Document Image Processing: The Next Revolution?

BUTLER COX FOUNDATION

A Paper by Tony Gunton
September 1987

BUTLER COX FOUNDATION

Document Image Processing: The Next Revolution?

A Paper by Tony Gunton
September 1987

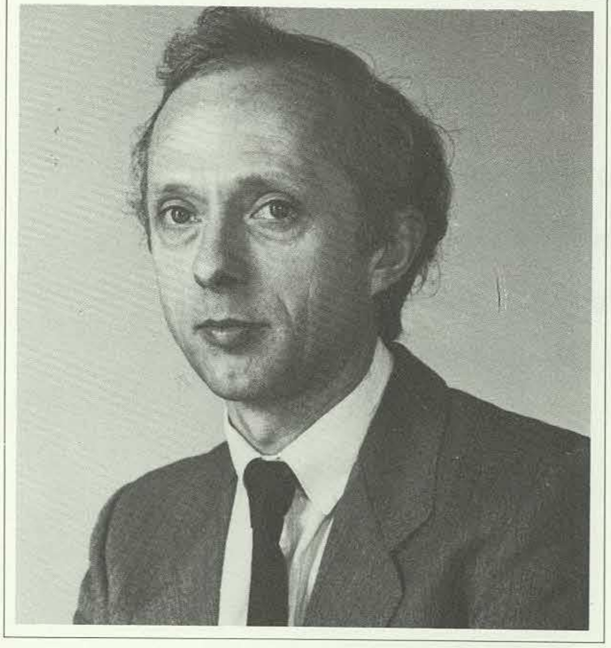

Tony Gunton was one of the founders of Butler Cox and is now an Associate of the company. He specialises in the implications for management of the move of information technology out into the workplace. His expertise is based on his own management experience and some 15 years of consulting and research for the most successful companies in Europe and North America, most of it with Butler Cox. His first book *Business Information Technology: End-User Focus* will be published by Prentice-Hall in 1987. He is also working on training and decision support tools based on expert systems technology.

Document Image Processing (DIP) is the latest manifestation of the move of information technology out into the workplace. Like personal computers before it, it offers both a promise and a threat. The promise is that company information systems will extend their reach into paper filing systems that are still a severe constraint on many aspects of business operations. The threat is that information systems people will be caught unprepared.

DIP systems meet four criteria that suggest they will indeed have a major impact on information handling. They promise to deliver substantial business benefits; they are in the mainstream of evolution of information systems; they have the hallmarks of a successful technological solution to the business problems they address; and technology is evolving rapidly to remove the cost/performance limitations of current equipment.

What should information systems managers do about DIP? Above all, bear in mind the central lesson of experience with end-user systems to date — *no gain without pain*. DIP will provoke organisational change, and will take systems into relatively unknown territory — the territory of document management, better known to organisation and methods specialists than to most information systems people. To choose the right targets, do not forget that the application must be right; people are the key constraints; and the technology has limitations as well as strengths, which must be respected.

Published by Butler Cox & Partners Limited
Butler Cox House
12 Bloomsbury Square
London WC1A 2LL
England

Photoset and printed in Great Britain by Flexiprint Ltd., Lancing, Sussex.

BUTLER COX FOUNDATION

Document Image Processing: The Next Revolution?

A Paper by Tony Gunton

September 1987

Contents

BUTLER COX FOUNDATION

Document Image Processing: The Next Revolution?

As each new technology comes along, someone can be relied upon to proclaim it the next revolution in information processing. What I call Document Image Processing in this paper is no exception. John Connell, for example, wrote in the IMC Journal last year:

> "Image processing is the next major breakthrough in advanced office technology and it is destined to have a significant impact on office operations and office-based personnel."

Writing in Computerworld, an Arthur Andersen consultant echoes the theme, somewhat more cautiously:

> "Image processing may be the next major type of information handling that will fundamentally change the way information is stored, accessed, transmitted, and analysed."

By Document Image Processing (or DIP for short) I mean the capture and storage of information as images rather than as coded data or text. In addition, identifying information is associated with the stored images, so that they can be identified and retrieved later. DIP systems do not permit the detailed manipulation of images associated with computer-aided design applications, but they normally allow users to make changes that do not involve manipulation of the images, such as by annotating them or by 'stapling' several of them together electronically. They also assist with comparison and analysis, by enabling users to 'cut-and-paste' portions of images onto different parts of the display screen. Sometimes images can also be drawn down into desktop publishing routines for editing prior to incorporation into documents. By way of illustration, Figure 1 overleaf is a block diagram of a composite DIP system, showing the kinds of features and facilities that might be included at each stage of the document-handling process.

Applications of DIP include existing paper archives which have to be retained for long periods for legal or other reasons, and for which manual rekeying would take far too much time and be much too expensive. They also include current documents, generated internally or received from outside, which contain signatures, drawings, or pictures, and which cannot be keyed in at all. Before the arrival of DIP, information of this kind, estimated to comprise nearly two-thirds of the paperwork used in business, was partially or completely outside the reach of digital information systems.

Two of the earliest products on the market — FileNet Corporation's Document Image Processor (seen during the Foundation's US Study Tour in 1984), aimed at records-management applications, and Philips' MegaDoc , intended for high-volume business applications — use special-purpose workstations with high-resolution windowed displays. Other suppliers, such as Xionics and Rank Cintel, are developing DIP systems that can be accessed from or can coexist with personal computers. Figure 2 on page 3 shows Rank's Videomicrographics system sharing a local area network with PCs.

A number of technological developments underlie the arrival of products such as these. WORM (Write Once Read Many times) optical-disc technology provides low-cost, high-capacity storage for large files of images, at a cost per unit at least an order of magnitude lower than magnetic disc; data compression techniques reduce storage and transmission demands; image-enhancement techniques improve the quality of stored images; and laser printing provides high-resolution hard copy.

So far, most applications of DIP (or at least those that have been publicised) have been implemented to meet specialised needs. For example:

- The US Library of Congress started a pilot project in 1982, aimed at storing the 80 million items it holds, including print, photographic, and audio material, on optical disc.
- In the United Kingdom, the Hertfordshire Police Force has installed a computerised system for storing and retrieving full-colour photographs of convicted criminals, again using optical disc.
- Patent offices in North America and in Europe are experimenting with DIP systems to automate patent searches.
- In the United States, the National Archives and the Internal Revenue Service between them

Figure 1 A composite document image processing system

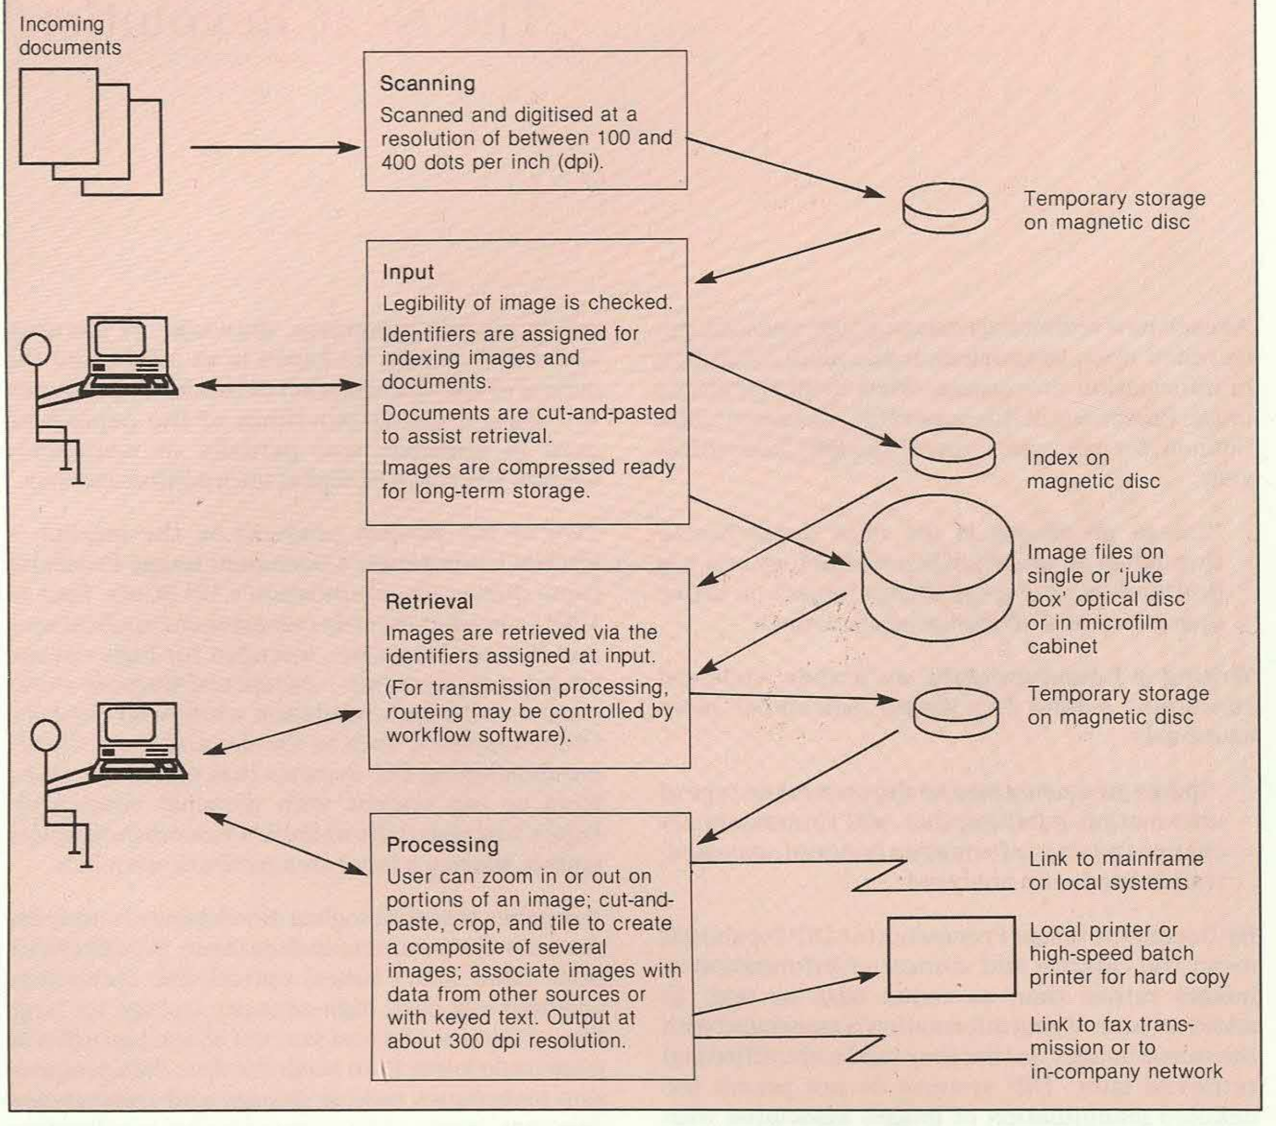

store tax returns taking up five million cubic feet of space, and the Files Archival Image Storage and Retrieval Optical Image Project *(sic!)* aims to reduce the $33 million spent per year on storage and transport.

But there are already signs that DIP is of interest to a wider community — not just libraries, government paper factories, and police forces. What is more, it could be another technology, like the personal computer, that takes the user world by storm, threatening to bypass the information systems function completely. Indeed, most of the DIP activity and interest to date has come from end users and organisations and methods people, not from systems specialists.

British Petroleum (BP) has done more than most business organisations to assess the potential of DIP, looking particularly at applications for low-cost systems, and it has recently initiated a joint venture with Xionics to produce a system linked with personal computers. Peter Taylor, who represents BP's interest in the project, is quoted as saying:

> "Unless management services get ahead and know where they are going, we will have problems holding off user demands."

So, is this a real threat or another false alarm? Will the promised Document Image Processing revolution really happen, or will it, like other strongly tipped technologies before it — voice/data integration, facsimile, and voice processing, for example — never quite make the big breakthrough? Is DIP of interest principally to the paper-handling specialists, such as libraries or patent offices or internal revenue services, or is it something all information systems managers should be taking notice of?

My perspective on this question is that of a long-time observer and analyst of best management practice in the field of information systems, and particularly

Figure 2 Rank Cintel's Videomicrographics system sharing a network with PCs

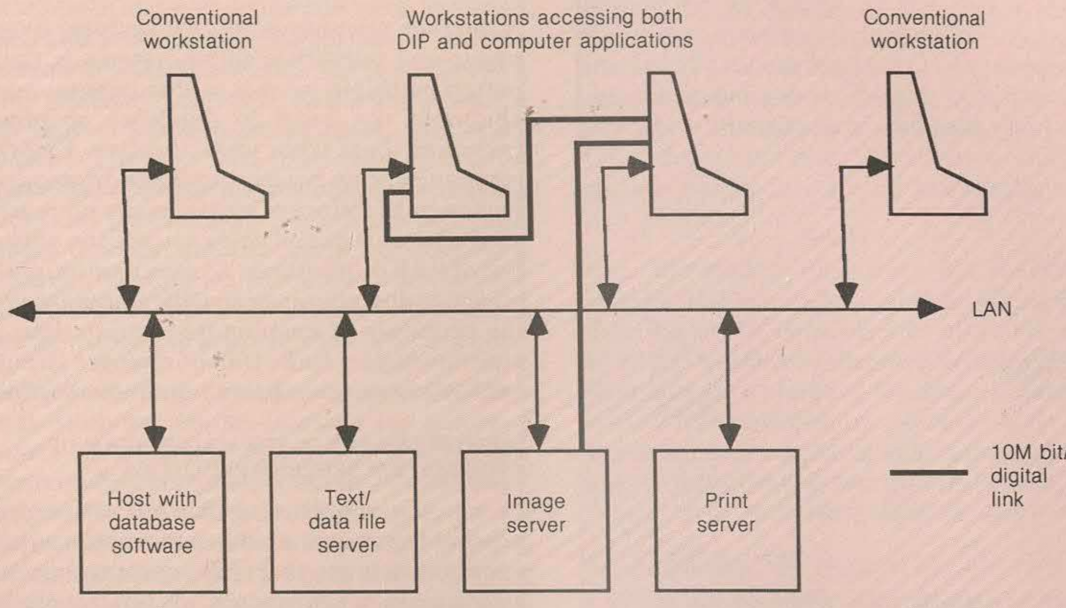

of those technologies that move systems out of the control of the specialist and onto the end-user's desk. My answer to the question is that Document Image Processing is not just for libraries and patent offices. We should all be taking notice. It will bring about a revolution, but probably a quiet and gradual one, just as 'office automation' is doing, and that means it must be approached with care and with determination.

In the remainder of this paper I explain how I came to that conclusion, and then elaborate on how managers responsible for information systems must prepare themselves.

WHY DOCUMENT IMAGE PROCESSING WILL FIND ITS PLACE

DIP systems have four characteristics that suggest they will indeed have a major impact on information handling in the modern business. In the first place, by tackling the considerable problem of working with paper files in today's office, they promise to deliver substantial business benefits. Second, they are taking up a position right in the mainstream of the evolution of information systems. Third, they have the hallmarks of a successful technological solution to the business problems they address. And fourth, technology is moving rapidly to remove the cost/performance limitations of current equipment.

BUSINESS BENEFITS ARE THERE FOR THE TAKING

DIP systems address the problem of working with large volumes of information that cannot satisfactorily be coded and captured as data or text. At present, this type of information is normally held in four-drawer filing cabinets, in filing racks, or in dusty basement archives. Paper filing systems are labour-intensive, and the problems of working with them are obvious — they include high risk of misfiling or loss; lengthy retrieval times; and physical bulk. Sometimes, to save space, paper files are microfilmed. Microfilm systems cope with the space problem but, because of their manual nature, tend to exacerbate the problem of retrieval, because they reduce the number of access points to the information.

Seen within the domain of information storage systems as a whole, DIP occupies a space between, on the one hand, the paper and micrographics systems suitable for information that is rarely retrieved and, on the other, the transaction-processing systems that can handle high retrieval rates but only limited storage volume. Figure 3 on page 4, showing how FileNet Corporation saw the market for its own DIP product, illustrates this point.

The spread of technology out into the workplace has highlighted the paper-handling problem. I have come across several managers and professionals who key (or have their support staff key) selected information from incoming correspondence into their desktop systems for later analysis or retrieval. I know many more who would dearly like to bridge the gap between their existing data/text-handling systems and incoming or stored correspondence, if only they knew of a practical, affordable way to do it.

Information systems people, I believe, are much less aware of the magnitude of this problem than their colleagues who specialise in records management and than end users themselves.

Figure 3 The market for FileNet's product

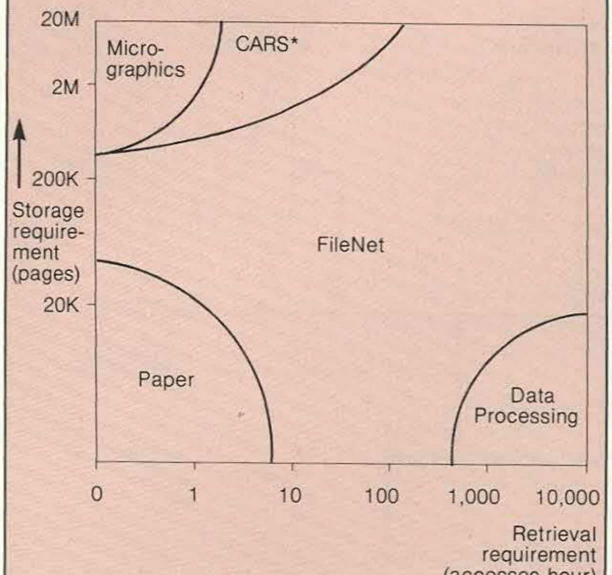

*CARS = Computer aided retrieval system

Hence there is a risk that they will underestimate the potential of DIP systems.

DIP systems can contribute to corporate performance in two ways. The most obvious is by directly attacking the costs and the drawbacks of paper-based filing systems. In the United States, according to the Association of Information Systems Professionals, 70 per cent of the costs of office filing are incurred on salaries, 20 per cent on space, five per cent on equipment, and five per cent on supplies. A number of companies have done studies that show that staffing and storage costs could be drastically reduced by introducing the new DIP technology. Engineering departments can typically achieve a two-year payback on DIP systems installed to handle drawings, while for normal business documents, systems would pay for themselves in three to four years. DIP systems, clearly, will also minimise the problems of loss and misfiling to which paper-based filing systems are prone.

But DIP systems are much more than a means of reducing the costs of paper filing systems. They also improve access to information, and it is from this, I believe, that the most valuable benefits will derive. By eliminating the delays associated with manual filing systems, they give a competitive edge to companies, such as those in banking or insurance, that rely on the speed and quality of their response to customer demands. Using a DIP system allowed a trust department in a US financial services company to double the number of customer enquiries it could handle before the income tax filing deadline, and the international division of a bank reduced the float associated with sending documents to correspondent banks throughout the world.

A further advantage of DIP systems over both paper and microfilm filing systems is that they permit multiple access to documents, and thus eliminate the need to duplicate or physically transport them from place to place. Engineering companies often need to duplicate drawings so that engineers in different locations can have access to them. Extra cost and delay are not the only penalties of such duplication. As soon as multiple copies of design documents or specifications are created, the problems of keeping them up to date and of maintaining an audit trail of changes (frequently needed to meet regulatory requirements) multiply.

DIP SYSTEMS ARE IN THE MAINSTREAM OF INFORMATION SYSTEMS EVOLUTION

To put the opportunity that DIP offers in more general terms and at the same time support my second point — that DIP systems are in the mainstream of information systems evolution — I need to review, briefly, the history of what was originally termed 'office automation', but which I prefer to call 'end-user systems'. By end-user systems I mean all information systems where end users (or their managers) have considerable discretion about whether and how to exploit the computing power at their disposal. These include personal computers and work-group systems (defined below), as well as office automation. Many types of DIP system will fit into this category as well. End-user systems may be contrasted with operational systems, which process business transactions or control production processes, where use is clearly not discretionary.

The evolution of end-user systems so far mirrors the changing role, or perhaps our changing perception, of the office. Initially, office automation equipment was aimed at assembly-line tasks such as typing. It focused heavily on support staff, whose activities could be analysed and systematised (it was thought) in the same way as factory processes can. Later, with the realisation that three-quarters of office costs were accounted for by professionals and managers, office automation moved to eliminate their so-called 'unproductive' activities. Results, however, were disappointing. Numbers of support staff were reduced, but neither easily (more often than not, they had to be reorganised into fewer, but larger, support units) nor dramatically. The big productivity gains for professionals and managers did not often materialise, and, even where they did, it was not clear that the office was contributing substantially more to business success than it had in the past.

In several telling cases, however, the really worthwhile gains came from quite unexpected quarters. A UK engineering company provides a typical example. The firm installed 32 advanced workstations to improve communication between directors and senior management and to contain administration costs. In practice, the first objective was only partially achieved, and administration savings only covered running costs. The main business benefit was one that had not been anticipated at all. Engineers responsible for making estimates and preparing tenders began using the equipment directly, and cut tendering times by a large margin. Equipment installed to streamline administrative processes had in the event enabled knowledge workers to improve their effectiveness.

Since 'office automation', we have experienced the personal computer boom. Demand for personal computers was driven mainly by knowledge workers requiring personal productivity tools. Now companies are installing work-group systems, combining office services such as word processing and messaging with support for personal computing and for departmental applications. Sometimes, these work-group systemd do give a clear payoff in straight cost terms, but more often the gains are less tangible. Lead times to produce documents or respond to service requests are cut, sometimes dramatically; staff morale rises and staff retention is improved; crises are handled much more effectively; and so on.

The fact is that the view of the office principally as a factory that deals with business transactions (rather than products) is misleading. This aspect of its role has been and is still being displaced by data processing systems and by trading networks. The crucial role of the office is to act as a link between the unpredictable, volatile demands of customers, competitors, suppliers, and regulatory bodies and the more stable production and administrative processes of the organisation. Rather than regularity and rigid control, the office needs flexibility and spontaneity to react intelligently to the demands placed on it. In short, it needs to act as an exchange for the knowledge on which the success of a business depends, rather than as a transaction factory.

End-user systems are powerful catalysts in changing the role of the office. There is now growing evidence that the real payoff from end-user systems is at departmental or, more accurately, work group level. (I prefer to talk of 'work groups' rather than 'departments', because the latter term describes organisational structures which information technology is calling into question.) It arises principally from integrating work patterns that have become fragmented. Over the years, principles of organisation derived from the factory, such as the division of labour, have been applied in the office. The resulting specialisation in narrowly defined tasks has created enormous problems of communication and coordination in large companies. In *In Search of Excellence,* Peters and Waterman give the example of a company that needed 223 formal linkages between organisational units to launch a new product. Taken by itself, they say, each linkage made perfectly good sense, but the outcome did not: "Needless to say, the company is hardly first to the marketplace with any new product."

Data processing, of course, has recently driven fragmentation a stage further by taking structured data out of the control of its original owners and storing it on shared computer systems. The task we face now is to preserve the gains of computerisation, while undoing the damage both of partial computerisation and of poor job design. Success will lead to gains in business responsiveness, and in the productivity and the quality of office work.

So where do DIP systems fit into this scenario? Potentially at least, they bring the domains of external correspondence, of technical documents, and of paper archives within the reach of information systems in the workplace. How important it is to integrate these record systems with existing information systems depends on the nature of the business operations. In some organisations, the pressures for better management of information of this type come from customers or trading partners for whom paper is the preferred, or the legally required, means of communication. Many organisations are under growing regulatory pressure to improve their internal records management. Others will need to improve internal coordination by providing easier access to such information as markets grow more demanding. For one or more of these reasons, DIP systems will play a central role in the process of re-integration on which business organisations are now embarking.

DIP SYSTEMS HAVE THE HALLMARK OF A SUCCESSFUL SOLUTION

DIP systems also have the hallmarks of a successful technological solution to the business problems they address. As the cost/performance ratio of equipment continues to improve, the rate at which organisations can assimilate information technology and gain benefits from it is no longer determined principally by whether or not the relevant equipment and software are available. The ability of users to come to terms with and gain value from the technology becomes the key limiting factor. From this perspective, DIP systems have two signal advantages.

First, DIP systems are like 'analogue software', in that they present the user with an analogue of operations with which he or she is already familiar. Spreadsheet packages are perhaps the most successful example of analogue software, mimicking the drawing of a table of figures on a sheet of paper. Word processing, similarly, gets as close as it can to typewriting, but, of course, with all kinds of added advantages. And windowing software, mimicking on the screen the shuffling of papers on a desk, is fast becoming a 'must have' for personal computers.

Analogue software is not just a frill. Research has shown that many people find it difficult to come to terms with an information system unless they can construct in their heads some kind of model of what is going on, unseen, behind the screen. Many non-specialists cannot easily construct such a model in the same terms that computer specialists can. DIP systems, because they provide an electronic analogue of paper documents and paper filing systems, will be easily accessible to a much wider range of endusers than the systems based on structured data and text that are prevalent today.

The second advantage that DIP systems have is that, as well as streamlining business processes, they enhance human capabilities, enabling their users to search for and associate information directly and much more powerfully than in the past. Based on my observation of a large number of work-group systems, of which DIP systems are a particular type, I conclude that the processing functions that they provide are only taken up and assimilated into work patterns where they add real value for the individual concerned. Merely to provide a lower-cost substitute for existing procedures or services, without adding extra value, simply is not enough to get most users — any but the 15-20 per cent minority of enthusiasts — over the initial learning hump and hooked on regular use. DIP systems are obviously capable of providing that essential added value, if targeted correctly, and because of this they will be welcomed enthusiastically by a wide range of users.

COST/PERFORMANCE IS SET TO IMPROVE RAPIDLY

Today's DIP systems have limitations, both in cost/performance and functionality, but these limitations are being addressed and will quickly be overcome.

Costs are already coming down rapidly as suppliers move up the experience curve with optical storage technology. The price of Toshiba's image file server (the Tosfile 550 based on compact disc technology with a capacity of up to 80,000 A4 pages), for example, originally \$50,000, has halved from the first to the second generation and is expected to decline further to \$15,000.

Volume production will bring further cost reductions, and here the attempts by a number of suppliers to link DIP into standard PCs will be significant. High-resolution workstations are still about an order of magnitude more expensive than standard PCs, but the gap is beginning to close. Desktop publishing has already created a market for add-on high-resolution displays, and the next generation of PCs with 32-bit processor chips will be capable of driving these without strain.

As far as functionality is concerned, the great weakness of DIP compared with paper systems (and also compared with free-text-retrieval systems) is that information identifying each item has to be explicitly defined, so that it can be held in a separate index, normally on magnetic disc, which can be searched in the conventional way. This, as we see below, is the Achilles' heel of DIP systems, since the utility of any filing system is determined by its retrieval mechanism. Scanning technology that includes an OCR capability will soon come to the rescue here. Scanner products capable of recognising over 100 fonts, aimed at the desktop publishing market, are already available at prices little over \$4,500. Japanese manufacturers are working on incorporating this technology into DIP systems, so that filing clerks can easily enter information identifying pages, items, or illustrations into indexes, while the hardware embeds this same information, in machine-readable form, into the stored images themselves.

We can also expect future systems to provide much more help with retrieval, using expert system techniques to interrogate the user about his information requirements, and then translating the response into the appropriate search commands.

'NO GAIN WITHOUT PAIN'

An important characteristic of DIP systems is that they cannot be introduced piecemeal, unlike personal computers or word processors. Inherently, they are a work-group technology, because their *raison d'être* is to enable people to share document images. The lesson of experience with work-group systems to date can be summed up in the phrase 'no gain without pain'. Invariably, they bring about major changes in procedures and working practices. Usually these changes take place in many small and gradual steps, as end users come to appreciate the capabilities of the technology and adapt their ways of working to take advantage of it; more rarely, they are driven by an individual's vision of a new organisational design that the technology has made possible.

In the case of DIP systems, the pain also derives from a second cause, apart from that of organisational change. This is the pain of learning about a new technology with novel properties, which demands new skills and new expertise from the specialists who must help to introduce it.

MANAGING ORGANISATIONAL CHANGE

Organisational change, as is now well recognised, must be prepared for and then managed. Unless analysts and/or user management build a new vision of the opportunity that DIP opens up for the business, and communicate that vision to the decision makers, the potential of the technology is unlikely to be realised. One manager responsible for introducing DIP into his organisation told me that his senior managers saw the potential of the technology solely in terms of current procedures. They were particularly reluctant to contemplate changes in the main areas of the business, because of the risks involved. This is, of course, office automation revisited. Failure to face up to this difficulty is likely to mean that DIP is limited to marginal applications where its real potential is not apparent.

From the point of view of an analyst responsible for the introduction of DIP, or a user 'champion' keen to exploit it, it is important not to take existing procedures as read. Of course you will not ignore the opportunities to save space or to streamline existing document-management procedures, which are likely to provide a necessary foundation for a business case. But existing procedures are only the starting point, and analysts must also ask more searching questions, such as:

- Can we gain a marketing advantage by improving our speed of response to incoming documents, such as customer inquiries?
- Can we reduce the number of steps in the document-management process, or improve our control over it?
- Can we make high-volume documents more accessible to decision makers and make it easier for them to select the ones they want?

Answers to questions such as these will point to the real business gains, which will in turn elicit management commitment rather than mere acquiescence.

As a work-group technology, dependent on organisational change to realise benefits, there is no low-risk entry point. One way to express this is to say that you cannot introduce a trial system. Your first try must be designed, and referred to, as a prototype. Only if users see the system as a preliminary version of the real thing, rather than just a practice run, do they put in the effort necessary to come to terms with change and make it work. In fact, more than one manager who has managed a successful project of this type has told me that a prerequisite for success is that users should come to rely heavily on the system at an early stage.

UNDERSTANDING DOCUMENT MANAGEMENT

It is very important to recognise that document management and data management require different skills. A number of organisations already have experience of shared filing of documents on office systems. This experience has revealed clear differences between document filing and retrieval (that is, document management) on the one hand, and the much more familiar techniques appropriate for data management on the other. DIP systems will extend and accelerate the move of information systems into document management, making it essential that information systems specialists acquire or develop the necessary skills.

The difference is apparent if you consider the nature of access to data files, as compared with access to document files. A data-retrieval system directly answers an enquirer's question, for example what were total sales of widgets in Brazil last month? Document retrieval, by contrast, is often more indirect; for example what reports do we have that discuss our main competitor's marketing posture? The system responds by finding those documents that *might* satisfy the enquirer. The differences are summed up in Figure 4, taken from a paper by an academic specialist in document management.

Figure 4 Data versus document retrieval

Data retrieval	Document retrieval
Direct retrieval that answers the enquirer's question: typical query is specific ("I want to know X").	Indirect retrieval that provides or refers to a set of documents that may contain what the enquirer wants: typical query is general or topical ("I want to know about X").
Necessary relation between the request and the correct answer (hence, data retrieval systems are deterministic).	Probabilistic relation between the request and a satisfactory answer (hence, document retrieval systems are nondeterministic).
Criterion of successful retrieval: Correctness of answer (Objective: "Does the system answer the enquirer's question correctly?")	Criterion of successful retrieval: Utility (Subjective: "Does the system answer the enquirer's need?")

(Reprinted from 'The management of information: basic distinctions' by David C. Blair, Sloan Management Review, Fall 1984, PP13-23, by permission of the publisher. Copyright 1984 by the Sloan Management Review. All rights reserved.)

In text-based systems, associative and free-text retrieval techniques help to meet the novel requirements posed by document management. Nevertheless, the analyst must still think through (and find out at the prototype stage) just how users will want to retrieve information, so that the most helpful information can be included in the title and other primary indexing fields used to identify documents. In image systems, because all the identifying information must be captured explicitly, this task is vital. The paper referred to above gives an example that illustrates the point:

"... a document retrieval system was developed to keep track of the substantial number of documents (engineering drawings, purchase orders, subcontracts, correspondence, receipts, etc) which were generated during the course of a large construction project. Since the major documents (drawings, orders, receipts, and subcontracts) all had unique numbers associated with each of them, the system designers felt that these numbers should be the primary access points to the documents on the database. Unfortunately, after the system was built, it was discovered that the users could rarely remember (or find) the exact number associated with a desired drawing, order, receipt, or subcontract. In fact, over 80 per cent of the searches were based on subject descriptions — an access point not well developed in the system."

Document-management skills needed to deal with design issues such as these are often possessed by records-management specialists dispersed throughout the business. But these are people who have been brought up in a paper-oriented world. Both their skills and those of information technology specialists must be brought to bear on DIP projects. An educational programme is needed to upgrade skills on both sides and bring them together.

CHOOSING THE RIGHT TARGETS

With advanced office systems, of which DIP systems are a particular type, success rarely depends on getting any single thing exactly right. But it always depends on getting somewhere near the right answer for all three of the elements that make up any information system — information, technology, and people. In other words, the application must be appropriate; the human factors must be taken into account; and the technology chosen must be capable of doing the job. This is a statement of the obvious, but it does provide a useful framework for looking at how targets for DIP systems should be chosen.

APPLICATIONS ARE AS DIVERSE AS INFORMATION NEEDS

Applications of DIP systems vary considerably, both in terms of overall system goals and at a detailed level. Many of us made the mistake of seeing 'office automation' in far too general terms in its early stages. We should not make the same mistake with DIP systems of treating them as general-purpose solutions to paper-shuffling problems. It is significant that system suppliers are devoting much of their present efforts to making their systems easier for users to customise. Philips, in fact, customises all Megadoc installations to individual requirements.

At the system level, we can distinguish several different ways of exploiting the new capabilities that DIP has put on the information systems application menu. These include:

- *Transaction processing* — such as loan approval, where applications for loans would be scanned and stored for access at various stages of the approval process.
- *Archival storage for converted paper files* — such as a registry, or a library of technical documents such as patents or specifications.
- *Networked distribution of incoming printed material* — the BP library is experimenting with a system that will distribute electronically a daily news digest circulated to directors in the UK and overseas.
- *Work-group shared filing* — as well as sharing data and text files, the members of a work group would also be able to share stored images, for example of incoming correspondence or technical drawings.

The detailed differences in applications are illustrated by the use of DIP systems to conduct patent searches in a Patent Office. Depending on the type of patent, the searchers look for different things within the patent document. If it is a chemical patent, for example, the searcher looks for chemical formulae. This affects how the document is prepared for storage — what items are cut-and-pasted onto the first page; what key information is entered; and also the search strategy used to retrieve images for examination during a search.

The wide variety of information needs means that care must be taken to decide which information should be stored and for what purpose, just as with a data management system. I once heard an IBM marketing man argue that all information should be archived as a matter of course, because electronic storage was cheaper than the manpower needed to decide what to store and what to discard. He was speaking about data and text, and now DIP systems also make it possible to store image information electronically with relatively little human effort. Leaving aside the point that, as an equipment supplier, IBM has a vested interest, the argument is fallacious because it ignores the fact that information is only stored in order that it can be retrieved. Unselective storage makes retrieval more difficult, since it obliges searchers to fight their way through redundant information to find what they want.

Selectivity and accessibility of information are the keys to the business benefits of DIP systems. If we are not careful, we will make this the information retention age, rather than the information age.

PEOPLE ARE THE KEY CONSTRAINT

More and more, the constraints on the success of end-user systems, including DIP systems, are human — how well these systems fit into existing human systems; how usable and accessible they are; how easily people affected by them can cope with change, and so on.

While the detailed 'human' qualities of a system are important, the *sine qua non* is that the system should add real value as far as the user is concerned — make the job easier or, better still, enable him or her to do it better. It is this that (combined with effective project management — too big a topic to take on here) generates the commitment on the part of users without which the real gains cannot be secured.

To justify DIP systems you will certainly need to show that they can displace costs, but this should not be the prime criterion by which you select targets. You should look first and foremost for areas of the business that are not cost-sensitive and whose paper filing procedures are inadequate, so that end users experience positive benefits from their use of the system. They will certainly lose flexibility compared with paper-based filing, a fact that they will be acutely aware of in the early stages of implementation, when the disadvantages of the new system will be more apparent than its advantages. The new system must give them enough added value to sustain their interest and commitment while they are learning about those advantages.

While the main business benefits are likely to derive from faster or more effective retrieval of stored information, you cannot afford to forget that everything depends on information being captured in the first place. If the capture procedures fail to provide the right key information, either because requirements have not been identified or because the staff responsible get bored, the system will fail. At ICI Mond there is a successful work-group system based on a Xionics workstation network. The only application to fail among a number of successes was an attempt to automate the registry. One of the prime reasons for the failure was that document-capture procedures were unwieldy and came to be seen as an unpleasant chore. In fact, DIP systems can and should take the drudgery out of filing procedures, by providing the type of sophisticated cut-and-paste features available in desktop publishing packages.

Some aspects of the task can be eliminated altogether — for example, where the system supports the equivalent of 'style sheets' provided in advanced word processing packages. The style sheet for a particular type of document would specify exactly how it was to be manipulated prior to filing. The filing clerk then need only indicate what type of document is being scanned, and the system does the rest.

RESPECT THE LIMITS OF TECHNOLOGY

Paper is good for imprecise manipulation of information, which is typical of much of the knowledge work carried out in the office today, and that is why people like it. DIP systems will bring more power of retrieval and of association to the knowledge worker's elbow, but they also take away some of the flexibility of working with paper. While seeking to exploit the new power which the technology makes available, you must also be aware of its limitations.

As indicated above, the image-capture stage, which sets up the identifying information used for searching, is potentially the Achilles' heel of the entire system. Paper file-search strategies, honed over many years of practical experience, can be very slick indeed (and, at the other extreme, can also be hopelessly cumbersome). Unless the system can come close to matching the filing and retrieval processes it displaces, it is unlikely to gain support from users, whatever the compensating advantages it can offer for those who do succeed in finding what they want.

At the retrieval and processing stage, DIP systems, as emphasised earlier, have the advantage of providing a familiar analogue of paper documents and of paper filing systems. But the completeness and accuracy of the analogue is sharply limited by the size and resolution limits of the display. At most, this will constrain the user to a limited view of perhaps three or four documents at any one time, far fewer than can be spread out for examination on a desktop. It is also much more comfortable physically to switch the eyes from individual items on a page to the context in which they are set, than it is to zoom in and out as some DIP systems permit. Except where high-resolution displays are being used, zoom is of dubious help because screen quality is so much lower than print quality. Compare the 1,800 x 2,400 pixels needed to display an A4 page at a resolution of 200 dots per inch with the 640 x 480 pixels provided on an upmarket personal computer such as the Apple Macintosh II.

Storage presents the familiar problem of backup and security, especially if DIP systems are distributed. (As touched on next, the cost of transmission of image information between image file servers and workstations will favour distribution of storage.) The backup problem is magnified by the very large volumes of information that can be held on optical media — roughly 600M bits on a 12-cm compact disc.

ARCHITECTURE — THE 'CHINESE WALLS' POLICY

Any newly arrived technology brings with it the threat of future incompatibility as standards evolve, and DIP is no exception. At this stage, there are no universal industry standards for document formats or for document interchange. Facsimile standards are already being adopted for DIP. IBM's influence cannot be discounted — DCA and ODA standards envisage compound documents including image. The X.400 'open systems' standard for electronic messaging is also relevant.

What is more, DIP systems will certainly provoke a rethink on network structures, because of the enormous additional traffic load they can be expected to create. Using the best compression techniques currently available, a one-page image document occupies 250,000 bits at a typical scanning density. This is at least an order of magnitude greater than a page generated by a word processor (see Figure 5). To provide the speed of response users are likely to expect from an electronic system (much faster, of course, than they put up with when their secretaries retrieve the same documents by hand), line speeds in excess of those now normally provided are certain to be needed. LANs will likely be needed on the site, and high bandwidth links between the sites. And transmission is only part of the problem. Experiments at BP to transfer images from an IBM mainframe to PCs via the corporate network, using standard IRMA cards, ran into serious performance problems — 15 minutes to transfer one page.

Figure 5 Size comparison of text and image traffic

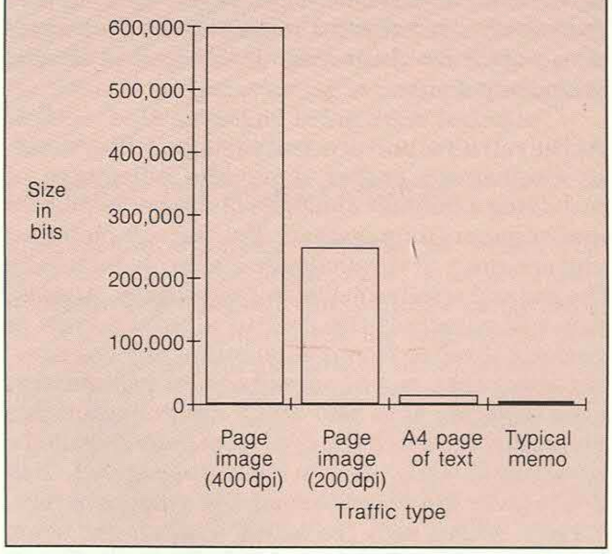

Unfortunately, despite these problems, you may have no choice but to proceed, or risk losing the initiative to the users. Offering the advice, "We think you should wait until the standards firm up," or "We need a year or two to get the network up to scratch" is likely to cut little ice with a manager with a serious records-management problem who has just seen an impressive demonstration.

It will certainly be some time before the full implications of the arrival of Document Image Processing are apparent. Some hard thinking will be needed to decide where image files should be located, to obtain the best trade-off between cost, performance, and security. I share the view of Mike Bevan, managing director of Xionics, that, for a work group sharing files, small distributed clusters of high-capacity storage devices are likely to prove the most attractive solution, serving work groups or departments, possibly with backups organised from a central point. (Figure 6, opposite, taken from his paper, illustrates one way of achieving this.) Meanwhile, how can you find a sensible way forward?

A certain amount of muddling through is unavoidable, but what I call the Chinese walls policy helps to limit the muddle. (Chinese walls are imaginary barriers between departments in financial services firms, erected to avoid potential conflicts of interest between business activities.) What you seek to do is to separate image traffic from existing network traffic, until you are in a position to integrate those traffic streams properly. You can do this, for example, by carrying image traffic on separate lines, which converge at the workstation (see Figure 2 on page 3).

You also keep the coupling between DIP systems and other systems as loose as possible, for example by dumping data files that are accessed from DIP systems out into the DIP system or a separate server, rather than handling data-access requests online. With users able to buy inexpensive add-ons to their PCs, this policy may be difficult to enforce completely, but it is considerably better than nothing.

CONCLUSION

With the arrival of Document Image Processing, information systems managers once again find themselves caught between a rock and a hard place, as they did a few years ago with personal computers. Pour cold water on users' enthusiasm for a new and attractive technology, and you reinforce the reputation of the information systems group as inward-looking and conservative. Encourage them, and you put further strain on your skilled people and risk creating new support and interconnection problems further downstream.

What is different now is that we have been through the personal computer experience and understand better how end-user systems should be managed, and how the partnership between end users and

Figure 6 A possible architecture based on small distributed clusters

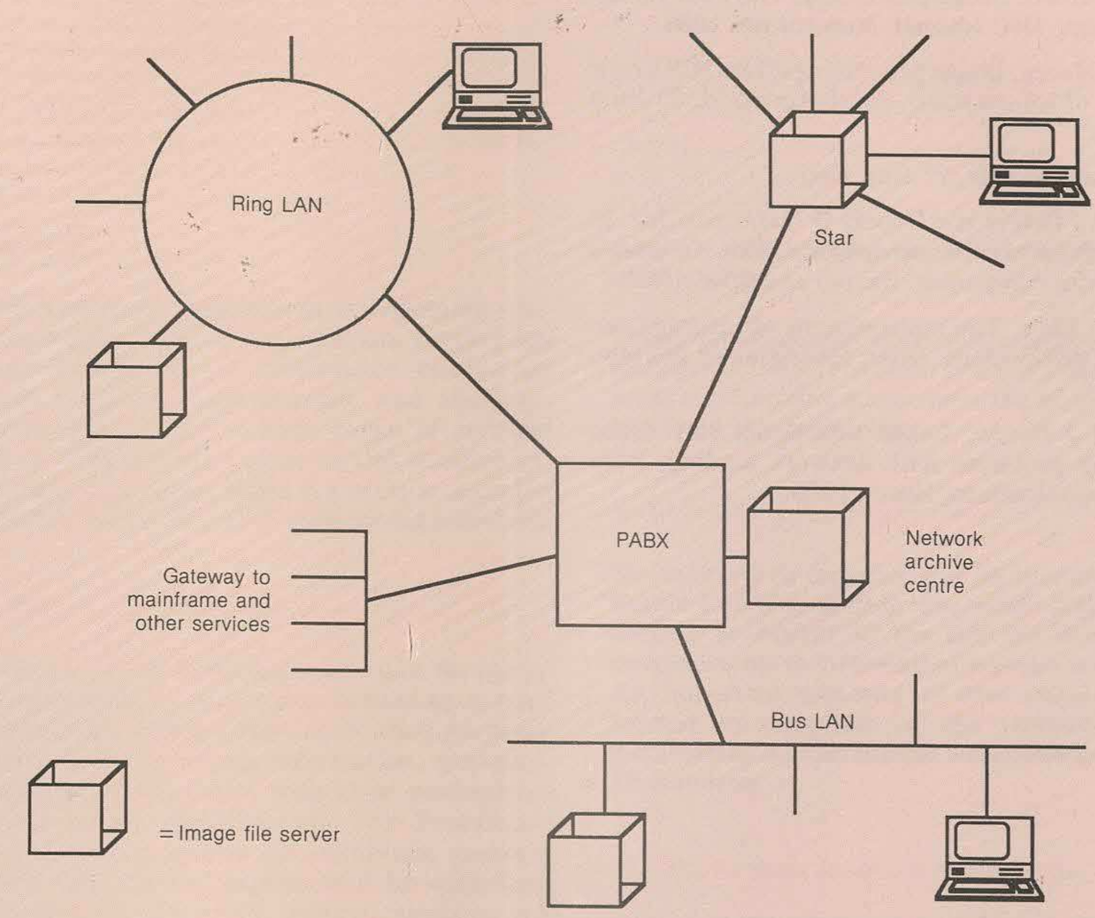

(Source: Bevan)

information systems specialists should be organised. In this short paper, I have tried to put across the main lessons arising from the experience with end-user systems of the past seven or eight years, as these apply to DIP systems. None of these lessons are new, but nonetheless they are more often honoured by being ignored, perhaps in the heat of battle, than by being observed.

To join the battle and then fight poorly is the worst possible course to take. Information systems directors must make a judgement, and jump one way or the other. Either, because of the nature of your business and your current systems strategy, DIP is important, or it is not. If it is important, you must: invest in the skills which you will need to succeed; go out and choose the targets before they choose themselves; and manage the changes which will deliver success. In short, be positive and proactive in ensuring that your organisation grasps the opportunity.

Alternatively, you may come to the conclusion that DIP is not important enough to displace other priority tasks, and that your scarce skills would be better deployed on those. In this case, your best option is to make clear to end users that they must take the initiative and the risk themselves, specifying what level of support they can expect from the specialists. In addition, of course, they must be warned of the dangers and given guidelines designed to keep future support and integration problems to a minimum.

If you are not in a position to pursue either of those alternative courses of action, then you need a wiser person than me to advise you. Pushed by the technology and pulled by the growing difficulties of working with all those paper files, DIP is unstoppable.

Bibliography

John Connell. Image processing: The next breakthrough. IMC Journal, March/April 1986.

Connie Moore. Image processing offers MIS a new view of information. Computerworld, 23 June 1986.

Computer Weekly, 17 July 1987.

Thomas J Peters and Robert H Waterman Jnr. In search of excellence: lessons from America's best-run companies. Harper and Row, 1982.

David C Blair. The management of information: basic distinctions. Sloan Management Review, 26 (1).

Michael J Bevan. Image processing may cause future problems with network loading. Data Communications, March 1986.

BUTLER COX FOUNDATION

Butler Cox

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

Objectives of the Foundation

The Butler Cox Foundation is a service for managers responsible for information technology in large organisations. It helps them to do their job more effectively by providing information, guidance, stimulation, and contact with their professional peers in other organisations. The Foundation provides this assistance by publishing research reports and position papers, and by organising international conferences, national meetings, and study tours.

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

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

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

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

Membership of the Foundation

The majority of organisations participating in the Butler Cox Foundation are large organisations seeking to exploit to the full the most recent developments in information systems technology. An important minority of the membership is formed by suppliers of the technology. The membership is international with participants from 19 countries.

The Foundation Research Programme

The research programme is planned by Butler Cox in consultation with the member organisations.

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

Position Papers

In addition to the research reports, the Foundation also publishes position papers. These papers are based on the views and personal research of their individual authors. Typically, four position papers are published each year.

Butler Cox & Partners Limited
Butler Cox House, 12 Bloomsbury Square,
London WC1A 2LL, England
☎ (01) 831 0101, Telex 8813717 BUTCOX G
Fax (01) 831 6250

Belgium and the Netherlands
Butler Cox BV
Burg Hogguerstraat 791
1064 EB Amsterdam
☎ (020) 139955, Fax (020) 131157

France
Butler Cox SARL
Tour Akzo, 164 Rue Ambroise Croizat,
93204 St Denis-Cedex 1, France
☎ (161) 48.20.61.64, Fax (161) 48.20.72.58

Germany (FR)
Butler Cox Deutschland Ltd.
Richard-Wagner-Str. 13
8000 München 2
☎ (089) 5 23 40 01, Fax (089) 5 23 35 15

United States of America
Butler Cox Inc.
150 East 58th Street, New York, NY 10155, USA
☎ (212) 486 1760

Australia
Mr J Cooper
Consultants (Computer and Financial) plc Australia
Level 5, 303 Pitt Street, Sydney 2000, Australia
☎ (02) 2870400, Fax (02) 2870450

Ireland
SD Consulting
72 Merrion Square, Dublin 2, Ireland
☎ (01) 766088/762501, Telex 31077 EI,
Fax (01) 767945

Italy
SISDO
20123 Milano, Via Caradosso 7, Italy
☎ (02) 498 4651, Telex 350309, Fax (02) 481 8842

The Nordic Region
Statskonsult AB
Stortorget 9, S-21122 Malmo, Sweden
☎ (040) 1030 40, Telex 12754 SINTABS

Spain
Associated Management Consultants Spain SA
Rosalía de Castro, 84-2°D, 20835 Madrid, Spain
☎ (91) 723 0995