

# Collaborative Tools and the Practicalities of Professional Work at the International Monetary Fund

*Richard Harper<sup>1</sup> & Abigail Sellen<sup>1,2</sup>*

<sup>1</sup>Rank Xerox Research Centre (EuroPARC)  
61 Regent St.  
Cambridge, CB2 1AB, U.K.  
<surname>@europarc.xerox.com

<sup>2</sup>MRC Applied Psychology Unit,  
15 Chaucer Rd.  
Cambridge, CB2 2EF, U.K.

## ABSTRACT

We show how an ethnographic examination of the International Monetary Fund in Washington, D.C. has implications for the design of tools to support collaborative work. First, it reports how information that requires a high degree of professional judgement in its production is unsuited for most current groupware tools. This is contrasted with the shareability of information which can 'stand-alone'. Second, it reports how effective re-use of documents will necessarily involve paper, or 'paper-like' equivalents. Both issues emphasise the need to take into account social processes in the sharing of certain kinds of information.

## KEYWORDS

CSCW, work practice, ethnography, paper documents, groupware, professional work, International Monetary Fund

## INTRODUCTION

Rank Xerox's Research Centre in Cambridge has a long standing history of examining work settings for the purpose of reasoning about the design of tools and technology in support of collaborative work. These tools and artefacts range from sophisticated, computerised 'groupware' to the more mundane use of paper documents.

One set of these studies has looked at the implications of electronic replacement of paper documents [12]. Others take a broader organisational perspective focusing on such things as the organisational context of office equipment design [1], the willingness of professional staff in research laboratories to allow technology to alter their hierarchical working relations [7], and the social factors underpinning the introduction of a network and associated technologies in the British civil service [2].

This paper reports on an ethnographic study of the organisational practices of professional workers at the International Monetary Fund (IMF). One goal of the study was to understand the role of tools and technology in their current work practices in order to understand what changes might be brought about by introducing new kinds of tools and technologies. We believe that these findings have more general implications for the design of technologies in the support of collaborative work. As will become clear, the collaborative tools we will discuss range from tools and artefacts (including paper documents) which support interaction in face-to-face meetings, to those which are intended to provide access to shared information amongst individuals who are not co-present.

## The Nature of Professional Work

Professional workers can be contrasted with secretaries, clerical workers, administrators and technicians, amongst others. They are people who are paid to organise their own work, make judgements and valuations, and who maintain an element of creative control over their own projects. They are of interest to us for at least two reasons.

First, there is the importance of professional workers to organisational activity. It is professionals who are key to decision-making, who have central roles in information production and use, and whose activities are influential in organisational effectiveness.

Second, despite their importance, very little is known about professional workers from a systems design perspective. What is known varies considerably in depth and quality. There is for example, Kidd's remarkable 'The Marks are on the Knowledge Worker' [10] (where knowledge worker is an operationalised term for professional). There, the claim is made that it is the process of making notes and jottings that is key to professional work, rather than re-use of those notes afterwards. This has obvious implications for a whole range of computer devices designed to support or facilitate note-taking. Kidd's study may be contrasted with one which observes that workers 'stack and pile documents on their desks' [13]. In between, one finds the view expressed by the likes of Zuboff [18] and Drucker [4] that manual work is being entirely replaced by professional work. Their claim derives from the belief that information

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of ACM. To copy otherwise, or to republish, requires a fee and/or specific permission.

CHI '95, Denver, Colorado, USA

© 1995 ACM 0-89791-694-8/95/0005...\$3.50

technology enables everyone to access information that they can then process, evaluate and act upon in skilled, that is to say, professional ways. This somewhat gnostic hope is as delightful to contemplate as it is, to date, empirically unverified.

Taken as a whole, it is clear that the work practices of professionals need much more investigation. Theoretical developments from other disciplines, most especially sociology, need to be empirically corroborated, and a great range of empirical material needs to be brought to bear on system design issues.

Needless to say there are many places in which one can examine professional work. We chose the IMF for several reasons: First, nearly a third of its staff have a professional role. Second, these staff are involved in producing highly complex and analytical reports. This work is quintessentially professional. Third, the IMF was willing to allow researchers to examine all aspects of professional work, including the most confidential. It was also willing to allow these examinations over extensive periods of time, enabling the researchers to better understand the nature of professional activity. Thus, in short, the IMF provided an opportunity for a thorough ethnographic examination of professional work. We shall say more about this in the description of the methodological approach.

### The Organisational Setting

The IMF, based in Washington D.C., is a financial 'club' whose members consist of most of the countries of the world. Member countries contribute to a pool of resources which can then be used to provide low interest, multi-currency loans should a member find itself facing balance of payments problems.

The IMF has some 3,000 staff, of which 900 are professional economists. These economists analyse economic policies and developments — especially in the macroeconomic arena. They have particular interest in the circumstances surrounding the emergence of financial imbalances (including those that lead to a balance of payments crisis), the policies to overcome such imbalances, and the corrective policy criteria for making loans. This involves going on 'missions' to the country in question. The resulting assessments and criteria of member countries are contained in documents called 'staff reports' which are used by the organisation's executive board for its decision-making.

### Methodological Approach

The ethnographic research, carried out by Richard Harper, is ongoing. The first stage consisted of six months field work. The purpose of the field work was to understand what ethnomethodologists, following Garfinkel [5], call 'practical reasoning'. Understanding of this, in turn, enables specification of what may be called the 'logics' of organisational action. In this first stage, the concern was the practical reasoning of the IMF's economists: those skills, methods, techniques, and rules of thumb that

enable IMF economists to produce documents which the organisation itself views as adequate. Because these economists are extensively supported in their activities, the study also included examination of the work and practical reasoning of associated administrative, clerical, and support staff.

The field work centred around the 'life cycle' of staff reports, from the first draft (what is called the 'briefing paper' prepared before a mission commences), through the mission process itself, to the post mission review, and then to translation, printing, and circulation processes. This was accomplished by:

- (1) Following a hypothetical staff report around the organisation and interviewing parties that would be involved in its life cycle. Interviews were conducted with desk officers and chiefs who author staff reports, with secretarial and research assistant staff who help in the composition of staff reports, with participants in the review process, including junior economists, and with Front Office chiefs and senior managers (including the deputy managing director). Staff involved in the post authoring and review stages were also interviewed, including the clerical staff who issue and release staff reports once they have been 'cleared', with those who copy and print staff reports, with translators, and finally with archivists.

In all, 138 personnel, including 90 economists were interviewed. These interviews were informal, but consisted of a systematic process of questioning and clarification, whereby the field worker gradually developed a picture of what 'practical reasoning' consisted of in any particular job.

- (2) Observing an IMF 'mission' and its allied document production practices. The field worker observed meetings between the mission team before the mission commenced, between the member authorities and the team during the mission, and observed post mission meetings. All related documentation was made available to the field worker. Interviews were also undertaken whenever possible in which observed parties were asked to explain what they were about.
- (3) Subsequent to the field work, a set of descriptions were generated and circulated around the IMF. Written comments were gathered and discussed with a number of key 'informants', ranging from senior staff, through to junior economists. A final report was made available to the Fund in 1995.

The results of this research are much more extensive than can be adequately covered here. However, we will focus on two interrelated sets of findings which concern the use of tools in professional collaborative work. We believe these findings have implications for the understanding of activities in other settings and for the generic design and evaluation of collaborative tools, which we will describe.

### SHAREABILITY, JUDGEMENT, AND THE DESIGN OF GROUPWARE

Since the invention of organisations, certain types of information have been shared, jointly processed, and used in professional work. With the increasing sophistication and 'user-friendliness' of computer technology, and most especially with the introduction of desktop devices and tools, the capacity for information-sharing was expected to be revolutionised. Not only would organisations be able to provide technological support for what was essentially large scale information-sharing, but now they would be able to support and encourage the sharing of information between individuals and their own, smaller scale, more local information production activities. Thus was born 'groupware' and its associated 'middleware' tools and technologies.

Much of the software originally designed to support groups was designed to provide asynchronous access to the information of others. In other words, it was designed to give access to the work of individuals who, for whatever reason, were not working together in the same physical space at the same time. Examples of this include shared databases and spreadsheets, asynchronous co-authoring tools, and electronic meeting schedulers. Henceforth, it is these kinds of tools we have in mind when we use the term 'groupware' although the term can be and has been used more generally.

Evidence for the success of this kind of groupware has sometimes been contradictory. For example, *Lotus Notes* has been taken up by a great many organisations with the expectation that professional workers would use this software to collaborate more effectively. Some research has found that the use of *Lotus Notes* has not enhanced collaborative work [15], whilst other research has [11]. However, the weight of the evidence suggests that groupware tools frequently fail [e.g., 6]. Numerous explanations have been offered for this. Most have to do with what may be described as social factors: the claim is made that professionals want to hold onto their 'own stuff' so as to preserve an advantage over their colleagues [15]; or that they are unwilling to alter their time-honoured work practices [2].

Our observation of work practice at the IMF, however, shows that differences in the utility of various types of groupware has to do with the type of information intended to be shared within these technologies. Though social factors may be imposed upon it, the fundamental issue is whether information can be shared or whether it cannot. This itself turns on how much individual judgement is used in the creation and management of information.

This can be illustrated by looking at two of the major tasks that professional economists at the IMF carry out: the production of staff reports, and the production of data for a statistical database.

### The Production of Staff Reports

Professional workers at the IMF use a variety of tools in their information work. Nearly all use *Lotus 1-2-3* as their spreadsheet tool, an in-house package called *AREMOS* for time series data, and *Word Perfect* for text generation. A number of in-house applications enable the construction of compound documents. Together, these tools are used to generate information that ultimately works its way into what are called *staff reports*. These staff reports are collaboratively authored, and are used by the management of the organisation to make most consequential decisions.

The introduction of a new computer network at the IMF was expected to encourage professional workers to share more of their information throughout the organisation. It was believed that email, the sharing of data files on servers, and the capacity for remote access to information stores would provide new opportunities for information use that previously existing information access and delivery procedures — namely paper mail systems — made difficult. Such broader sharing of information was expected to help create new forms of group collaboration, ones that transcended currently existing group structures. But, as has been found with professional workers in other organisations [2], such changes do not appear to be happening. This can be explained by looking at the practical reasoning, and hence organisational logic, of what professional staff at the IMF do.

Professional economists at the IMF are each responsible for maintaining data on a member country. On an annual basis, or when a request for financial assistance has been made, these individuals are joined by four or five other economists who will work together on the production of a staff report on that particular country. This is known as the mission process. Work on the mission process is divided amongst economists so that one will collect data for, say, the balance of payments, and another for industrial production.

The collection and analysis of data for *Lotus* spreadsheets or *AREMOS* time series is not simply a clerical matter, however. Economists working on a mission are required to make professional judgements about where to fill in missing figures, inconsistencies or to clarify areas where the data seem muddled. This is a natural feature of current economic information. It is never complete, never certain, and always subject to revisions and amendment (sometimes years after the period in question).

Because this work involves judgement, there are checks and balances to ensure that the judgements are correct. These checks and balances consist of the social process that is mission work. For this process involves economists 'working up' their individual data sets iteratively, by corroborating them against the data 'worked up' by their colleagues. Gradually, a commonly specified set of interpretations is agreed upon by the group. This is used to compile the staff reports.

A corollary of this is that data stored on any individual workstation or PC consist of a mixture of judgement and agreed fact whose meaning reflects that stage in the social process of agreement and iteration. Hence at some point in the mission process, the data are rough and incomplete; the social process of figuring out the data being only begun. At another, later stage, the data are more complete, more effectively understood and developed, the social process of which it is part being nearer completion.

We can begin to see the consequences of this by putting things another way. When a team starts work on a mission, when data has just begun to be gathered, and the first meetings have occurred, data are too rough to be shared amongst the team, although each member will have some understanding of other team members' data. Toward the end of the mission cycle, data can be more readily shared, since by that time data will have been more thoroughly assessed and cross validated.

Therefore, at any specific moment in time, the adequacy of data is only visible to participants in a mission team, since it is only they who are aware of what stage of completion the data have reached, and who understands the exact boundaries between judgement and fact. All others, outside the mission team and outside the mission cycle, will find those data opaque and unsuited for use.

A further corollary of this is that outside of the mission process, when economists are working on their own, the data they store will be very difficult for anyone to use but themselves. And it needs to be remembered that these non mission periods are extensive. (As noted, missions typically only occur once a year). For the rest of the time, individual economists are effectively on their own, even though they work closely with their chiefs (and occasionally other colleagues) on the preparation of a variety of small scale data analyses and commentaries.

Taken as whole, these work practices have important implications for groupware. First, when individual economists are working on their own data stores, their data are unsuited for sharing and general use. Those data have not been through the social processes of validation and assessment. Second, when the data have been through such processes, only those within the coterie of a mission team will be able to know at when the data are usable. In addition, only they will know the ins and outs of the data, the difference between judgement and hard fact.

It is for these reasons that professionals at the IMF have not used the network to share information in new ways. This failure to cross information boundaries is not due to physical problems of information distribution, but because there would be no organisational logic in doing so. In summary, then, there are two main findings:

*Information involving a high degree of judgement in its production is best interpreted by the producer of that judgement.*

At the IMF, individual economists do not use the network to access data created by colleagues working on different countries or in different mission teams. They view those data as being 'individualised' and having a temporally located provenance, therefore being unusable for anyone else's purposes.

*Collaborative processes are required to check the judgements used in the production of professionally assessed information.*

The IMF's staff reports go through extremely complex and elaborate review, revision, discussion and checking procedures before they are used by the management. These procedures involve numerous personnel and several distinct departments. This is because staff reports contain a great deal of information that derives from individual professional assessment. These individual assessments need to be checked by others.

### The Production of Statistics

At the IMF a great deal of effort is put into collecting and publishing statistical data. These data are used for historical analysis but not for management decision-making. Currently, the data are entered into a database known as the 'Economic Information System' running on an IBM 3090 mainframe. This system is about to be replaced with an as yet unspecified server-based environment running on a Novell network.

Irrespective of the system chosen, it is recognised that users will be able and will want to share and access data. They can do so because it is *in the nature of the information itself that it is shareable*. This is because the information that composes the statistical database is strictly only that which derives from standard methods. If any judgement is required to determine vagueness or inconsistency, then those numbers are not added to the database. One consequence of this is that there are numerous omissions in the database. A second is that data can sometimes take years for figures to be agreed and added to the data base. This is particularly troublesome for those involved in policy work and the mission process, since they need up-to-date data. A third consequence, and one we are concerned with here, is that whatever is in this database can be used by anybody. Unlike data which involve judgement, these data are objective and therefore shareable.

Thus, the production of statistics at the IMF stands in contrast to the production of staff reports. For, we find that:

*Information that does not require judgement in its production can be more easily shared than information which does.*

A straightforward illustration of this is the fact that statistical data on the Economic Information System at the IMF are accessed by individuals other than those who created that information.

Also in contrast to the work practices of the producers of staff reports, we find that:

*Social interaction is not as crucial to the sharing of objective information as it is to the sharing of interpreted information.*

Professional staff at the IMF will use statistical data without reference to the producers of those data or to other colleagues. Furthermore, documents that derive from the use of those data (and those data alone) are not subject to the same review and assessment procedures that exist in relation to staff reports. This is because there is no need to discuss, check, review, revise and iterate information. The information in question can stand alone and separate from social processes.

### Implications for Design

That there is potentially only a small role for information-sharing groupware tools in the production of staff reports, and that there could be greater use for them in the production of the statistical database can be attributed to the extent of judgement used in the production of the two different classes of data. The more judgement used in its production, the less likely that conventional groupware for the asynchronous sharing of information will be useful. The less judgement, the more likely. This relationship can be represented by the schematic graph presented in Figure 1.




Figure 1. A schematic diagram showing the relationship between the nature of shared information and the suitability of groupware tools.

This is not to say that information which requires high levels of judgement cannot be supported by technological tools. However, the implications of these findings are that the tools for supporting this kind of process must also support the social processes of collaboration. This rules

out many kinds of conventional groupware technologies which are aimed at information exchange of the work products of others, not the work processes.

However, some technologies are aimed at encouraging and facilitating social processes. Tools to supplement meetings are one class which might be suitable. Professional staff at the IMF spend a great deal of time in meetings and in discussion, exchanging drafts and reviewers' comments.

Media space technologies are another, (i.e., audio-video links in conjunction with access to shared documents). A media space would be suited for those involved in the production of staff reports, but not necessarily useful for those involved in the production of statistics. One wants to support social processes when they are part and parcel of getting the work done, but not necessarily when they are not. (In fact, for a case where a media space was misapplied in this way, see [8])

Finally, it might be objected that the situation one finds at the IMF is unique. In particular, it might be claimed that the information professional workers store and produce does not always involve judgement. Though it is certainly true that the amount of judgement incorporated in the process of creating information will vary amongst different organisations, it must necessarily be the case that all professional work involves some judgement. Otherwise the data collection and production tasks would be a clerical one, and there would be no justification for having highly trained, expert professionals. It is in the nature of all professional work that it involves judgement of one kind or another. Further, the social processes of review can be seen across a variety of organisations: legal briefs are reviewed, medical assessments checked, and auditors' evaluations corroborated.

### THE NEED FOR PAPER DOCUMENTS

One of the great hopes of office information systems designers was that their designs would lead to the paperless office. To date the attainment of that goal has not been achieved.

One reason for this is technical: these systems have been characterised by problems of incompatibility, incomplete WYSIWYG editors, and inflexibility, amongst other things. Observation of work practices at the IMF confirms the existence of such problems with their office information systems. However, we can expect that the development of integrated compound document technology (text plus graphics), middleware, and 'open' systems will eventually go some way towards solving these problems.

Another set of difficulties relates to what is sometimes called the cultural preference for paper. It is argued that people like paper, it is familiar to them and they can see no reason why they should stop using it and turn to electronic document forms. This is the cultural inertia argument.

However, our observations suggest that the primary reason for the persistence of paper is that electronic office information systems do not effectively replace the functionalities or affordances of paper. We will argue that users prefer to use paper at certain points in the document life cycle not as an issue of cultural preference, but because these functionalities interact with organisational logic. We will describe an example from the IMF in which the need to re-use documents for ad hoc and unpredicted purposes requires that professional staff re-specify the 'recipient design' of those documents. This process is one which is well supported by the use of paper.

#### Document Re-Use at the IMF

Professional staff of the IMF produce a whole range of documents with substantive information content. All of these are composed with specific reference to the ultimate user(s) of those documents. Thus, staff reports include all the information that is necessary for IMF executives to make their decisions; briefing papers for a managing director prior to a meeting include only those issues that are pertinent to that meeting; documents for public consumption are designed for a general audience, where the issues of concern are neither politically sensitive or still subject to analytical investigation. Thus, all important documents are 'recipient designed'.

This is important in a setting where there is a great proliferation of documents, each designed for particular sets of purposes and/or audiences. At the IMF, over 70,000,000 pages of documents are copied a year. The executive board receives over 4500 documents, subdivided into a dozen main categories, including approximately 350 staff reports.

However, an important characteristic of the IMF's work practice is that it can also react to unexpected circumstances, and it can do so in part because it can re-use the documents designed for its standard document processes for these unexpected needs.

This effective re-use is achieved by authors redesigning their documents for the new purposes or for the new recipient(s) of their documents. This redesign is especially important in relation to professionally authored documents, since these are likely to contain information that derives from professional judgement. Only the authors will know which information is entirely separate from their assessments and which is not, and thus will be able to assess which of the information is suitable for the unexpected use.

So for example, when members of the IMF's research department wanted to investigate the ratio of military expenditure to private sector growth in underdeveloped economies, they were able to do so, in part, by re-using staff reports. But for this, research department staff needed to get the authors of these reports to explain which of the figures contained in the staff reports derived from calculation and which from professional judgement. These

explanations effectively 'redesigned' the reports for the new, unplanned use.

Furthermore, to ensure that these explanations were provided, members of the research department did not access the staff reports electronically but arranged meetings with the respective authors. During these meetings, they were able to outline what they already knew and what they had to find out, and, on this basis, the authors were able to offer the appropriate guidance on using the materials in their reports.

These document-related work practices lead us to three important implications:

*For documents to be re-used for unexpected purposes, the authors often need to be 'in the loop', or directly involved in the process of document modification.*

At the IMF when staff reports are used for purposes other than executive board activities, such as for gathering research material, the authors of those reports prefer to meet the new recipients. This provides them with an opportunity to learn about what the new recipients need, and to offer them guidance on how to use the information contained in the reports. Thus, this process of document re-use is often essentially social.

In being social another implication is made clear:

*To support document re-use, paper documents are preferable to electronic ones. Paper documents support the social mechanisms of document redesign.*

Paper documents can be the focus of a face-to-face meeting, can be placed on a desk in view of all parties and each page discussed in turn, and paper documents can be ritually exchanged once an agreement as to its interpretation has been made. In other words, as Luff et al. have noted [12], paper has an 'ecological flexibility' which allows it to be used as a focus for discussion, and for the co-ordination of social interaction. Luff et al. also point out that paper can be more easily interweaved into ongoing collaborative activity, as opposed to screen-based documents which cause interaction to be more localised and fragmented.

So, for example, at the IMF, when the research department wanted to re-use staff reports, meetings were arranged with the area departments that produced those reports. During these meetings, the respective staff sat around a desk and discussed the paper versions of the reports placed in front of them. They made notes on these documents, flicked through them, and drew attention to certain parts. When the authors believed that the research staff had fully understood what they could use the staff reports for, the paper versions were handed over to them.

Accordingly, the third implication is:

*There will always be a critical role for paper (or technologies with paper-like qualities) in organisational work practice.*

It is important to realise that the above findings will not hold true for all documents. A great many documents can be used again and again for unexpected purposes without the need for authors to be involved. Moreover, certain electronic document tagging applications such as those which use Structured Generalised Mark-up Languages (SGML) can enable users to create their own individually tailored documents. This can be especially useful in relation to such things as technical manuals.

However, for those documents which embody high degrees of professional judgement, the utility of SGML-type applications is reduced. Further, it is difficult to see how such tagging applications could be used to mark out the parts of a document which incorporate judgement from those parts which incorporate just the hard facts. Inevitably these different kinds of information are intertwined so as to convey a cohesive interpretation. This brings us back to the need for keeping the author in the document re-use loop. Only in so doing can correct re-interpretations be devised and agreed.

#### **Implications for Design**

A large proportion of the document life cycle in professional work can be carried out electronically. It is largely assumed that it is electronic tools which enable effective document re-use in allowing documents to be easily modified and reproduced. Indeed, electronic tools do support much of standard document-related work practice. However, for an organisation to be flexible, it must be able to respond to unexpected circumstances, and it must be able to modify existing documents. In the face of unpredictability, it is best that authors be in the loop of re-use — a process which we have argued is most effectively supported by the use of paper.

This has important implications for designers. Above all, it means that designers should not try to obliterate the use of paper altogether but should attempt to understand and preserve some of its important functionalities.

One way to do this is to integrate or 'interface' the paper and electronic worlds. In this way the need for paper and the benefits of electronic document forms can co-exist. Thus far, technological advancements have mostly benefited moving from the electronic to the paper world. By contrast, moving from the paper to electronic forms is still a cumbersome process in its technological infancy, (albeit with some exciting possibilities [14,16]). However, new developments such as 'glyphs' [9] are representative of advancements made in this direction. Glyphs provide support for the effective moving from paper to electronic and back to paper forms, enabling electronic systems to access (or, more accurately, recognise) important formatting, style, and other information. This can go some way to avoiding the pitfalls of some OCR technology which often destroys

just that information that may be crucial to interpreting and hence using some document.

Another possibility is to consider the use of technologies which possess paper-like qualities. So far, most claims to be developing 'paper user interfaces' refer to the ability to use pen input on flat panel, portable displays, e.g., [3,17]. These technologies have the advantage over paper in being re-imageable and perhaps most interestingly, in providing some interactive capacities, enabling producers of text to manipulate and edit on the 'page' itself.

However, the very fact that the displays are dynamic and reimageable may fundamentally alter the ability for these tools to support the kinds of social processes we have been describing. For example, this may detract from the ability for a group of discussants to easily 'walk through' a document together and gain at-a-glance information by spreading it out on a table. Thus the new affordances offered by alternative technologies need to be set against those offered by paper. The ability for new paper-like technologies to support processes such as sharing, talking over, and exchanging documents between one professional worker and another has yet to be demonstrated.

#### **CONCLUSION**

We want to make clear that the preceding discussions and design recommendations are not based on analysis of what is often called 'organisational culture' or, more loosely, 'social factors'. The evidence we have provided turns on the claim that the activities we have described have a fundamental *organisational logic* to them. Professionals prefer paper documents for certain aspects of their work not because they are used to it, but because they afford certain advantages for the achievement of their practical ends. Cultural preferences and other social factors will be superimposed upon this logic, making it obscure to the casual observer. It is only through in-depth, extensive and thorough ethnographic examination that these logics can be uncovered.

Needless to say these 'logics of practice' are considerably more complex and broad than we have been able to convey in this short paper. But those aspects we have presented, relating to the nature of information and the extent to which tools support the social processes involved in the sharing of that information, are we think of fundamental importance in any organisational setting where professional work is undertaken. Therefore, we believe that the recommendations we have offered will have considerable general applicability. If designers take them on board, then we can be confident that we have gone some way towards ensuring that the tools and technologies professional workers will have at hand in the future will be appropriate for their practical requirements.

#### **ACKNOWLEDGEMENTS**

We greatly appreciate the help and co-operation of the IMF in allowing the field work to be undertaken and findings to be published. There were a great number of individuals who deserve special thanks. These include

Chris Yandle, Saleh Nsouli, John Hicklin, Alan Wright, Gertrude Long, and Terry Hill. We also thank William Newman and Paul Luff for guidance and comments on this paper.

#### REFERENCES

1. Anderson, B., Button, G., & Sharrock, W. (1993). Supporting the design process within an organisational context. *Proceedings of ECSCW '93*, (13-17 Sept., Milan), Netherlands: Kluwer, 47-59.
2. Bowers, J. (1994). The work to make a network work: Studying CSCW in action. *Proceedings of CSCW '94*, (22-26 Oct., Chapel Hill, North Carolina), New York: ACM, 287-298.
3. Brocklehurst, E.R., (1991). The NPL electronic paper project. *Int J. Man-Machine Studies*, 34, 69-95.
4. Drucker, P. F. (1973). *Management: Tasks, responsibilities, and practices*. New York: Harper & Row.
5. Garfinkel, H. (1967). *Studies in ethnomethodology*. Englewood Cliffs, NJ: Prentice-Hall, Inc.
6. Grudin, J. (1989). Why groupware applications fail: Problems in the design and evaluation. *Office: Technology and People*, 4(3), 245-264.
7. Harper, R. H. R. (1992). Looking at ourselves: An examination of the social organisation of two research laboratories. *Proceedings of CSCW '92*, (31 Oct. - 4 Nov., Toronto), New York: ACM, 330-337.
8. Harper, R. H. R., & Carter, K. (1994). Keeping people apart: A research note. *CSCW*, 2, 199-207.
9. Johnson, W., Jellinek, H., Klotz, L., Rao, R., & Card, S. (1993). Bridging the paper and electronic worlds: The paper user interface. In the *Proceedings of INTERCHI '93*, (24-29 April, Amsterdam), 507-512.
10. Kidd, A. (1994). The marks are on the knowledge worker. *Proceedings of CHI '94*, (24-28 April, Boston), 186-191.
11. Korpela, E. (1994). Path to Notes: A networked company choosing its information systems solution. In R. Baskerville, S. Smithson, O. Ngwenyama, & J. DeGross (Eds.), *Transforming organizations with information technology*, North-Holland: Elsevier, 219-242.
12. Luff, P., Heath, C. C. and Greatbatch, D. (1992). Tasks-in-interaction: Paper and screen based documentation in collaborative activity. *Proceedings of CSCW '92*, (31 Oct. - 4 Nov., Toronto), New York: ACM, 163-170.
13. Malone, T. (1983). How do people organize their desks? Implications for the design of office information systems. *ACM Trans. Office Info. Systems*, 1(1), 99-112.
14. Newman, W., & Wellner, P. (1992). A desk supporting computer-based interaction with paper documents. *Proceedings of CHI '92*, (3-7 May, Monterey, CA.), 587-592.
15. Orlikowski, W. (1992). Learning from Notes: Organisational issues in groupware implementation'. *Proceedings of CSCW '92*, (31 Oct. - 4 Nov., Toronto), New York: ACM, 362-369.
16. Wellner, P. (1993). *Interacting with paper on the DigitalDesk*. Unpublished Ph.D. dissertation, University of Cambridge, Computer Science Dept.
17. Wolf, C. G., Rhyne, J. R., & Ellozy, H. A. (1989). The paper-like interface. In G. Salvendy & M. J. Smith (Eds.), *Designing and using human computer interface and knowledge-based systems*, Amsterdam: Elsevier, 494-501.
18. Zuboff, S. (1988). *In the age of the smart machine*. New York: Basic Books.