing plan. I call it a plan somewhat advisedly; I would say that it is more of a direction. We recognised when we started that it would evolve and that there was nothing in our environment that was stable enough to indicate that we could maintain any type of clear plan in place for anything approaching a five-year period, much less even a one to two-year period of time. So it was basically a directional kind of move.

The primary objectives of this move as we saw them at that time were, first, to take an approach that would better serve the needs of our then current users of the primary operational, traditional financial, sales users of our systems. Secondly, it would provide an information systems capability that had the flexibility to adjust to what we were already seeing as a very dynamic environment, and as we looked to the future we could see it becoming even more dynamic. Finally, we felt that in our particular situation we would be able to flatten the data processing cost curve, and at the same time increase the functionality of that capacity without incurring any further cost.

IMPROVE SERVICE

1. MORE RELIABLE DP OPERATION

2. INCREASED ABILITY TO HANDLE SYSTEMS THAT PARALLEL THE "REAL WORLD" ENVIRONMENT (I.E., REAL TIME OR INTERACTIVE SYSTEMS).

3. GREATER USER CONTROL OVER OPERATING ENVIRONMENT OF THEIR APPLICATION.

In terms of improved service, the first thing we saw — and one of the primary objectives as we looked at distributed processing — was that we felt that we could move to a more reliable type of DP operation. The simplicity inherent in a distributed minicomputer network where you have systems dedicated to specific applications, we felt, was an inherently more reliable mode of operation than we were into with a large, multi-purpose mainframe environment.

Speaking to Mike's point about the difference between a minicomputer and a mainframe — as we looked at the equipment the major distinction that we made was not so much in the size and power of the equipment but rather in the evolutionary history of that equipment; mini-computers having been built from day 1 to be on-line, interactive type systems, with all of their operating systems geared to that kind of environment, the large mainframes having been built and having their operating systems initially designed primarily for batch business operations, and having layers and layers of software added to that core in order to accomplish the on-line types of applications that were becoming increasingly important in our environment.

That also speaks to the second point: we did see that, in order to service our current users, we had to move toward systems that had a more real time component that would allow us to match the response time frames of our systems to the time frames of the user activities that we were servicing. This was particularly true on our plant operations where, although not quite as critical as the bakery situation that we saw this morning, we certainly have some critical time dimensions.

Finally, it was our objective to provide and allow the users to have greater control over their own operating environment for their applications to not be part of a large central operation.

INCREASE INFORMATION SYSTEMS FLEXIBILITY

1. DEDICATE DISCRETE RESOURCE UNITS TO MAJOR BUSINESS ACTIVITIES ALLOWING USERS TO MAKE GROWTH AND COST/BENEFIT DECISIONS BASED ON TRUE INCREMENTAL COSTS.

2. PROVIDE A SYSTEM NETWORK AND STANDARD INTERFACING TECHNIQUES TO ALLOW NEW BUSINESS VENTURES TO SELECTIVELY SHARE EXISTING BUSINESS SYSTEMS OR DEVELOP NEW ONES.

3. CONTOUR THE DATA PROCESSING ENVIRONMENT TO CORRESPOND TO THE BUSINESS ORGANIZATION CHART.

Looking at flexibility, again we saw that the ability to dedicate minicomputers to specific major business activities would allow us to match these capabilities closely to the growth and needs of those particular functional activities, and that we would be making much better decisions relative to those systems than we were as we tried to run them and tried to analyse the needs and resource requirements of these applications within our traditional environment.

Secondly, we saw that in this environment we would be able, through a network and standard interfacing techniques, to allow new business ventures to either piggy-back on our existing systems in an easier way, or to allow new business ventures the option of not having to piggy-back on our existing systems but to be able to develop new ones where there really was a need to do so.

General Mills' food businesses are basically in a market that is not in total growing that fast — the population of the US is not increasing at a very high rate and growth of the food industry is basically coupled to population. So General Mills' growth strategy is very heavily oriented towards new venture activities in a wide range of new products. The latest example of this was our entry into the yoghourt market, where we have set up a new company under our food group umbrella to market yoghourt. This activity is one of the first where we have used this strategy effectively, and are implementing

a stand-alone system to support that business on a Hewlett-Packard.

Our real attempt was to contour the data processing environment to correspond more closely to the business environment that we are trying to service.

FLATTEN DATA PROCESSING COST CURVE

1. ADDITION OF PROCESSING POWER IN SMALL INCREMENTS FACILITATES USE OF LOWER COST TECHNOLOGY.

2. THE HARDWARE BEST SUITED TO DO A PARTICULAR JOB CAN BE ACQUIRED AND TUNED FOR MAXIMUM PERFORMANCE.

3. DIRECT CHARGE BACK TO USER DEPARTMENTS CAN BE MADE ON ACTUAL HARDWARE USED BASIS.

Finally, looking at the data processing costs, we felt that by working in smaller incremental units, as we were with minicomputers, we could stay a little closer to the edge of the technology in terms of price performance. Typically, from Hewlett-Packard we are dealing with a 90-day lead time on delivery of new 3000s. They have a pretty good track record over the last two years of matching and passing on the decreases both in their costs and also matching competitive price moves, so that we have seen a constant reduction in the cost of this equipment over that period of time that we were immediately able to capitalise on as we moved through our plans.

Secondly, we felt that we could match hardware to a particular job to maximise performance, and that we were not straitjacketed by very large decisions on mainframe systems , and we could tune them to particular applications; and thirdly, that we could get around some of the hidden cost problems inherent in a typical large operation where it is very difficult to analyse the costs of any particular application and that, in trying to make users feel more responsible both for the operation and the cost of their systems, this was much easier to do in a minicomputer environment.

ELEMENTS OF DISTRIBUTED PLAN

1. OFF-LOAD MAJOR, DAILY SYSTEMS WHICH HAVE A SIGNIFICANT REAL-TIME COMPONENT TO DEDICATED "MINI" COMPUTERS.

2. INSTALL MINI'S FOR INVENTORY CONTROL AT 6 PACKAGED FOOD PLANTS.

3. RETAIN LARGE MAINFRAME(S) FOR MAINTENANCE OF REFERENCE DATA AND LARGE HISTORICAL/STATISTICAL DATA BASES.

4. LINK ALL COMPUTERS INTO A NETWORK.

5. EACH INDIVIDUAL PLAN ELEMENT MUST BE JUSTIFIED USING ROI AND NORMAL APPROVAL PROCESS.

Just briefly, and I will go back to this in a little more detail later, our distributed plan was composed of these key elements:

- first, to off-load our major daily systems from our central system to dedicated minicomputers, dedicated to a specific function;
- secondly, to install minis not just for inventory control but for a total warehouse management function at our six packaged food plants;
- thirdly, to retain the large mainframes for maintenance of what we call 'reference data' — and I will get back to that later — and also our large, historical/statistical databases;
- fourthly, to link all the computers into a network;
- fifthly — and I put this in to underscore the fact that this plan was really a directional plan and that as we moved ahead to implement the phases of it we would be looking at each phase — having to justify our move to minicomputer on an ROI basis using the same process that we had used on our central systems.

ENVIRONMENT

1. LARGE CENTRAL - MACHINE ORIENTATION

2. DIVERSE REMOTE FACILITIES

   A. 6 PACKAGED FOODS MANUFACTURING PLANTS

   B. 6 RELIEF WARESHOUSES

   C. 20 PACKAGED FOODS SALES OFFICES

   D. 6 FLOUR MILLS

   E. 7 GRAIN BUYING OFFICES

To try to give a little background on the kind of environment that we were coming from and why we made some of the decisions in terms of the basic direction and strategy, we had a large central machine orientation, a very long history of that type of operation. We had a diverse number and types of remote installations: our six packaged food plants; six relief warehouses; 20 sales offices; our flour mills; and our grain operations.

As you can see from this environment, we are a little bit different for a $1½ billion operation; we are quite a bit different looking from probably a typical British operation. We are almost entirely dependent on rail for the transport of our product. We service the whole of the United States out of essentially 12 distribution sites for our grocery products, making very little use of truck transport at this point, although we do see some need to move in that direction because of problems with rail services.

ENVIRONMENT

3. EARLY USE OF REMOTE DATA ENTRY AND REPORTING

   A. INTELLIGENT DISKETTE TERMINALS AT 35 LOCATIONS

   B. DIAL-UP, BATCH TRANSMISSION

   C. MINIMAL CENTRAL DATA ENTRY

4. LONG TERM COMMITMENT TO DATA BASE MANAGEMENT

   A. HEAVILY COBOL ORIENTED

   B. SOME ON-LINE INQUIRY

   C. NO INTERACTIVE DATA BASE UPDATING

From a DP point of view, we were very early into remote data entry and reporting, using intelligent diskette terminals, Datapoint equipment, in a dial-up, batch transmission mode. We moved to a point where we had minimal central data entry operations.

We had a long term commitment to database management, primarily in a Cobol mode on our Burroughs equipment; a limited amount of on-line inquiry; and almost no interactive updating in our existing systems.

I want to qualify this database management comment, particularly after some of Carl's comments this morning. Our commitment to database management here is really not a philosophical one in terms of a large, centrally controlled, very integrated single database, but really an operational kind of a commitment to database management systems as a solid foundation for systems development. So our system's databases are typically oriented towards specific application systems; they are not oriented towards any large architecture or overall plan, with the exception of this reference data system that I will speak about later.

ENVIRONMENT

5. TREND TOWARDS GREATER USER PARTICIPATION IN DP ACTIVITIES

   A. OPERATIONAL CONTROL OF SYSTEMS MOVED TO USERS

   B. DATA CAPTURE AND ENTRY MOVED TO USERS

   C. SYSTEM DEVELOPMENT AND IMPLEMENTATION REMAINS CENTRAL

   D TWO PILOT "DISTRIBUTED" INSTALLATIONS.

Further in the environment, a key thing has been a strong trend towards greater user participation in DP activities. About two years prior to the start of this plan, as best we could we bundled up the existing batch types of systems and moved both a large number of people and the responsibility for the systems back to the various functional divisions, again with the objective of having the users feel a greater accountability for the systems that they were running.

At the same time, we changed to a direct charging procedure. I think that is a very common thing for many companies to have gone through. As part of that, data capture and entry was moved to the users. On the other hand, we had maintained systems development and implementation capability on a central basis.

Finally, we had experience with two pilot distributed installations. We installed a Burroughs 1700 at one of our plant locations, hired a DP manager, a programmer, and tried to get them to work on what the people in Minneapolis thought they should work on. That did not work out very satisfactorily. We learned a lot from the situation, getting both some insight into the kinds of control problems that exist in a distributed environment, and also from our point of view we got a very good understanding of how critical reliability was in a remote installation. We had some reliability problems with that equipment and it became very clear to us that reliability factors that were adequate on a central basis were not adequate when looking at a remote installation.

Our second pilot installation was one that was in Minneapolis and related to controlling coupon fraud activities on retail coupons that people send back, sometimes when they buy the product and sometimes when they do not. That was a much more successful operation, primarily because it was in Minneapolis and we did maintain good, central control over that operation.

INITIATION OF DISTRIBUTED PROCESSING EFFORT

1. DRIVEN BY SPECIFIC NEED (INVENTORY CONTROL)

2. EQUIPMENT EVALUATION FOR INITIAL APPLICATION PROVIDED INSIGHT INTO "MINI" CAPABILITIES

3. OVERALL DISTRIBUTED PLAN DEVELOPED BY DP.

4. GMI MANGEMENT REVIEWED AND APPROVED FIVE YEAR RESOURCE PLAN.

So looking now at how we really got into the distributed processing area, in our particular case it came about by looking at a specific need. We installed this Burroughs 1700 in one of our plants, for the purpose of beginning to develop a plant-based inventory system. When it became clear that that venture would not produce an inventory system that would be usable at all of our plant locations, we pulled back and began to look for a way that we could implement a distributed system at all plants.

We went through a very extensive evaluation of equipment to meet that particular need. It was as a result of that analysis that we developed or evolved this total distributed processing strategy. As we looked at the economics and the operating characteristics of the minicomputer equipment that we were evaluating for the plant, we began to relate that to the overall DP plans that we had for the next five years, and really to a large extent backed into the distributed plan that we are now pursuing based on that evaluation.

This plan was developed primarily by DP. We did not at that time have a great deal of user involvement. Again, it was primarily a directional statement and the review and approval on it was a top level management review and approval, not an application by application as we probably should have done it.

SUMMARY OF PLAN

1. TRANSFER SELECTED SYSTEMS FROM LARGE CENTRAL COMPUTERS TO SMALLER "MINI" MACHINES DEDICATED TO SPECIFIC APPLICATION AREAS.
   - SELECTIVE USE WHERE ECONOMIC AND PERFORMANCE CHARACTERISTICS OF MINI-COMPUTERS ARE PARTICULARLY ATTRACTIVE.
   - NOT A WHOLESALE CONVERSION.
2. DEVELOP ALL MAJOR NEW "OPERATIONAL" SYSTEMS ON APPROPRIATE MINI EQUIPMENT.
3. DECREASE THE SIZE OF THE BURROUGHS'S CENTRAL MACHINE CONFIGURATION AS WORKLOAD PERMITS.
4. CONTINUE TO EMPLOY THE NEWEST, MOST COST-EFFECTIVE MACHINES IN BOTH THE LARGE-SCALE AND MINI CLASSES.

This is another summary of the plan. This is basically in terms of presenting the plan that we followed to General Mills' management, the directional kind of statements that we made, that we were going to transfer selected systems from our large central computers to minis. Again, we were very careful to say that this was selective where the economic and performance characteristics of minicomputers were attractive; it was not a wholesale conversion.

Again, our company had been through a large conversion from one vendor to another of its main systems, five years prior to this, and that was not something that anybody wanted to go through again. We had to be very careful in laying out this plan to make it clear that we were not talking about that kind of effort. That in the future we would develop all our operational systems — by which we mean systems that support the daily operations of the company — on minis; that we would decrease the size of our Burroughs' central machine configuration as the workload permitted.

In fact we have moved from a three-CPU environment on Burroughs 6800 class machines down to two Burroughs' machines of that class. We did make a commitment to continue to use the most cost-effective equipment in both the large scale and mini classes, since again there was concern on management's part that we should not back ourselves into a corner with obsolete equipment as a result of any plans in which we engaged.

EQUIPMENT EVALUATION CRITERIA

A. PRIMARY

1.	RELIABILITY/MAINTAINABILITY	40%
2.	EXPANDABILITY	30%
3.	PRICE/PERFORMANCE	15%
4.	SOFTWARE	10%
5.	COMMUNICATIONS	5%

Let me take a minute now to talk about the criteria that were used for evaluating the minis. This was done primarily looking at the plant-based systems that we were concerned about, not the ones that would be located in Minneapolis specialised to specific functions.

Reliability and maintainability was our first consideration, again based on our experience with the 1700 in our Toledo plant. We were also very concerned about expandability. We did not want to under-buy on the equipment; we wanted to make sure that there was a large safety margin in the capability of the equipment versus what we could immediately see as the needs on that equipment.

Price/performance was of some consideration, but it was really a fairly minor item. Software we rated as a minor item, although I think it carried a much heavier weight in the final decision than we really admitted.

Finally, we rated communications lowest of all. If I were going to do this today, based on our experience to date and where we are having problems, I would say that communications capability would be the second item, right behind reliability. This is the place where we have had the most problems in accomplishing our objectives and in having a smooth transition to a distributed approach.

EQUIPMENT EVALUATION CRITERIA

B. SECONDARY

1. PRODUCT MATURITY
2. SUPPORT ORGANIZATION
3. SOFTWARE STABILITY
4. SYSTEM "FRIENDLINESS"

Also very important in our selection of Hewlett-Packard were some secondary considerations: the maturity of the product with which we were dealing. We did not want to be the first ones in with a new product hot off the press. We did want a support organisation that was as close as possible to what we were used to from a mainframe vendor. We saw in many minicomputer organisations very skeleton kinds of support organisations that we did not feel we would know how to work with or we did not have the internal resources to take up the slack. We were looking for software stability. We wanted an operating system that had been around for a while, and we were looking for a system that had some usability characteristics. We will use Hewlett-Packard's term since this is one of the things that they push on their systems: it was basically a "friendly" system for us to be using.

Looking a minute now at what our basic overall plan was and what we have accomplished to date in this plan: the circle at the top represents our 6800 capability; each of the boxes represents a Hewlett-Packard 3000 minicomputer. These are fairly large scale minicomputers; all of them are at least 512K machines. The typical plant to configuration includes a 512K machine, about eight terminals, and about

Figure A

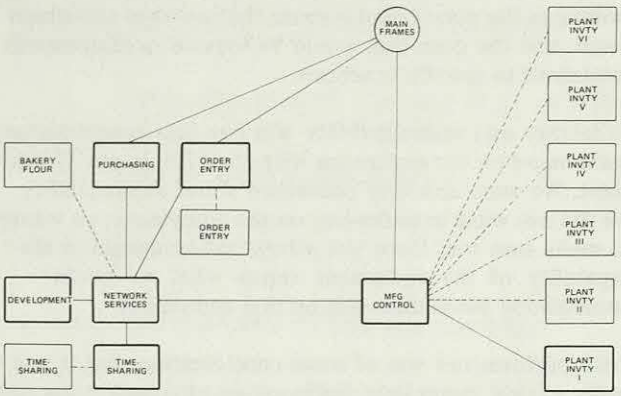

120 megabytes of disc, plus some printers and assorted other equipment.

Again, we did not try to cut it real close on the hardware. Our objective was to make sure that we had the capacity that we needed to handle the various functions that we were trying to support and to give ourselves enough headroom to cover any kinds of problems or inefficiencies that we could not anticipate at the time we went into this.

The boxes that are in dark outline are the part of the network that we currently have installed. We have one of the six plants installed. This has occurred this spring. Starting in September, we will be rolling the plants out, one every three months so that the rest of the plant side of this system will be moving ahead rapidly.

Time sharing is basically a capacity situation. We will add a second time sharing machine as soon as we need it. Order entry we are still in a pilot mode, servicing only part of the country. We will be adding an additional machine as we roll that system out to service the total country.

From a distributed processing philosophy point of view, there are three kinds of distributed machines here. There are geographically distributed systems represented by the plants. There are functionally distributed systems represented by the purchasing and order processing systems. Then there are business distributed systems represented by the bakery flour system and the yoghourt system that I indicated will be coming on line. They were not part of our original plan but are examples of systems that are distributed on the basis of specialisation to a specific sub-business in which we are involved.

SYSTEM DESIGN APPROACH

1. RETAIN COBOL DATA BASE MANAGEMENT ORIENTATION

2. MINIMIZE INTER-MACHINE DEPENDENCIES
   - ALLOW STORAGE REDUNDANCY IF NECESSARY

3. OPERATION OF APPLICATION SYSTEMS AS AUTOMATED AS POSSIBLE

   A. USER INITIATES ALL PRODUCTION RUNS

   B. COMPUTER ROOM OPERATOR DOES ONLY DUMPS, OFF-LINE REPORTS

   C. NO CARD EQUIPMENT ON MINIS

Looking at the design approach that we used in making this major change from a central environment to a distributed environment, basically we have tried to hold as much of that environment the same as it was on our Burroughs' equipment. We have stayed with Cobol. We are continuing to use database management systems as a primary tool in the development of applications; in this case it is the Hewlett-Packard system as opposed to the Burroughs system.

Secondly, we have specifically tried to minimise inter-machine dependencies to the point of allowing a significant amount of storage redundancy. Again, the cost of storage is no longer a significant ecomonic factor; the real issue in storage redundancy is a control issue, not a cost issue. We have approached these systems on that basis as we have designed the applications.

Thirdly, operation of these application systems was to be as automated as possible, with the user to initiate all production runs. Try to minimise the involvement of the computer room operator. We have not been as successful on that score as we would like. There is no card equipment involved.

SYSTEM DESIGN APPROACH

4. STRIVE FOR HIGHLY INTERACTIVE SYSTEMS

   A. ON-LINE UPDATING

   B. MINIMIZE PAPER REPORTING

   C. GET THE DATA INTO THE COMPUTER AT EARLIEST POSSIBLE POINT

Finally, as I have already pointed out, we were looking exclusively at on-line types of systems for minis, trying to move the terminals out in all cases to the actual data entry point in the business system, not necessarily to the traditional entry point in our previous systems.

MANAGEMENT STRATEGIES

1. USER CONTROL OF SYSTEM OPERATION

2. RETAIN STRICT CENTRAL CONTROL OF REFERENCE DATA

3. RETAIN CENTRAL CONTROL OF PROGRAMMING - DEVELOPMENT AND MAINTENANCE

From a management strategies point of view in terms of how we saw these systems being managed, again user control of the operation. However, strict central control of reference data. I should now tell you what reference data is. Dating way back to our very earliest database planning, which came in the late '60s and early '70s, we created a set of what were then files, and eventually evolved into a single database system of reference data that contained all of the key customer code, product code, pricing, particular promotional kinds of activities and when they are going to occur and what the specific promotion terms are, which is a very critical part of the business — this was all maintained as part of this reference data.

It was our decision, at least for the time being, to retain this central reference data system that we will control from

our operation and refresh copies out to the various minis that require this data. Again, this data, because of its nature, is needed in almost every system tht we are running, so we will be refreshing and maintaining multiple copies of this information on most of our systems.

Thirdly, again for the time being, we see the ability and the desirability of maintaining central control of our programming development activities. As with the bakery plan this morning, we will be down line loading all of the code to our plant locations; there will not be compilers on those systems. So, at least for the time being, we will be maintaining a good control over the code that is running on all of our systems.

MANAGEMENT STRATEGIES

4. WHERE POSSIBLE PHYSICALLY LOCATE MINIS IN CENTRAL MACHINE ROOM

5. CENTRAL MONITORING AND CONTROL OF FLOW OF OPERATIONAL AND REFERENCE DATA BETWEEN MACHINES

6. JOINT USER/DP BUDGETING, DEVELOPMENT PLANNING

7. HEAVY USER PARTICIPATIONG ON DEVELOPMENT PROJECT TEAMS

Further, this is a point where I have had several people tell me, "Well, if you're going to put them in the computer room, that's not distributed processing," yet we have found that physical location — particularly for the functionally specialised machines — was not a critical issue. We had tried with the coupon fraud control system that I mantioned earlier to put that machine physically out with the user, and found that that was not a particularly relevant thing to do. It is much more economical to keep them in a central equipment environment and you are still able to maintain the same level of user control and responsibility for the system that you were by having it physically located out with the user. So for the systems that are supporting Minneapolis-based functions we continue to service those out of our data centre.

We saw a need to centrally monitor and control the flow of operating and reference data between machines; that is, we saw the need to assume a network management function that was concerned not just with the operation of the network but also with the data that was flowing on the network. The other items here primarily relate to maintaining a strong user focus in the total planning and implementation process, wherever possible, to shift the primary burden and responsibility for systems development away from DP and back to the user organisations.

Like most of you, when looking at distributed processing one of the biggest issues and concerns was the control issue. As we looked at it at the time we were going into it, these were the kinds of control issues that we identified. We did not necessarily have a solution for them, but we could see that we were greatly increasing the number of access points to the data on our systems; that we had multiple site security problems; that we had this transmission management problem; that while we were not concerned about redundant data from an economics point of view, we did have to be very concerned about it from a data integrity point of view, ensuring that we could keep the various systems in sync

CONTROL IMPLICATIONS OF DISTRIBUTED PROCESSING

1. INCREASED NUMBER OF POINTS OF ACCESS TO INFORMATION

   A. INQUIRY-RETRIEVAL
   B. TRANSACTION INPUT

2. NEED FOR SITE SECURITY AT MULTIPLE INSTALLATIONS

3. MANAGEMENT OF TRANSMISSIONS BETWEEN COMPUTERS

4. SYNCHRONIZATION OF REDUNDANT DATA AT MULTIPLE SITES

5. UNIFORMITY OF CONTROL AND OPERATING PROCEDURES AT MULTIPLE SITES

in terms of the data that they had. Finally, in order for us to provide some level of central support we would have to have some uniformity of control and operating procedures at remote sites and could not let each site evolve its own control and operating procedures, otherwise we would not be able to provide them with the kind of service that we were planning.

Just to summarise this part of my presentation, I should like to go back and say where we are now, what we think we have accomplished, and what kinds of problems we see with the system as it now stands.

If we look at my three initial objectives: service, flexibility and cost, on the service side we do think that we have achieved part of our objective there in terms of user reliability, but I think that as part of this plan we have introduced another element that I do not think we anticipated would have the impact that it does — and that is the network environment.

As we tried to link these machines together into an effective communications network, we have found that the network software and the network reliability issues have, to some extent, offset — and in a couple of cases more than offset — the gains that we have made in the reliability of the individual piece of equipment and the individual application that was running on that. This continues to be the area where we are putting in most effort.

We are using the Hewlett-Packard DS 3000 networking capability to link the 3000s together. That was a new product announcement at the time that we began to install the Hewlett-Packards; that was the one area where we knew that we were using software that was not fully tested, and it has been an area where we have continued to experience some problems, although it is continuing to get better.

From a flexibility point of view, which was the second objective, we certainly achieved that objective. In moving to distributed processing and getting the users heavily involved, they began to see the flexibility and get a much better understanding of what the distributed approach could mean to them and their needs. Really the problem we have had here is trying to control this process and trying to keep ourselves focused on the intial set of applications that we had identified because, as I think you have all begun to experience, as the users get involved it is very difficult to

control the process at all and there is an immediate awareness level increase as to the abilities that they have.

We also found ourselves within the DP operation getting a better feeling for what could be done in a multiple machine environment. We have also branched out from our initial set of objectives. We have now, for the first time, installed an IBM mainframe in our data centre — something that we probably would never have done while we were still in a single vendor environment; but having had a relatively good experience with the Hewlett-Packards we had an opportunity to get some packaged software that would meet the needs of one of our users very well, but it was only available on an IBM system. Because of our experience we were able to move ahead and make that decision to base that decision not on hardware considerations, but on which software would best meet the user's needs.

That has been a very successful implementation. We are putting in a totally new general ledger system for the General Mills' food group, and that operation is scheduled to go into production this month. That project is on schedule; it is on cost. It has been one of the smoothest implementations of any project that we have ever accomplished; and we really did not start working on it until last November/December.

Cost. Again, we are meeting our objectives in the cost area. We have had to beef up the configuration of some of the systems more than we had originally anticipated, but we have been baled out from a cost point of view by IBM's aggressive pricing strategy over the past couple of years and by Burroughs' continued aggressive response to that pricing strategy. So our total hardware costs for the past two years have stayed flat as we have gone through this development process.

So far as the basic applications that we are involved in are concerned, again we have had relatively good success with the exception of these networking problems, and with the exception of our purchasing system which has been an absolute disaster from a project point of view. It is way over schedule; it is way over cost; the user is not happy; my project team is not happy.

I guess this all goes to show that even in the best of all possible worlds you can still have a project that simply does not come through as you would like it to; that you still need the same kind of leadership and planning going into a project that you needed in a central basis and, if it is not there, we still run into the same kinds of problems that we have traditionally experienced. But the other activities, the order processing and manufacturing projects, have been and continue to be very successful.

Time sharing is another area that has been very successful. I should like to go through briefly what we are doing on time sharing because I think that as all of you move into a distributed environment, one of the types of capability that becomes very easy to implement is a time sharing or what we call an ad hoc systems capability. It is an approach to time sharing that provides both you and the users with a great deal of experience, a very cheap kind of experience as to the type of data processing environment that they will be operating in increasingly over the next five to ten years.

Ad hoc is defined as "for a specific case or situation"; basically throwaway kinds of systems, or traditional time

AD HOC "FOR A SPECIFIC, CASE, OR SITUATION"

- DEDICATED "FAST RESPONSE" STAFF
- DEDICATED TIMESHARING SYSTEM
- SPECIALIZED SOFTWARE

sharing type systems where the life cycle and the stability of the system is very uncertain. Essentially we have created a separate business within my department to run this operation on a zero cost basis and are allowed to go out and compete for time sharing business against the outside time sharing vendors who are still allowed to come in and do their thing.

Basically we have a dedicated staff, we have dedicated a Hewlett-Packard 3000 to time sharing, and we have a fairly large portfolio of specialised software.

USES

- DECISION SUPPORT SYSTEMS
- WORKING TOOLS FOR SOPHISTICATED USERS - MARKETING RESEARCH, ENGINEERING, FINANCIAL ANALYSIS,
- BREADBOARDING NEW SYSTEMS
- ONE-TIME SPECIALS
- USER PROGRAMMED SYSTEMS

The uses that we are seeing for this system are decision support systems. Again, I think that decision support systems are notoriously unstable and difficult to focus in on the design requirements. We find that this is a much better environment to try to develop those kinds of systems than part of a large, sophisticated application system on a mainframe where we try to piggy-back a decision system on top of an existing application. By isolating it off on a separate piece of hardware, we seem to be able to focus it much more quickly on key issues.

It is a working tool for what I call 'sophisticated' users; everyone else has been calling them 'naive' users. I think that they are naive from a data processing point of view; that is, they do not worry about controls, back-up and disasters, they just want to get their job done and they are willing to take some risks and some responsibility to do that job.

I guess that is why I referred to them as sophisticated users. They are not looking for somebody to hold their hand, they just want to be given some capability, and that is what we are trying to do here. We try to make sure they understand what they are doing and that if they design a system that they will want to live with for a couple of years, there will

be some problems and additional costs to transfer that system from an ad hoc basis to an ongoing operational system.

In fact we are drafting now a document that we are going to make these users sign, which is essentially a limited warranty, so that they understand exactly what they are getting into and also so that we can keep the auditors off our backs.

The auditors are not particularly happy with our providing this kind of service. We have not been able to convince them that if we do not provide it it will be provided in other ways that are totally out of their control. I think through at least having this limited warranty approach we will be able to shift the focus of control back to the user and out of the data processing department. That is, again, the long-range direction.

I think that auditors have now got the feeling that they can control the key systems of a company by developing a very good understanding of data processing and by keeping good control over the data processing organisation. I think that they will suddenly find out that the horse is no longer in that barn and we will have to go through another whole period of the auditors coming up with a new approach, which is really the old approach of going back out to the user departments and to develop adequate control systems at the user department level.

We use it to breadboard new systems; to do one-time special items; and to the extent that we have knowledgeable users, to do user programmed systems.

ADVANTAGES OVER OUTSIDE SERVICES

- LOWER COST
- EASIER ACCESS TO GMI DATA
- GREATER SECURITY
- SIMPLIER MIGRATION OF APPLICATIONS TO FULL PRODUCTION
- BROADER KNOWLEDGE BASE FOR INFORMATION SYSTEM PLANNING

Basically, the reasons that we brought this type of capability in house are that it clearly did have a lower cost in our analysis. It provided easier access to GMI data, even if we are doing some fairly crude kinds of transfer of data from our central databases over to the time sharing system. It provided greater security than outside services, and simpler migration of applications to full production. If a system does become a stable operation, it is a little easier for us to move it over and take full production responsibility if it has been done in house.

Finally, from our point of view, it gives us a much better feel for where users are or what their needs are, which we would not get if they were going outside to a time sharing service during the prototype stages of a system and coming to us later in the game.

Software capabilities that we have installed on this system

SOFTWARE CAPABILITIES

- DATA MGMT SYSTEM
- QUERY LANGUAGE(S)
- REPORT GENERATOR
- FINANCIAL MODELING LANGUAGE
- STATISTICAL PACKAGE
- GRAPHICS
- DELIVERIES DATA BASE

include a data management system, several query languages, one of which is specifically oriented to our marketing intelligence system. It is some proprietary software that we operate. There is a report generator that is again tied to this marketing intelligence system; generalised financial modelling language; statistical package; graphics; and a fairly small but very important deliveries database, coupled with some of the marketing intelligence databases that we purchase from the outside. On this system we have a large percentage of the data that many of our key decision makers, particularly marketing decision makers, need in terms of the kinds of applications with which they are involved.

SUMMARY

THE INFORMATION SYSTEMS SOLUTION PORTFOLIO WILL INCREASE WITH OR WITHOUT THE PARTICIPATION OF THE TRADITIONAL DP/MIS ACTIVITY.

DP/MIS'S PRIMARY OBJECTIVE IN THE EARLY 80'S SHOULD BE TO MINIMIZE THE COST OF THS CHANGE TO THE ORGANIZATION.

THE BEST STRATEGY FOR ACCOMPLISHING THIS OBJECTIVE IS TO AGGRESSIVELY LEAD IN THE RIGHT DIRECTION.

Just to summarise both the ad hoc activity and our whole experience with distributed processing, it has been our observation over the last two years as we have gone through this process that what we got into here we thought was distributed processing.

We currently have a centralised DP operation and we have made a specific decision to distribute that processing load out either geographically or functionally. What we are clearly seeing — as we go through this process and as we see the impact of the technology on our corporation — is that the illusion that we are in control of that process, that we are making the decisions to distribute that workload out to a variety of minicomputers and ultimately microcomputers and whatever else, is really an illusion. This process is going to take place. It will take place whether we participate in it or not.

The total portfolio of solutions that is available for solving information systems problems in a corporation will increase very dramatically over these next few years. It will increase whether we participate in that process or not. The companies that are still sitting back, trying to figure out how to control

this process before they begin to implement it, I am afraid will wait too long and not have any impact at all on the direction that distributing of processing will take in their companies.

I think that the primary objective of any DP shop in a large company over the next five years has to be to look at strategies for that particular corporation that will minimise the cost of that change to the organisation, or perhaps in a more positive way even to maximise the potential benefit of that change and the potential benefit of this technology to the corporation.

I think that the only way that a DP shop can really be effective over this period of time will be to provide leadership; not to try to provide control and not to try to provide the total direction and total work and try to keep it all as part of their operation.

The only thing that you can do is to identify the key strategy, try to see what the needs and the specific requirements of your business are; and then to make sure that you have the type of staff and overall resources that you need to provide the leadership.

COLIN: Gary, thank you very much. On behalf of all of us here today, I should like to thank you for talking to us about your experiences. I am sure that you are really more in control than you pretend you are.

For me, the highlight was not so much the fact that I have come across an on-line yoghourt system for the first time, it is the fact that — looking round — you struck so much accord with the audience here today. Obviously people have similar problems and they want to face them in the same way that you have done. Thank you very much.

CASE STUDY OF A DISTRIBUTED PROCESSING SYSTEM USING MICRO COMPUTERS

John Jones
Southern Railway System

From 1951 to 1957 Mr Jones was in the United States Air Force and served with the USAF Comptroller in Data Processing and assisted in assembly and check-out of the first three UNIVAC 1s. He was in charge of the Engineering Division Computer Center of the Chrysler Corporation from 1957 to 1959, and concurrently he was a Management Consultant in Data Processing to the Air Force Logistics Command in 1958 and 1959. In 1959 he became a full-time civilian employee with the Air Force Logistics Command, responsible Command-wide for programming systems and standards, EDP equipment evaluation and selection, management of installed equipment and data systems research.

In 1963 he became Assistant Vice President of the Southern Railway Company with responsibility for all corporate data processing activities. In 1969 Mr Jones was appointed Vice President of the newly established Management Information Services department responsible for all corporate systems and data processing activities, including operations research and industrial engineering.

Mr Jones is Chairman of the Executive Committee of CODASYL and a member of the General Committee of the Data Systems Division of the Association of American Railroads.

COX: Gentlemen, we continue looking at the experience of others in this area. It gives me particular pleasure to welcome our next speaker. I first heard Jack Jones speak and met him some eight or nine years ago. I was struck not only by the kind of systems that he was describing and how advanced and successful they were at the time, but with a real sense that here was a very hard-nosed, practical practitioner in an area then largely peopled by technical enthusiasts. I was therefore very pleased, when I saw the agenda being compiled for this conference, to find that Jack Jones had agreed to come and give us the benefit of his experience as it relates to today.

JONES: Thank you very much. I should start by saying that I thought Mr. Specker's talk was an extremely interesting one. I am going to dip you in a very different flavour of system. His talk gave me two distinct mixed emotions. He talked about their use of Cobol and their continuing use of Cobol, and as chairman of the organisation known as CODASYL — which sponsors the work of Cobol — it gave me a warm glow in my heart. Then he also talked about the fact that, while they ship very heavily by rail, they are now starting to use some trucks, which gave me a pain in my pocket book! So I had those mixed emotions about his talk.

I would like to mention that next week, on Monday and Tuesday, CODASYL is having its 20th anniversary meeting in Washington, D.C. I mention this just in case anyone is interested in that meeting. It will be an interesting two-day programme. We will have some very well-known speakers there, such as the Controller General of the United States and the Secretary of the Navy; John Cullinane of Cullinane, who markets IDMS. Tom Neece of Cincom Systems, who markets TOTAL, which is a non-CODASYL database spec., will talk. We will have progress reports, and try to get some response from people as to the work that CODASYL might do that would be interesting in the future. I mention that as a matter of passing interest.

To give you some perspective on Southern Railway Company; I suspect everybody knows what a railroad is, but to make it somewhat more identifiable for you we are about the same physical size as British Rail — our one company. We have many differences, however, with British Rail. One of them is that we are a privately owned company as opposed to a nationalised company.

We are really in the freight business, not a mixture of freight and passenger business. As a matter of fact, passenger is less than one third of 1% of our gross. Our gross last year was about $1.4 billion. One other way in which we are very different from British Rail is that our net income last year was about $137 million. We also paid about $125 to $130 million in federal, state and local taxes, which is usually not the case in a nationalised organisation. We are by no means the biggest railroad in the United States, by no means the smallest; we are ranked seventh, eighth, ninth, somewhere in that range.

I am going to talk to you first about some of the basic principles that we have tried to apply in our use of mini and micro processors. The theme of that part of the talk is that "the more things change the more they are the same".

Secondly, I will do a quick case study look at what we are doing in Southern Railway Company. I want to emphasise that I do this only to try to give you an example of what one company is doing. I would not begin to pretend that the way that we are approaching the problem, the techniques that we are using, or anything else is applicable to any one other given situation than that of Southern Railway Company. But I hope that by annunicating some of the principles that we have tried to follow and describing some

of the things that we have done, trying to talk a bit about some of our problems and some of our successes, and what we think the bottom line of all this will be, will be interesting and helpful to you. But it will be left as an exercise for the student to draw those lessons which are appropriate.

I was a little concerned about the title of my talk when I realised that it said "using micro computers" because, while we do use a large number of the micro processors, it is not a system only of micro processors; it is a mixture of central processors, minicomputers and micro processors. What it is will become clear as we go along. I will try to emphasise the micro processor role in this. But today's micros are tomorrow's minis, or next week's maxis. A name is only a fleeting thing, particularly when it comes to trying to describe the size of a computer.

I make a great distinction between the idea of a distributed system and a decentralised system. To me, a decentralised system implies some form of local autonomy or authority, and in the case of the network that I am going to talk about, that is not true with the Southern Railway Company. In our case, the distributed processing puts the processing capability and the storage of local data out locally, but there is no local autonomy in terms of what is done out there. As a matter of fact, all of the computers out there run the same programs.

To put it even stronger, the reason we are installing this network is not to give the field flexibility, but to take it away; so that all the railroad yards will operate on a similar basis — hopefully a similar efficient basis. The old Mom and Pop grocery store syndrome which exists in some rail yards will disappear.

The railroad is a very interesting business in that it is a wonderful example of how you can sub-optimise yourself into bankruptcy. Every element of the railroad must work as a part of the whole and not necessarily optimise its own local operation. It is interesting because that is difficult to explain sometimes, because all of us are under certain pressures as to cost control and other measures, and then when it turns out that some of the things that we have to do really do not benefit us, but benefit somebody down the track 50 miles, or 100 miles, or even on another railroad, that is sometimes hard to swallow if it costs us anything more to do that.

The final point that I want to make is something with which I know you are very familiar, but I just want to put it in simple terms: just because the minis and micros are cheap and small, they aint simple. They provide us with all the opportunities to make the same mistakes that we made in the 1960s with what we now call the big machines. Just because they are small and cheap, they aint simple.

Let us talk about some of these principles that I said I would try to list. I want to point out again that, first, as with any set of basic ideas that you try to pattern yourself with, they are not necessarily mutually exclusive. Sometimes you have to pick and choose in a given situation as to which principles you are really going to follow, which ones on which you are going to compromise a little bit, and which ones you are not going to pay any attention to in this case.

Secondly, the order in which I talk about them is of no particular significance. One of the things that is not changed is the importance of having top management understanding and participation in what you are doing with these minis and micros. I am sure that that is no different message than you have been getting from the other speakers, but it certainly is true. There is a temptation for it to happen because these machines are so small and cheap. So many times a division manager is within his dollar discretionary authority to buy one of these things, but there are some real opportunities for problems.

Another principle that has not really changed from the big machines to the small machines is that of common sense. I always say, "If you don't make common sense, you don't make any sense." There is no more magic in the minis and the micros than there was in the big machines. There is nothing magic about them. The systems that you design and implement must make common sense, just as they have to make common sense on the big machines.

There are no pat solutions. We are great in this business for having pat solutions. I am so long in the tooth in this field that I can remember back to the days, in the mid-50s, when we thought the assembly language systems would be the solution. Then it was input/output control systems. Then it was operating systems. Then it was databases. Now it is distributed processing. Everybody is doing it. It is the answer for all problems. And nobody knows what it means.

There are no pat solutions. One must design and tailor to the situation in which you find yourself. Your organisation has a certain managerial style. My approach probably would be a disaster in General Mills. General Mills' approach would not work in Southern Railway Company, because we are structured so differently.

I emphasise that because it is so easy sometimes to feel that there are pat solutions. You can go out and buy one. You can do something that somebody else has done and it will work for you. It is not true. Management style is different; capabilities are different; techniques are different; your people are different; and your problems are different.

Just as I always have, I believe that the straightforward approach is the best approach. Give me a good old simple solution any time, as opposed to a sophisticated, fancy, tricky, or maybe aesthetically pleasing from a technical point of view, solution. Straightforward is better.

Modular is better. I believe in straightforward modules. I have never been an advocate of the total, optimum, integrated, everything-for-everybody kind of system. I very much believe that you are better off to start out with modules, to attack the basic problem, to leave the fancier things for later; get the basic system in and working in a very straightforward way and build from there. The problem is, of course, that sometimes you have to go back and unbolt one of those modules; you have got to re-bend it a little and bolt it back down. The nice thing about it is that you can do it; often in a very integrated or interleaved system it is impossible because you have something to fix here, and when you do that something up here changes. After all these years, these principles in my mind have not really changed.

I very much believe that whenever possible you ought to use off-the-shelf hardware and software. I will give you an example of a case where we did not do that, but we knew fairly well what the risks were. We certainly did not know

what the bottom line benefit was going to be. It was a very large project which came in on schedule and under budget, which was wonderful for those of us whose career pattern was involved in it. But you must recognise the problem and manage the risks involved.

I still believe, as I have for many years, that the earliest Christians get the hungriest lions. There is just no sense when you are working on an information system, or a business system which aims at the thrust of your corporate organisational function, to take unnecessary risks whenever you do not need to.

Just like with the big machines, you certainly do need standards and conventions and higher level languages to the greatest extent. Every once in a while, one of my young technicians will come in and say, "You simply don't understand. These machines are too small to have a lot of standards and conventions or to use a higher level language," or, "They're too slow," or, "They're too cheap," or anything so that they can go out there and fiddle with the bits like they used to do many years ago. What they do not realise is that there is almost no machine that you can buy nowadays which is smaller than the old Univac I that I helped to assemble in the early 1950s. Even a hand calculator can calculate a square root faster than the Univac I could. So having grown up there with some compilers, their arguments do not get very far with me. We think we do have some very good standards and conventions, and we do use Cobol in programming our machines.

We very much believe in a pilot approach to the design and implementation of these kinds of systems. You will not hear me talking much about a plan, because in general we do not have one. That does not mean that we do not think a little bit about where we are going, but we do not have any formalised or structuralised plan. We tend to identify areas where it seems clear to us that there is something to be done which is worthwhile as far as the railroad is concerned, and we attack that as a pilot. We do not put it anywhere else until we have solved that problem.

I will step through a specific case here to demonstrate how we go about doing that. That is not necessarily the way for everybody to do it, maybe we have so many things to do in the railroad that we can clearly see what they are; but in any event, that is the thing.

The other thing that I have learned over the years, and particularly on the railroad, which is just as true with the minis and the micros as with the big machines, is that everything you plan to do, you had better have a retreat position; because no matter how small or insignificant it seems, how easy it is going to be, how impossible that this one little instruction change could cause a bug, it is just as true today on the minis and the micros as it was on the big machines.

There is one added complication as far as I am concerned, and that is we are using the minis and the micros out in the field, helping to run the railroad. Now they are not out there doing accounting work, doing anything else but helping run the railroad. When there is a problem out there with a machine, or a problem with a program, the first thing that happens if it is in the daytime, or even in the middle of the night, is that my telephone rings. I have learned that when you pick up a telephone you always ought to start with it well away from your ear, because if it is the president or the chief operating officer he is probably screaming! So the minis and the micros, at least in our situation, have all the opportunities to give you a bad day that the big machines do.

The railroad problem is really a very simple one. We have freight cars which are either empty and need to be taken somewhere where they can get a load, or they are loaded and they have come from somewhere; we have got to know where and we have to track them to make sure that they are moving correctly, we have got to know where they are going, what is in them, and how much they weigh; any special instructions. If it is a load of cannon balls you handle them one way; if it is a load of coal you do it another.

So cars and all the information about those cars and their characteristics — and there are hundreds of different kinds of freight cars. You cannot take a car that has just hauled a load of raw hides to a tannery to a tobacco manufacturer and have him load cigarettes in it — it just does not work, so there are all kinds of data that needs to be captured about the cars and the movements of those cars and the contents and characteristics of those contents.

The second thing that we have to deal with is trains. Everybody knows what a train is: it is an engine pulling a bunch of cars. That sounds like a right simple thing to do, just hook all those cars together. But it is quite a bit more complicated than that, particularly when you get into the aspects of dispatching.

We are using minis to computerise some of the dispatching functions, help the dispatchers decide how to meet and pass trains. In the US this is quite a problem, because we routinely run freight trains that are 150 to 200 freight cars long and average 50 to 60 feet in length, so a freight train on the US railroad can easily be a mile and a half or two miles long. There are certain little problems in handling those trains and dispatching them, and knowing precisely at all times where the train is, what the status is, any delays he has had, and so on.

A third piece of our problem is the railroad yards. A railroad yard is a place where the cars sit when they are not on the train. The purpose of the yard is very simple. When the train comes in, the cars are in the train in some specific sequence. In general, when a train goes into the yard, those cars need to be re-sorted into some other sequence so that the cars that are going to the same place are together. That is so that when the train gets to this future place, he can stop or set off or take off all the cars that belong there, he might pick up some other cars and go on, but to minimise the time of the train going over the road.

So those three things are the basic things we have to work with. Unlike General Mills, which has yoghourt companies and all the other various entities, we are a one-product company. We manufacture transportation. It is a highly perishable product, because either the car is on the train or it is not. If it is not, that transportation product which we manufactured, nobody uses, and it is gone forever. We have got to manufacture some more for the next time that car is going to move.

The basic cycle of a freight car is that if it is empty it gets moved to a customer. The customer puts a load in it and the car is brought back to the yard. The yard classifies it or

switches it into the proper sequence. It is put on a train, moved to some other yard or yards where it is finally switched to an industry movement, out to a different customer, unloaded, brought back empty, and the cycle starts over again. That is all there is to the railroad business.

Unfortunately, with about 75,000 cars on our line at any one time, 235 or 240 places where cars can be exchanged physically with other railroad companies — we have a lot of places where the cars can get off the railroad, but only those 235 or 240 where we like them to get off the railroad — our problem really is that all this activity is going on in the 13 south eastern states in the United States, at hundreds of locations, and all those freight cars, by some plan or some action, have got to be moved somewhere. It is not like an airline passenger. An airline reservation system is a very different kind of problem. They have tremendous volume problems, which is not particularly our problem, but at least the passenger walks up to the desks and presents himself to be boarded. The freight car is happy to sit there as long as nobody bothers him.

So our problem is to capture the data, get it in and get it processed in time to make it available for somebody in the field to make a decision. I am not talking about the president, or the vice presidents, or the general managers; all of us kind of guys may feel as though we make decisions all day long, but actually the guy who is making you or breaking you and making decisions is standing right out there in the ground, deciding, "Should I do this? Should I do that? Should I hold for this train? Should I get that car now? Should I get it later?" Those are the people who need information.

We call that a real time system, and it is. It is not real time like a system that is controlling a missile shot, where everything is happening in thousandths of a second; but in the case of a missile, in a few thousandths of a second it has gone a long way. In the case of a railroad car, in a few minutes or even a few hours, it maybe has not gone anywhere. So your real time is different. It is real time in the sense that it is decision making information in time to make decisions.

When we think of real time we often think of it as being instantaneous. We have instantaneous responses and so on, but our timing problem is not quite the same as some other, more popular real time types of systems.

We have had a system running in Atlanta since June 1965, keeping track of the cars and trains. That was a very highly centralised system. In fact not only did it run all on central computers, but it was very centralised in the entire data collection process. It is a unique and interesting story in itself, but that is not the story that I need to tell you.

The accuracy, the completeness and the timeliness of that data was very "good" — and I use the word good in quotes because it was good for the time and the use that was made of the data in the late 1960s and early 1970s.

In late 1971, we began studying the design of a new automatic classification yard to be built in north west Alabama. In that design we concluded that we wanted to use what now would certainly be called distributed processing. I am not sure that those words were used at the time. What we are doing at that yard in today's knowledge and technology would not be considered to be very outstanding. If you can put yourself back to what you knew about minicomputers in the early 1970s, you might realise that in fact it was a fairly interesting decision and approach at that time.

We decided to build a railroad yard using a network of five minicomputers in the railroad yard, but linked by communication to the Atlanta computer. In the Atlanta computer we have what we call the 'consist' or the list of every train that moves. We know precisely what locomotives are on that train, what cars are on that train, the sequence of those cars on that train, and everything about those cars.

The concept of this yard is very simple, that is before the train gets to the yard the Atlanta computer transmits to the minicomputer out there what the consist of that train is. When the train comes into the yard, somebody sits there and watches closed circuit television on a CRT to make sure that everything is in step — because sometimes they are not — sometimes there is a ghost. A ghost is a car that the computer thinks is there, and when he looks at the train it is not there. Sometimes there is a stranger; he looks at the train and there is a car and the computer does not know about it. So he has to get rid of these ghosts and strangers. But by and large, he sits there — he has the exalted title of Inbound Clerk because he watches the inbound movements — and watches closed circuit television and the CRT and the computer.

Assuming that things are straight, or he makes them straight, the next process is to classify those cars. To do that, in this particular yard we have very cleverly scraped up a little hill or a hump — we call it a hump yard. We go out there and put a locomotive at the end of that string of cars and push them over that hill. As one might imagine, as those cars go over the hill, we use the considerable help of gravity, and uncouple every car just as it goes over the crest of the hill. The cars are coupled together and, in order to uncouple them, there is a little pin that goes through the coupler and that has to be pulled by a lever on the side of the car.

As you will detect, we are very clever with titles, and the title of the fellow who does that job is Pin Puller. Every car that goes over, he pulls that pin and the car rolls down the hill. Then through radar and the speed of the wind, speed detectors and scales and so on, we control the speed of that car and classify or switch it into the proper track. It is a fairly complicated process, requiring about 600 inputs and 300 outputs to the physical yard. There are certain requirements as to the speed that those cars couple together and how far they have to roll.

For example, little things mean a lot in such a situation. The car is rolling down the track and we try to have it coupled to the car that is sitting there. It might have to roll a couple of thousand feet or a couple of hundred feet, but we try to control the speed of the car so that it couples at less than four miles an hour. It might not seem too much of a difference if it couples at six miles an hour, but if you think about the efficiency of that steel wheel, rolling on that steel or rail with 150 tons, and the irresistible force crashing into an immoveable object, at six miles an hour versus four miles an hour the potential damage is considerable.

So it is a fairly complicated process control problem, even including computer control of the locomotive that is out there at the end of those cars, pushing them over the hill at a very specific speed.

The yard also has a very specific inventory control problem. The cars come into a long, thin yard that can hold the entire train. We call it the receiving yard for the simple reason that that is where we receive the train. We inspect it; check brakes and so on; switch it and it rolls over the hump into a yard which we call the 'class' yard because that is where we classify, or more properly re-classify, the cars.

Finally, we go to the other end of those classification tracks. We go in there with a crew that we call our 'pull-back' crew because they go in and pull back the re-sequenced cars and shove them into the forwarding yard, and assemble a new train where we forward the train on to the next destination.

That whole process at this new yard is done entirely under the control of the minicomputers, using the data received from Atlanta, with some additional local information which is constant at the yard. All the people are doing in that yard is watching closed circuit television and CRT and making sure that the world is in step with what the computer thinks is going on.

There is no data input at that yard, other than to correct for ghosts and strangers. The tracking of the cars, everything about it is automatic. When the train goes out, the yard master keys into his CRT and says, "Train Number So-and-So departing at such-and-such a time," and the minicomputer, having now completed its job there at this yard, transmits back to Atlanta the consist of the outbound train, which is the first time that the Atlanta computer has heard about any of these cars since they went into the yard.

I have a little schematic that I will show you to make sure you have the picture, but I go through that to point out a couple of things. First, the process at the yard itself is a distributed process in that there are three minicomputers on line and two in hot standby at all times, as opposed to a single large computer doing the local process.

But the second aspect, which is even more interesting, is that the Atlanta computer which has a bunch of information about the cars which are unimportant to this yard — where they came from and only the next place they are going — they do not know what is in the car and they could not care less out there. It transmits to the yard only that information that it needs to properly handle that car. I neglected to tell you that our central processors are in Atlanta, Georgia. However, once the car is in that yard, other than knowing when it got there and what train it got there on, the Atlanta computer could not care less whether that train went into track 3, track 7, whether it has been switched, whether it has not been switched, or anything about it. It does not care that they are putting Birmingham cars in track 17 today, and when that gets full they will turn them to track 23. All that stuff is local information that is terribly interesting, terribly concerning to the local people. The central processor could not care less.

Once the train goes out of that local yard, the local minicomputer transmits the outbound list back to Atlanta, and from that point on it could not care less. The local people could not care less about those cars. Once that train goes out, they all wave goodbye and that is the last they ever want to see or hear of those cars again. So it is very much a distributed processing of the data — one logical system. When the guy out in the yard hits that CRT or gets a printout, he does not know and he does not care whether it came from the central site or the local processors. It is one logical system.

Having completed that project successfully in June 1973, we then went on to a couple of other aspects which first of all involved the question of whether or not we could perform this same sort of function at what we call a 'flat' yard. After hearing all this terminology, it may come as no surprise to you that a 'flat' yard is flat, as opposed to a 'hump' yard, which has a hump! We have many more flat yards. They are older types of yard. You actually switch the cars by using the locomotive and giving the cars a little shove. There is no automation; the switches are manually thrown. The question was: could we achieve the same kind of breakthrough in the technology of operating that yard — not the computer technology but the yard technology — that we had with this new hump yard?

The reason that was so interesting to us was because we had achieved a productivity increase of 40% per employee at this new switching yard. The railroad is a very labour intensive business, and that was very important to us.

So the next place we went to was Savanna, Georgia, which is a very large, complicated flat yard. In fact, instead of having one place from which you switch, you switch from both ends of the yard and from four leads into the yard. It is a very complicated yard. We began that project in 1974, and in 1976 had successfully completed that project.

I might say that we knew it was successful because we were able to reduce eight employees out of the yard office, and that is the biggest criterion to make sure that you have got a productivity increase and have taken advantage of it. So many times when we study things we say, "This will reduce the clerical work so many hours, and then that clerk can do this." That is not the way we do it in the railroad; we flat cut off the jobs when we have eliminated the work. So we knew we were successful because we got eight payroll authorisations.

I should point out something that maybe does not come to mind right away about this situation. Clearly, this is a very great help to the local people in the yard because they have great information as to exactly what is going on; they can do better planning to the extent that they do planning in terms of what cars to handle next and so on. But it does one other thing for us: everything they do in that yard has to be done on the CRT; that is the only way they can do it. If they want to put out an instruction to do something, they have to hit that keyboard. If they want to prepare some document, they have to hit that keyboard. If they want to get the train out of the yard or a movement out to industry, they have to hit that keyboard.

Now the computer puts out work orders and things like that, but it does one other thing that they do not even realise: it captures all those key strokes. So without any further effort on the part of the people out there, without doing anything that is not a normal part of their job, we have captured what is going on out there. This is what we call 'source data capture'. I make a great distinction between the idea of data entry and source data capture. All data entry in Southern Railway Company is done centrally. Data entry is any situation where somebody sits down and keys information off an existing document. That is data entry to me. We have not had punched cards for ten years in

Southern Railway Company, but it is a key punch type of thing. Source data capture is when the person who is creating that document hits those keys.

The reason that is so important is that we discovered that as we use this information on our real time system more and more in rail operations, it was slowly becoming less accurate, less complete, and less timely than what we needed. The data was not degrading over what we had had in the late '60s and early '70s, but the ability of the operating people, the transportation people, to use it was getting better and better. They were using it more and more.

Every morning in our company we put out a morning report. On that morning report is every activity of the past 24 hours that anybody in the company needs to know about in terms of delays, problems and so on. You would be stunned at the detail in that morning report. I am part of that morning report. The chief operating officer called me, six or eight weeks ago, and he said, "You know, one of the problems, Jones, is that I've now got to read your damned morning report before I read my damned morning report." The processing has become so involved in rail operations that it really has become an integral aspect of it.

So it was clear to us that we had to improve the accuracy, completeness and timeliness. The way you do that is to get a person to do that input for you who has an incentive to do it accurately, completely and timely. Let me tell you that that is not a key punch operator; that is not a data entry person. By and large, they do not know what the data is that they are entering. They do not care whether it is payroll or expenses. They have a desire to get it complete to the extent that they do not get disciplined. They have a desire to do it timely to the extent that their boss does not get on to them and chew them out for goofing off. But that is the extent of their commitment.

What you have to do is to get to the person who has something at stake, like the yard master. When the train goes out of the yard, it goes out because the yard master says that it will go out. Why does he want that train out of the yard? Because if he does not get that train out of the yard on schedule, he will hear about it. It will be his telephone ringing instead of mine from the chief operating officer. Timeliness, completeness and accuracy. You bet. He understands that. That is part of his job.

The yard clerk correcting strangers and ghosts — can you imagine what happens if you shove 150 cars over the hump, and the third one is a stranger and all the rest of them are out of step? Guess where that yard clerk would be tomorrow? Well, you do not know where he would be, but he would not be on Southern Railway property! Accuracy, timeliness and completeness: he has an incentive. So source data capture is a very key and critical element of what we are doing here.

There is one aspect of this which is associated with the yard, but really not part of it. That is what we call a 'way-bill'. That way-bill is a single, 8½ x 11 sheet of paper that controls the movement of a car. That piece of paper is so important that even if it is wrong it is right, because that is what is going to happen. Whatever it says on that piece of paper. The piece of paper may say, "This car is empty. Take it to Reynolds Tobacco, in Durham, North Carolina for a load of cigarettes." They get it there and open it up and it is full of tyres, made in Michelin plant in Greenville, South Carolina, destined for California. That is not what the piece of paper that they get on that car said.

That way-bill is terribly important. So the other aspect that we are in the process of capturing is the creation of that piece of paper by our agent, who is actually our agent dealing with the customer, creating that information.

So these two things, the yard systems and the way-billing systems, are the key to the source data capture getting the accuracy, timeliness and completeness that we need on our data.

We went through a series of steps, designing and implementing these systems — the yard system, the way-billing system, the flat yard system which is actually the identical set of programs as it turns out — keeping track of all the cars in the terminal area. Around the yard there may be hundreds of customers and thousands of places to put freight cars which have to be kept track of around the yard — at different yards and at different places. Then we finally brought them all together and brought them up at Savanna, running on a single set of minicomputers. We completed that successfully in October 1977, and then began the rather massive installation of these systems.

We were lined up installing these yard systems at 39 major yards. We will install the stand-alone, way-billing part of it, for which we use the same computer. We use all Data General S130s for the main minicomputers. I always say that any minicomputer is good as long as it is a Data General S130. That is not because I really believe that, other than in the sense that I do not believe that you ought to mix minicomputers. If you pick out a machine to do a job for you, then you ought to use that same machine everywhere, because every one has its own architecture, its own operating system, its own programming language, and it does not make sense to have a mixture. Just because they are small and cheap, they aint simple.

One other aspect of this that I need to drag in here is the concept of the micro processor. We had at one time about 90 terminals out in the field. They were Univac DCT 1000s, just dumb polled terminals. We wanted to swap those devices and get out there with a programmable type terminal and get to a standard protocol, because Univac was not a standard protocol. It was about this time that Data General announced the MicroNova — a 16-bit micro processor.

The MicroNova had two outstanding advantages. One was that it was Data General. Data General is beginning to understand our problem, because when you run 24 hours a day, seven days a week, your maintenance requirements are not the same as they are in other companies. Sometimes it is hard to find a technician at three o'clock on New Year's Day, to go out to a railroad yard and fix something, unless you have previously made some arrangements. We have worked that out. So that is one great advantage to us, just the fact that it was Data General.

The second big advantage of that MicroNova is that it is a 16-bit micro processor, compatible with the rest of the Data General hardware. It uses the same operating system; the same language. I was terribly familiar at this point with the Data General stuff. So we took that micro processor and programmed it to imitate the Univac terminals, put it out in

the field, and basically plugged together by buying the micro processor, writing a little program, and renting a terminal — actually now we have purchased a Diablo HiType 2s — plug them together and we have basically made our own intelligent terminal and put it out there in the field. We got those 90 out there about a year and a half ago.

Now we are finding something very interesting, which is that that micro processor, by upgrading memory and a few little things like that, still coming up with a device that is only about $6,500, including a CRT and synchronous and asynchronous communications and 64K of memory and some read-only memory and so on, we can program that thing to do the way-billing function. It is amazing how many places you can put a device that only costs $6,000 or $7,000.

So we are in the midst of installing these micro processors. We will wind up with a couple of hundred of them. They obviously do not have the disc storage that a yard system would have, but not because they could not. In fact, we are playing around with a micro processor driving a 10-megabyte disc to see if that fits in the scheme anywhere. It blows your mind to have a computer on a quarter-inch chip and 10 megabytes of memory on it. Well, it does mine, but then I am an old fellow here.

We currently have installed about 120 of the micro processors. I would like to go to some slides now, and run through them quickly to give you an idea of the magnitude of this system. Again, I emphasise that even though we are going to wind up with 240 or 250 — or maybe a few more than that — Data General computer installations in Southern Railway Company, they are all the one system. It all operates like one logical system. They are all linked together by communications to the central site. At the central site we have four 370/158s, two of which drive the real time network. We have always used one processor for tele-processing and a separate processor for the handling of the databases. People always say, "You could go to one big machine," and I say, "Well, that's a wonderful chance for somebody."

We have made some progress on operators being involved. We have one operator who operates the consoles of all four /158s; we have two operators who handle the 30 tape drives; and we have one operator who handles the three printers, the microfiche machine and the plotter. So we do operate with four operators per shift. They do not have time to be too involved; they stay so busy that they do not mess up the system too much for us.

The railroad yard at Sheffield, Alabama, is about four miles long and about half a mile wide. The receiving yard is at this end. The class yard is a short, fat yard. In Sheffield it has 32 tracks. On 11 June we are opening a new yard, designed just like this one except that it is a $48 million yard up in North Carolina. On the same principle, that yard will have 48 tracks in the class yard. Some of them will have 65, but the class yard is designed to take a large number of groupings of cars. The forwarding yard is down along the side.

This is the flat yard at Savanna. There are two class yards there, which is how we switch from four places. There were two railroad yards there, side by side — one called the Central of Georgia, and the other called the Savanna in Atlanta. We own both of them and we put the yard together. That was a very difficult yard to work with.

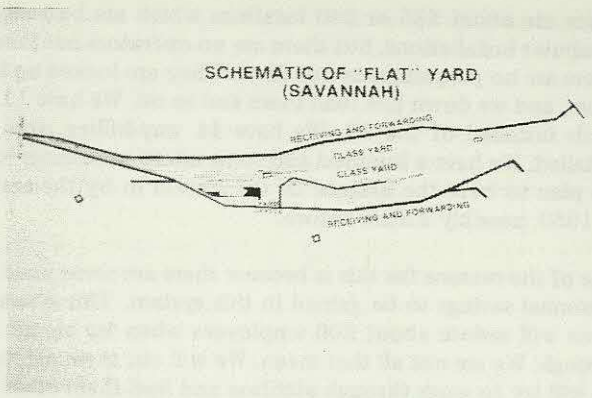

I have another slide which has the same basic functions but it just does not have the process control, with switches and things like that. You have one opportunity at a hump yard

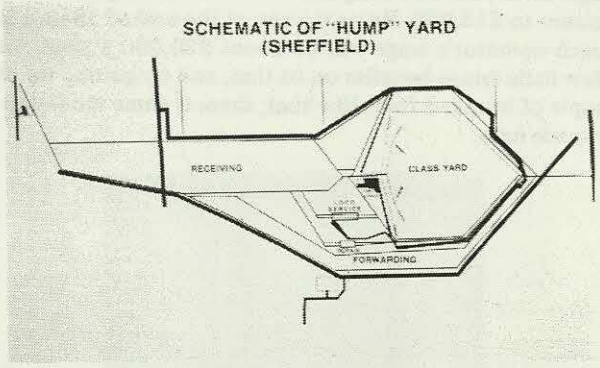

with your computers that you do not have in a flat yard. The process control part of the computer is sitting there throwing switches, routing the cars properly. You have the opportunity to throw a switch at a bad time, which would be right under the car, so that the front wheels would go over one track . . . At least you do not have that opportunity here at the flat yard.

SOUTHERN'S NETWORK

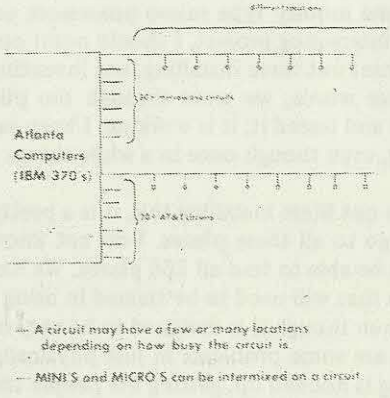

Our network is a very simple, straightforward, stay-type network. Fundamentally, it is the central site with 2400 speed buffered network. We have a very large private microwave system, the third largest microwave system in the United States. AT&T and General Telephone have bigger microwave plants than we do, and we have the third largest one. So we use a lot of our own microwave circuits. Every device out there, though, is a computer. There are no dumb terminals in this network anywhere. They are all either minis or micros.

There are about 235 or 240 locations which are basically computer installations, but there are no operators out there. There are no programmers out there. They are locked up in a room, and we down line load them and so on. We have 11 yards installed of the 39. We have 44 way-billing sites installed. We have a hundred and some micro processors in. We plan to have the balance of the system in by the end of 1980, possibly a little before.

One of the reasons for this is because there are some huge personnel savings to be gained in this system. This system alone will reduce about 200 employees when we are all through. We are not all that mean. We will cut them off, but we will try to work through attrition and find them other jobs. But some people will be walking the street when this is all said and done. But in any event, railroad wages in the United States are so high. This year, a key punch operator's wages alone, not counting the 25% to 30% fringe benefits, amount to $18,000. By next year, at the end of 1980, a key punch operator's wage will be about $20,000 a year. Add a few little fringe benefits on to that, and by getting rid of a couple of hundred folks like that, there is some money to be made here.

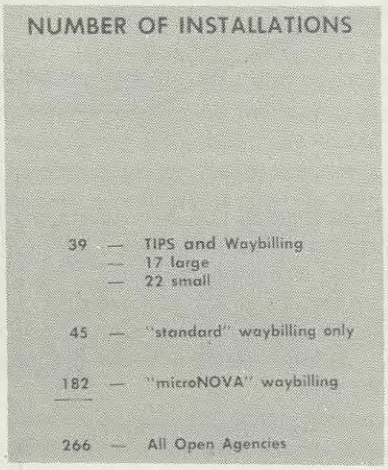

This gives you the run down of the locations. There are a bunch of micro processors not included here in that we still will have some inquiry type micro processors out there. It is a fairly interesting project. I should point out that we are at this point out there installing, not inventing anything new. In other words, we went through the pilot; we developed it and tested it; it is working. I keep saying that it is working, even though once in a while things go wrong.

While we are out there installing this, it is a pretty significant problem to go to all these places. I do not know whether we will even be able to find all 266 places. We have a lot more people that will need to be trained in using this stuff out there. Even though it is designed to be very natural for them, there are some problems in just physically getting there, getting it hooked up, getting the people trained and running with it, and so on.

We already know a bunch of things that we can do, having put this capability out there, in terms of other things that we we can capture. We could do it right now. But that is not today's problem. Today's problem is not inventing something new any more, it is getting the stuff installed and running. Then we will get back out there and add to it.

I make that point just to emphasise that it is a great temptation to say, "Ah, we've got these computers out there, we'll add this program and that one." The only thing you get with that is unreliability. You just put more things out there to give you a problem.

To put this in perspective, we have in our company a very unusual degree of understanding and participation by the president and the vice presidents, due to another story with which I will not bother you. But in any event, for 13 years we have had a very high level of participation by our senior management. They understand what is going on here very thoroughly.

Our president made an interesting comment about this project. He said, "The impact of this project on our business and our company is going to be second only to the invention of the diesel locomotive." I think that sums up the impact that this kind of thing can have on a business if done right. I hope that it is done right. I still have some years to go until I care to retire. The opportunities are there. One must carefully pick and choose so that you have some chance of success instead of a bigger chance of failure.

Let me put it another way. A project like this gives you a lot of opportunities to take a lot of small, difficult problems and make them all into one big, unsolvable problem. That is what you must watch out for.

QUESTION: Are you using Cobol on the MicroNova?

JONES: We are not using Cobol on the MicroNovas. The reason for that is that the Data General Cobol for the MicroNova, which does exist, does not at this point have a well-tested, concurrent communications capability. In other words, you cannot do communications and processing concurrently. We could not wait for that to come about in order to get on with this project. There are certain timing pressures on us in terms of some labour considerations, so we simply could not wait. We intend at some point to go back and re-work those programs in Cobol. That sounds like something one always says one is going to do and never does, but we have done it once on the minis. When we started out with the minis in 1972 there was no Cobol for that either, and we wrote those in assembler language and went back and re-wrote them in Cobol. All the minicomputers' programs are running in Cobol.

QUESTION: May I ask what the total data processing cost is?

JONES: First of all, all of my central hardware is purchased so I do not have any cost in there except for maintenance for the central site. I do not have any communications cost, other than AT&T circuits. In our company we do not work with charge back systems. The way that we work is that everybody is aware of what is going on and if the company decides they will provide a service, then you get a budget provided for that service. I provide the computer service and the communications department provides communications, so I do not have that cost.

I also have operations research, industrial engineering, all the corporate printing, all the corporate photocopiers, and a bunch of junk like that. Operations research and industrial engineering are not so bad, that is a lot of fun, but some of this other stuff . . . But the total annual budget is about $15 million. An awful lot of that cost is the key punch

people. The data processing staff, what most people would call programmers and analysts — we do not make that distinct, we call them all programmers — would number about 60, including supervision. We sure work hard, though.

COX: On that note we must close the session. Jack, you have thoroughly engaged the audience, as you can tell. I am sure that, like me, they appreciate your incredible ability to preach a very hard-nosed and disciplined approach to systems with such relaxed charm. It has been a pleasure to have you here.

CASE STUDY OF A DISTRIBUTED PROCESSING SYSTEM USING MAIN FRAME AND MINI COMPUTERS

Steven Rowe
Bank of America

Steven Rowe is the Chief Analyst in charge of Bank of America's distributed system. He is based at the Bank's headquarters in San Francisco, California.

COX: We now come to our final presentation of the day's proceedings. And in this case we have gone to the financial sector for our case study. The Bank of America is the world's largest bank and appropriately enough, we have brought along Steven Rowe, who is their largest analyst. Steve will describe what they have been doing in their very extensive distributed systems.

ROWE: What I'd like to do today is to cover four areas. I want to describe the problem to which we adopted a distributed solution. Secondly, I want to describe why we chose a distributed solution. Thirdly, I'd like to describe that distributed solution and, fourth, I'd like to give you the lessons we have learned and discuss the mistakes we made.

To begin with, the problem faced by the Bank of America, and California, is largely a function of our size. The State of California has approximately 200 banks. There are 3,600 branches of those banks in the State of California. The Bank of America has roughly a third of that market; we have about eleven hundred branches.

The way we perceive the competitive environment in California is that our success in the commercial market place is going to be determined by our ability to provide a high level of customer service.

This thinking was going on in the late '60s and early '70s, and approximately in that time frame a decision was made to implement a State-wide on-line system to provide timely information to the Bank staff dealing with the customers themselves.

During this time, salaries were again a major consideration and the cost of providing responsive customer service was increasing every year, and again — as in the railroad industry — labour costs were a very big concern. The Bank has about 70,000 employees in California and a large number of those employees are dealing with the back-office operation of the Bank — the paper work, the accounting, these kinds of function. So any kind of system that we could implement, in these eleven hundred and some odd branches, that would reduce the number of people dealing with the clerical and accounting functions and increase the number of people dealing with the customers would serve both ends.

The slide I have got on the screen here documents our solution to this basic problem. The distributed computing facility is our central processing facility. Community on-line system is really the composite of the central facilities, the terminals in the branches and the network that go to provide this on-line system.

DISTRIBUTIVE
COMPUTING
FACILITY

COMMUNITY
OFFICE
ON
LINE
SYSTEM

At the current time, our distributed system provides those services to our branch offices. I am not sure how familiar you are with the terminology — demand deposits, saving deposits etc. Demand deposits are essentially just checking accounts.

COOLS APPLICATION FUNCTIONS

* INQUIRY AND UPDATE
  * Demand Deposit Memo Accounting
  * Savings Deposit Memo Accounting
* ONLINE DATA ENTRY
  * Non Dollar Data Entry
  * Replacement for OCR Entry Mechanism
  * Online Interactive Validation and Correction
  * Capture of Holds, Stops, Warnings
* TRAINING FACILITIES

Savings accounts with on-line data capture are new account set-ups and changes to old account set-ups. Training is the training of our teller personnel in the use of the facility.

JUSTIFICATION FOR COOLS/TELLER APPLICATION

* REDUCE TELLER TIME PER CUSTOMER
* REDUCE INTER OFFICE PHONE EXPENSE
* REDUCE AMOUNT OF PRINTED REPORTS SENT TO OFFICES
* REDUCE CHECK CASHING FRAUD LOSS

The system I am talking about is a twenty million dollar system and it was justified on these hard dollar savings. With twenty thousand tellers in California, we can afford to spend quite a bit of money to reduce the amount of time each teller spends with a customer. The way the Bank of America is structured is each office or branch functions as its own profit centre. It has its own reports, computer print-out reports, account lists etc., and when a branch is dealing with a customer from another branch, before the system was implemented, in order to provide that customer service it had to call the other branch. We had a very large telephone bill.

On the last item, I have heard that the savings have been so significant that we could have justified the system on that factor alone.

The Bank basically divides California into thirteen regions. And it's my understanding that on any day of the year, in each one of those thirteen regions, there is some sort of fraud activity going on against the Bank.

COOLS SCOPE

HIGH TRANSACTION VOLUMES

* 50 Transactions Per Second Statewide
* 25 Transactions Per Second Each Center

SEVERE RESPONSE TIME CONSTRAINTS

* Average Response Time to Terminal – 3 seconds
* 95% of Responses Within 6 Seconds

HIGH RELIABILITY REQUIREMENT

* 98.5% Central Site Uptime
* 9 Minutes of Central Site Downtime Per 10 Hour Day

HIGH GROWTH POTENTIAL

* Transaction Volumes for Current Services Will Continue to Increase as the Bank's Customer Base Increases
* Transaction Volumes for Current Services Could Increase through Policy Changes (e.g. Check Cashing Floor Limits of Zero Would More Than Double Transaction Volume)
* New COOLS Services Could Greatly Increase Total Transaction Volume

Here is some general scope of the problem that we were dealing with in developing on-line systems to serve our branches in California. Basically, our data processing is done in two data centres — one in San Francisco and one in Los Angeles. They are essentially parallel installations in that we have divided the State evenly between the two data centres.

We felt we had to develop a system that was divisible between the two data centres and that the total State-wide volume for the system had to be at least 50 messages per second, in terms of capacity. We obviously defined a number of database accesses per transaction, a profile of a typical transaction etc.; our response time objective was that 90% of them would be within three seconds, 95% within six seconds.

Because once this system was implemented the Bank would become totally dependent upon this system, and there would be no back-up system in its place that was practical, we felt we had to have a very high reliability factor. In fact we have achieved better than the 98.5 indicated there.

COOLS TERMINAL NETWORK

* 9200 TERMINALS
* 1120 INTELLIGENT MINICOMPUTER-BASED TERMINAL CONTROLLERS
* TERMINAL CONTROLLER FUNCTIONS
  * Local Input Editing
  * Standard Message Formatting
  * Terminal Security Facility
  * Off-Line Calculator Function
* TERMINAL CONTROLLER SOFTWARE LOADED FROM CENTRAL SITE OVER COMMUNICATION LINES
* 111 COMMUNICATION LINES
* TERMINAL CONTROLLERS MULTIDROPPED; UP TO 14 PER LINE
* 2400 BPS SDLC FULL DUPLEX.

This system is perceived as the Bank's entry into electronic funds transfer. We are automating the banking process. And the bottom line in banking operation is, basically, reducing the paper costs. Banks, and particularly the Bank of America, deal with a tremendous volume of cheques. We process on the average one cheque per account per day, which works out that we are processing approximately through our operations somewhere between six and 12 million pieces of paper every evening. So anything that is going to reduce that paper flow is encouraged.

Given that paper is our biggest problem, then any move to electronic capture of data or transfer of funds is very desirable. We perceive that we have to create a foundation environment to facilitate the movement to an electronic funds transfer environment. And this system therefore had to be generalised to support very large potential volumes, or to be expandable to support large volumes, and we wanted it to be sufficiently general that we could support new and quite different applications on it.

This is the system as it's currently implemented. We have now over ten thousand terminals. These are dumb terminals, each tied in to a minicomputer which is located in each branch. The minicomputer is a Bunker Ramo, programmeable control unit, which is very much like a DEC PDP 103. It essentially performs the functions listed there, plus it provides some system level interfaces into our central side facilities.

The 110 communication lines are all leased lines from the telephone company, and on average we have 14 PCUs per line. The lines operate at 2400 BPS, full Duplex, plus DLC.

Figure B Basic Single System Approach for a Single Center

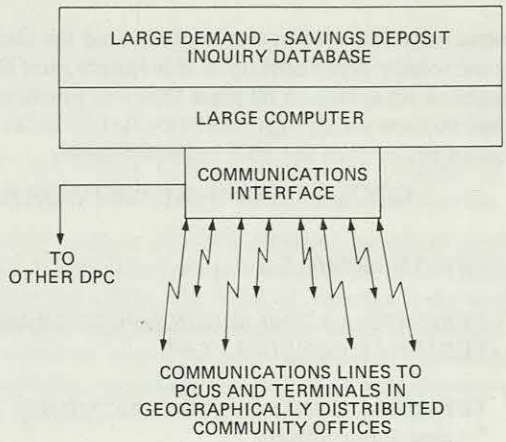

When this system was being considered our first approach was to build a large mainframe base system. And, in fact, we signed and went into a joint venture with the United Airlines to develop a system based around their PARS derived system. In fact we signed a contract and worked with them for over a year on development of this system. Approximately six to 12 months away from our initial implementation we made the decision to drop everything, back out of our contract with United and go to a distributive approach. I guess the fundamental overriding reason for that decision was the concern that a single mainframe CPU would not be adequate to handle our growth and expandability requirements. Subsequent experience of United Airlines, I believe, has substantiated that view because they have recently come face to face with the problem that a 360/195 simply wasn't adequate for their reservations system any longer and I think they have now gone into a joint venture with Eastern Airlines.

Figure E Simplified Distributive Computing Facility Processing Module

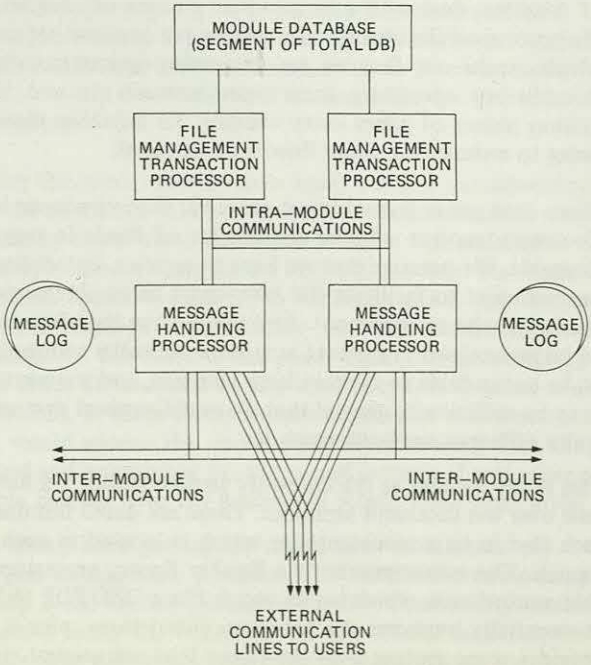

The slide refers to central site processing functions and this is basically the alternative that we took to the single mainframe approach.

DISTRIBUTIVE SOLUTION CHOSEN BECAUSE

* HIGH TRANSACTION VOLUMES
  * Concern That Volume Growth Could Exceed Processing Capacity of Most Powerful Mainframes
* RELIABILITY REQUIREMENTS
  * Redundant Hardware and Automatic Switching of Function to Backup Component Needed to Meet Requirements
  * Isolation of Applications in Their Own Set of Computers to Localize Impact of Instability in a Particular Application
* CENTRAL SITE HARDWARE COST
  * 10 Million Dollar Cost Savings through Use of Many Minicomputers as Opposed to Duplexed 370/168s in Each Center
  * Less Expensive Increments in Processing Power
    – Increasing Processing Capacity for Given Service
    – Adding New Service
* EXPAND ABILITY
  * Ability to Add New Applications Services Without Disrupting Existing Services

There's a more formal statement of the reasons behind our decision to go with a distributive approach. As we perceive the situation in California at the current time, once again the bottom line for the banking industry is the number of cheques you process. At the current time about 10% of the total items come across our teller windows; 90% are through institution transfers or transfers from corporations and that kind of thing. But the potential exists that with an EFT system, an on-line system, a real time system, a considerable portion of the 90% that now comes in as paper can be replaced with electronic systems or electronic transfer of some kind, with a consequence that any system we develop would have to be expandable to handle a much higher volume than even our branch system would now provide.

Once again, the presence of an on-line system would really eliminate the possibility of a back-up, so our system would have to be more reliable than we felt the current mainframe hardware could be made to be.

Central side hardware costs is an argument that was valid at the time the decision was made in 1974. It is somewhat questionable now. The general decline in hardware costs across the board has made this a questionable decision.

Expandability, the last item — the ability to add new applications without disrupting existing services — is really an implementation issue.

MODULARITY AND RELIABILITY

* MODULAR ARCHITECTURE ALLOWS EASY EXPANSION FOR CAPACITY INCREASE AND FUNCTIONAL ADDITIONS
* FLEXIBILITY IN EXPANDING BEYOND TWO PHYSICAL DATA CENTERS
* PARALLEL HARDWARE ARCHITECTURE WITH AUTOMATIC SWITCHING FOR RELIABILITY
* SECURITY THROUGH
  * TERMINAL ACCESS CONTROL
  * PASSWORD SECURITY

These really represent part of our design goals for our distributed system. We wanted a modular architecture, and the objectives in modularity were twofold:

- one, to allow us to match our hardware with our capacity requirements, and
- two, to modularise the software into maintainable, discreet components. How that comes about, I'll show you in a second.

Flexibility is spanning beyond two physical data centres. This was important to us, as I mentioned before, as we have this basic structure which divides the State of California into 13 regions. At one time we had a system based around centralised processing on a regional basis and it's been a concept that has not yet been replaced. In fact, we may go back to it. So any system that we implemented had to be open-ended in terms of its ability to divide or expand to 13 regions.

We needed a hardware architecture with an automatic switching, with automatic and quick switching for reliability. We wanted to have a facility that was failsafe to the greatest extent possible. Obviously we needed security on the funds transfer system.

Figure D DCF System Functional Block Diagram

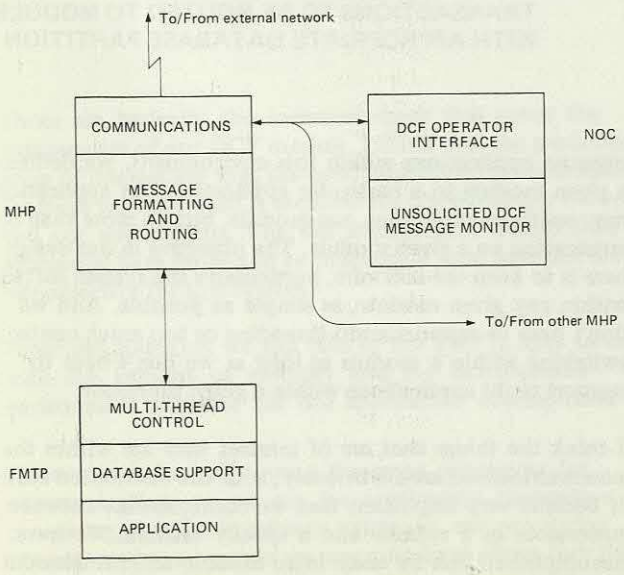

This represents the functional breakdown, the way we divided the functional components of our distributive system. It so happened that it is also the way we have physically sited those facilities in that those functions each reside in a separate processor and not in a distributed network. There is a single processor dedicated to the message formatting and routing function. There is a single processor dedicated to the application and database support area. There is a third processor that controls the network and controls the distributed processors at the central side.

DCF PROCESSING MODULE HIGHLIGHTS

* BASIC SYSTEM BUILDING BLOCK
* ALL COMPONENTS DUPLEXED
* TWO MESSAGE HANDLING PROCESSORS (MHP)
  * PROVIDE COMMUNICATIONS FRONT END FUNCTIONS
  * LOG ALL INCOMING AND OUTGOING MESSAGES ON TAPE
  * INSPECT MESSAGE TEXT AND ROUTE MESSAGE TO THE APPROPRIATE MODULE FOR PROCESSING
* TWO FILE MANAGEMENT TRANSACTION PROCESSORS (FMTP)
  * CONTAINS SYSTEM SOFTWARE NECESSARY TO PROVIDE APPLICATION MESSAGE PROCESSING ENVIRONMENT
  * CONTAINS APPLICATION SOFTWARE
  * CONTAINS SYSTEM SOFTWARE TO PROVIDE ACCESS TO APPLICATION DATABASE
* IF AN MHP OR FMTP FAILS, ITS COMPANION (BACKUP) AUTOMATICALLY ASSUMES THE PERIPHERALS AND THE WORKLOAD

The central side facilities on our network were provided by a distributed network and minicomputers. These minicomputers are organised into clusters, and typically there is one cluster per data centre, and right now we have two data centres. Then within a cluster they are further divided into modules. The module is the basic building unit. This represents a module. It consists of four GA 16440 minicomputers organised in parallel paths. In other words, within each module there are two communications processors called MHPs, there are two file transactions processors called FMTPs. Each processor within a module is connected to every other processor within the module. Each partner backs up, or each half of a module backs up the other half, meaning that if any given MHP fails the system will switch to the other MHP and that MHP assumes the load for both.

Figure F Bank of America Distributive Computing Facility Processing Module Detail

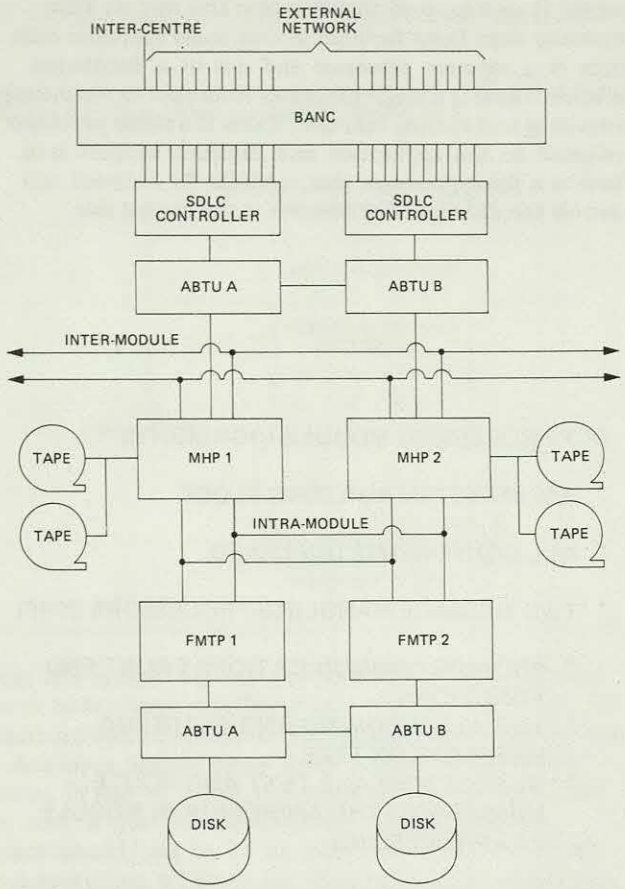

Here's a slightly different breakdown of a module. Starting from the top, there is our network control facility with 110 communications lines that have a diagnostic and network

facility. Then the STLD controllers. Their role is purely that of communications lines controller. ABTU is an automatic bus transfer unit. It is a piece of switching hardware which senses the state of its host and its partner's CPU, and if the situation is appropriate we will switch the peripherals to the other processor.

There's the MHP in the tape drive, which is used for logging each message as it comes across to the network. Then we have communications lines connecting each MHP to each MFTP as well as across the same communications bus to other FMTP.

Finally, we have a processor that does the actual application processing and transaction processing connected to its database via ABTUs, again with the same functions as before.

This is a high-level diagram of how the whole thing ties together. In each data centre we have one to n modules connected by a communication line to the other data centre as well as to the external PCUs and the branch offices themselves.

Figure G Network Block Diagram

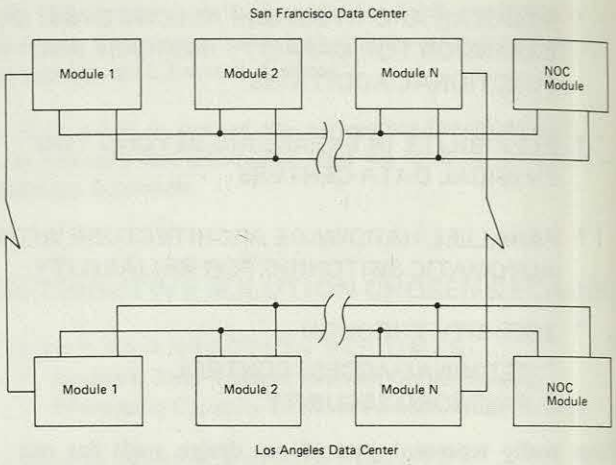

Here's another diagram showing the same thing, the major difference being that it indicates the way in which we

MODULE USAGE

* ONE APPLICATION PER MODULE
  - PREVENT INSTABILITY OF ONE APPLICATION FROM JEOPARDISING OTHERS
  - ALLOW PERIPHERAL CONFIGURATION OF MODULE TO BE TAILORED FOR APPLICATION
* IF PROCESSING REQUIREMENTS OF AN APPLICATION EXCEED CAPACITY OF ONE MODULE
  - PARTITION DATA ACROSS MULTIPLE MODULES
  - REPLICATE APPLICATION SOFTWARE IN MODULES
  - MAKE ROUTING TABLE CHANGES TO CAUSE TRANSACTIONS TO BE ROUTED TO MODULE WITH APPROPRIATE DATABASE PARTITION

organise applications within this environment. We dedicate a given module to a particular application. An application may reside on more than one module, but no more than one application on a given module. The objective in our design here is to keep the software, particularly the system software within any given module, as simple as possible. And we don't have to support multi-threading or too much contact switching within a module as long as we don't have to support multi applications within a given processor.

I think the things that are of interest here are within the communications area. Obviously, with the distributed system, it became very important that we communicate between processors in a reliable and a speedy fashion. We have accomplished this by these intro-module and intra-module links, which run at 2.4 megabits. These use an SDLC-like

NETWORK CONTROL

* CENTRAL OPERATIONAL CONTROL
  * DCF OPERATOR CONSOLE
  * REMOTE IPL OF ALL PROCESSORS
  * STATUS INDICATIONS AND MESSAGES
* NETWORK DIAGNOSTIC FACILITIES
  * FULL DIAGNOSTICS ON TERMINAL NETWORK
  * MINICOMPUTER CONTROLLED DIAGNOSTIC FACILITY
  * USE OF SIDEBAND ON COMMUNICATION BANDWIDTH
  * DETECTION OF MODEM AND LINE PROBLEMS

protocol. They are communications links. They are not bus connections. Although it's probably a little impractical, we still have the option of coupling in to other systems if we choose.

DCF SDLC COMMUNICATION PROTOCOLS

EXTERNAL LINES

FULL DUPLEX PRIMARY
2400 BITS PER SECOND
MULTI-DROP POLLED
FULL OR HALF DUPLEX SECONDARIES

INTRA-MODULE AND INTER-MODULE LINKS

HALF DUPLEX
2.4 M BITS PER SECOND
MULTI-DROP CONTENTION

INTER-CENTER LINK

FULL DUPLEX
9600 BITS PER SECOND
POINT-TO-POINT
PRIMARY/PRIMARY

Those are basically the technical specs that cover the components of our DCF module. I think that the particular areas of interest are that the communications processors have far more capacity than is required for our particular network and our applications. The file management transactions processors have proved to be inadequate, and have not met our objective for this system.

Particular restraints we run into in the FMPTs have to do with the amount of memory and the development environment available for our application development.

Since we are basically using the same processors for communications as we are for application transaction processing, a basic imbalance exists. In our processing capacity within a module we currently have about four times the processing capacity in the front end of our modules as

Figure J Two Center Distributive Processing Logical Concept

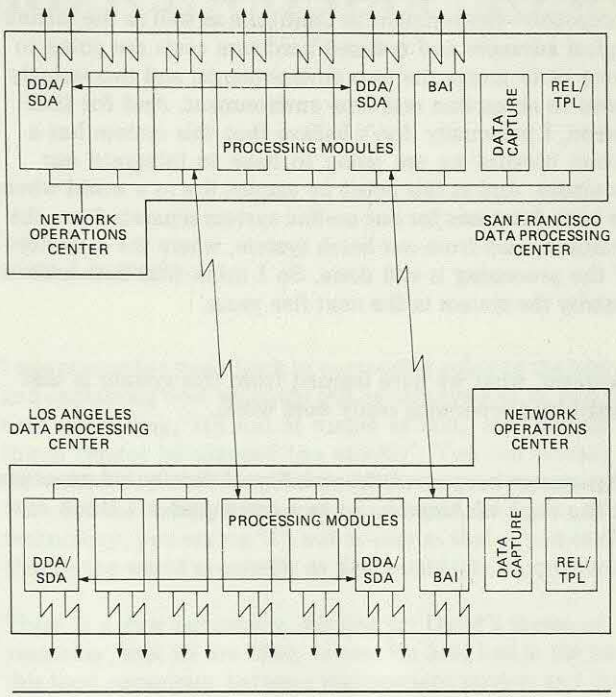

we do in the back end, where the actual work is done. So if we have to do it over again we would quite clearly go to a different architecture — something like 2X4 or 2X6.

I guess, to sum up our experience in data and distributive data processing, several conclusions can be reached. One is, Mr. Jones' conclusion that minicomputers are not simple is very valid. Our experience has been that the system was originally sold to senior management on the basis that minicomputers were simpler than large mainframes. This is not proven to be true and we have had considerably more development and software effort in our minicomputer system that we did in a large mainframe based system. And that our future efforts and developments will be directed towards systems built around standard software, standard hardware, as opposed to custom-tailored systems such as this one.

The system did meet its immediate objective. We are currently meeting our response time objectives easily. In other words, over 90% of our transactions are being responded to within three seconds or less. And our uptime exceeds the 98.5% objective that we set for this system. We are actually running over 99% and that includes all systems components, central site as well as network and branch programme control units and terminals.

Where this system failed, I think, falls into two areas. One is that it was conceived to be a building block for the Bank of America to build on for future applications to complete automation of our branch offices. I think experience has told us that it will not do that, and the basic reason why it will not is that the development environment is too unfriendly. All of our applications are written in a macro-assembler language. We have a great deal of effort and time invested in these applications.

A key point here is the way we perceive the future direction of data processing at the Bank of America. Up to this point

we have regarded teleprocessing as a separate and smaller portion of the overall data processing pie. However, the competitive environment in California as well as the technological advances and reduced hardware costs are going to force us to merge the two environments, and increasingly move to an on-line real-time environment. And for that reason, I personally don't believe that this system has a future because we are going to have to integrate our databases. And at this point we cannot live in a world where we have databases for our on-line system separate from the databases used from our batch system, where the majority of the processing is still done. So I think that that will destroy the system in the next five years.

However, what we have learned from this system is that distributed processing really does work.

This system has proved the viability of distributed processing at the Bank of America, so as a consequence I think our follow-on and our add-on systems are going to be distributed. They will probably be distributed using more standard system software and different kinds of hardware. But nonetheless, the idea of breaking the problem into small components, into modules that you can deal with, I think has been proven to be valid.

QUESTION: What is the average cost per transaction of the system?

ROWE: A real rough one is about three to four cents a transaction. Considerably less than our walkthroughs with MIS.

COX: I won't even ask how that compares with Samuel Montagu.

Steven, we are very grateful to you for coming along here and describing your design philosphy with us and sharing your experience with us. Thank you very much indeed.

SUMMARY AND CONFERENCE CONCLUSIONS

G. E. Cox
Butler Cox & Partners Limited

COX: It now falls to me to undertake the very difficult task of summing up. I realise now that I am rather ill-equipped for this. I don't have that much knowledge of classical Greek history; I don't even speak the language. Therefore I am going to take a very recent and short-term view of developments.

It also feels a bit like trying to sum up a football match just before half time. So it can only be an interim summing up; I think a lot of the value of this week is still to come with our visits.

We have devoted the conference to two areas of fast-moving development — office technology and distributed processing. These are going to usher in major, widespread changes, not just in business systems, but I believe in our lives in general.

For people here it represents not only a need to master new technology, which I think we'd expect, but also the need to develop new skills in our organisations and new methodologies. What we are on the brink of are pervasive systems that interweave with the way people work. Now I know that many of us have had interactive systems in our business for several years, but if you actually think about it, the number of people in our organisations in contact with such systems is very limited and, by and large, they are people who were specially trained for that role.

The systems technology that we are now discussing becomes an everyday part of everyone's job. That's a fundamental change. I find it exciting, but I think it's going to call for a fast response to this situation on our side. That is, if we are actually to direct and influence these developments, rather than just react to them, and to anticipate the demands of the users and the moves of the suppliers.

It's happening awfully fast. What we are going to experience I think is far greater user 'pull' for systems than we have ever experienced before. You can see this starting to happen. People being titillated by little bits and pieces of cheap hardware, cheap devices that they hear about. If you have been watching television in the UK in the last few weeks, Wordplex have been advertising word processors in the commercial breaks — and I think this type of thing will start to happen with some quite advanced processing systems.

Something struck me that Randy said yesterday: that people being regarded as reactionaries in a business was not a function of age but a function of how long they had been in a role. I might add that people in this room have been in the role of controlling systems in organisations for some time. And I suspect that you can recognise in yourself reactionary attitudes towards a lot of these developments.

I can remember years back in computing going to the user and explaining how we could put certain systems in, and the user responding, 'It's not as simple as that,' and 'certain things cannot be changed too quickly'. You can actually see examples nowadays of exactly the reverse happening. With the user asking for a system, asking to use a piece of technology, you saying 'it's not as easy as that — you can't change the world as quickly as you would like to, out there'.

There is a new conspiracy, picking up David's theme of yesterday, that we are likely to see. We have had in the past this tacit conspiracy between management services and the supplier. I believe there is a new conspiracy that we can see: that of the user and the supplier.

I believe that the right reaction to this is to appreciate that the user is only going to continue to hold us in the esteem that we currently enjoy if:

- we match his enthusiasm, and
- we demonstrate skills that are still welcome and make an essential contribution to successful systems.

And I think certainly these skills are going to be required not just in a technical area, but increasingly in the area of the interface between people and equipment. As many of you know, we acquired onto our staff in Butler Cox at the end of last year Tom Stewart, from Loughborough University: not an information technician like most of us, but a specialist in ergonomics. And one of the assignments he had recently, I found most intriguing. A large company had installed a very large word processing system. They have been running this now for about four or five years — it really was an early one — and they had no end of problems. In not a very happy industrial relations environment they had all sorts of problems.

For example the women would complain about their working conditions, and the company would throw money at these problems. They would say, 'Here's a book of wallpaper, pick out whatever you like'. And they did, and then another problem would come up and the women would say, 'It's too hot (or cold) in here,' the lighting was bad, they didn't like the supervisor . . . So they changed the supervisor and the lighting. The problems were manifesting themselves in many forms: regular absenteeism, high degree of errors, and so on. And this eventually reached a head, just after Christmas, when the union concerned just went to management and said, 'These machines — we've realised they make our people sick.' And the people responsible for the system said, 'That just can't be the case. Believe me, these

machines don't make people sick.' And they said, 'Our women get sick a lot and it's undoubtedly the machines.' And after a little bit of discussion the union put its case more strongly and said, 'We are not going to work them any more.'

At that point they turned to us, and Tom Stewart went in as referee. To read his report is fascinating. It documents all the incidents that have happened and you can see how the successive reactions of the management of this company — well-intentioned as they often were — actually exacerbated the problem.

I think that you have got a number of problems that have come to light. I think really we have got far-reaching decisions to make in terms of systems strategy, and the architecture we use. I tend to think we'll deal quite confidently with those. When we recognise a nice solid problem, the cause for careful technical evaluation, I think most organisations here tend to approach it thoroughly.

However, when you get some of the more nebulous areas, such as our methods of analysis, our methods of evaluating whether we should do a project, our control of projects, our interface with the user, most of us still have problems.

And I think they are exacerbated by some of these developments. It was interesting to notice some of the comments on cost justification. I think this becomes an interesting area particularly when the cost, the known cost of what we are doing in some of these systems, is a very small part of the total cost, i.e. the cost of equipment, the one really tangible thing you know about, becomes a less and less consequential part. And how you actually justify some of these systems — particularly when the savings are not in terms of displacement, not always measurable, not on a small scale — is quite fascinating. It caused me to think of an example in our own organisation.

I have a very sound approach to requests for new facilities in the company. A filter mechanism. What I do is wait until the people asking for them become quite apoplectic about their absence. At that point I take it that there's a prima facie case for a preliminary look at it. And this has happened over the last two or three years with word processing, and it got to the point where people would ask me to intervene, and say, 'This report has got to be out by Friday and here it is,' and they would send a little note saying, 'If we had a word processor, this would not be a problem.'

Eventually I succumbed to this and said that we would investigate word processing and how we should use it. Now you may think that specialists turning attention on themselves is no problem. Don't you believe it!

I commissioned one of our experts in the office technology field to do a report for me on whether we really did need a word processor and what the benefits would be and so on, and he went for what he thought would be the weak spots in my eyes. One was, straightforward cost saving — we would undoubtedly save on girls; and the other was, it would undoubtedly remove bottlenecks in preparing reports. I didn't believe this, but I accepted the recommendation. We put in a word processing system. It certainly hasn't saved any staff at all, and it has created a bottleneck.

Yet I am delighted with what we have done. The reason that I'm delighted and we are actually looking now for extending the system, is that those savings are really immaterial. The real attraction is in the quality of the final document; the fact that you can afford redrafting and redrafting. Now that is a matter of policy. It's impossible actually to cost justify that. What one could have said in the first place is, if you want this, this is what you will pay for it. It really is not a simple case of quantified pros and cons.

Furthermore, I think you find that when you put in a system which really gets down to the way a person does their job, it actually alters the problem. If you look at the Xerox machine, the copies you take now are not the kind of things you'd ever have thought of doing before Xerox machines were available. It becomes a tool just because you have got it. An interesting figure was given yesterday: the average number of copies per document. That wasn't the way people originally thought of using high speed copiers.

That then is the situation. I am not advocating, incidentally, a loose approach to assessing the rationale for new systems or a loose approach to quantifying benefits. I am just saying that in certain areas the benefits are not apparent at this stage, particularly when you talk of small, cheap, pervasive technology, which is an integral part of the way the person works from hour to hour.

Now as Gary Specker said, we will have a dramatically expanded portfolio of solutions to our information problems. Now that's a message that you probably, like me, get again and again, but I think to take advantage of it requires one or two important changes.

Firstly, it requires a subtle change in our position within the business, the management services function no longer having an absolute right to control data processing merely because we control the processing equipment.

Secondly, it requires an extension of our skills in a number of areas. The necessary extension of our skills in certain areas like telecommunications is obvious. The extension of our skills in actually analysing how people work and how they use equipment is not yet so apparent. It also requires an addition to some of our methodologies, and it requires an addition to our systems tool-kit. There are new skills in systems analysis and design, which I am quite sure many of the people who have been designing systems now don't have, and as yet they are not being taught — areas like data analysis, still a term loosely used but not understood, ergonomic aspects of systems design, still something we really haven't had to encounter too much before.

So over the past few days, to help us in assessing this changing situation, we have had some general presentations to give us perspective and we have had some insight into the views and the experience of others. In closing, I'd like to thank on your behalf all the speakers, to whom we are extremely grateful, for coming along and sharing their views with us. I trust it's given you some food for thought as, I can assure you, it has me. Thank you very much.

The Butler Cox Foundation

Butler Cox & Partners Limited
Morley House, 26-30 Holborn Viaduct, London EC1A 2BP
Tel 01-353 1138, Telex 8813717-GARFLD

SISDOCONSULT
20123 Milano – Via Caradosso 7 – Italy
Tel 86.53.55/87.62.27

Akzo Systems B.V.
Velperweg 76, Arnhem, The Netherlands
Tel 85-662629

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

Printed in England