COMIC

ESPRIT BASIC RESEARCH PROJECT 6225

Field Studies and CSCW

Abstract

The thrust of this Strand of the COMIC Project is mainly methodological in investigating aspects of the relationship between the social analysis of work settings and the systems development context. It has been assumed throughout the work that real world CSCW systems will be large scale and, thus, need be produced by appropriate software engineering processes. In addition, it has also recognised that CSCW raises special problems regarding the effective analysis of the social organisation of work settings and that among the more promising methods in this regard, namely ethnography, does not sit easily with many of the current methods of requirements elicitation and system development.

Document ID	D2.2
Status	Final
Type	Deliverable
Version	3.0
Date	October, 1994
Editors	Lancaster University and Manchester University
Task	2.2

Project coordinator:
Tom Rodden
Computing Department
University of Lancaster
Lancaster LA1 4YR
United Kingdom
Phone: (+44) 524 593 823
Fax: (+44) 524 593 608
Email: tom@comp.lancs.ac.uk

The COMIC project comprises the following institutions:
Gesellschaft für Mathematik und Datenverarbeitung (GMD), Bonn, Germany
Risø National Laboratory, Roskilde, Denmark
Royal Institute of Technology (KTH), Stockholm, Sweden
Swedish Institute for Computer Science (SICS), Stockholm, Sweden
Universitat Politècnica de Catalunya (UPC), Barcelona, Spain
University of Amsterdam, Amsterdam, The Netherlands
University of Lancaster, Lancaster, United Kingdom (Coordinating Partner)
University of Manchester, Manchester, United Kingdom
University of Milano, Milano, Italy
University of Nottingham, Nottingham, United Kingdom
University of Oulu, Oulu, Finland

Authors of this report:

John A. Hughes
Department of Sociology
Lancaster University
Lancaster LA1 4YL
United Kingdom
Phone: (+44) 524 594 174
Fax: (+44) 524 594256
Email: j.hughes@lancaster.ac.uk

Wes Sharrock
Faculty of Economic and Social Studies
University of Manchester
Manchester M13 9PL
United Kingdom
Phone: (+44) 275 2510
Fax: (+44) 275 2514
Email: msrssws@cms.mcc.ac.uk

Tom Rodden
Department of Computing
Lancaster University
Lancaster LA1 4YR
United Kingdom
Phone: (+44) 524 593 823
Fax: (+44) 524 593 608
Email: tom@comp.lancs.ac.uk

Jon O'Brien
Department of Sociology
Lancaster University
Lancaster LA1 4YL
United Kingdom
Phone: (+44) 524 594 186
Fax: (+44) 524 594256

Mark Rouncefield
Department of Sociology
Lancaster University
Lancaster LA1 4YL
United Kingdom
Phone: (+44) 524 594 186
Fax: (+44) 524 594256

Dave Calvey
Faculty of Economic and Social Studies
University of Manchester
Manchester M13 9PL
United Kingdom
Phone: (+44) 275 2510
Fax: (+44) 275 2514

Contributions from the work of:

Hans Andersen
Risø National Laboratory
Cognitive Systems Group

Liam Bannon
Department of Computer Science
University of Limerick

Richard Bentley
GMD
Bonn

Peter Carstensen
Risø National Laboratory
Cognitive Systems Group

Val King
Department of Sociology
Lancaster University

Kari Kuutti
Department of Information
University of Oulu

James Pycock
Department of Psychology
University of Manchester

Dave Randall
Department of Interdisciplinary Studies
Manchester Metropolitan University

Kjeld Schmidt
Risø National Laboratory
Cognitive Systems Group

Dan Shapiro
Department of Sociology
Lancaster University

Carsten Sørensen
Risø National Laboratory
Cognitive Systems Group

Mike Twidale
Department of Computing
Lancaster University

Stephen Viller
Department of Computing
Lancaster University

ISBN 0-901800-54-6
Lancaster University, 1994
This report is available via anonymous FTP from ftp.comp.lancs.ac.uk.
Further information on COMIC available from htttp://www.lancs.ac.uk/computing/research/cseg/comic/

Table of Contents

Section 1
Restatement of Themes and Problems 3

The Requirements Problem and CSCW 5
Viewpoints 13
This Year's Work 15

Section 2
The Role of Ethnographic Fieldwork Methods in CSCW 17

The Case for Ethnography in CSCW 19
Moving Beyond Research 21
Ethnography in system design 24
Summary and Lessons Learned 34

Section 3
Fieldwork Studies
Toward Corpus, Codification and Framework 39

Introduction 41
The Presentation of Fieldwork Studies 42
Air Traffic Control 44
Small Office in Training Centre 55
Branch Office in Building Society 62
Manufacturing Engineering 71
The Production of Technical Documentation 80
Technology Centre 88
Software Production 104
Branch Office and Specialised Services Banking 111
Summation 123

Section 4
Perspectives on the Social Organisation of Work 129

Introductory considerations 129
Locality of social organisation 131
Distributed Coordination 137
Awareness of Work 148
Summation 159

Section 5
Developing a Framework for the Presentation of Field Studies 129

Introductory considerations 163
An Overview of the DNP 169
Viewpoints in practice 173
Summary 185

Section 6
Concluding Remarks 191

Bibliography 197

Appendix A
Use, Design, and Evaluation 205

Appendix B
Situated evaluation for cooperative systems 223

Preface

This Deliverable continues the work presented in Deliverable D2.1 of the Comic Project. Deliverable D2.1 was a major ground clearing exercise in order to see more clearly some of the methodological issues in CSCW design, particularly as these impinge on requirements. In particular, we discussed some of the important implications of the 'turn to the social' in system design that CSCW represents, presented an extensive review of the major methods of requirements elicitation and gave a preliminary consideration of ethnography in CSCW.

Our major position throughout the research is to take a pragmatic stance toward the issues of CSCW design, rather than regarding the 'turn to the social' as a panacea for the numerous and varied problems of system design. We are in no doubt that at least some sociological approaches, along with ethnography, have much to offer CSCW, but that much needs to be done to turn these into effective tools for CSCW design.

The major theme of Strand 2 is to explore, with some thoroughness, the problems involved in relating ethnographic fieldwork method to the system design process but not necessarily as the *only* method for supporting CSCW design requirements and system development. Our objective is to align the method with others and, as part of this, examine the place of ethnography within the CSCW design process.

This Deliverable builds upon D2.1 by focusing more directly on some of the lessons and implications of ethnographic fieldwork studies. We also begin a discussion of some elements of a framework for system design as the basis of work to be reported in the next Deliverable. The third theme of this Deliverable is to report on progress made with the development of the Designer's Note Pad (DNP), a support tool for representing and analysing fieldwork materials for CSCW requirements elicitation and system design.

Structure of the Deliverable

Section 1: *Restatement of Themes and Problems*

This section constitutes a brief review of some of the salient issues arising out of the discussion contained in Deliverable D2.1, and how it relates to the research reported in this Deliverable. We also lay the basis for the viewpoints approach used in the rationale for the design of DNP.

Section 2: *The Role of Ethnographic Fieldwork Methods in CSCW*

In this section we report on some suggested roles for ethnographic fieldwork methods in CSCW based on our own and others' experience of fieldwork in a number of projects, some of which are reported in Section 3.

Section 3: *Fieldwork Studies: Toward Corpus, Codification and Framework*

This section is a compendium of shortened accounts of studies, some of which are continuing, related to the COMIC programme. It is material such as this which is the basis for the development of a framework for work analysis in CSCW. It is also intended to support the work of Strand 1, 'Organisational Context', and complement that of Strand 3, 'Mechanisms of Interaction and Notation'.

Section 4: *Perspectives on the Social Organisation of Work*

This section identifies some important themes of the social analysis of work for CSCW. It forms the basis for a more systematic development of a framework for work analysis which we intend to develop in the coming year.

Section 5: *Developing a Framework for the Presentation of Field Studies*

The ongoing work using the DNP is reviewed as a tool for the organisation and presentation of field work materials for CSCW design. In particular, we discuss its potential contribution to systems development within the framework of a viewpoint oriented analysis.

Section 6: *Concluding Remarks*

This section presents a short review of the research reported in this Deliverable and identifies the research to be undertaken in the coming year. In addition, we identify some longer term challenges for field studies in CSCW.

Section 1
Restatement of Themes and Problems

Section 1
Restatement of Themes and Problems

In Deliverable D2.1, 'Informing CSCW System Requirements' we addressed the following topics:

- identification of the problems of requirements elicitation in CSCW design;
- an examination of some of the consequences of the 'sociological turn' for conceptions of system design;
- a review and evaluation of current methods of requirements elicitation;
- an examination of the rationale for ethnography in CSCW design;
- identification of the main problems in relating ethnographic field work studies to CSCW design.

The Requirements Problem and CSCW

In general terms, the requirements problem for any system design is easy to state: how are designers and developers to understand what is to be built and how may it be done? However, and as is exemplified by the tremendous variety of methods and approaches available to system design, the problems are less than straightforward. In Comic Deliverable D2.1 we identified the following general issues in design:

- understanding the domain of application;
- conveying this understanding to designers;
- managing the design and development process.

Understanding the domain of application

This aspect of the design process is typically referred to as the 'requirements capture' or 'requirements elicitation' process. It is a process which is variably conceived, and much debated, but it is that aspect of the design process which is concerned with achieving an understanding and characterisation of the domain of application. In general terms its concern is with identifying the functions that the system should deliver, how these may be displayed to users, what parameters of the human-computer equation should be satisfied, and so on.

CSCW has forcibly reminded the design community that work is social and, in this way, CSCW proposes to pull HCI away from its traditional preoccupation

with cognitive science and, thus, provide an opportunity for extending an invitation towards sociology and anthropology to play a role in the design process. The major benefit of such a move was seen, perhaps over optimistically, in the ability of these disciplines to inform system design, and CSCW system design especially, through studies of the 'real world' character of work, a character which is pre-eminently social.

In the review of most of the salient methods and approaches in Deliverable D2.1 it was concluded that a major problem in evaluating the various methods was that each has evolved to tackle a particular set of problems and concerns, so that it is hazardous, to say the least, to even contemplate any one of them as *the* method of requirements elicitation. Methods are massively heterogeneous as manifested in their origins and presuppositions, the problems they address, the extent to which they focus on detail, which phase(s) of the design process they emphasise, the community the method is intended to serve, and more.

A major problem identified with respect to many of the methods currently available is their unsatisfactory attempts to bridge the "great divide" (Bowker et al, 1993; Kuutti, 1993; Jirotka and Goguen, 1994) between the concerns of engineering and those of the human sciences. Although this is often spoken of as a difference in 'mentalities', it is, more specifically, a difference in practices and ways of dealing with their respective problems. Briefly, system designers require that problems be broken down into manageable proportions, that they be thoroughly analysed and specified, that the process of design and development be organised as methodically as possible, that the problems be simplified and systematically laid out, and so on. In the human sciences, it is difficult to simplify the problems in ways that engineering requires, and even difficult to secure a clear and more or less agreed conception of what the problems are in the first place. Hirsch et al (1987) draw a contrast between formal disciplines (their own example is that of economics but it could apply equally to engineering) and the prized realism of disciplines such as sociology. In the former, the "style is characterised by the development of models based on deliberately, vigorously, and rigidly simplified assumptions". The "elegance of the models", along with their economy, are highly regarded, and their aim is to be predictive. By contrast, disciplines such as sociology tend to "value description or explanation over prediction" and "the realism of the concepts and propositions used, their resemblance to the perceptions and meanings of participants is highly valued". The clear implication is that one type of discipline cannot be simply transformed into the other because each defines its problems, poses its questions, evaluates results and designs research in entirely different ways.

It is important to note that this is an argument against a simple transformation of one type of discipline into another not one against such disciplines collaborating. However, it does mean that this collaboration will have to be on the basis of recognising the differences in approach and working out ways in which, and at what points, each can make contributions to understanding the problem. In the context of system design it is working out where to draw the line between

engineering problems and those which are better tackled by the social sciences. As we argued in Deliverable D2.1, the problems which the human sciences need to address are not solved by trying to use methods and descriptive techniques, such as those of many of the methods used in requirements elicitation, which owe more to engineering than to trying to understand the 'real world, real time' features of socially organised domains.

However, the opportunity for the social sciences to take part in system design is a mixed blessing. For one thing, these sciences are singularly ill-equipped to meet the important needs of system builders and developers because they lack any widespread agreement as to their own methods and approaches with the result that arguments within the social sciences have now been brought to issues in system design. In other words, whether or not sociology is capable of informing CSCW design effectively is a promissory note rather than an achievement.

Our own approach to these issues within social science is not so much to regard them as matters of principle that need to be disputed in relation to issues of system design, but to take a more pragmatic stance by exploring as thoroughly as possible one promising method of acquiring information about application domains, namely, ethnographic fieldwork, asking how and in what forms it may be deployed in CSCW system design and development. This work is reported in, along with the background arguments only sketched here, in Deliverable D2.1.

Relating the understanding to designers

Coming to understand work as a 'real world', socially organised domain is, for us, an essential precondition of CSCW design.[1] However, it is also acknowledged that this understanding must be effectively related to the design and development process. To the extent that design and development is likely be concerned with large scale systems and, accordingly, involve a large number of people of different specialisms, this issue becomes a much more explicit one. The issue of how far and in what ways users should be involved is also pertinent.

While recognising that many of the conventional approaches to requirements are inadequate for CSCW, we also accept that 'human-centred' approaches also lack the ability to deliver clear and systematic system specifications. One objective of the work in Strand 2 is to investigate how one such 'human-centred' method, ethnography, may be integrated with at least some of the conventional approaches to system requirements engineering.

Fieldwork methods involving ethnography are capable of providing much rich material and analyses on the 'real world' character of the social organisation of work. As a method of social investigation, it has a long pedigree in sociology and especially anthropology, though not, it is important to add, one which receives universal assent within the discipline. Be this as it may, bringing ethnography into

1 This does not imply, as will become apparent later in the Deliverable, that this should necessarily be a first or even a single stage in the design process.

the CSCW design process raises important issues concerning the agenda for such fieldwork as well as issues to do with relating the outcomes of fieldwork to systems engineers.

Changing the agenda of fieldwork and ethnography

One of the hallmarks of ethnography is its claim to 'presuppositionless inquiry'; that is, to a mode of inquiry which makes as few assumptions as possible about the socially organised character of the domain under investigation.[2] It does, of course, make the presumption, not unreasonably, that the domain is socially organised and its task is to find out how, and relate this finding to sociological issues. In the context of system design, however, other agendas, not necessarily antithetical to those of ethnography, are likely to be in force. For example, in industrial contexts decisions about technical innovations are rarely matters for any design team. They are typically given a remit for design, and have perhaps some, but a variable, and usually small, role in influencing the shape of that remit, which they are then required to implement. Whatever one's views about this way of organising industrial innovation, it is a common experience and one which, for the foreseeable future, is likely to be the 'real world' of system design. Many of the obstacles to change are to do with the arrangement of powers and responsibilities in the organisations which undertake large scale system design and redesign, and the position of design teams in relation to these. This means that ethnographic fieldwork methods, if they are to be used in design, will have to accommodate in some way to such constraints as these.

However, it is important that in doing so, the essential virtues of the method are not lost; that is, its ability to show the 'real world' character of work and its activities. Nevertheless, compared to its use in sociology and anthropology, in system design the fieldwork agenda will need to be more targeted in terms of the design remit. The most immediate practical consequence of this is likely to be a reduced period of fieldwork, making it a matter of a couple of months at best, rather than the couple of years which is not untypical of fieldwork in sociology and anthropology. This alone is likely to place some restrictions on what the method is capable of delivering to designers.

In Section 2 of this Deliverable we discuss some of the roles which, in our experience, ethnography could play in the design process.

Frameworks, theories and conceptual schemes

As indicated earlier, ethnography is not especially theory-driven in its approach, preferring to operate initially within a loosely formulated, loosely defined, loosely confined framework of ideas. It is not, of course, without its theoretical-cum-

---

2 It is not, of course, without presuppositions. No form of inquiry could be without these. The phrase in connection with ethnography is intended to draw a contrast between the more formally structured methods of social research, such as surveys and interviews, with the much more open-ended style of fieldwork. It also contrasts with methods of social research which begin with a theory, conceptual scheme or framework which is imposed upon real world activities.

methodological tradition which justifies this stance as an appropriate one for sociological inquiry.[3] However, it is a stance which contrasts not only with other traditions of inquiry in sociology, but also to strongly established traditions in requirements elicitation and, more generally, HCI. While some of the arguments pertinent here are reviewed in Section 2 of this Deliverable and some of them more fully in Deliverable D2.1, it is worth noting the strategy which is distinctive to Strand 2.

An important aspiration of much of human science since its beginnings has been to develop a set of well-grounded theories, frameworks or conceptual schemes as a common vocabulary for describing and explaining human behaviour. If realisable, such an aspiration would have tremendous benefits to engineers in that, by using the common vocabulary, by using the theories, their design work would be more successful. Ergonomics, HCI, task analytic approaches, among many others, in significant respects reflect such an aspiration (See, for example, Bannon, 1991). Even though it is held with less optimism these days, our view about this aspiration is not that it is in principle unrealisable but that it is, at least, premature.

Our own approach is to emphasise the requirement of CSCW design that it be informed about the sociality of work and that, currently, ethnographic fieldwork method offers much promise in this regard. However, its distinctive attitude to what it regards as premature theorising means that it cannot be in the business of setting up analytic frameworks in advance of inquiry. The important phrase here is 'in advance of inquiry'. In other words, the objection is not so much to theories or conceptual frameworks *per se* but to a strategy which sees these as necessary preliminaries to inquiry.

It is this which presents some of the problems of relating ethnographic analyses and materials to CSCW design. As we have said elsewhere, ethnographic analyses are typically textually discursive and, on their own, offer little guidance as to system requirements. Software engineers, as noted earlier, because of the nature of their work require a mode of presentation and tools of analysis which are comparatively systematic. They need to decompose their problems into manageable segments, achieve a good measure of abstraction in their descriptions of the operations of existing or planned systems, clearly lay out the flows of information, data, channels of communication, and so on, in order that the complex engineering problems and the software requirements can be clearly exhibited in graphical forms. This way of thinking and working has to be understood not as some arbitrary and whimsical choice but one which is intimately a 'tool of the trade' and which has evolved to meet the problems posed by this kind of work.[4] Bridging the gap between a mode of presentation which is

---

3 See, for example, Ackroyd and Hughes (1992) for an overview of these arguments. Also Hughes et al, (1993) and Section 2 of this Deliverable.

4 It is recognised that there are many notations, etc. used in software design and development, but the point remains the same.

essentially textual to one which is predominantly graphical is one of the important challenges in formatting ethnographical materials for use in design.

As we keep on stressing, our approach to problems such as these is pragmatic. In other words, we accept the 'gap', so to speak, and treat the problems that it poses as practical problems to explore. In this particular case it is thinking about how to effect means of adequate translation between one mode of presentation and another.

There are two main tasks which are initiated in this Deliverable.

First, to begin the development of a framework for characterising the social organisation of work. This framework, as will be explained more fully later, is derived from a corpus of fieldwork studies done either directly as part of COMIC or as part of research projects closely associated with COMIC. In this respect, it is a framework which is 'inductively derived', though we treat this phrase with some caution. Its aim is to present a set of guidelines to sensitise CSCW design to the properties of socially organised work and its activities. The preliminary work toward this framework is presented in Section 3 and 4.

Second, through the development of Designer's Note Pad, we hope to evolve a mode of presentation and of representation to provide computer support for the design process. This is discussed in Section 5. Our aim in this part of the research in Strand 2 is to use material and analyses from ethnographic studies to think about different ways in which this material can be represented to display features of the social organisation of work for the varying needs of the design process, including those of computer scientists.

Managing the design and development process: the organisational context and design

Any large scale system will need to be built within the familiar constraints of time, manpower and cost. In addition, it will need to be a systematically organised process involving a division of labour among possibly changing personnel of varying specialisms. These are features which need support.

There is also the very important issue of the organisational context in which design is done as well as those organisations who are the recipients of the design. Many of the issues here are the direct concern of Strand 1 of the COMIC Project, nevertheless, they impinge on the research in this Strand in a number of ways particularly related to methodology. Among those that we particularly want to identify are, first, the issue of scale, and, second, how organisations are to be conceptualised. Clearly, these issues are closely bound up with each other but in complex ways which are difficult to see.

Scale

This is not a single issue but a complexity of topics and problems. The most obvious way in which scale appears in connection with ethnography in system

design is that it is, *par excellence,* the method for the investigation of small-scale worlds such as the small community, the control room or the work group, rather than for the investigation of large-scale processes and organisations. As was found in one of the studies reported later in this Deliverable, a lone fieldworker trying to investigate a large software team in an aircraft company proved to have a much more difficult task compared to one studying a single control suite in London Air Traffic Control Centre. In the design team study many of the work activities happened simultaneously making it difficult, if not impossible, to follow them as fully as required. Similar problems emerge in respect of distributed systems where users may be widely geographically dispersed or in respect of processes which can transcend the boundaries of work team or even an organisation. In other words, the interdependencies of work activities and processes can extend beyond the small-scale worlds that ethnography typically studies. This was, for example, the case in the banking study where, though many work activities were routinely sequential, following through the complete process proved difficult.

The issue of scale also appears in the management of the production of large-scale software systems. Not surprisingly, there have been a number of voices within CSCW urging the importance of studying the software production process itself in order to devise appropriate computer support for what is a co-operative activity. As Brooks (1975) points out in his now classic study, the production of large scale software systems requires the organisation and management of a large number of people and this becomes a major overhead, not least because of the increased need to devise an appropriate division of labour, plan and co-ordinate the work, assess its progress, check its quality, and so forth. In other words, the production of large scale software systems becomes more and more like production engineering and, as a result, accumulates the inescapable features of such work, including plans, specifications, records, management and control procedures, along with the 'natural' troubles and problems these give rise to. Among these, as we discussed in Deliverable D2.1, are not only the technical and engineering problems but also those which arise from the wider organisation, such as those to do with constraints on investment, on staffing, on recruitment, the economic health, or otherwise, of the company, changes in marketing policy, and more, all of which are typically outside of the control of a design team.

Closely bound up with the issue of scale is that of time. Once again this can surface in a number of ways. In the context of ethnographic fieldwork it arises in particular when trying to follow work processes which are temporally extended. Similarly, there are many settings in which the work follows 'seasonal' rhythms. In branch banking there are regular peaks and troughs in the flow of customers, in manufacturing there can be ebbs and flows of work and in air traffic control the density of the traffic is closely related to times of the year.

Time, spatial extension and work processes are related in complex ways.[5] However, for our part, we want to emphasise the importance of seeing these as features of the social organisation of work itself.

As we have indicated the issue of scale is intimately bound up with the conceptions of organisation, and in this respect assumes the character that it does because of the picture of organisation which they encapsulate.

Conceptualising organisation

Another guise in which the issue of scale appears, arises from the way in which organisations are conceptualised as a context or framework for human activities. As we noted earlier, the issue of 'organisational context' is of particular concern to the work of Strand 1 and is reported in its Deliverables. However, it also impinges on the methodological research in this Strand in that how we choose to conceptualise organisations has ramifications for the choice of methods to study organisational activities, processes and their context, as well as for issues of scale.[6] As we indicated earlier, if we think of a bank or a manufacturing company as a large scale entity then, methods such as ethnography, given their focus on small scale, clearly bounded interactions, are at a disadvantage. However, we would argue that the issue arises in this form because of the analytical assumptions which underlie the characterisation of a bank or a manufacturing company as a 'large scale entity' with the strong implication that what we are dealing with here is an entity which is 'greater than the sum of its parts'.

This view, of course, draws upon familiar conceptions within sociology, reflected in distinctions such as 'macro-micro', 'individual-society', 'structure-agency', as well as in concepts such as 'social context', 'social structure', 'culture', the 'institutional order', and more. What such conceptions do is propose the idea that there are two sorts of entities - the 'micro' ones of individuals, actions, activities, 'face-to-face situations', etc., and the 'macro' phenomena of 'organisation', 'situation', 'social structure', 'groups', 'encompassing social order', etc. - which analysis must bring into alignment, either through some method of empirical correlation and/or through theoretical specification. In either case, the dominating idea is what Coulter (1982) refers to as the "container model" of social structure, or of organisations, in which the actions and activities of persons are seen as constrained within the wider, supra-individual frame, rather like a liquid is contained in a bottle.[7]

The methodological strategy of this Deliverable is broadly nominalist in that we take the names of organisational entities as summary descriptions referring to the actions and activities of the individual members of those organisations acting

---

5 One can plausibly see much of technology from the earliest days as artifacts designed to 'reduce the cost' of time and space on human endeavour.

6 The literature on organisational paradigms is voluminous. See Bowers (1994) for a review of the main variants. Most of the current wisdom in organisational studies tends to take the view that organisation theory is in a state of crisis due to the bewildering array of paradigms.

7 The literature on the issues here is voluminous, but see, for example,

in some organisational capacity.[8] Although this is not the place to develop these arguments in any greater detail, the picture it offers of an organisation is as a connection of locally organised activities, sometimes widely distributed but organised moment-to-moment by its members through the course of organising the organisation's activities. The nature of these 'connections of locally organised activities' is clearly an important issue but from the point of view of ethnography they become investigable as 'real world, real time' features of the work activities themselves. This, we hasten to add, does not ignore, play down, or trivialise, the kind of issues which organisation theories place at the centre of their concerns, including those of hierarchy, power, authority, and so forth. What it does is propose that these be examined in a different way as features exhibited in and through the locally organised activities of an organisation's members.

This strategy does not, of course, avoid the issue of scale but reconfigures it in a way which makes it more manageable as a feature of 'real world, real time' work activities. As will be shown in the studies reported in Section 3 of this Deliverable, organisational issues are the constant practical considerations of persons within the work settings and it is through the local organisation of the work that what we earlier referred to as 'the connections' in organisations are created, reproduced and maintained.

Viewpoints

As we reviewed in Deliverable D2.1 recent methods of requirements engineering have adopted the notion of viewpoints (Moyse, 1991; Kotonya and Sommerville, 1992). This explicitly recognises rather more than other methods that there can be different interests and perspectives on the requirements for a system and that the requirements for a system can be derived from many different sources. An additional bonus for CSCW is that many proponents recognise the importance of co-operation in determining requirements rather than simply the abstracted resolution of conflicting needs and inconsistencies (Finkelstein and Fuks, 1989). However, despite this, to date such approaches tend to focus on the recording of the detailed relationship between the system and its users rather than trying to understand the wider context of work within which a system must find its place. Further, many frameworks for implementing viewpoints into a method for system design over schematise the procedure which tend to fix, for example, the number of roles, viewpoints and so on of the organisation and the system underdevelopment, so running the risk of imposing a structure on 'real world, real time' activities in advance of inquiry.

Accordingly, for the purposes of this Deliverable and in keeping with our pragmatic stance, we feel it worthwhile to distinguish two broad and related categories of viewpoints: the methodological/representational and the socially located viewpoints.

---

8 This position is consistent with the one adopted by Weber (1949).

Methodological and representational viewpoints

By this category we refer to those techniques which have been developed, and are being developed, to implement the viewpoints approach to systems engineering. These are intended to structure the exploitation of the viewpoint notion for the purposes of requirements specification. Many of these methods were reviewed in Deliverable D2.1.

In our own case, DNP incorporates the viewpoint notion based on the different rationales of ethnographic fieldwork and systems design; rationales which, in CSCW design especially, should be made complementary. In this respect, the techniques incorporated and being developed in DNP are intended to meet the needs of handling ethnographic fieldwork materials for system design. In effect, what they facilitate is the handling of work analysis data for system design using the familiar computer science tools of graphical representation. It is in this respect that the tool exploits the different but complementary representational viewpoints on system design and, in particular, identifies incompatibilities in design decisions and specifications. We anticipate that one of the virtues of the system is that it makes such incompatibilities, both in engineering and in organisational terms, more visible.

Socially generated viewpoints

This is a more sociological rendering of the viewpoint idea and, in this respect, draws upon fairly standard sociological ideas - recognised in the origins of the notion of viewpoints in requirements engineering - that the members of some socially organised domain may well have various interests and perspectives on aspects of their social lives both within and without the domain.[9]

One of our main interests is with the way in which different perspectives and viewpoints can arise from particular roles and positions within a division of labour. The traditional exploitation of this idea in much of social science is that between management and labour within hierarchical organisations. However, and perhaps more prosaically, from the point of view of ethnographically based work analysis the concern is much more to do with viewpoints as products of the location within a working division of activities.

For example, many work processes extend beyond the activities of a person or a team and one of the objectives of establishing a division of labour is that persons performing activities within such an arrangement only need to know about their tasks not those of others. The interdependence of the tasks is achieved via their organisation into a division of labour.[10] No one is required to ‘know the whole

---

9 It is an idea which has been developed in, for example, the Marxian notion of socially generated interests and ideology, the symbolic interactionist conception of ‘the actor’s point of view’, attitude studies, and many more perspectives within sociology.

10 There is clearly more to it than this but, nonetheless, this does characterise the core idea of the division of labour as a socially organised arrangement which facilitates the organising of large scale, complex cooperative endeavours. We also recognise that, organisationally, the division of labour can take a variety of forms.

process' except perhaps as a plan or as a schematic. As a 'real world, real time' phenomenon, as a 'working division of labour', the above brief analytic specification of the division of labour parallels the experience of those working within a division of labour. As Anderson et al (1989) argue, from the point of view of a person embedded within a division of labour, work activities seem to be organised around an "egological principle" and "horizons of relevance", that is and to summarise, as 'tasks which I do' and 'tasks which others do', 'tasks which are my concern' and 'tasks which are the concerns of others', and so on.

One important implication arising from this approach is that the division of labour itself becomes an object of empirical investigation as a feature of the social organised work activities and, as a result, places co-operation and co-ordination as themselves matters to be worked out in and through the work itself. Part of these processes is also working these matters out from the viewpoints which arise out of a person's location within a division of labour.

There is no implication here that such viewpoints are inherently in conflict. What we are referring to are more like vantage points from within a socially organised set of activities. Of course, it may be that out of different vantage points different requirements emerge which must then be reconciled. However, for us this is a matter of inquiry rather than *a priori* stipulation.

This Year's Work

In the conclusion to Deliverable D2.1, we made the claim that social science had much to offer CSCW and that, within this, ethnographic fieldwork methods hold considerable promise. However, in order to realise these much more research was needed to clarify the role that ethnography could play in the system development process.

As part of this endeavour, it was also felt vital to develop a greater understanding of the social organisational features of work itself through a series of case studies which would form a corpus of materials to serve as the basis for the construction of a framework for the analysis of work for CSCW design. Indeed, much of this year's effort has been devoted to organising, conducting and reflecting on these field studies. This effort complements some of the research interests reflected in Strand 3 which seeks to develop a notation for the representation of the properties of co-operative work as specified in a particular conceptual scheme. As we have indicated earlier, our approach is more inductive in intent but it is envisaged that toward the end of the project closer connections will be made between the two approaches.

Section 2
The Role of Ethnographic Fieldwork Methods in CSCW

Section 2
The Role of Ethnographic Fieldwork Methods in CSCW[11]

In this section we intend a retrospective look at our own experience of using ethnography in a series of research studies, and suggest some roles which ethnography can play as a contributor to CSCW system design. Though we are strong aficionados of the method we do not regard it as a panacea for the problems of system design which are complex and "wicked" (Rittel and Webber, 1973). In other words, if ethnography is to take a more regarded place in CSCW design, then it is important to appraise not only its virtues but also its vices. Here we identify four uses of ethnography in various phases of the design cycle as a contribution to an evaluation of the method.

We also briefly examine the arguments which have motivated the introduction of ethnography into systems design. We then reflect on our particular experiences in the use of ethnography across a number of projects and present some more general implications arising from them.

The Case for Ethnography in CSCW

Although it is not our intention here to review the impetus behind the rise of CSCW which has been effectively done elsewhere (See, for example, Grudin, 1991; Bannon and Schmidt, 1991) it is worth briefly reflecting on the rationale for ethnography in CSCW systems design. Two trends have strongly motivated the prominence ethnography currently enjoys:

- The growing plausibility of the diagnosis that the reason why many systems fail is due to the fact that their design pays insufficient attention to the social context of work; a failure often attributed to the inadequacy of existing methods of requirements elicitation and work analysis (Schmidt, 1993).
- A growing awareness with the emergence of low-cost technology that the ubiquitous nature of networked and distributed computing pose new problems for design which require the development of new methods which analyse the collaborative, hence social, character of work and its activities.

The tentative incorporation in system design of a social perspective emerges from these two trends and the insistence that the computer moves into the world of

---

11 A version of this section was presented at CSCW'94, Chapel Hill, North Carolina. See Hughes et al, 1994.

work and organisation (Grudin, 1990). Given this 'turn to the social' and the need to study the 'real world' character of work, drifting toward sociology through ethnography is almost a natural inclination. Thus, in the way that HCI has previously looked to psychology for an understanding of human behaviour, CSCW turns to sociology and in particular ethnography to provide insight into the social nature of work. The expectation is that requirements elicitation is to be informed by an analysis of the 'real world' circumstances of work and its organisation (Goguen, 1993).

This is reflected more generally in a growing awareness within the software engineering community that understanding the 'social' real world is an important factor in software design and development itself (Potts, 1993). There is, not surprisingly, some equivocation about just what the 'social' means in this connection. The influences have come from a number of different directions, including the sociology of technology, the sociology of industry and the sociology of organisations among them. Not all of these have a direct interest in sociology as an input to system design. (See Quintas, 1993 for a selection of papers).

The main virtue of ethnography is its ability to make visible the 'real world, real time' sociality of a setting. As a mode of social research it is concerned with producing detailed descriptions of the 'workaday' activities of social actors within specific contexts (Hughes *et al*, 1993a; Hughes *et al*, 1993b). It is a naturalistic method relying upon material drawn from the first-hand experience of a fieldworker in some setting. It seeks to present a portrait of life as seen and understood by those who live and work within the domain concerned. It this objective which is the rationale behind the method's insistence on the direct involvement of the researcher in the setting under investigation. The intention of ethnography is to see activities as social actions embedded within a socially organised domain and accomplished in and through the day-to-day activities of participants. It is this which provides access to the everyday ways in which participants understand and conduct their working lives.

It is the ability of ethnography to understand a social setting as it is perceived by those involved in that setting (the archetypal 'users') that underpins its appeal to developers. However, it is not without its problems. As has already been noted in the previous Section, there are, for example, those problems to do with presenting the results of ethnography in a form which is readily assimilable by designers. For many software engineers ethnography seems far too unsystematic a method, its results presented in a discursive form, design options are not clearly stated and do not attend sufficiently to engineering needs. Its virtues, in other words, become vices.

Against this is the argument that what is wrong with many of the traditional methods of system design is that they owe far too much to the needs of engineering with the result that crucial aspects of the 'real world' of work are obscured, misrepresented or never properly treated (Schmidt, 1993). It is in this respect that 'analytic approaches', Task Analysis, Office Automation for example, which focus on the flow of data within a domain, are found wanting (Shapiro,

1993; Suchman, 1983). While it is accepted that a balance needs to be found between the requirements of engineering and the need to adequately characterise the domain of application, such methods are an intrusion of the 'engineering mentality' into areas where it is inappropriate. The result is, so it is argued, that essential aspects of the socially organised character of the domain concerned are obscured or, worse, misrepresented. More specifically, the analytic deconstruction of work activities into ever more finely grained components removes the essential 'real world' features which make them practices within a socially organised setting. This complaint attacks the individualistic slant of the cognitivism which underlies 'analytic approaches' by acknowledging the implications of the observation that work is, typically, collaborative. Though the activities constituting work are done by individuals, they are performed within an organised environment composed of other individuals and it is this which gives shape to the activities as 'real world' activities. Thus, the focus of ethnography is on the social practices which enable the very processes which 'analytic methods' identify but which they decontextualise. It is through social practices that processes are established and, accordingly, rooted in socially achieved sets of arrangements.

There are, of course, many aspects to these kinds of arguments, some of which involve a critique of the nature of work in modern society and how current methods of design instantiate the dehumanising rationality of modernism[12] Again as we have emphasised earlier, our own arguments for ethnography are more pragmatic in nature. If we accept that CSCW design needs to attend to the sociality of work, then any method must respect the nature of this phenomenon. However, many of the existing methods fail to sufficiently recognise the social nature of work. This is not a call for the wholesale abandonment of more formal methods; they, like ethnography, will need to find an appropriate place in design.[13]

Accordingly, although there is a case for ethnography in CSCW system design, at the present time it is a promissory note rather than a claim based on substantial achievement. Its main use has been in research and mainly field sites which are small scale involving highly focused interactions, such as control rooms. Accordingly, if it is to substantiate its case as a method of system design, it will need to go beyond these and, in addition, face up to the problems of large scale system development.

Moving Beyond Research

For our part we accept that there are very real problems in the design and development of large scale systems, problems which have been well-rehearsed in

12 See, for example, COMIC Deliverable 2.1 (1993) for a review of some of these arguments.

13 We do recognise that matters are not quite so easily resolved. Many of the arguments about methods do involve critiques of their underlying presumptions. An example of this is the attack on the cognitivism which underlies many of the task analytic approaches.

the literature (Deliverable D2.1). Briefly, these have to do with obtaining adequate knowledge of the relevant domain, communicating this to designers and organising the process of system building. In commercial contexts these problems are deeply infused with the familiar commercial constraints of budgets, time and resources.[14] This means that methods such as ethnography must service a number of demands if they are to be widely accepted in industry. Without this acceptance the use of ethnography in systems design runs the risk of becoming a research curiosity and, thus, devalued as a tool to support effective CSCW design.

As a number of studies have shown software engineers typically work under some pressure (Curtis, et al, 1988; Bansler and Bødker, 1993); a pressure which is, in part, determined by market factors. However, the familiar moan that most system development projects are 'over time, over budget' cannot be entirely laid at the door of market pressures. Building large scale systems is a complex and difficult business. Many of them are 'one off' with little in the way of past experience to serve as a guide. It was problems such as these which provoked the development of software design methodologies to systematise and manage the process of design and development so that systems had a reasonable chance of meeting both technical and commercial targets. These pressures still hold true and apply equally well to ethnography.

On the face of it ethnography does not accommodate easily to the pressures of development. A set of tensions become apparent when we examine ethnography in the light of systems design and it is important that the role of ethnography is considered within this context. These tensions include the familiar pressures of scale and time and place new demands on ethnography in system design.

The problem of scale

This has been discussed above: to date the main use of ethnography has not only been within research settings but also confined to relatively small scale and relatively confined environments, such as control rooms and other micro interactional contexts. In such settings there tends to be a clear focus of attention for the participants, who are typically few in number, and in which there is a relatively clearly visible differentiation of tasks at one work site. Though 'team ethnography' is a possibility, this is likely to be an ineffective and impractically expensive strategy which could well create as many problems as it would solve. The realistic possibility is that ethnography will continue to be undertaken by a single investigator or, perhaps in exceptional cases, by a couple of them. For the lone fieldworker the sites described above are ideal. They minimise travel and communication problems, and all that the fieldworker needs to see is there in one place and can be gathered with a minimum of disruption. Scaling such inquiries up to the organisational level or to processes distributed in time and space is a much more daunting prospect in raising issues of depth and representativeness.

---

14 This is not to say that research contexts do not have their constraints of budget, time and other resources, only that commercial software development has different ones.

The pressure of time

As one of our computer science colleagues expressed it, ethnography is a 'prolonged activity' and in the context of social research can last a number of years, certainly time scales which would be considered a joke in software engineering. Added to this are the problems, noted earlier, of communicating ethnographic findings to designers. The output of ethnographic analyses are typically discursive and lengthy, looking nothing like the blueprint diagrams which are *de rigeur* in systems engineering.

The role of the ethnographer

Moving out of the research setting into a more commercial one also raises different sets of ethical responsibilities as well as making access to sites more vulnerable to the contingencies of the commercial and industrial world. Ethnography insists that its inquiries be conducted in a non-disruptive and non-interventionist manner, principles which can be compromised given that much of the motivation for IT is to reorganise work and, as part of this, often seek to displace labour. Less dramatically, but important nonetheless, fieldworkers not only require access to relevant sites but also need acceptance on the part of those who work in them. Protecting the identity of people, respecting the fact that the fieldworker is like a guest within their lives, and so on, become much harder to sustain in applied work of this kind. It is also too easy for the fieldworker to get drawn into intra-organisational politics especially when dependent upon 'facilitators' within organisations who have their own agendas to pursue.

Of course, few of these issues are easily solved. However, it is important not to be too ambitious for any method, least of all in software engineering where new methods follow one another with monotonous regularity. The 'wickedness' of its problems means that design is, at best, a 'satisficing' activity and a matter of doing the best one can with the resources available. Nevertheless, if it is accepted that designers should be informed about the social character of work, and that ethnography is an important means of gaining such knowledge, then serious attention needs to be given to the variety of ways in which ethnographic studies can be used by designers. What follows is an attempt to specify some of these ways using, in the main, our own experiences over four years of collaborative and interdisciplinary research. We do not offer these in anything other than the spirit of what can be done now. Research on ethnography and system design is continuing in a number of quarters and it may well be that in a few years the picture will be very different. For us, the important issue at the present time is to sensitise CSCW system designers to the sociality of work as systematically and as effectively as possible.

Ethnography in system design

The wish to incorporate ethnography into the already diverse collection of methods, tools and techniques used in system design must be viewed with some trepidation. While we accept the need for the inclusion of a social perspective on design we must be careful to avoid seeing ethnography as a ready-made solution. The experiences of ethnography within systems design are limited and, as pointed out earlier, mainly confined to small-scale settings and of highly focused activities. (See, for example, Heath and Luff, 1992; Heath and Luff, 1994; Hughes et al, 1992; Harper and Hughes, 1992; Harper et al, 1991; Heath et al, 1993 and, of course, Suchman, 1987.)

However, ethnography is a much richer method than these previous studies and the reports of design experiences might perhaps suggest. It is important that existing studies are complemented by a consideration of the variety of different ways in which ethnography can influence systems design. Our aim in this Section is to propose some different uses of ethnography within the design process. These are based, to repeat, on our own experiences gleaned from studies we have undertaken over the last four years. In this respect, the categories presented below represent a codification of the lessons we have learned rather than setting out the strategies which originally directly informed the particular studies we use to illustrate the approach.

The different uses of ethnography within design we identify include:

- Concurrent ethnography: in which design is influence by an on-going ethnographic study taking place at the same time as systems development.
- Quick and dirty ethnography: in which brief ethnographic studies are undertaken to provide a general but informed sense of the setting for designers.
- Evaluative ethnography: in which an ethnographic study is undertaken to verify or validate a set of already formulated design decisions.
- Re-examination of previous studies: in which previous studies are re-examined to inform initial design thinking.

In the following we suggest what each of these has to offer design and also identify some of the problems that could arise. These categories should not be read as if they were mutually exclusive ways of using ethnography in system design. As we will suggest, some of the uses could be harnessed together and the differences between them seen as differences of emphasis rather than sharp demarcations.

Design, as in so much else, is a matter of responding to contingencies of various kinds. What is also important to note is that the schema recognises that design objectives are themselves various and that this will have a bearing on the role of ethnography. In other words, while not necessarily buying into the picture

of the design process as a series of discrete, clearly delineated and phased steps, it undoubtedly has different objectives at different stages and, accordingly, implications for how design needs to be informed by relevant information about the domain.

Concurrent ethnography

This use is perhaps the one most commonly associated with design and the one most commentated on (See, for example, Heath *et al,* 1993; Hughes *et al,* 1993). It is a sequenced process in which the ethnographic investigation of a domain precedes the design development of the system. This is the method we followed in the design of a tool for the rapid prototyping of interfaces for controlling (Bentley *et al,* 1992). In this case a period of some four weeks ethnography in the London Air Traffic Control Centre (LATTC) was followed by a lengthy debriefing session involving both the fieldworker and the designers. Meanwhile, a first prototype was constructed. The process of fieldwork > debriefing > prototype iteration > fieldwork was repeated about four times until the team was satisfied that little more could be usefully gained by more fieldwork. The penultimate version of the system was then evaluated using working controllers. The process was a directed one in that each stage of the fieldwork was intended to target issues raised by the designers during the debriefings, although the first phase was more concerned with the very important task of the fieldworker familiarising himself with the setting and the work of the controllers.

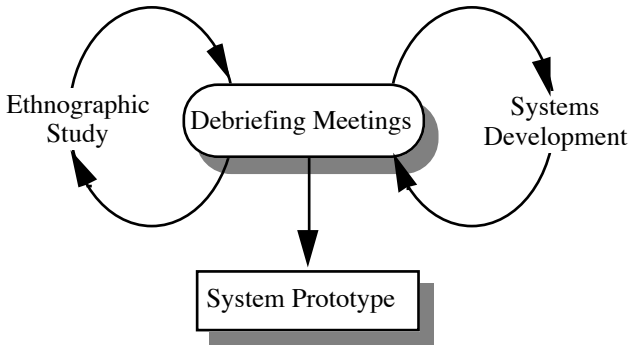

Figure 2.1 The use of Concurrent Ethnography

It is important to note that the aim of the project was research rather than the development of a system to be used in the ‘front line’ of controlling. Thus, we did not have the problems which would have arisen in implementing the tool. The research team was small so that much of the communication between the sociologists and the computer scientists could be done informally. There was little need for a requirements document or for a process model since the development work was done through rapid prototyping.

What the ethnography especially provided was a thorough insight into the subtleties involved in controlling work and in the routine interactions among the

members of the controlling team around the suite; subtleties which were rooted in the sociality of the work and its organisation. The vital moment-by-moment mutual checking of 'what was going on' by the various members of the team had been missed by earlier cognitive and task analytic approaches to describing controlling work. What also became clear is that any new interface system would have to keep the controller 'geared into' the work by not automating, for example, the ordering of the screen-based flight strips. In other words, we felt it important to retain at least some of the functionalities of the current paper flight strips while, at the same time, being in a position to evaluate what information the controller needs, what is less important but needs to be 'ready to hand', and what was inessential.

We also learned that there was a declining rate of utility for the fieldwork contribution to the design. This is not to say that there was not more to learn or that we could not have learned more sociologically from further study of the control room, only that in terms of the project the 'fine tuning' of the design needed to be informed by experts actually using it. The intricacies of tool use in air traffic control typically require a background of several years of work experience for their mastery, a grasp of the ways in which the employment of the tool is capable of being effortlessly used within the manifold requirements and exigencies of safe controlling: levels of sensitivity which cannot be matched by the ethnographer on the basis of a few weeks acquaintance with the work. In other words, although there is always more to learn, the payoffs for design, at least in this case, came relatively quickly in comparison with social research uses of ethnography.

Quick and dirty ethnography

This characterisation of a role for ethnography is, in many respects, a rationalisation of the experiences of a project which did not go quite as intended, but which, nonetheless, provided valuable insights not only into the use of ethnography but also about the character of 'real world' software engineering design and, through this, some of the limiting conditions affecting the provision of computer support in CSCW contexts.

The principal distinction between this project and our previous experiences within ATC was the larger scale of the work setting. The ATC suite provided a natural focus and location for the work taking place. However, in the case of software development both the location and focus of work was considerably less apparent to both the developers and ethnographers and the issue of scale needed to be directly addressed.

Large scale settings

We have already noted some of the problems of scaling up ethnography beyond the confines of such as control rooms. In the case we use for illustration, the project was concerned to use an ethnographic investigation of software engineers

at work in order to inform the design of a support tool; a tool which would, hopefully, enable designers to display the rationale behind their design choices and, through this, improve the quality of the system and the maintainability of the software. The aim was to develop a tool which more adequately reflected the collaborative and interdependent character of 'real world' design work. We planned to follow the pattern of the study mentioned previously; that is, a first period of familiarisation fieldwork while, at the same time, building the basic prototype, to be followed through by a series of iterations of debriefing, more directed fieldwork and prototype iterations. Although we had ready access to various sites, and to colleagues working in the same area, it was difficult to find projects we could study which were starting as opposed to those which were already some way along their development trajectories. Nevertheless, we felt that we would still gain a great deal for our purposes.

We realised from the beginning, and this was one of the purposes of the study, that the fieldwork would represent new challenges in involving a much less 'confined' field site than the control suite at LATCC. For one, the development engineers in both of the sites we eventually looked at, were working in industrial environments and, accordingly, subjected to a wider range of contingencies, events and policies which impacted more directly on their work. For example, one of the projects at the first field site was cancelled and access to another project within the same company proved more difficult due, to put it diplomatically, to one of the 'gate keepers', a team manager, being less than enthusiastic about a fieldworker studying a team under considerable pressure.

While we may have been unlucky in this case and more fortunate in the case of LATCC, it does highlight an important feature of ethnographic research, namely, its reliance on being accepted in the setting and, even if this is forthcoming, being subject to the range of contingencies that are capable of afflicting all 'real world' organisations. Among these, of course, are those to do with, for want of a better phrase, the 'local politics' of the organisation.

In addition to these were the problems arising from asking a fieldworker to cover what proved to be a much larger task than we had anticipated. Software development is a complex business and tracking through its unfamiliar complexities, understanding the management of its components, seeing how the teams worked together, trying to figure out how the integration of the various components was achieved, and more, all proved to be a much more immense task than we envisaged originally.

Nevertheless, and despite less than ideal circumstances such as those noted, one can always learn something from ethnography. Indeed, seeing how the kind of contingencies we have reviewed can impact on design and development is important and, of course, illustrative of the argument CSCW makes about the necessity of studying the 'real world' circumstances of work to inform system design. In this case, we learnt sufficient about the design process as a 'real world' phenomenon to indicate that the tool as originally envisaged was, in significant respects, wrongly conceived. Briefly, it would only be effective if it was

consistently used by members of project teams. However, in the conditions in which they typically worked, this would represent a considerable overhead. Also, given the personal and company investment in CASE tools of various kinds, persuading engineers to learn and use "yet another bloody tool" when they were already less than enthusiastic about their current ones, would have been a mammoth task.

In the second site many of the problems indicated above also emerged. The development involved, at the time of the fieldwork, approximately a hundred software engineers working on an avionics systems for a new version of an aircraft.[15] The work was organised according to a strict Process Model which was highly constrained, document driven and implemented under very tight budgetary constraints. This again provided insights into the 'real world' of design, particularly on the impact of management styles, the importance of professional pride the engineers exhibited in 'their craft', and a better understanding of the relationship of the Process Model used to organise the work to what actually goes on (Rodden *et al,* 1994). As far as the last point is concerned, in some respects the implementation of the work plan was so constraining that the engineers frequently made recourse to 'fixes' of various kinds in order to get the work done at all. Indeed, a surprising finding was the extent to which 'social and interactional issues' were constantly addressed with the aim of improving the efficiency and the quality of the work. For example, during the fieldwork the project team was reorganised to improve communication, the sharing of experience and skill, and various 'team building' exercises were arranged by management, though the latter did not attend them.

The phrase 'quick and dirty' does not refer simply to a short period of fieldwork but signals its duration relative to the size of the task. The use of ethnographic study in this category not only seeks relevant information as quickly as possible but accepts at the outset the impossibility of gathering a complete and detailed understanding of the setting at hand. Rather the focus is informing strategic decision making to select those aspects of the work setting of particular importance in informing design.

There are two points of comparison with what we have called 'concurrent ethnography' that are worth noting. First, compared to the much more focused attention of 'concurrent ethnography', and this emerged in the example we have used out of the problems of access and those of finding a clear focus for the study, 'quick and dirty' ethnography is capable of providing much valuable knowledge of the social organisation of work of a relatively large scale work setting in a relatively short space of time, and this includes what we were able to learn from the organisational problems that arose when trying to establish the research site. Indeed, it can be argued that the 'pay off' of the 'quick and dirty' ethnography is greater in that for time expended on fieldwork a great deal is learned. Second, such knowledge can be built upon for a more focused examination of the detailed

15 A review of this research is set out in Section 3 of this Deliverable.

aspects of the work which is more typical of ‘concurrent ethnography’. What the ‘quick and dirty’ fieldwork provides is the important broad understanding which is capable of sensitising designers particularly to issues which have a bearing on the acceptability and usability of an envisaged system rather than on the specifics of design. Both aspects, of course, are important.

The research also raised the problem of communicating the findings from the ethnographic study to designers, mainly because of the increased scale of the setting and the problems of finding a clear design focus. While the fieldworker learned a great deal in the study just discussed, certainly much that is useful for a sociological study, it proved difficult to hang this onto clearly formulated design objectives. In spite of this, even if used with this limitation ‘quick and dirty’ ethnography is capable of providing an informed sense of what the work is like in a way that can be useful for designers in scoping their design. In other words, although in our own case the research raised important questions about the initial design objectives, and this is not a pointless finding by any means, it did suggest useful ways in which ethnography could be used to provide designers with a better sense of the setting and its work activities.

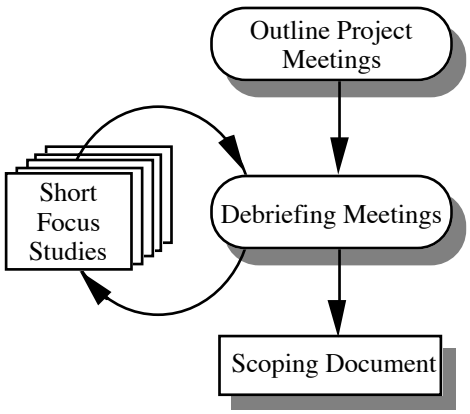

Figure 2.2: Quick and Dirty Ethnography

Evaluative ethnography

The third category can be considered as a more focused version of the ‘quick and dirty’ in that while it does not necessarily involve a prolonged period of fieldwork, it is directed at a ‘sanity check’ of an already formulated design proposal; that is, it is used in evaluating a design.

The example we draw on to illustrate was research which involved approximately three weeks of fieldwork in two branch offices of a building

society.[16] It was commissioned by a computer company to check out, using ethnography, some aspects of a model the company was interested in using for IT developments in the financial sector. In particular, we were asked to investigate customer relations at the front desk and mortgage processing.

In the relatively short period of fieldwork of about two weeks, it became clear that the model on offer had almost wholly ignored the character of 'front desk work' in branch offices, representing it as a series of information flows and tasks which could be unproblematically instantiated in the 'real world' conditions of branch work. Again in brief, much of this work was customer driven in the sense that the routine but essential work of processing the immense amount of paper that was generated was persistently interrupted by the need to serve customers or respond to their enquiries. Though, as said, much of the work was routine, including much of that with customers, there was an unpredictability to it in that cashiers did not know in advance what any particular customer wanted. Transactions with customers could be straightforward or involve complications of various kinds, neither of which was predictable. But, and customer satisfaction was an important element in the public face the company was anxious to promote, speedy and efficient service was important in the conduct of the branch's personnel. Among many other things, this meant that queues should not be allowed to build, a task which was difficult on particular days of the week due to specific local conditions. All of this was compounded by particularly unforgiving static screens which, given the complexities of even routine transactions, meant a laborious scrolling through of screens, a lack of confidence in much of the information displayed, and more, all of which occasioned considerable 'demeanour work' by cashiers in an effort to maintain "customer satisfaction".

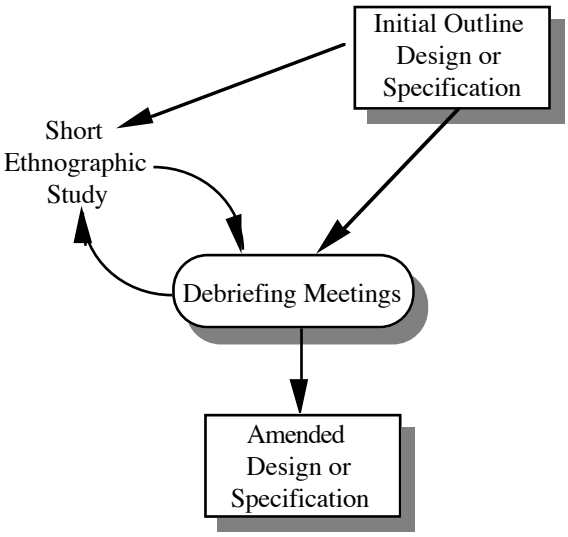

Figure 2.3: The role of evaluative ethnography

16 An outline of this study is to be found Section 3.

Another important aspect of the work, and one which aficionados of the 'paperless office' need to note, is that while there was already considerable IT technology in the branch, including help screens of various kinds, displays, remote terminals, and so on, the office space was full of 'personalised' items used by the staff in their work. Addresses of local solicitors, hints on who to call in case of particular problems, notes, and more. There was also a constant sharing of skill and experience during the work.

The above is, of course, again only a brief characterisation of the results of the ethnography. Nonetheless, they were sufficient to suggest that the model was, in significant respects, deficient. Such a conclusion is not necessarily of much comfort to designers who have, no doubt, spent many hundreds of 'person hours' developing the model. However, although in this case it reinforced the computer company's initial doubts, so much so that they withdrew from the negotiations to purchase the model, it is not difficult to see this use of ethnography in a more positive light. Independently of the commercial pressures which surrounded this project, the approach identified here could well be used to develop and improve system development. It is no part of our proposals about ethnography that it is a suitable replacement for other methods of software development or that the very real requirements of engineering are ignored in some cultic embrace of ethnography as the panacea to all the problems of systems engineering. As we have already indicated, and an important aspect of the rationale of this section, the problems are to do with incorporating ethnography into the system design process in order to improve system design while recognising that this is, again as indicated earlier, a satisficing activity and one, moreover, governed and influenced, as it should be, by an interplay of political, moral as well as technological considerations.

However, important as these matters are, the immediate point we want to make is that this use of ethnography as evaluation could be developed as a systematic means of monitoring systems in their use. Although human beings have an extraordinary ability to 'make do' with the technology with which they are provided, ethnographic studies could be useful in 'tweaking' existing systems and/or to inform the design of the next generation of systems. The first of these is, we suspect, of major relevance to many of the organisational contexts of IT use in which nothing is ever ideal. Investing considerable amounts of money in a new system is not an option for every organisation, and those who do often live to regret it. However, this proposal of continuous but modest redesign through periodic ethnographic field studies of system use may have considerable benefits if appropriately managed. And, again in support of one of the main tenets of CSCW, ethnography's focus places particular emphasis on the social context of innovation rather than simply allowing the technology to drive the innovation. It is in this context that the proposal for 'continuous but modest redesign', other things being equal, allows for persons using a system to make contributions to its

evolution and contribute their skill and experience to the next generation of systems.[17]

As an example of this process we can cite an ongoing field study of a technical documentation unit within a manufacturing company.[18] Although the research objective was to identify and analyse the characteristics of co-operative work arrangements, an important part of the study was to produce recommendations for changing the computer system, work organisation, and for training. What became clear was that the existing CAD system, despite its lauded functionalities and the support of system management, could not satisfactorily handle much of the material that came through the unit. Part of the problem was that the drawings retrieved from the CAD database served other purposes, were often inconsistently layered, contained details not required for 'customer friendly' documentation, and more. Many drawings were not available in the system but existed as paper copies in a drawing archive. In other words, in order to get the work done the CAD users had to depart from the company's prescriptions and internal standards in order to produce a variety of effects which were not well supported by the existing CAD system.

It is studies such as this which, by closely attending to the details of the work and its arrangements rather than to idealised conceptions, can inform a process of design and redesign which does not restrict itself to the computer system.

Re-examination of previous studies

One of the major problems that arises when new approaches, new methods, new systems are proposed is that not only do they challenge existing methods and approaches but there is also a lack experience and a corpus of case studies, examples, exemplars, etc., which can be used either as sensitising material concerning a domain or, in some cases, informing preliminary design. Although ethnography is relatively new in systems design, it is, as we have pointed out, a method which has been used for many decades in sociology and anthropology. Many of these studies have been related to work and occupations and while not all them have been studies of work of interest to system design, nonetheless, they can be informative.

In our own case, we have returned to previous ethnographic studies to inform the preliminary design of a Shared Object Service (SOS) platform which, among other things, is intended to handle documents in a wide variety of domains. The explicit aim of the project is to provide a set of services which allow objects to be shared by a community of users. The distinguishing feature from existing multi-user storage facilities is the focus on co-operative sharing across a group of users

17 There are a whole host of sensitivities in this which go beyond the confines of system design, namely, the willingness of people to participate given that one of the options of such investigations is to use systems to dispense with labour. This issue, of course, is not peculiar to ethnography.

18 This study was led by Risø as part of the COMIC Project and other aspects of it are discussed in Section 3.

and the provision of mechanisms which support the management of this sharing. It is intended that the shared object service provide a set of facilities for a group of users which abstracts from the properties of the underlying infrastructure to provide a well defined set of co-operative facilities (Hughes *et al,* 1993).

Though ideally 'concurrent ethnography' would have been an appropriate method to use because of the objectives of the research and the time scales involved, it was felt that much could be learned, and at relatively low cost, by using available studies even though they had not been carried out with system design in mind. What we were looking for were domains which would exhibit some of the varieties of document production, management and use as socially organised features of the work. To this end we chose studies of social work, police work and invoice processing in a multi-site fast food company (Ackroyd et al, 1992; Wattam, 1992; Anderson et al, 1989) These represented a range of domains which, though not in any sense exhaustive of document use, enabled us to identify sufficient similarities and differences to inform the preliminary design of the SOS.

The use of a range of studies also holds the promise of uncovering some properties that generally hold true and which a common service should support. For this reason, we suggest that this use of ethnographic materials is especially useful where obtaining sight of general infrastructural CSCW principles is the prime goal. In this respect, we learned the importance of history and record of use within the information store. Somewhat in contrast to current research trends non real-time interaction through access to common shared documents was prominent within our studies. The need for effective and dynamic management of access to shared information was a central demand on the shared object service. Also of considerable importance was the need to manage considerable heterogeneity as part of the shared object service and to provide facilities that maintained links between electronic and paper records.

Of course, not all ethnographies easily lend themselves to system design objectives. Ethnographic researchers, like any other researchers, have their own objectives in mind which may not always, or even very often, accommodate to the specific interests of a particular system design problem. Indeed, we had to discard a number of excellent ethnographies for this reason. However, there is another important consideration here. Unlike many of the natural sciences and engineering disciplines, the social sciences, on the whole, have failed to produce a cumulative corpus of findings to underpin any application of their knowledge. Although this feature can be overdrawn even in the natural and the applied sciences, the situation is that the multi-paradigmatic character of social research makes it very difficult, not to say hazardous, to presume that there is an available bedrock of findings which designers can consult. There is little doubt, however, that designers would find such a corpus extremely useful, though it would need to be used with due caution. In other words, re-analysing ethnographic studies could well be a useful way of sensitising designers to the social organisation character of a considerable variety of settings. This is not a substitute for the more directed

uses of ethnography when there are specific design issues to address but, depending on the design objectives, can perform a useful role in making designers aware of what to avoid and what the more specific issues might be.

Indeed, developing a corpus of such materials is one of the objectives of this Strand, and some of the preliminary work for this is set out in Section 3.

Summary and Lessons Learned

What we have tried to do in this Section is review some lessons from our experience of using ethnography in system design and proposed a number of uses to which the method can be put. As we said in its introduction, these are not to be taken as mutually exclusive, strictly demarcated methods but ones which, in significant respects, shade into one another. The motivation for this is part of a longer term attempt to place ethnography within the broader methodological context of system design in light of the focus of CSCW on 'real world' contexts of use. For CSCW it is vital that designers understand the work setting as a socially organised setting as a preliminary to design, and it is in this respect that ethnography has a role to play. In other words, the prime objective is not so much ethnography as such, but ethnography as a means of uncovering the 'real world' character of work, and it is by this test that ethnography needs to be judged in system design. Thus, it is a matter of looking at the method in light of the varied circumstances of system design, including those that arise in industrial and commercial systems development.

	Concurrent Ethnography	**Quick & Dirty Ethnography**	**Evaluative Ethnography**	**Re-assessment of previous studies**
Detail of Work	Dependant on focus of study	Outline understanding	Dependant on initial design/model	Dependant on Previous studies
Type of design information	Informing prototype through different stages of development	Overview of domain of work to inform initial design.	Check implications of design from initial model	Motivation and scope of design.
Duration of Study	12-14 Months, Balanced use of study/ debriefing	2-3 Weeks of study prior to analysis.	Analysis of original model, 2-4 Weeks of study prior to re-assessment	No fieldwork but costs of reanalysis
Influence of Field site	Strong and unpredictable	Greater ability to select field	Dependant on field site and previous model	----------
Design / Study Relation	Driven by Study	Driven by Study	Driven by initial design/ model	Driven by outline design
Form of system	Interactive workplace systems with emphasis on detail of work	Interactive workplace systems and overall system structure	Interactive workplace systems and overall system structure	General platforms to support a range of different applications

Table 1 : Outline features of the different roles of ethnography in design

Of course, the judgement about ethnography, as about any method, is a long term business and one, moreover, which will be influenced not only by its results but also by what can only be described as 'methodological fashion'. This is not to say that methodological issues are only matters of fashion, but this is one of the factors which play their part in its acceptability in at least the short term. Ethnography is currently fashionable in CSCW but if it is to survive this kind of attention then it is important that the method finds an effective voice rather than remaining content with ephemeral celebrity. In other words, we do not propose the above categories as solutions to the complex problems of system design; all we claim is that they have a useful contribution to make given the requirement for CSCW design to become sensitised to the social organisational context of work. Table 1 summarises some of the main points associated with the different roles ethnography can play in design.

There are a number of specific lessons we think worth emphasising from our varied experiences of ethnography across a number of projects over the last four years. Rather than adopting a particular stance we merely present these in summarised form below as a collection of our own pragmatic insights. Lessons of particular note include:

A variety of roles for ethnography in design

Designers require different information at different phases of the process: a point which has more than just a passing bearing on the role of any design method. While not wishing to become embroiled in the question, 'what is design?', it is a process which involves a number of skills though not all equally throughout. Indeed, some methods are intended for particular phases of system design. Process Models, for example, have their greatest utility when it comes to integrating the output of a large number of software engineers. What we have suggested for ethnography is that it has a role to play in various phases of system design and makes different contributions to them. Further, and again this is a virtue which is dependent upon the point of the fieldwork, fieldworkers can be extremely flexible in their response to the various contingencies that can arise, and deal with them as they occur. The very engagement of a fieldworker within a 'real world' work setting presents opportunities to learn much about that setting which is of relevance to design.

Responding to the pressure of time and budget

A charge often levelled at ethnography is that it is a 'prolonged activity'. As we have suggested, this is not quite the problem that it is imagined to be. Depending on the purposes of the design, much can be learned from relatively short periods of fieldwork. Indeed, within the context of design, and we emphasise this, diminishing returns to fieldwork set in relatively quickly. In other words, fieldwork of prolonged duration is not always necessary in that it would be more

effective to direct that effort in accordance with design objectives once an effective understanding of the setting of the work and its characteristics have been obtained. This point is also relevant to the claim that ethnography is an expensive process in person effort expended, a critical issue in commercial environments where, often, the contract has been gained by cutting costs to the bone. These are, of course, complicated and disputatious matters but we can say that the problem is not as severe as it might seem.[19] In any event, if the argument for ethnography is right, ignoring its value could be much more costly in inadequate systems and dissatisfied customers.[20]

The importance of focus

A major determinant in the successful undertaking of projects involving ethnography was the question of focus. In our study of air traffic control and the subsequent development of the flight strips systems a natural focus was provided by the setting for the study. Work was oriented toward a control suite which was placed within a control room with the explicit intent of making work publicly available and accessible through manipulation of flight strips. In contrast, no single location or set of work activities existed which provided a complete insight into the work of software development. Much of the effort of ethnography was in determining this focus through a series of 'quick and dirty' ethnographic studies. An existing focus was also provided by the initial design intentions within the shared object service and the existence of a previous specification within the building society.

The importance of previous studies

One of the major problems of CSCW design, and one of the reasons for the turn to ethnography and studies of the social organisation of work, is that it represents a set of new challenges. This means that, to a degree, it lacks experience and a corpus of findings to draw upon. CSCW systems are likely to be if not quite the first of their kind, at least sufficiently innovative to pose challenges in which previous systems are likely to prove of little help. This also represents a challenge to ethnography and the contribution it can make, through an accumulation of its studies, toward informing 'good practise' in CSCW design. This means paying attention to the ways in which a corpus of studies can be made available to software engineering, and others involved in the design process, ways which while enunciating general features of the social organisation of work, also display the variety of ways in which these become instantiated in 'real world' contexts. An initial attempt at such a formulation is set out in Section 4.

---

19 Of course, at present ethnographers are, typically, academics whose salaries are relatively low. If ethnography ever became a consultancy practice, the costs would be much greater.

20 Though, of course, there can be no question of guarantees here, but nor can there be with respect to any method.

System and work design

Finally, and this is to reiterate one of the main virtues of ethnography, system design is work design. This is, we would suggest, an unalterable fact about system design, let alone CSCW system design, and one which is too rarely given the emphasis it deserves. Ethnography, by its nature, has to attend to this aspect even though its studies will be concerned with 'work as it is currently done'. Thus, even though design may be concerned with developing a completely new system, understanding the context, the people, the skills they possess, what kind of work redesign may be involved, and more, are all important matters for designers to reflect upon. It is also more capable than most methods of requirements elicitation, as it ought to be, in highlighting those 'human factors' which most closely pertain to system usage, factors which are not always just about good interface design but include training, ease of use in work contexts full of contingencies which are not the remit of system design, and more. It is in respect of these considerations that ethnography is especially useful in CSCW design.

Section 3
Fieldwork Studies
Toward Corpus, Codification and Framework

Section 3
Fieldwork Studies
Toward Corpus, Codification and Framework

Introduction

This Section presents summaries of a number of fieldwork studies done under the auspices of COMIC and related CSCW projects. The studies were done at various times, conducted by various researchers and done under varying auspices. Not all of them are directly related to the design and development of CSCW systems in the sense, for example, indicated by some of the illustrations used in the previous section. Nor do they reflect the systematic application of a conceptual framework though they are deeply informed by the need to inform CSCW design through the investigation of the 'real world, real time' social dimensions of work. They are not representative of the totality of work environments so are not a sample in any strict sense of this term.

They are very much a 'contingent corpus' assembled out of developing interests in CSCW on the part of a number of researchers in the COMIC consortium and to this extent created within a context where issues of CSCW design predominated. In other words, they are studies accumulated in close associated with design issues even when these are not yet directly addressed.[21] They are studies done with the longer term aim of obtaining a better understanding of the varieties of 'real world, real time' work within their 'real world' organisational context as a basic requirement for the development of CSCW systems. In this respect, the collection represents the beginning of the development of a corpus of such materials which can be used to identify some of the main features of co-operative work in a variety of work settings.

This work is relevant to Strand 2 in a number of respects.

First, the studies have provided valuable and varied experience in the use of ethnographic fieldwork methods in the study of work for CSCW. Some of the more methodological conclusions emerging from this have been presented in Section 2 of this Deliverable. In addition, and as we have pointed out, such studies make a contribution to the long term understanding of work in its organisational context as a requirement for the development of CSCW systems. In this respect, the work reported here makes a contribution to Strand 1 of the COMIC Project.

---

21 One of the problems in adapting ethnography to the needs of designers is precisely that of assembling a corpus of studies which can be interrelated in a systematic way given the time consuming, individualised and localised circumstances of the method.

Second, a valuable resource for design is a corpus of case studies which embody lessons, principles, experiences, and such like, which add to what might be called, the 'lore of design'. More specifically, such a corpus can serve as a resource of materials about the social character of a variety of work settings which can be used as a baseline for design. In this way it can sensitise designers to the importance of understanding the social nature of work and, through this, make an important contribution toward requirement elicitation and design processes in CSCW.

Third, and a more systematic rendition of the preceding, the corpus can constitute a source of materials for formulating a framework for the analysis of the sociality of work to inform CSCW design. That is, using the corpus as the source materials for the construction of a framework which, among other things, could help codify the main properties of co-operative work as a 'real world, real time' social phenomenon. This approach usefully contrasts with the more theoretical driven approach to developing a framework which is also represented in COMIC, and which constitutes a substantial part of the research of Strand 3.[22] As we say, these are contrasts but useful ones in that being able to explore two approaches to the same problem, we may be able to learn more.

The kind of framework we envisage, and a preliminary formulation of this will be presented in Section 4, would present some generic features of the sociality of work and be closely based on the materials gathered through the kind of fieldwork represented through these studies. Its purpose would be to formulate a set of properties which characterise the complexities of co-operative work, the interdependencies of activities, the ecologies of awareness, and more, in a form which can guide CSCW designers toward those features of any work setting which need to be considered. Such a framework cannot be a substitute for a more direct investigation of particular settings but, by offering guidelines for such investigations, may help underpin the various roles for ethnographic fieldwork that we discussed in the previous section.

Such a framework could also inform the development of tool support of the kind represented by the DNP and, accordingly, one means at least of practically trying to solve the problem of integrating fieldwork studies more fully into the design and development process by developing a means of communication between fieldworkers and developers. This stage of the research, however, is only just beginning.

The Presentation of Fieldwork Studies

As we said earlier, the fieldwork studies we present here are a series of studies done over the last few years and done under varying auspices, at different times

22 See Schmidt (1991); Schmidt and Carstensen (1993) and Kuutti, (1993).

and places and by different researchers, although all were motivated by CSCW concerns.

The presentation of each study consists of:

- a table giving the date of the study, the approximate duration of the fieldwork and some relevant references, where available, from which a fuller discussion can be found;
- a summary account of the study;
- a discussion of some salient features of the social organisation of work illustrated in the study.

The summary account of each study is an edited version gleaned from various reports, working papers and, where available, published materials. Although every attempt has been made to keep close to the original texts, and checked with the relevant authors where this has been possible, the selective accounts provided here are very much the responsibility of the authors of this Deliverable. The presentation of the studies is not comparative in the sense of using a set of common dimensions across the studies. This would have required more knowledge about the social organisation of work than is currently available.

From the point of view of the strategy being followed in this Strand, the first step is to understand the particularities of work first, and then try to develop more generic analytic tools; analytical tools which are generally applicable but which, at the same time, are sensitive to the particularities of work within different domains. Indeed, and as we briefly suggested in the previous section in connection with reusing previous studies, gaining experience and ways of working with a corpus such as is represented here, is an important aspect of the work of this Strand.

We have attempted to provide as much detail as is necessary to convey an adequate sense of ‘what the work is like’ as well as highlighting those aspects which illuminate features of the social organisation of work. These are more fully dealt with in sections following the summary account of each study, ‘Emergent features of the social organisation of work’. These sections draw out features of the social organisation of work illustrated by the study. No claim is made that these are unique to the particular study concerned, only that they are suitably illustrated by it. Indeed, some of the features could be illustrated by a number of the studies. Nor are the features we identify necessarily ones which the original researchers identified.

Finally, it is important to stress that the emphasis of this and the following section is on the social organisation of work. Although the point of such studies is to inform CSCW design, design implications are featured rather less. This is not because we think them unimportant: our main concern here is with building a fuller understanding of the social organisation work activities.

Air Traffic Control[23]

Date	Approx. Duration of Fieldwork	Relevant References
1990-3	2 Months	Hughes et al 1988; Harper et al 1989; Shapiro et al, 1991; Harper et al 1991; Harper and Hughes, 1993; Hughes et al, 1993.

Objectives of the Study

The study was directed toward using ethnographic investigation of 'controlling work' to inform the development of a prototype design tool for the rapid prototyping of interfaces. The study reported here was closely built upon a previous ethnographic investigation of air traffic controllers completed in 1989 (Hughes et al, 1988).

Duration and Pattern of Fieldwork

In terms of the categorisation presented in the previous section this is 'concurrent ethnography' in which periods of fieldwork were systematically interspersed with debriefing sessions which served to inform the development of the prototype system. The fieldwork consisted of the close observation of the staff on a controlling suite, shift by shift over a total period of some 14 months but broken up into fieldwork periods of approximately two weeks followed by prolonged debriefings to inform the development of the prototype.

The Setting

Currently, the London Air Traffic Control Centre (LATCC) at West Drayton is responsible for co-ordinating all flights in UK airspace.[24] The control room consists of 8 control suites, each dealing with 2 sectors of airspace: that is, blocks of airspace defined laterally and vertically and layered into flight levels.[25] The airspace of England and Wales is divided into 16 sectors. As an aircraft flies through controlled air space along static routes it is passed from sector to sector under the direction of Air Traffic Control Officers (ATCOs) whose main task is to ensure the safe separation of aircraft as well as the expeditious flow of traffic. Co-ordination between sectors is the core task of controlling.

---

23 The study, 'Social Analysis of Control Systems for HCI Design', was funded by the Tri-Council Initiative on HCI..

24 There are other control centres at Manchester and Prestwick in Scotland. The former is responsible for a small sector of the airpace around Manchester Airport, and the latter for Scottish airspace.

25 Under certain traffic conditions sectors can either be merged or split further.

As indicated two sectors are controlled from a single suite. Fig. 3.1 shows a simplified arrangement of the Pole Hill suite which supports the 2 controllers who direct traffic in the adjoining Pole Hill and Irish Sea sectors.

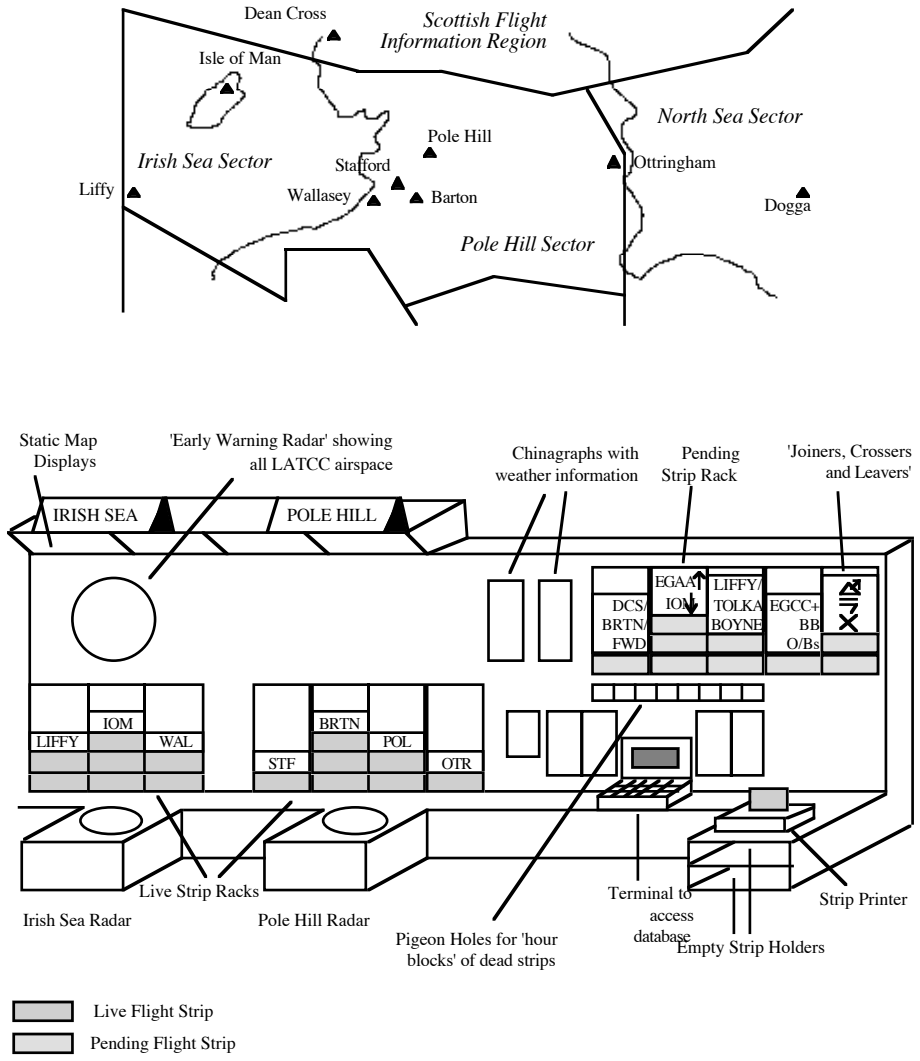

Fig 3.1: Simplified Arrangement of Pole Hill Suite and Map of Pole Hill Sector
The suite arrangement is simplified and not to scale. In particular, Designator Strips are in fact all the same size (the size of a flight strip and the width of strip rack ‘bays’ is uniform. The map is included to show how Designator Strip positions correspond to geographical locations of the reporting points they represent.

The controlling team around a suite responsible for a sector normally consists of the radar controller for a sector, a ‘wingman’, or assistant controller, and a Chief. The active controller is responsible for the moment-to-moment control of aircraft. The assistant controller is generally responsible for collecting and checking strips as they are produced from the printer, and occasionally liaising with adjacent sectors. The Chief is responsible for non-standard flights and generally overseeing the suite.

Radar controllers have 3 main instruments to support them. First, the radar display which shows a trail of blips representing a particular flight, with a data block alongside showing flight number and flight level. Secondly, RT links allow controllers to talk to pilots, radar controllers on other suites and to neighbouring airspaces. Thirdly, paper flight progress strips contain more detailed information about particular flights. These are printed from a data base of filed flight plans and show more detailed information about flights. (See below and Fig 3.2).

The Work Process

ATC is a command and control system in the sense that virtually all the UK airspace is controlled from one command centre, LATCC. However, within the centre itself there is no centre of co-ordination. Co-ordination is achieved by controllers at the suites applying the rules and procedures of controlling using the technology 'to hand'. Since, typically, an aircraft will fly through more than one sector, a great deal of co-ordination has to be done between sectors through sector controllers. Much of this can be done by 'silent handovers' when the relevant aircraft is at a bearing and height as specified in the formal procedures. Otherwise contact between the respective sector controllers is required.

At its simplest and most general, the controller's problem is to organise the stream of air traffic as and when it enters the segment of air space for which he/she is responsible and thread it together into an orderly stream before handing over to the next sector. This scheduling has to be achieved in and through making the traffic flow. It is the 'strips' which are the crucial elements in the melding of the information available to the controller which also includes radar and RT.

The controller has to apply the rules as set out in the *Manual of Air Traffic Services* which lays down, sector by sector, route by route, the procedures to which aircraft have to conform as well as defining flight and co-ordination levels for transfer between sectors.

The 'strips'

These are formatted strips of paper containing boxes of information relating to individual flights including the next reporting point on the route, time due to pass that point, call sign of the aircraft, its 'squawk code', aircraft type, planned flight path, requested cruising height, departure and destination points. Strips are printed anything up to 40 minutes before an aircraft is due at a reporting point within the sector. Assistants tear them off the printer and place them above Designator Strips which signify the reporting point corresponding to the bay. The live strips are placed under these designators for each aircraft due to pass over the corresponding reporting points, usually on a top to bottom principle so that the strip corresponding to an aircraft moves down the rack as it moves through the sector until, once it is passed on, the strip is discarded. In addition to the live strips is a rack of pending strips maintained by the assistant to allow early co-ordination of future traffic.

As instructions are given to aircraft, controllers write these down on the strips using a standard set of symbols and different coloured pens corresponding to their role within the controlling team.

Working the strips

Strips are not a real time record of the progress of an aircraft until they accumulate the notes written on them by the controller.

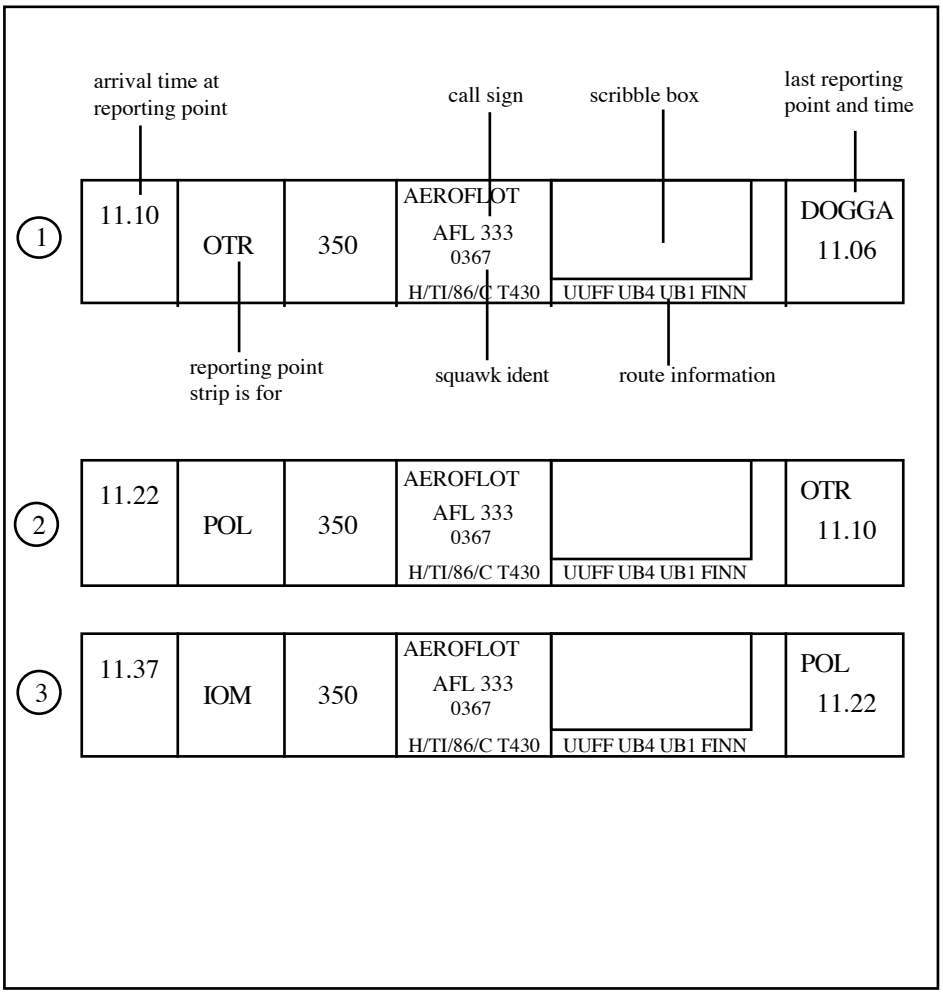

Fig 3.2 Examples of Live Flight Strips printed for a single aircraft.
These strips are not based on actual flight strips but are included as examples of strip arrangement. Strips 1 and 2 are for the Ottringham and Pole Hill reporting points respectively, and would be placed in front of the Pole Hill sector controller. Strip 3 is for the Isle of Man reporting point and would be placed in front of the Irish Sea sector controller.

The local culture describes a controller's orientation to the configuration of traffic in the sector and the problems it poses as "getting the picture" and refers to various activities involved in controlling work, not least the regular habit of the incoming controller spending anything up to 10 minutes watching over the shoulder of his or her predecessor before taking over the position in order to

"build up the picture". An important part of this is "looking about 5 to 10 minutes ahead all the time...sometimes a little further than that", in order to get an idea of what actions need to be taken 'now'.

Strips are not some detached record but a visible and vital instrument of the work. 'Working the strips' is a means of achieving a solution to the scheduling problem. Once the strips are placed on the racks in front of the controller, they can be used as a resource, among others, to organise what the controller does. The next step is to order the strips in ways that reflect the work that needs to be done. Planes fly through a sector in a sequence which is reflected in the estimation of successive times at which an aircraft will pass the reporting points. These estimates are used by the controller to order the strips on the racks so that the next plane due at any point in time is, say, at the top of the rack, the last to arrive at the bottom.[26]

The strips are not related to the real-time events in the sky as represented on the radar screens since, at the time they are first printed, the aircraft the strips refer to are not yet the business of the sector controller; the aircraft is 'pending'. The ordering of the strips at this stage is very much preparatory. Though the ordering may come to be seen, in as yet unpredictable ways, as in need of revision, supplementation or, indeed, correct, the business of the controller at this stage is ordering this resource so that what he/she may require from them is available in useful ways. It is making the strips 'come to be at hand'. Typically, the bulk of the time estimates on the strip are fairly accurate and, as a consequence, so is ordering the strips into a time sequence. This means that a controller will know immediately where a strip is, what special features need to be taken into account, and so on.

The point of this preparatory ordering is that it enables the controller to get a clear idea of what decisions will need to be made in the near future. It shapes the controller's attention in terms of what is likely to be happening in the sector, what is special, what is routine, what needs to be anticipated, etc..[27] Although some of these may be of no immediate concern, controllers mark out such problems by slightly lifting the strips out of the rack - "cocking them out" - so that when the strips become live, the controller will already have prepared them so that what they indicate is available not only to the controller but to others in the team.

'Working the strips', making them 'at hand', continues once they become 'live' on receipt of a radio message from the aircraft when it enters the sector or nears the reporting point. The controller selects the appropriate strip and moves it down to the live strip section on the rack. However, the now live strip is not just placed anywhere on the rack but is placed in sequential order - for example, the latest at the top, old and finished strips exiting from the bottom. The order reflects, and helps organise, the fact that decision-making is a sequential matter.

---

26 Practices do vary among controllers as to whether the ordering is top to bottom or vice versa, on the rack.

27 Such problems include such things as 2 aircraft estimated to reach the same reporting point at the same time and same height, non-standard traffic, faster aircraft catching slower aircraft, and so on.

As indicated earlier, the strip is used to display information which is not part of the computer output. Updating is provided by the members of the team using the strip as a notepad in ways described earlier. For example, as controllers instruct pilots to ascend or descend, or follow particular headings, etc. these instructions are written on the relevant strip, as are the pilot's acknowledgement and the attainment of the instruction. Other information may also be written on the strips, such as unusual routes, symbols designating 'cross-overs', unusual destinations, and so on. Strips come to embody the control history of a particular aircraft.

When an aircraft has crossed the relevant navigation-reporting point, a cross is put through it as a marker that the controller's work has been completed.

The way in which the controller organises the strips, organises information about the state of the sector as well as organising the work that will need to be done and, through this, organising the traffic. Strips play a key role in enabling controllers to use the radar quickly and effectively and achieve 'good technique'. The organisation of the strips is a 'proxy orderliness' for the configuration of the traffic flow. While the radar is a computer generated 2-dimensional representation of the relationship of the elements of the traffic, the strips are the means by which the patterns on the screen, and thus the sky, can be seen as the patterns that they are.

The resources available to the controller - strips, RT, and radar - are mutually determinative for what it is they indicate in that each is 'made sense of' by what the others indicate; a process of 'making sense' that is, typically, 'at a glance'.

The picture as a sequence of working tasks

Earlier what was described, in terms of the local culture, as "getting the picture" refers, among other things, to the controller's capacity to "keep it all together"; to see and give coherence to the patterns of aircraft movements under varying conditions. It is also important in working out the organisation of a set of tasks. Controlling work is not like an assembly line in which a recurrent sequence of steps has to be followed through, but is one in which the work consists of putting the tasks to be done into a sequence of steps that can be followed through. Interactions with pilots and with other members of the controlling team are not isolated contacts but part of a developing sequence whereby current steps are built upon previous ones and shaping subsequent ones, and their sense derived from this location within a stream of actions. The picture, from this point of view, is seen as a display of a set of task requirements. This is how the controller is looking at the information presented in the strips, the radar, from the RT: to see what needs doing 'now', 'in a moment', 'sometime later', and so on.

Using the rules

The rules of air traffic control, which are the responsibility of the controller to apply, are contained in the constantly updated *Manual and Air Traffic Services*. This specifies in detail the 'rules of the road' including the routes, procedures, and

the regulations to be followed by aircraft in controlled airspace and by controllers in directing the traffic. But the rules do not stand in some disembodied relation to the work of controlling, but they are integral to its activities. The rules have to be imbibed as activities that constitute controlling as a situated and artful skill. They furnish instructions for controllers to see the information 'at hand', much of this information provided by the technology.

The rules are, however, discretionary in that the rules have to be applied within the situation of the unfolding details of air traffic in real time rather than following a plan which determines in advance what a controller should do next. There are always inevitable and contingent factors, including technical troubles, of various kinds, which have to be dealt with.

The team work character of controlling

Although much of the moment-to-moment burden of controlling is the responsibility of the radar controller, it is a team which operates the sector suite. The formal specification of the respective duties of the suite personnel are set out in the *Manual of Air Traffic Services* which lays down that under normal traffic loads a suite will consist of 2 controllers, a Chief, and Assistant Controller, or 'wingman', and an assistant. This pattern is designed to service the radar controllers. However, there is more to the team work than this formal division of labour. A 'working division of labour' is closely related to the interdependence of activities around the suite and the 'at a glance' availability of information on the strips. A great deal of mutual checking goes on, again typically 'at a glance', and with a minimum of verbal interaction.

Emergent Features of the Social Organisation of Work

Real Time Decision Making, Rules and Plans

Within a sector a controller's task is to make the traffic flow safely and expeditiously using the procedures of air traffic control and the information 'to hand' provided by the strips, radar and RT. As said earlier, the controller's task is a scheduling one, but this is not according to some predetermined centrally specified plan. The procedures are a resource through which a sequenced flow of traffic is achieved in real time using the information available. Although there is no central plan to ATC, the controller's work at the console is intensely planful in the sense that he/she has to make decisions in real time, that is, making decisions 'now' as part of the unfolding implementation of the rules, and sometimes doing

so with 'what if...' contingency plans in mind should some proposed course of action prove, in the event, unfeasible.[28]

It appears that controllers 'build a picture' of the state of the sector by relating traffic to a schematic of that sector's characteristics. For example, in a sector where two routes cross each other, building up the picture means looking for the alignment of traffic in a 'single file' relative to the main axes - though not simply so since aircraft have different performances and some may be 'catching up' while others 'fall behind'. This feature is both a characteristic of the sector and a condition which the controller needs to actively create. It is a schematic in that not all flights will conform, or be capable of being directed to conform, straightforwardly to this pattern. Some will turn from one airway onto another, which may involve them crossing over flights coming in the opposite direction. Others may be completely non-standard, such as military flights crossing airlanes.[29] 'Getting the picture', the planfulness of controlling, means working out what proportion of flights fit the scheme, and with what variation, and what proportion are exceptions; assessing these for the kind of action required and planning for them within the rules and procedures.

Air traffic is constantly on the move. Even when in holding patterns they are moving with respect to other aircraft and need to be attended to. A sequence of traffic has to be achieved in and through the flow of traffic in an interplay with a sequence of decision-making. Again as indicated earlier, ATC is a discretionary system in that there is no central administration of a plan for the traffic flow and it is up to the skilled judgement of controllers, in co-operation with the controllers of other sectors, to apply the rules and procedures to the actual flow of air traffic. It is knowing the rules and procedures which enables controllers to see the patterns of traffic and their implications as a sequences of tasks which need to be done.

One of the chief characteristics of the radar controller's work is 'buying time' in order that more of it can be devoted to the difficult problems. In other words, controllers are not merely looking to schedule the distribution of their attention but are as much concerned with the kind and amount of attention that they will need to give to the aircraft under their control. Part of achieving this is using standard flight profiles for the majority of the aircraft under control at any one point in time so that more deliberation can be given to other aircraft within the overall configuration which are more of a problem. As was pointed out earlier, unlike an assembly line, the flow of air traffic is a developing sequence - an

---

28 Of course, sometimes the traffic can develop to such a pitch that the controller has no time to plan in this way but is engaged in 'head down controlling', that is, his/her attention is entirely and intensely focused upon the radar screen as flights are manouevred in close formation and where decision-making is almost moment-by-moment. Such situations are normally ones which controllers can anticipate and, to some extent prepare for. In cases like these, it is the Chief who tends to take over a more strategic planning function.

29 Military controlling in the UK is a separate activity though both civilian and military ATC are based in LATCC. Liaising between the two arms of ATC constitutes a considerable amount of controlling effort.

'entrained' sequence (McGrath, 1990) - rather than a strict one-step after one-step at a time in a predetermined sequence of steps.

In other words, the relationship of the rules and procedures of controlling to the actual situation and activities of controlling is not one of determination but one of making judgements in the light of the constantly unfolding circumstances. Occasionally, situations arise, often as a consequence of the system itself, where the rules need local amendments in order to conform to their spirit.[30] For brief example, occasionally aircraft have to be put on conflicting headings in order to deal with a particular problem, knowing that 'in a few minutes' the first problem will be returned to and resolved. Rules cannot determine their own application but need to be 'filled in' by those who use them, since it is only on such occasions that the practicalities of what the rule or the procedure involves can be made relevant for the situation to hand. It is as actions that the rules and the procedures have their voice within a pattern of work. It is the understanding rules and procedures, such as those used in air traffic control, as features of the practical actions done within the setting which, in this case, constitutes the skill of controlling. Acting-in-conformity-with-the-rules is to exercise competence in knowing the 'just whats' which will satisfy the rules.

Working Division of Labour and Team work

There is a division of labour around the suite which is specified in the *Manuals of Air Traffic Services* and these specifications are closely related to the various tasks each member of the team performs. It is a division of labour which is oriented to by members of the team and used by them as a resource for organising the sense and the activities of the working environment. The notion of 'working division of labour' is intended to direct attention not so much to the division of labour as an abstract specification, but to the actual course of activities as they are socially organised and understood by parties to the work (Anderson et al, 1989). In this case, the notion brings out a vital element of the team work around the suite in that although formally each member of the suite team has their separate tasks and responsibilities, they also attend to the activities of others in the team.

The notion of a 'working division of labour' is intended to capture the sense of the work as experienced - as mundane, ordinary and moment-to-moment - by parties to that division of labour. From this point of view, work is not encountered as a set of discrete tasks involving chunks of data and distinct procedures, but is more like a free flowing gestalt contexture with activities overlapping, moving through foreground and background in response to, and constituting, the situation as it develops. It is through a collectively developed, negotiated and evolving knowledge and practice - not, of course, perfectly matched across all members - that the work of controlling is achieved. The orientation of the individual within the social organisation of the work is not primarily to the work as a whole, but

30 See Zimmerman (1970) for an analysis of rule following in the intake section of a social work agency which makes the same point.

rather to the tissue of connections and separations as they fan out from the particular position which s/he occupies. The orientation to the working division of labour is an "egological" one (Anderson et al., 1989), concerned with engaging one's own activities, the boundaries with the activities of others, and the technology and other resources at hand, in terms of 'decisions-that-I-can-make' and 'actions-that-I-can-take', as against those that others deal with. Working through the stream of tasks means 'doing-what-I-can-do', and passing tasks on to others so that they can do what they do. Though often involving extremes of skill, judgement and co-ordination, and thus intensive 'work', its smooth accomplishment can render the working division of labour silent and virtually invisible. "Gearing into" the work (Gurwitsch, 1979) is about having the resources of knowledge and practice with which to understand these elements and relations, and inserting oneself into the flow.

Thus, and for example, much of the interaction of the team members around the suite is unspoken, routinely limited to gestures such as pointing to strips or to the radar, 'cocking strips out', reviewing them, and so on. Strips are 'glanced at', 'taken heed of', 'ignored for now', etc., not just when they arrive but continuously so making their information 'at hand' and 'in hand'.[31] The activities within the 'working division of labour' are governed by structures of relevance in that what is taken heed of, what is necessary for the task in hand, what is ignored for now, has very much to do with each person's responsibilities though not in any predetermined fashion, but rather in a way which is responsive to the requirements of the moment. What a controller feels he or she needs to know, what the screens or the strips indicate, is worked out on a moment to moment basis as the situation unfolds.

As an additional example, a 'wingman's responsibilities are mainly to do with checking strips as they are output from the computer; checking the information on them and making them ready for when needed. By attending to the controller's current work and the display of 'live' and 'pending' strips, 'wingmen' can work out appropriate headings, levels, routings, and so on, to facilitate the tasks of the radar controller. It is knowing what others in the team are doing and how one's own work can be made to fit into that co-operative effort. As indicated, much of this intra-suite co-ordination work is unspoken, relying on the 'know how' and the 'know what' and, importantly, the trust among members of the team that they 'know what they are about'. 'Cocking out' strips is another way in which mutual monitoring goes on. What this does is bring a flight to the attention of the controller, though this is also something the controller sometimes does as a reminder. Such activities, and there are many more instances that could be cited as illustration, are all aspects of the implicit social organisation of work, learned, understood and continually worked at and relied upon by members of the team. Much of this is taken-for-granted, but it is essential to the performance of the

---

31 It should be pointed out that the cooperation across suites is often very explicit. Nor are we claiming that cooperation among members of the suire team is always silent. There is often continuous talk but much of the crucial elements in the work remain tacit.

work as a seamless sequence of tasks that have evolved as a means of smoothly integrating separate responsibilities and tasks into an overall flow of work done by a team.

Team work is something which needs to be actively recreated in response to the constantly changing flow of traffic. Thus, for the radar controller to focus 'single-mindedly' on the job it is important that s/he knows that the strip delivered by the assistant has already been scrutinised by someone with immediate knowledge and experience of what is going on; that particular co-ordination problems and non-standard flights can be passed to the chief to handle; that the active management of the strips is available for all to see so that they can work up 'the picture' for themselves without having to ask for explanations, and so can be expected to handle tasks delegated to them competently; that puzzling situations or apparent errors can be discussed with other members of the team; and that they are independently and competently on the 'qui vive' for problems and looming catastrophe. Hence it is a working division of labour**,** which is social through-and-through, which supports safety.

Representation and 'at a glance' availability of the work

One of the features of the work which facilitates smooth team work is the ways in which the work activities are made available 'at a glance'. Routinely, mutual monitoring of the state of the sector and what this indicates in the way of work activities, is done without undue inquiry or other verbal interaction; it is done 'at a glance'.[32] This is not just a property of the radar controller's attention, but also one which is important for the activities of other members of the team. The 'at a glance' availability of the work is embodied in the public readiness of the strip as a worksite of controlling and a function of its placement on the suite.[33] They come to embody a history of the flight and, together with other strips, a recent history of the state of the sector and can be used as a resource for others to make sense of what is going on and 'where we are now'.

The suite is also rich in representations of all kinds, including the strips, the radar, maps, weather indications, many if not most of them standardised and conventional displays of information. Indeed, the work of ordering the strips, and its relationship to the radar display and information from RT, stand as a proxy orderliness for the 'order in the skies' created through the controlling work. As we say this is achieved through the moment-to-moment practical accomplishment.

---

32 Sudnow (1972) in 'Temporal Parameters of Interpersonal Observation.' notes 'at a glance' availability as one of the basic features of mundane interaction which is provided for and expected within public domains as a means of determing 'what others are doing'. That is, there seems to be a routine sufficiency for information gleaned and provided for 'at a glance'.

33 This feature seems to have been an inadvertent feature of the suite design, Nonetheless, its importance for mutual checking of 'what's going on' should not be underrated.

Small Office in Training Centre[34]

Date	Approx. Duration of Fieldwork	Relevant References
1994	5 days	Rouncefield et al, 1994

Objectives of Study

The objective of the study of a Training Centre Office (TCO) attached to a hotel and jointly run by a university and a hotel company, was the hope that a more effective use of IT might achieve a better balance between the various elements of the work, especially in terms of what was seen as excessive paperwork. It was also stressed that any such innovations had to have a positive impact on the staff in increasing the variety of their work rather than having to spend most of the time at a terminal. Currently, while one member of the staff co-ordinated University business, one external business, and the other involved in both, the intention was to train staff to an equivalent level so that they could deal with any business and not only make the work more interesting but also build in more flexibility. The manager was also keen to develop a "proactive customer culture" that was more able to anticipate customer needs and did not treat them as a problem. This idea of being proactive rather than reactive was a feature of all the staff comments about the expected benefits that IT could bring. By releasing extra time and improving output, IT would enable them to "crack on".

The fieldworker's remit, however, was to undertake a short period of fieldwork in the TCO, the results to be fed back to the Manager of the office who, in consultation with the other staff, determine any changes in IT strategy for the office.

Duration and Pattern of Fieldwork

The fieldwork encompassed five days in total and designed to include both 'quiet' and 'busy' days. It also included observation of a staff meeting when the coming week's business was discussed as well as spending some time with each of the workers getting them to describe the work they were doing, examining how co-ordination was done, and tracking the process of document production and distribution. Informal interviews were conducted in order to clarify issues that had arisen during the fieldwork.

As indicated in Section 2, this study is an example of 'quick and dirty' ethnography.

---

34 A version of this study was be presented at CSCW94, Chapel Hill, North Carolina, USA.

The Setting

As stated earlier, this research was an opportunistic study which grew out of a request by the manager of Management Training Centre Office (TCO) attached to a Hotel to study their work routines. The complex is managed jointly by a university and a hotel company. The TCO is responsible for co-ordinating and managing training sessions, workshops and conferences booked both through the university and external clients. The hotel side of the complex is looked after by staff employed by the hotel company. The TCO staff are employed by the university.

The TCO has a manager and three other employees who book the facilities, collate the information and disseminate it to the various units within the complex. They also provide clerical services, do promotional work and even attend to the physical layout of the training facility and its resources. Computer facilities in the office include PCs whose use is primarily confined to word processing and the production of templates and memos. Although other packages are available, such as an accounts programme, these are not used mainly because the staff do not have time to learn them.

The TCO is a ‘typical’ small office. Despite some obvious individual and idiosyncratic features, it contains many of the items which would be recognisable to anyone used to working in a small office. It is a medium sized room with access to a ‘front desk’ in the training centre where customers request services they may require. Three secretaries work in the office each with their own PC linked through a network. Besides a photocopier and a fax machine there is also a Hotel Bookings terminal which is used to check the availability and take-up of accommodation. A filing cabinet contains information about external clients, and card indexes and files above and around the desks contain the paperwork which is used in the work. Another cabinet, the Date File, contains future bookings while another, the Move Forward File, acts as both a reminder of work to be done and an indicator of the progress of a booking. Finally, on the table near the Hotel Bookings terminal rests the Diary, or “Bible” as it is referred to, which is an outline of confirmed and provisional bookings, along with customer requirements for the present and future weeks. (Figure 3.3 is a plan diagram of the office layout).

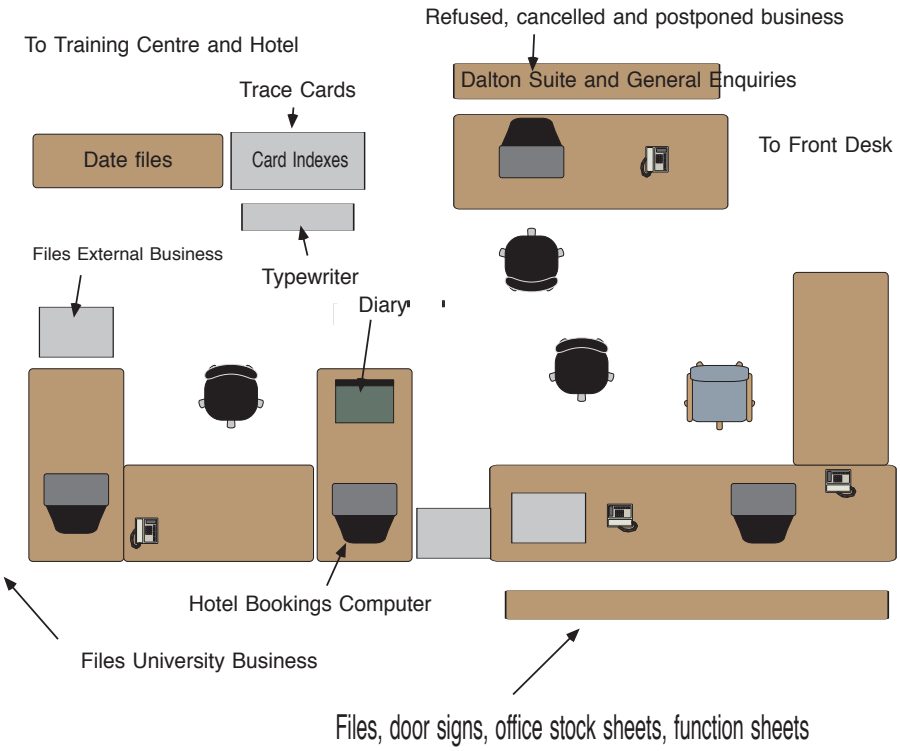

Fig. 3.3 TCO Office Layout

A key feature of the work is the need to attend to details which is reflected in the amount of paper work: details of accommodation requirements, conference needs, meals, cleaning, and so on. At the same time, the staff need to respond to client requests which, it is constantly stressed, is a major part of their responsibilities. The work, accordingly, can be described as balancing 'computer work', 'paper work' and 'people work'.

The Work Process

A division of labour organised the work into University and External business. A secretary co-ordinated each of these while the responsibilities of the other staff member was split between the types of business. In significant respects however the work was very similar. Nevertheless, the division of labour between the respective origins of the business, and the amount of work that was involved, meant that 'helping each other out' was not always as smoothly and as 'professionally' done as the staff wished. By taking responsibility for each side of the business, the member of staff was able to get a sense of the flow of work, the time scale of needed actions from the paperwork that appeared on her desk, from the 'post it' notes on Enquiry Forms, and so on. The organisational of work activities reflected the "egological principle" (Anderson et al, 1989) in which the "horizons of relevance" of each member of staff were shaped by their respective

responsibilities. This was practically exhibited in flows of information, tasks to do, things to check, co-ordinating between the two sides of the business, all of which tended to be specific to the workaday experiences of each member of staff.

Much of the work, seen as a series of information processing activities, consisting in responding to enquiries and dealing with their implications for the TCO and the Hotel.

Dealing with enquiries

Enquiries are dealt with by recording name, address, group, phone number, details of requirements, etc. on pink Enquiry sheet (Conference or meetings) or a blue Dalton Suite sheet (weddings). The available hotel accommodation is checked from the Hotel terminal. Training Centre rooms are checked for their availability from the Diary. A checks of which rates, if any, have been quoted in the past are made using Trace Card. This may require consulting the External Business File. If it is just an enquiry another memo is written to ring back, put into Move Forward File or Lost Business File. If a booking is made the Reservation Form is filled out and sent to the Hotel Reception. The Provisional Booking Form template is called up on the computer, details completed and then sent to client with confirmation booking form taken from files. These details are also entered in the Provisional Booking File and into the Move Forward File which is looked at daily by all three staff. Entered into Diary, Call Sheet and Trace Sheet which are used by the Hotel for marketing and sales analysis.

The next set of processes follow the return of the confirmation booking form. The Daily Timetable, or Function Sheet, becomes the central document in the next chain of paperwork. The original information on the Enquiry Form is copied onto the Function Sheet and other subsets of this information transcribed onto other forms and sheets. The Confirmation Booking Form initiates a memo and daily timetable accommodation requirements form and an accounts sheet. If there is a cancellation a Cancellation and Lost Business form is completed. The Dairy is completed and a copy of the Function sheet placed in the Date files.

The Business Summary Sheet, known as “the Bible”, is drawn up 1-2 weeks in advance from the Diary and the Function sheets placed in the Date File. Copies are sent to all departments and this dictates the work of those departments during the relevant weeks. The Function sheets are used for drawing up instructions to caterers, cleaners, room arrangements, coffee times, door signs, and so on.

At the weekly staff meeting the Function sheets are gone through, the bookings for next week looked at, requirements determined, etc., and instructions passed on to the relevant member of staff.

The step by step movement of information from sheet to sheet is a ‘modal transformation’ whereby one set of information is turned into another set with different procedural relevances for the actions of others.[35]

---

35 For a discussion of the notion of ‘modal transformation’ see Anderson et al (1989).

Getting the paperwork done

'Doing work' in the office was primarily seen in terms of the production and processing of paperwork - forms, diaries, schedules, records, etc.. Even selling, which the staff enjoyed and sought to develop further, was manifested in and through the production of paper. A major characteristic of 'getting the paperwork done' was its relentless repetition, much of it completed by hand. Most forms contained information which was principally, sometimes solely, derived from other forms, and many forms shared essentially similar information derived either from the original Enquiry Sheet or the Daily Timetable.

Working with interruptions

The above is an idealised version of the standard procedures for dealing with enquiries and generating the necessary paperwork in order to organise the various information and services that need to be put in place. In practice, the work rarely took this idealised form. A noticeable feature of the work, and one commented on repeatedly by the staff, was the "constant interruption" to the processing of the work as described above. These were typically high during the morning and least later in the day. They commonly took the form of the Front Desk bell, the telephone, unexpected arrivals and inquiries from colleagues from both the office and the hotel.

Paradoxically, these 'interruptions' consisted of aspects of the work that the staff said they most enjoyed, namely, contact with customers, and which 'interrupted' the work they least enjoyed, namely, the paperwork. Another irony is that the 'interruptions' were often events which initiated the paperwork in the first place, that is, enquiries about bookings and other arrangements.

Emergent Features of the Social Organisation of Work

Co-ordination through paperwork

Many of the work tasks consist of processing documents of various kinds as co-ordinating 'mechanisms of interaction' not only for the office but also as a means of communicating actions that need to be done by the hotel staff. This co-ordination was achieved through various 'transformation processes' in which a base of information was selectively drawn upon and recorded for the relevant files and ensuring actions. Typically, the documents which constituted these various processes were strictly formatted and organised into a set of routine processes.

This feature of co-ordination of actions through paperwork is a very ordinary one in bureaucratic organisations large and small. What this feature calls attention to is the ways in which documents, treated as elements within a socially organised pattern of work, are indicative of and implicative for a variety of actions (See

Hughes and King, 1992; King and Hughes, 1994). It is in this respect that the documents are organisational objects in that they are seen and understood as representing and having implications for organisational activities. They are not simply records of 'actions taken', though this is one of their important functions, but can also have implications for 'actions to do' and, more than this, 'actions taken' and 'actions to do' by relevant organisational personnel. That is, the documents embody instructions as to what is to be done, when, and how. It is this last feature which displays the fact that documents such as these are integrated within the normative structure of organisation. Documents can be 'traces' and, thus, not only constitute a record of 'what was done' but also who 'should have done it'.

Co-ordinating work in real time

Although the office is small, a considerable amount of the work deals with co-ordination of activities as information resulting from enquiries is modally transformed through the various processes. These transformations take place in real time and crucial aspects of them are dependent on the actions of others, such as a customer confirming a booking, presenting requests for facilities, and so on, as well as being 'interrupted' by other tasks which are unpredictable as to when they occur and what they might occasion. In other words, the timing of the sequences of activities was dependent upon events external to these sequences, some of these events necessary to the completion of the sequences, such as the confirmation of a booking, and others entirely contingent and unpredictable, such as the front desk bell or a telephone call.

Ecology of co-ordination and awareness

The organisation of the paperwork as visible in the arrangements of the office space: the files, filing cabinets, card index systems, computers, and so on. Much of what constitutes 'doing the work' consisted of movement around the office taking forms from one file to another, amending or copying the details of different forms and locating them elsewhere in the ecology of the office. The spatial organisation of the paperwork - where files are, which files contain which forms, etc. - constitutes a 'working map' for the workers in the office and provides the 'at a glance' availability of what people are doing, what stage they are at in the work, how quickly they are getting the work done, and so on.

In many respects the ecology of the office is a standardised one for those who use it. Its layout provides a sense of the organisation of the work which is constantly reproduced through its use. 'Knowing the job' involves not just knowing the office procedures but, as part of these, knowing the spatial organisation of the office in its details.

The ecology of the office is a socially constituted arrangement which facilitates the co-ordination and awareness of the work. For those who 'know the office', its arrangements offer affordances for 'seeing at a glance' where someone is in the

course of the work activities.[36] This is not just a matter of seeing 'in general' how the work is going, what more needs to be done, and so on, but is also a means of coping with 'interruptions' in the course of 'doing something else'. The 'present' state of the office space, the accumulation of paper documents on a desk, attached 'post it' notes, jottings, memos, and the like, are examples of 'territorial markings' used by staff to indicate where they are in the course of current work.

An important feature of this 'marking' is those affordances arising from paper. Not only are piles of paper visible but because of the paper driven character of the work, there is a commonplace relationship between, for example, the placement of the amount of paper on a desk and the amount of work done or 'yet to do'. Similarly, 'interruption' from the current processing of documents can be handled through a procedure, "begin where you left off".

The ecology is also a moral order in that given the 'working division of labour' which is also a division of responsibilities, 'missing items', 'forms not yet done' or 'delayed', 'work not yet completed', can be seen as 'someone's responsibility'. This is not necessarily an occasion for blame but one for 'seeing who needs help' when 'pulling together' is needed.

The ecology of the office is an achieved public space using the arrangements of its 'objects' to facilitate the shared awareness of the flow of work, including its 'interruptions' and, through this, affording a means of co-ordinating in real time the division of labour through the paper records.

Local Knowledge

Working in and through the system depends upon 'local knowledge' which is not simply an adjunct to the pattern of work activities, but essential to them. Such 'local knowledge' is knowledge of the particularities of the work as exhibited in its day-to-day routines. Although much of the paperwork deals with standardised formats, and much of it is routine and repetitive, the particularities of customers, and their requests, are noted and recalled by the staff, where files have 'gone', war stories, strange names, and all the multifarious ways in which experienced workers display their 'know how' and the 'real world' organisation of the work's activities. It is knowing how to use the system as an ordinary, taken-for-granted, commonplace organisation of work activities.

---

36 See Gibson (1979) for an examination of the notion of affordances. However, we agree with Sharrock and Anderson (1993) that such affordances are not so much cognitive as socially constructed.

Branch Office in Building Society[37]

Date	Approx. Duration of Fieldwork	Relevant References
1993	Two Weeks	Randall and Hughes, 1994

Objectives of the Study

This study arose out of an invitation by a financial services system developer for a 'sanity check' on a structured model of the transactions and information flows within a building society that was to be used as a possible basis for system development. Our brief was to look at the daily work of the front office in order to throw light on whether the presumptions made in the model about the nature of the work were consistent with the work as actually done of a day-to-day basis.

Duration and Pattern of Fieldwork

The ethnographic fieldwork lasted for approximately two continuous weeks in the 'front office' of a building society. Although this is not an 'in-depth' study, sufficient material about the character of the work was provided to suggest that in crucial respects the model poorly expressed important facets of financial transaction processing, particularly in respect of the work of the cashier, the 'interface' between the customer and the organisation.

The study can reasonably be described as an 'evaluative ethnography' since its main aim was to test the assumptions of an existing model.

The Setting

The work place was a typical office within a building society consisting of a counter behind which is a large office space containing desks, telephones, cupboards, filing cabinets, ledgers, typewriters, and so on. To hand the cashier has a menu driven computer screen which can be used to provide information about accounts, recent transactions, etc. Under the counter are relevant forms, such as paying in slips, cash drawers and stationery. Around the desk space is other relevant information including memos from Head Office as well as various 'post it' notes for more local use. Elsewhere in the building are offices for the manager and interview rooms for dealing with more prolonged customer business, such as mortgage applications.

Customers enter the branch and go to one of the counters to be served. If no counters are currently available, as regularly happens during busy times of the

37 In the UK building societies were financial organisations mainly concerned with savings and lending money for the purchase of homes. While this is still their main role, they are now beginning to branch out into other financial services on the model of banks.

day, then customers queue in roped off lanes and then wait until a counter is available and go to that on a 'first come, first served' principle.

The office investigation was in one of the urban branches of a relatively large regional building society.

The Work Process

The cashier is the 'interface' between the organisation and the customers of the building society. The cashier directly services the customer and passes on other information and work dealt with by others in the organisation. To hand is much of the equipment and information that the cashier will need in the matter of course treatment of customers and their requests. Thus a customer request or transaction is entered into the system by the cashier, and starts its journey through the organisation. It is coded and standardised as information and fed into bureaucratically organised channels. In many ways, the work process consists of classic office procedures in which activities are defined by rules which determine what should follow from what. In this way the system achieves a high level of predictability and efficiency of functioning.

The customer driven nature of the work

However, although the work of the cashier is procedurally driven in that, with respect to a given transactions, the steps are laid out, much of this work is customer driven. The flow of work begins with the customer and is customer driven in that it is what *this* customer wants *here and now* which will determine what work the cashier will have to do *here and now*. Although much of the work is routine, such as bringing ledgers up-to-date, filing, and all the other typical activities of the office, there are variably busy periods of the day and of the week, it is the customer who takes priority: a priority reinforced by the building society's policy of 'customer care'.

Customers structure their enquiries in a number of ways, including making a series of requests at the beginning or, alternatively, waiting until an initial request has been completed before making a second, or simply making a simple request such as for a withdrawal. However, the cashier does not know the nature of the request until the encounter is underway. It may be a single request or it may consist of multiple requests.

Transactions can, accordingly, become more or less problematic depending on the various ways in which cashiers and customers jointly orient to them. It is considerations such as these which make the notion of a 'flow of work' less than straightforward.

The flow of work is less than predictable depending on the nature of the customer request and includes not only the unpredictability of whether it is a single or multiple request, but also the various contingencies that can 'normally' arise in processing requests. These include problems generated by a customer's oversight, such as a passbook not signed, identification not provided, failure to

understand account rules as to when cheques are cleared, and more. Each customer represents an element of unpredictability. Although it is perfectly possible to provide aggregate descriptions of a typical day, such as how and when such transactions occurred, when such transactions are likely to be more or less frequent, what cannot be known in advance is what *any* particular customer transaction will be and what it will involve.

From the customer's point of view, and this is one of the main concerns of the building society's 'customer care' policy, satisfaction is an amalgam of a number of factors, not least the cashiers' ability to keep a queue moving so as to minimise waiting time while, at the same time, meeting the expectation that requests and enquiries will be adequately dealt with. Branch offices speak of the need to maintain "customer confidence" and a significant part of this is avoiding unnecessary waiting as well as dealing effectively with customer requests. Maintaining "customer confidence" is not a single, discrete activity but one which permeates throughout every activity which involves dealing with customers, from the interaction over the counter, to dealing with requests and inquiries promptly, not allowing queues to build up, and so on.

Customer enquiries

By far the most common source of complication, and the most difficult to deal with, is the customer enquiry. These may be in person or by telephone, and may be about a range of matters, including the status of an account, standing orders and direct debit payments, advice about insurance claims, advice about new accounts, what to do about a lost passbook, and so on.

Enquiries are often made in the course of other transactions and are nearly always time consuming, not least because they interfere with the smooth execution of the transaction itself. They are also unpredictable. Customers can, at any time, make what for them is a relevant enquiry regardless of how irrelevant it may be for the cashier's work-in-hand. The cashier is simply the first point of contact for enquiries which may, in the end, have to be redirected to sales, financial advisors or others in the organisation. Enquiries do, of course, have to be handled and as expeditiously as possible. The cashier needs to make a decision about how to deal with an enquiry bearing in mind the dual need to maintain customer confidence and, if the enquiry is in person, keeping the queue moving.

As we have suggested, keeping the customer satisfied and, as part of this, maintaining customer confidence is, for the cashier, a matter of juggling with a potentially conflicting set of demands: expediting the requests of *this particular customer* to his or her satisfaction, as well as maintaining the flow of customers through the process.

The technology to hand

Another ingredient in the balancing act that the cashier needs to perform is the technology to hand. Earlier it has been described how even the simplest of

transactions has the potential to become complicated and additional complications can arise from trying to weave the technology-to-hand into the interaction with a customer. For example, possibly the most time consuming aspect of the process of opening an account is inputting information into the Static Displays. Sometimes cashiers will leave this task until a customer has left rather than get them, and any others who may be in the queue, to wait for what can seem a long period. Technically, however, this is a breach of procedure.

The technology is based on a range of likely customer requests and designed to process these. Once the customer has requested a particular service, then the relevant screen display is interrogated to expedite that service. The system consists of proprietary software designed specifically to manage financial information. However, the current technology can create further troubles for customer-cashier encounters. The following is a lengthy extract from an interview that arose from a customer's discovery that he was in arrears on his mortgage payments and his subsequent inquiry as to how this arose. It portrays quite dramatically how both the operator's and the customer's confidence eroded due to 'inadequacies' in the Screens.

Customer: I came in and got told I was £700 in arrears and then, no, it was only £220 and I want to know why...we left our payments running at a higher level thinking we were paying the interest off, which we obviously weren't...and I came to in to find I was £220 in arrears and I said "No, it couldn't be..."

Interviewer[38]**:** Are you on Annual Review? (No) You've not been on arrears since the beginning. What we've got to try to do is find out where it's gone wrong...let's have a look...

INTERROGATES SCREEN

What it is, your insurance has just gone in, which is not on your mortgage...

Customer: I thought it was

Interviewer: No, only if you're on Annual Review

Customer: It doesn't account for the £700 I got told I was in arrears by first of all...that figure there (POINTS TO SCREEN)

Interviewer: No, but this column doesn't mean much, only the last figure...If you're not on Annual Review we'll have to work it out annually...It's the 29th September every year, it should be on your statement... (FIVE MINUTE DISCUSSION OF ANNUAL STATEMENTS).

Customer: At one point I was told I was £900 in arrears...where did they get that from?

Interviewer: I can only apologise...you shouldn't have been told that...your arrears balance has been reducing as you've been paying over the odds

Customer: But shouldn't there be another 8 months of £30 (overpayment)?

Interviewer: Have you ever been on Annual Review?

---

38 Note that the 'interviewer' here is one of the cashiers of the building society not the fieldworker. The interview was conducted in one of the rooms in the branch office set aside for this purpose.

Customer: I'm not trying to be awkward like...I just want to know...

Interviewer: Well, your arrears must have been getting worse and worse 'cos you've never paid your insurance.... there's four lots altogether since 1988...this column here is what we charge and this column is what you actually pay...(CALCULATES INSURANCE ARREARS) Now that comes to 31200...what we've got to do now is find out what you've overpaid...(POINTS)...that was the first overpayment you made...your £159.80...now what's this £274.46? (SPENDS SOME TIME TRYING TO FIGURE IT OUT)...ah...I know what it is, you've paid your first insurance payment and your mortgage together...there should be another £144...It might be on the next screen (ENTERS CODES)...there it is...we've charged it to you' cos you haven't paid it...the £144 is the additional premium for your contents 'cos you took out Buildings and Contents...can you see that? Do you agree with that?

Customer: Yes, yes...I can see that...where did that £510 come from?

(INTERVIEWER USES SCREEN AND CALCULATOR TO TOTAL SUMS)

Interviewer: Can you see what I'm doing? (CARRIES ON USING SCREEN)

So it comes down in the end to (EVIDENTLY REALISES THAT THE APPARENT ARREARS ARE STILL TOO HIGH) you must have paid one or two of those insurances...so I'll have a look (SCROLLS THROUGH SCREENS AGAIN)..nowhere near it.

(EVIDENTLY HAVING CONSIDERABLE DIFFICULTY)

Interviewer: You've never been on Annual Review?

(PHONES MORTGAGE ACCOUNTS)

Interviewer: Can you give me a hand on this one? Annual Review, he doesn't think he's every been on it but I'm wondering whether in fact he has....I've taken all that into account...Screen 09, is it?...Start of year....I've added all the insurance premiums together....they're coming to over a thousand...it's got to be at least two on the insurances, but usually it tells you the day they've come off, doesn't it, on an 07? (POINTS)....that was for your bikes...(ON PHONE AGAIN)...so where's the rest of the payments coming from?...'cos I'm still in arrears

Interviewer to Customer: She'll do a breakdown and send it out to you 'cos it's going to take some time.

Customer: I just can't understand it...they keep sending us letters if we want to reduce and at the same time they're letting the arrears build up (INTERVIEWER TRIES TO REASSURE CUSTOMER) Am I looking at £700 arrears or £200? It doesn't fill me with a great deal of confidence.

Interviewer: I imagine the computer will be right, but no, I understand...

Customer: Do you need any more information from me?

Interviewer: You can't recall making any payments over the counter? Cash sums? (No}...I'll have a word with my boss, 'cos she's very good at this sort of thing.

Customer: And when will I hear?

Interviewer: Oh....first thing next week?

Customer: Shouldn't I have been told...don't you think I should have been told?

Although the queuing problem is not relevant here in that it was an encounter in an interview room, it is a graphic example of the difficulties that can arise in weaving the technology into the interaction with a customer. The initial problem arose because a cashier who originally dealt with the enquiry was unable to establish the exact reason for arrears on a mortgage account and, indeed, gave the customer three different estimates of the amount concerned.[39] There were pressing reasons for providing a satisfactory and plausible explanation for the confusion. But this is not what occurred despite the fact that the officer approached the problem with the confidence associated with an experienced practitioner. At a very early stage, a likely cause for the confusion is identified and attempts are made to demonstrate the source of the error, namely, the 'failure' to pay an insurance premium. The problem gets worse as candidate explanations fail. The realisation that the problem is not at all clear from the information available on the screens prompts an appeal to further expertise via the telephone and, eventually, an apologetic admission that it will take further investigation to locate the source of the problem.

What was also clear, though not from the edited extract, was that interaction with the screens in the course of the interview caused a number of difficulties. Interrogating the screens is time consuming and officers report that they have considerable difficulty in conducting smoothly flowing conversations with their clients.[40]

One of the major problems with the system in use is that it is structured according to the flow of transactions not to the flow of enquiries. The display of an account, for example, will record every transaction, the amount and the overall balance. It will, however, give no clear indication of what payments are for, and to whom they are made. The orientation brought by a particular customer, on the other hand, will be driven by relevances which are the concern of *that* customer. It is the problem of taking these into account in an elegant way which often disrupts the flow of work and often generates further inquiries to other departments.

It is in light of these problems with the technology that in many instances cashiers prefer, instead, to rely on paper based information.

---

39 Cashiers in the building society concerned do not receive substantial training in the complexities of mortgage transactions which was why this enquiry was referred to an officer with over a year's experience of mortgage processing. The officer in our extended example was relatively practised in the use of the system. However, the requirement for an effective knowledge of information resources, such as the Enquiry Screens at the same time as trying to maintain the confidence and trust of the customer, is a daunting one especially in an organisation which, historically, has not placed a high priority on training. Explicit reference to the inadequacy of training was common and a typical example is as follows: "It's the training really...course people in the branches are keen to do this kind of thing, to learn about it, 'cos it makes their life easier. But they don't have the information in front of them...they have to look for it...even the basic information screens..."

40 It should be pointed out that the system, though not state of the art, is a standardised and widely used one for such offices and is reasonably serviceable for most routine transactions.

Behind the counter work

Dealing with customers at the front desk is not the only work that cashiers do. There is also 'behind the counter' work such as filing, completing ledgers, bringing manual records up-to-date, answering correspondence, and the myriad of other activities which constitute the work of a building society branch. Even this work, however, is interrupted by customers 'to be finished later'. It has to be 'put down' as customers appear or the telephone needs answering. At any moment of the day, the office will show tasks that are awaiting completion.

Much of the work is paperwork using ledgers, customer files, and other documents, records, forms and, just as in the cashier's work space at the counter, there are a number of local and personalised information resources pinned up on desks and on the walls.

The officers orient to each others' skills and local knowledge in getting the tasks done. Personal 'bibles' are used on a regular basis as resources for solving problems and making decisions. They draw on each others' work and help each other out. A considerable amount of time is spent 'file chasing', seeking the whereabouts of a relevant document - a problem which is compounded when the customer is present or on the telephone. Files can go to other departments but customer enquiries tend to come to the branch office.

Emergent Features of the Social Organisation of Work

The locality of organisation

One of the features we want to draw out of this study is the locality of organisation. By this we mean that although for differing purposes organisations can be treated in various ways, such as goal-oriented units, as brains, as organisms, as cultures, as political systems, and more (See Bowers, 1994, for a review of some of these metaphors), the point we want to emphasise here is that however else we may conceive of organisations, as identifiable schemes of socially organised activities their substance lies in the activities of their members. It is in the board rooms, the workshops, the laboratories, the offices, the branches, telephone calls, and so on, where the 'real world' of social organisational life is to be found.[41] In other words, the activities, the work, are locally organised.

There are three further points which need emphasising about this view.

First, 'local' in this context is not simply a geographical locality, though this is by no means an unimportant aspect to it, but to the 'here and now' where this is

41 We recognise that this remark avoids the vigorous and extensive debate about the character of organisations. See, for example, Burrell and Morgan (1979) and Morgan (1986) as well as the summary in Bowers (1994). However, while acknowledging this it is beyond the remit of this Deliverable to do any more than note this. It is, of course, a concern of Strand 1 of the COMIC Project and we will have further comments to make on these issue arising from subsequent work in this Strand.

understood to point to a collection of orderly activities done by persons and bound within particular contextures of relevancies (Lynch, 1993). What it resists is the attribution of some generic principle by which activities within a domain are organised. For example, the building society's Customer Care policy was formulated as a set of principles to be followed by all its employees. However, what that means in terms of 'real time, real world' activities has to be worked out moment-by-moment in dealing with all the contingencies that arise in the local circumstances of, in this case, the branch.

Second, in opposing the imposition of some generic principle of orderliness on socially organised activities, be these concepts such as the division of labour, information flow, cognitive structures, etc., the notion of 'local organisation' helps to draw attention to how co-operation and co-ordination within a 'working division of labour' are achieved as practical matters for the parties involved. Achieved, that is, as "situated actions" (Suchman, 1987). Accordingly, and to develop the point a little more, it helps make available for empirical inquiry the ways in which co-operation and co-ordination are achieved as real world, real time phenomena.[42]

Third, the notion also makes available for empirical investigation how the trans-situational properties of organisation are achieved and sustained through locally organised actions. In other words, the idea does not deny a sense in which organisations can be seen as holistic entities; what it does is propose a way in which these can be understood in terms of the local actions and relevancies.

The local management of conflicting demands

The above remarks about the locality of organisation are also relevant to understanding the shape that conflicting demands can have within a work domain. In this case, the Customer Care policy of the building society has very practical implications for such prosaic matters as "keeping the queue moving", "being polite on the telephone", "going out of one's way to help", and so on. However, these did not always smoothly resonate with other demands placed upon staff in the course of their daily activities. Prolonged customer enquiries, for example, had direct ramifications for 'keeping the queue flowing' and, hence, for contented customers.

The point is that these conflicting demands appeared in the local environment and had to be dealt with 'there and then' as part of the day's work.

Aggregate predictability and local contingency

Although much of the work within the branch is routine, and designed to be so, and although staff were able to predict with reasonable accuracy the rhythm of the work by time of day, days of the week, weeks of the year, and so on, particular interactions with customers were radically unpredictable even though many of

---

42 This is one rendering of the notions of 'mechanisms of interaction' and 'articulation work' which are concepts prominent in Strand 3 research.

them turn out to be routine and straightforward. As the study shows, customer requests were often a source of 'troubles' of various kinds in that while the customer 'held the floor', so to speak, the cashiers were required to respond to 'whatever else' the customer required. It was these interactions which often stretched the limits of the cashier's 'resources to hand'.

This loose relationship between aggregate predictability and local contingencies raises important issues about the ways in which work activities are represented. Although it is more than possible to provide adequate characterisations of the 'what' and the 'how' of the work, it is also important to know about the 'when' and, in this example, the 'when' can also introduce elements of unpredictability for the flow of work.

Demeanour work and 'getting round the technology'

The unpredictable nature of customer enquiries was often occasioned 'demeanour work' on the part of the counter staff especially when these created problems with the technology. As indicated earlier, the rationale for the design of the technology used was in terms of transaction flow rather than the flow of enquiries and because the system was relatively dated, was cumbersome to use with the flexibility required for processing customer enquires. Nonetheless, one of the important ways in which the staff coped with inflexibilities was their skilful demeanour work which not only helped to maintain customer confidence but also 'smoothed' the trouble.

Despite some problems with the technology used in the office, by and large the staff has evolved ways of working which 'got round' its limitations and which, in effect, contributed to it working more effectively than it might have done. Such activities included not only using notes, manuals, personal diaries, and such like, but also 'smoothing' the interactions between customer and staff.

Manufacturing Engineering[43]

Date	Approx. Duration of Fieldwork	Relevant References
1994	Ongoing	Borstrøm, Carstensen and Sørensen, 1994a; Borstrøm, Carstensen and Sørensen, 1994b; Borstrøm and Sørensen, 1994c; Carstensen, 1994; Carstensen and Sørensen, 1994; Sørensen, 1994a; Sørensen, 1994b; Sørensen, Carstensen and Borstrøm, 1994.

Objectives of the Study

The aim of this study was to investigate relatively large scale manufacturing to further the understanding of the ways in which large numbers of people with different specialisms cooperate. Its particular focus was to study the ways in which various artefacts were used to organise and manage the complexities arising from the interaction of a large number of people within a design process.

Duration and Pattern of Fieldwork

The research, which is ongoing, used a combination of qualitative interviews and ethnography to study the S4000 Project, a new instrument for testing raw milk. Over 20 long interviews were conducted with staff on the project as well as long periods observing the work and attending meetings of the design staff.

The Setting

Foss Electric is a company whose business is the design and manufacture of equipment for measuring quality parameters of agricultural products. It has few competitors and is among the largest in the world. Research and development is based in Denmark with subsidiary companies in the UK and in Germany and its products are sold all over the world. The holding company employs about 700 people. The products it manufactures are used for measuring the compositional quality of milk, such as its fat content, amount of protein, lactose, bacteria, etc., the composition and micro biological quality of food products and for measuring grain quality. The measurement technologies typically used include infra-red, fluorescence microscopy and various bacteriological techniques.

---

43 The research was partly funded by the COMIC Poject, and also by the CODEM project sponsored by Fisker og Nielsens Fond, Ib Henriksens Fond and the Danish Technical Research Council.

The company's customers are laboratories, slaughterhouses, diaries, etc., and due to the high market specialisation and the centralisation of laboratories, there are few competitors. Nevertheless one of the strategic goals of the company is to innovate toward new, better and faster measuring techniques. Accordingly, research and development are essential activities. The company has implemented concurrent engineering techniques in an effort to integrate manufacturing and development.[44] Typically projects include different specialisms including mechanical engineering, hardware and software design and, on some projects, specialists in optics and chemistry are also involved.

Work Process

Much of the work is structured around projects which are seen as following a number of steps:

- a product concept is formulated which defines the overall architecture and the technologies involved;
- a small number of functional models, or mock-ups, are built;
- between five to ten prototypes of the instrument are fabricated to verify detailed designs;
- a test series of five to ten instruments are constructed to test the manufacturability of the instrument.

This study concentrated on one of the company's large projects, System 4000, a new instrument for testing raw milk.

The S4000 project

For Foss Electric, this was the first 'system' instrument in that several instruments that could be plugged in and out were to be integrated in the one machine. It also introduced new measurement parameters, such as testing for urea and citric acid. The measurement speed was also enhanced compared with previous machines. It was also the first product with an Intel-based 486 PC built in so that configuration, control and operation of the instrument could be made via a Windows interface. The instrument consisted of approximately 8000 components grouped into a number of functional units, such as the cabinet, pipette unit, conveyor, PC, hardware, flow system and measurement unit. More than 50 people were involved in the project which lasted for approximately two and a half years. The core personnel included hardware designers, electronics engineers, software designers and chemists. In addition, draughtsmen and personnel from production, the model shop, quality assurance, quality control, marketing and management were also involved at various stages.

Figure 3.4 is a picture of the S4000.

---

44 On concurrent engineering see Harrington (1984) and Helander and Nagamachi (1992).

Fig 3.4 : The S4000 system being tested in the Quality Control Department

The S4000 was not only a large project for the company, it was also considerably more complex than previous ones. Not only was it an engineering challenge to integrate what were previously separate instruments in the one machine, but control and configuration required over 200,000 lines of software code as well as using a Windows application in which the software designers had no previous experience.

Organising for complexity

There were many sources of complexity in the S4000 project including those arising from the technicalities of the design itself, loosely formulated initial guidelines, conflicting requirements arising from the concurrent engineering process, co-ordinating large numbers of specialists, and so on. Foss Electric uses a matrix organisation for its projects in which projects are treated as organisational units with the project manager serving as ‘head of department’. All participants in the project are located in the same area of the building. To ensure that overall goals were met, scheduled project meetings, some weekly and some bi-monthly, were held. In the most intense phase of the project over 27 regular meetings, each involving from 6 to 26 participants, were held. In addition to these more formal meetings there were innumerable informal meetings of various members of the project team to deal with various problems as they arose.

Of special interest to the study were the various artefacts, some of which were developed during the course of the project, used to keep track of the ‘state of the

project' during its course; artefacts designed to schedule the relationships among the staff of the project and the interdependencies among the tasks.

The CEDAC Board

The Cause and Effect Diagram with the Addition of Cards (CEDAC) system was originally developed in Japan in the 1970s and subsequently adopted by many North American and European companies. The purpose of the CEDAC system (see Fig 3.5 below) is to reduce the number of manufacturing defects through continuous improvements by enabling team members to pool and make use of their accumulated knowledge and experience. Its main purpose in the project was to register defects, shortcomings and problems found in the production of the prototypes and the test series of instruments. Improvements in productivity and product quality was anticipated by involving all the employees in solving problems.

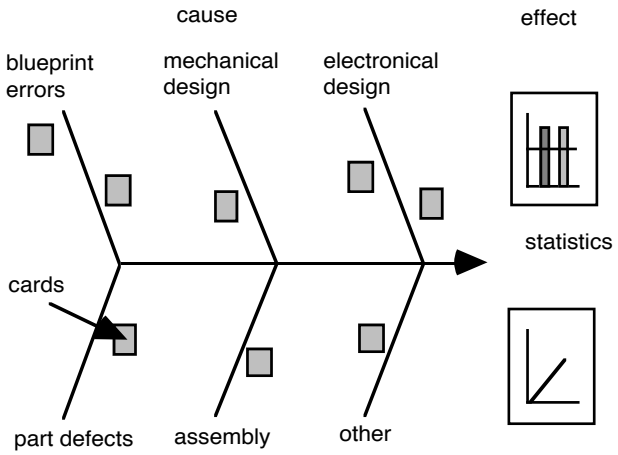

Fig 3.5: CEDAC Board as implemented in engineering design, process planning and production of parts for functional models and prototypes at Foss Electric

The system is instantiated on a 1 x 2 metre steel board placed on a wall in the shop floor. When a problem is encountered or a suggestion for improvement made, an observation card describing it is attached to the board with a magnet. Examples include poor and insufficient drawing and tool specifications, inadequate processes, poorly thought out manufacturing principles, and so on. People from the various aspects of the project - mechanical and electronic designers, draughtsmen, project leader, process planners, as well as shop floor workers - participated in weekly CEDAC meetings to discuss the problems and suggestions on the cards. At the CEDAC meetings, the observation cards are discussed, and for each observation card a suggestion card is added on the right-hand side of the 'fishbones'. Suggestion cards can be classified in the following categories:

1. Of interest
2. In progress
3. Proposal not possible
4. Being tested
5. Successfully tested
6. Test yielded poor result.

Statistics on the accumulation and processing of the problems are also placed on the board. At Foss Electric, CEDAC systems are implemented in the production, assembly and quality control departments.

Product Classification Scheme

One of the major organisation problems in any large project is managing and keeping track of the possibly thousands of components which may be used in manufacture. The Product Classification Scheme was a means for distributing, storing and retrieving CAD models in a data management system. The CAD system was introduced in 1990 at the same time as the S4000 project began. One of the advantages of using a CAD system is the opportunity it provides for reusing old components in new products, which is both time-saving as well as supporting the use of standardised components. To facilitate browsing the CAD system's database, the company developed a classification scheme for the components and units for all the instruments it produced.

See fig 3.6 for an example of a small section of the scheme. The different sub-categories belonging to the category 'valves', which are a further sub-divided into hydraulic and pneumatic categories.

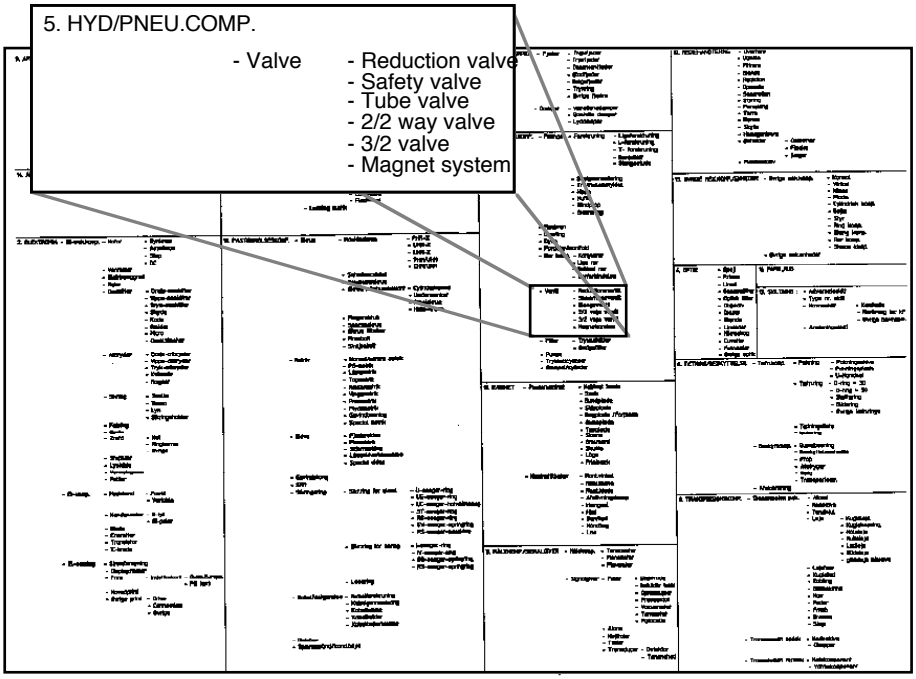

Fig 3.6: A small fraction of the product classification scheme

When a CAD model for a component has been identified, it is classified according to the scheme. Use of the classification scheme is not stipulated in any organisational procedure but is conventional. The classification is most often determined by consulting an alphabetically ordered print out from the database. The paper based system is organised into a tree structure with classes, categories, sub-categories, and sub-sub-categories. There are 16 classes and approximately 340 different categories, sub-categories, and sub-sub-categories.

The classes and categories used in the scheme are a mix of functional and geometric properties of the components and the units. Some of the categories reflect the practical problems of classifying irregular components.

From time to time, because of new designs and because the system is used by many different people in the company, the categories are modified and new ones added. These changes are the result of negotiations between designers and draughtsmen at designated meetings. What the scheme does is enforce a standard format for retrieving CAD models.

The Augmented Bill of Materials (ABOM)

Due to the large number of components used in the instrument, as well as the large number of people involved in the project, an augmented bill of materials, ABOM, was invented to manage the production of components. See figure 3.7. Each ABOM forms holds up to 19 components for the same unit in the same instrument.

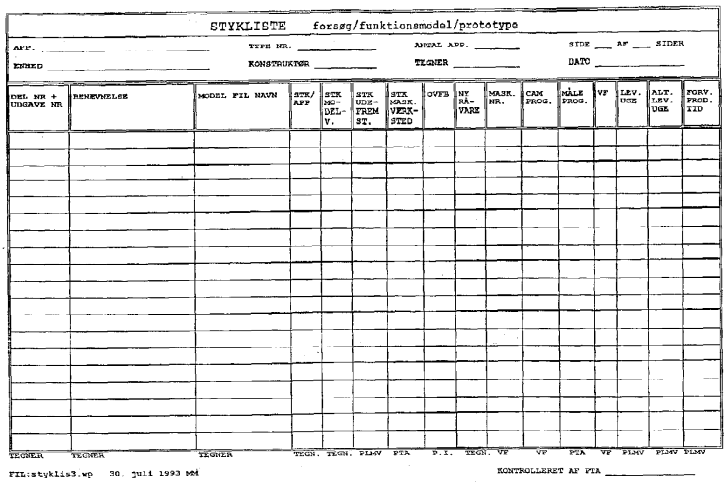

STYKLISTE forsøg/funktionsmodel/prototype

DEL NR + UDGAVE NR	BENEVNELSE	MODEL FIL NAVN	STK/ APP	STK NG-DEL-V.	STK UDE-FREM ST.	STK MASK. VÆRK-STED	OVFB	NY RÅ-VARE	MASK. NR.	CAM PROG.	MÅLE PROG.	VF	LEV. UGE	ALT. LEV. UGE	FORV. PROD. TID

TEGNER TEGNER TEGNER TEGN. TEGN. PLMV PTA P.I. TEGN. VF VP PTA VF PLMV PLMV PLMV

FIL:styklis3.wp 30. juli 1993 MH KONTROLLERET AF PTA _____

The ABOM as the artefact used at Foss Electric

Augmented bill of materials: Experiment/functional model/prototype

Instrument: *Name* Type ID: *ID number* Page of pages
Unit: *Unit name* Batch size *Integer* Date:
Designer: *Name* Draught person: *Name*

Categories	*1. component*	*2. comp.*	*19. comp.*	*Responsible*
Component ID	*ID + version*			Draught person
Description	*Text*			Draught person
Model name	*Model database ID*			Draught person
# Pr. instrument	*Integer*			Draught person
Model shop	*Check*			Draught person
Subcontractor	*Check*			Production planner
Production	*Check*			Process planner
Surface processing	*Check*			Project purchaser
New input materials	*Check*			Draught person
Machine ID	*ID number*			Foreman
CAM program	*Type of CAM prog.*			Foreman
Measure program	*Check*			Process planner
Foreman	*Initials of foreman*			Foreman
Delivery week	*Week number*			Production planner
Alt. delivery week	*Week number*			Production planner
Production time	*Estimated time*			Production planner

Controlled by process planner: *Signature*

Figure 3.7: The Augmented Bill of Materials (ABOM).
NOTE: Each ABOM form holds information of up to 19 components for the same unit in the same instrument. The original is inverted, i.e., the rows in the figure are columns in the original ABOM. The original has 19 rows, corresponding to 19 different components. These are illustrated in three columns. This inversion is made in order to depict the whole ABOM and not only part of it. All italics are our additions. For each category of field in the form, the type is indicated.

Earlier projects had used a simpler bill of materials allied with weekly meetings. The ABOM helped to reduce the complexities of handing over CAD models for components from engineering design to process planning and production. Organisational procedures stipulated who was responsible for which tasks when producing components whether for experiments, functional models or for prototypes.

The information in the header is completed by the designer and the draughtsman. When the latter has finished polishing the CAD model for a particular component, the model is sent to process planning via the server. The first 5 columns (rows on the figure) are filled in with information regarding each

component, its ID, version number, description, database ID, batch size, whether the component is to be produced by the model shop, and whether new input materials are to be used. The ABOM is then sent to the responsible process planner. Here further information is added by the production planner, process planner and foreman. They, among others, schedule machines, staff, type of CAM programme, measure programme for quality control, delivery and production time. In this way, and unlike the previous bills of materials, distributed tasks are supported as well as retaining a comprehensive record of materials, processes and responsibilities. ABOM made a number of meetings unnecessary.

The bugs form and the problems list

One of the early problems faced by the project team was the lack of overview and co-ordination of the software. This was mainly due to the familiar problem of a large number of people working on the project as well as the constantly changing design. Eventually, a tool was devised for registering, prioritising, co-ordinating the correction of the bugs and other software problems. To ensure that all registered bugs were diagnosed and corrected a related set of procedures was established to assign responsibility and checking.

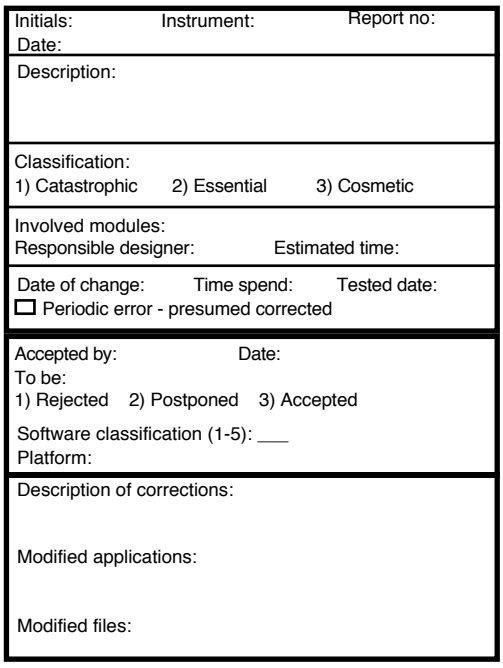

Fig 3.8: A simplified version of the 2 pages bugs form

The bugs form can be filled out by anyone involved in testing the software, including marketing people. The originator describes and classifies the problem - see fig. 3.8 - and then a team of three software designers (called the ‘spec team’) adds information about affected modules, designer responsible, platform period and a priority. All forms are filed in a central bugs forms file using the categories, non-corrected catastrophes, non-corrected semi-serious problems, non-corrected

cosmetic problems, postponed, rejected, corrected but not tested and corrected problems, the forms being successively refilled.

The list of not-yet-fixed problems are continuously produced and accessed by the software designers. Each designer is responsible for fixing the problem and reporting back to the platform engineer.

The purpose of the forms, the central file, the problems list, and the procedures for classifying, correcting and reporting on the problems was to ensure that all problems were registered, that responsibility was clear and visible to all the designers, and to make clear how the problem was dealt with. By making the list available to all the designers, they could become more aware of the state of the total system, and particularly aware of problems which might have a bearing on their own responsibilities.

Emergent Features of the Social Organisation of Work

Organising complexity

One of the features emerging from this study which we want to highlight is coping with complexity by the development of standardised artefacts of various kinds. In this case, the complexities mainly arise from organising not only a large collection of people but people with different specialisms and different responsibilities. Classically, this state of affairs exhibits both the virtue of the division of labour - increasing efficiency by combining specialisms - as well as its vice - the increased possibility of confusion and disorganisation. Again, classically, the solution to this is managerial organisation.[45] Or, to put it slightly more accurately, the enduring response to this is various forms of managerial organisation. This general issue is, of course, reflected in various modes of software project management with all the problems so classically elaborated by Brooks (1975).

One of these problems is the increased need for “articulation work” (Strauss, 1985, 1988; Schmidt, 1990; Schmidt, 1993) as projects become larger; that is, achieving co-ordination and co-operation among the various people involved in the project becomes a considerable and inescapable overhead as tasks become distributed and the costs of organisation increase. Such “articulation work” is practically exhibited in the round of meetings, both formal and informal, the constant reviews, planning, updates, and so on; in general terms, the need for the continual reviewing of ‘where the project is’, ‘what needs to be done’, ‘where

---

45 This, of course, is a truncated statement of the rationale of bureaucratic organisation both as a means of organisation large numbers of people and, importantly and as part of this, organising specialist skills and competences in a rational manner. See Weber (1947) for a classic sociological statement of the problem. It should go without saying, of course, that the whole history of organisation to the present day can be written in terms of a constant attempt to resolve this problem.

effort needs to be put now', 'what problems remain', and so on. Organisation itself becomes a vital ingredient in getting the project done (Anderson et al, 1993).

In this particular case, largely because such a large project was a new experience for Foss Electric, many of the systems used to organise the complexities of the project had to evolve or adapted from existing systems. The aim of these systems was to embody much of the "articulation work" within organisational artefacts in an effort to reduce the organisational overhead and, at the same time, improve co-ordination and co-operation within the project team.

Providing for mutual awareness

One of the main objectives of the kind of artefacts focused on in this study is the provision of mutual awareness among members of the project team in a systematic fashion by using standardised formats for the presentation of relevant information. The CEDAC Board is a good example of this. By being in a prominent place and available to all members of the team, and by organising its means of presentation in terms of problems, it not only served to acquaint them 'where the project is at' but also gave some direction as to where effort needed to be directed. The Bugs Form and Problems List, though specifically directed at the software engineers, operated in much the same way; that is, improving the overall awareness of others' work and, through this, supporting the co-ordination of tasks and co-operation between members of the team.

Honing tools

As has been pointed out, one of the features of the artefacts used to effect co-ordination within the team were developed from existing tools; the ABOM is a good example of this as is the Bugs Form. In other words, standardisation was leavened by adapting the tool to local circumstances. In effect, these tools were designed with local considerations and needs in mind. Thus, although standardisation is a key factor in reducing complexity, what we have in this case is the adaptation of an existing standard tool in order to standardise for the particularities of the local circumstances of the project.

The Production of Technical Documentation

Date	Approx. Duration of Fieldwork	Relevant References
1993	9 weeks	Andersen, 1993a, 1993b; Hughes at al, 1994

Objectives of the Study

This case study focuses on the production of technical documentation within a large international company with special reference to the co-operative work

arrangements involved in their production. In addition, the study was also intended to provide recommendations for improving co-operative work arrangements through computer support, work organisation and training.

Duration and Pattern of Fieldwork

The fieldwork took approximately 9 weeks and consisted of ethnographic observation and analysis, unstructured interviews with selected personnel of the technical documentation and other related departments, document analysis, video and some participation in aspects of the work.

The Setting

Although the technical documentation department - known as Technical Marketing within the organisation - is a distinct unit within the organisation, its main role is to support other units by producing, maintaining, updating, archiving, and distribution of technical documents about the company's products. Its main relationships are with marketing, production, and service departments as well as with installers and end-user organisations.

The main organisational units within the department are the development group, the product management group and the technical documentation group itself. Generally speaking the project management and development staff supervise the content of the technical documentation, mainly through scheduled meetings.

The technical staff are located on the same floor except for the document store itself which is in the cellar of the building. The product management department is located in several offices on the same floor close by, while development staff are distributed at various places around the factory, though at the time of the fieldwork there were plans to shortly relocate product management and development staff in one place and, following this, to bring all three groups closer together. Fig 3.9 shows the current layout of the office.

As can be seen from the diagram the marketing technicians are equipped with a computer terminal connected to a mainframe in which the document archives are stored. Also, standards, reference books, drawings, dictionaries, etc., are placed in the archives. Local archives are placed in desk folders.

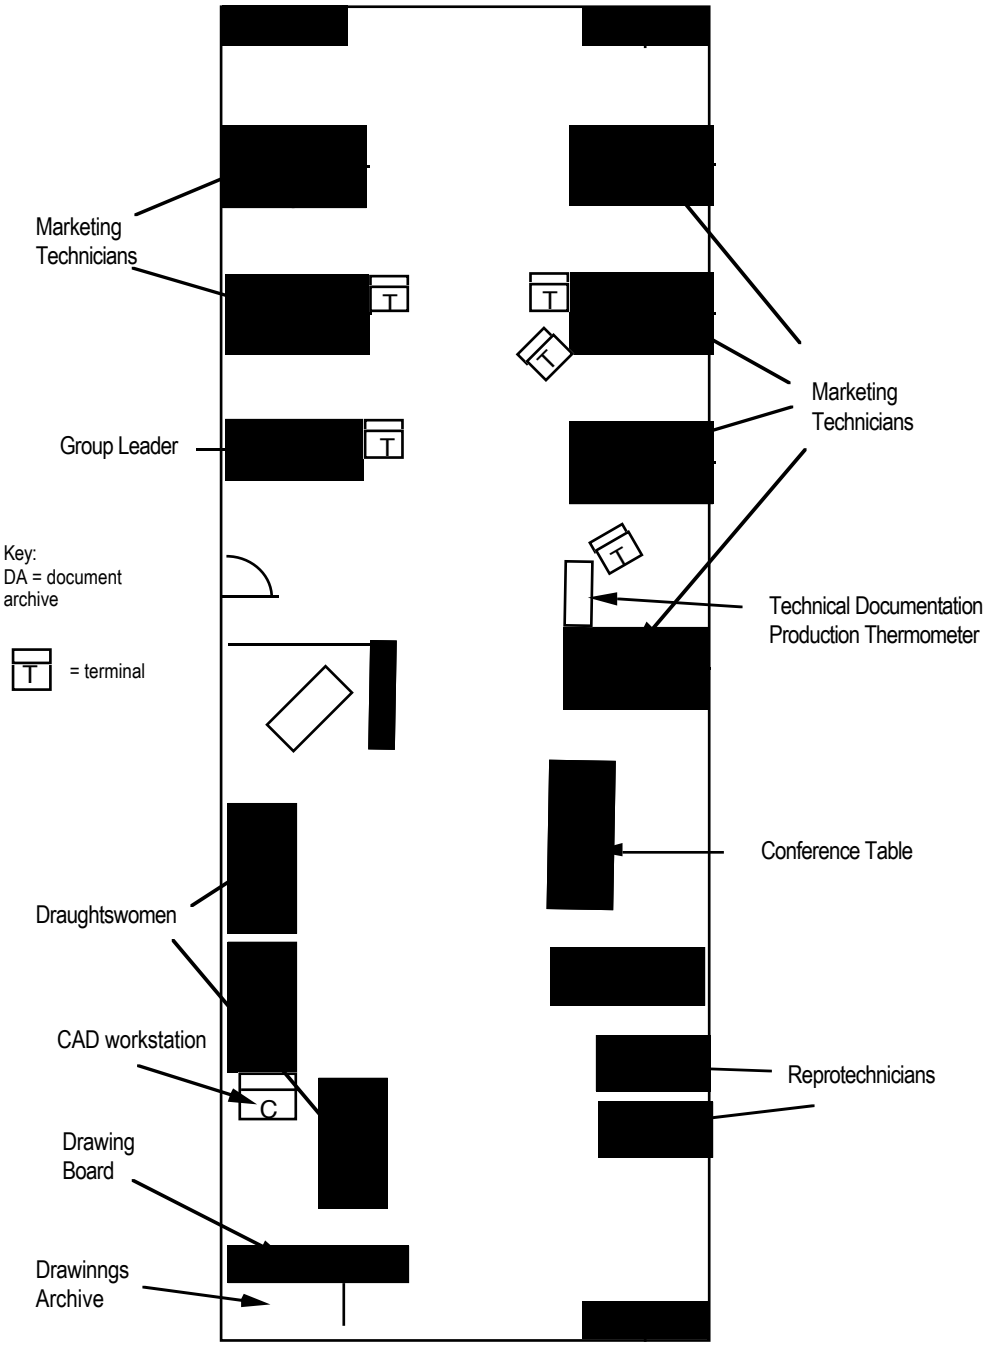

Fig 3.9: Layout of technical documentation department

Work Process

The technical documentation department is responsible for the production of technical documents for the company and is co-ordinated by project management. This involves making sure that the relevant technical documentation is available for when a product is marketed. Information is produced in the product development departments but the existing stock of documents in the department

functions as a document format and reuse pool. There are around 1500 document variants containing information on 25000 or more product versions.

Document producers are engaged in a series of parallel and non-discrete activities gathering, distributing and producing the documents. The work is extremely complicated and detailed and is by no means always a smooth process. Constant updating and checking is required as well as translating documents for the world-wide market that the company deals with. Legislation often affects the specification of products and this, too, has to be checked and reflected in the documentation. See Fig 3.10 for a schematic of the department's functions.

Product analysis

One of the first requirements is the detailed analysis of the product. Though field studies are often part of this process, if possible marketing technicians prefer to have the product to hand for this purpose. But one of the most important things to know is who knows about the product and this requires knowing who has been involved in its development. Interviewing the relevant people is one technique that is used. Another is to study the existing corpus of information about similar products to 'get a feeling' for what the product is like, using existing drawings, using reusable elements, and so on.

However, product analysis is not a phase in document production but a process that goes on through design, development and document preparation. Scrutiny meetings are, for example, one place where this goes on and are a forum for co-ordination, negotiation and the testing of the technical documentation about a product. This involves looking at what is new about a product, whether or not there are any changes, checking through the existing documentation to see what changes will need to be made, allocating responsibilities, and so on. Scrutiny meetings are typically attended by all those involved with a product, including marketing technicians, quality engineers, group leaders, product managers, and development engineers.

Tailoring documents in terms of customers

One of the major problems in producing technical documentation is categorising the users of the documents. The department has a number of internal and external customers for its products. Among the external customers are various end-users as well as installation and service specialists. This means that, to a large degree, the style as well as the information the documents contain, has to be geared to the varied interests and competencies of these different groups. The marketing group have evolved various categories - "this piece of documentation is meant for the ordinary person", "this is meant for the engineer like the one in the local heating plant", and so on - which are relevant to the character of the document while, at the same time, trying to ensure that the information contained in it is complete, correct and encourages safe usage. Installation and service specialists need information mainly in the form of procedure specifications for product installation

and maintenance. Service units are concerned about the availability of spare parts as well as details on how they may be installed.

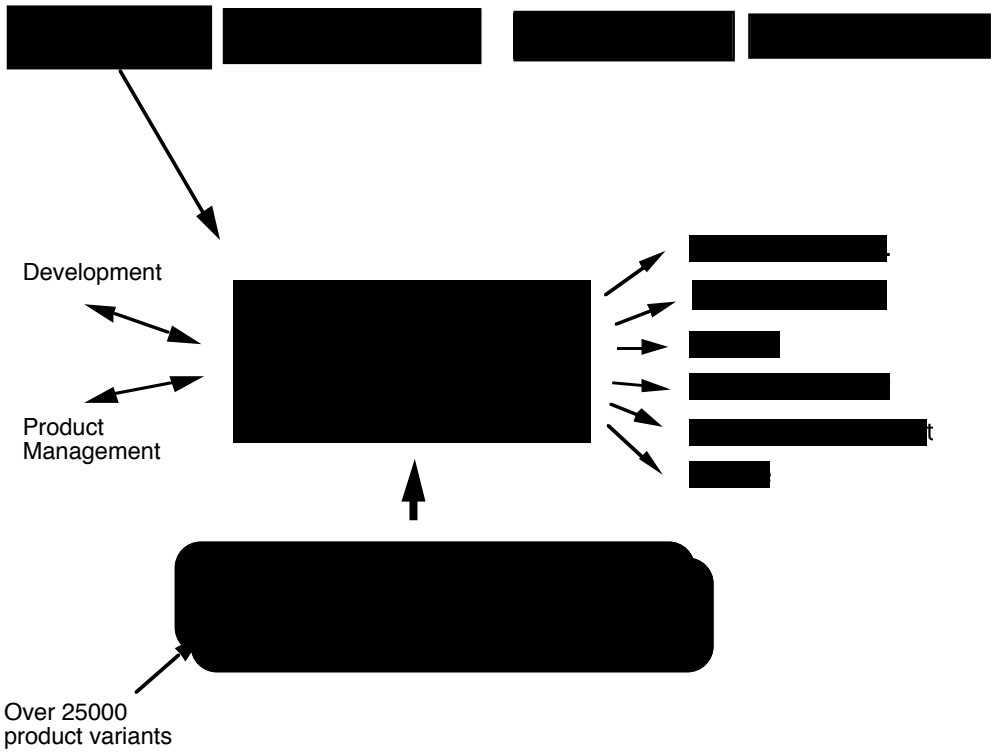

3.10 Schema of Technical Documentation Department's Functions

Internal customers again have different requirements. Sales, for example, need information so that they can market the products to customers who themselves are likely to constitute a varied clientele from technical specialists to high level management which, once again, need to be reflected in the technical information which is produced.

Keeping up-to-date with this varying pool of recipients of the documentation is no easy task and is generally a cause of complaint by the technical documentation staff. Scrutiny meetings are regularly held to test out and assess various documentary drafts in relation to various customer profiles.

Product managers, who determine release dates, are directly involved in the production of documentation for the products which, in principle, are not allowed to be released until the proper documentation is available for distribution. However, in fact products are often released before this since technical documentation tends to be regarded as an 'after sales service'. In any event, product managers are more than ready to pressure the department to get the documentation for a product ready in time for the product release, and in a form which reduces the need for what they often regard as 'trivial questions' by customers and users. However, because product managers are directly involved in the production of the relevant documentation, they are an important source of

information which is useful for producing documents appropriate to the intended users.

Using the technological resources

Not surprisingly, the department uses a range of technological tools in the course of the work. Most of the computer technology, including word-processing, databases, email, and a curve editor, are mainframe based. The draughtswomen are equipped with CAD workstations connected to a CAD mainframe which is often used not only for manipulating drawings but also to do graphical editing.

The repro-technicians are equipped with a PC and a mainframe terminal. The PC has a DTP software system and drawing tools but are not heavily used. They mainly use a mainframe table editor used in the production of part lists and the same mainframe word processing, database and email systems used by the marketing technicians.

In many respects the systems available are inadequate and the source of numerous problems and complaints. Word processing, for example, is available on the mainframe and based on the General Mark-up Language but for the inexperienced user it is difficult to use. The print preview facility particularly is expensive of CPU time which necessitates a large amount of draft printing. Its limits are also quickly reached and, in general, the system is cumbersome and made worse by having to handle character sets in different languages. The word processing system does not adequately support the integration of text, tables, graphics, drawings, photographs, etc., all of which are essential to technical documents. Often drafts are set up using cut and paste in paper, scissors, glue and photocopying.

The table editor, which is used for editing parts lists, is also cumbersome to use. It is command based and controlled by function keys and key combinations. Only five columns can be viewed at a time necessitating a great deal of scrolling. There is no print preview facility. Data input is received and retrieved from development by a technician who produces a hand-written draft and the repro-technicians do the actual editing from this draft. The worse problem is that the data cannot be transferred to other systems and it has to be retyped.

The CAD system is widely used for transforming drawings for the documents. It handles 2 and 3-dimensional drawings in solid and in wired formats as well as providing tools for their manipulation. The main problems attaching to this system are the inconsistent file naming conventions which makes search strategies often complicated, the inconsistent use of layered models in development and the fact that many drawings are not available to the system but exist as paper copies in the local drawings archive.

This makes the preferred ‘exploded’ 3-D drawings a very prolonged activity. One of the draughtswomen estimated that such a drawing from scratch could take between one to two person months. At the present time, the local drawings archive is used. This consists of drawings which are partly hand-drawn, and partly

CAD based and are largely the result of cut and paste using photocopier, repro-camera, drawing boards, etc.

The transformation and refinement of technical data

Producing a technical document requires the transforming of an enormous amount of diffuse, unstructured mixed media technical data. This is typically a combination of text, procedure specifications, graphics, diagrams, tables, drawings, illustrations, photographs, all presented in various formats and layouts depending on the document to be created. The form and layout of the product information, though formalised is dynamic and negotiable, the latter usually through scrutiny meetings with the product team. All the time the usability of the document by the relevant customers and end-users is the prime consideration. Also of concern are issues to do with consistency, readability, logical structure, table formats and the general appearance of the document. Important, too, are the checks made about directives, legislation, standards that need to be satisfied in respect of the various countries in which the product is being marketed.

A key process is that of proof reading the documents. This activity involves a wide range of people including translators, marketing technicians, development engineers, product management and marketing people in subsidiary companies in other countries. The translators spend approximately 75% of their time proof reading. Proof reading is highly iterative and documents will go through this process several times before being released for printing. It is also a process which goes on through much of the development of a product and its documentation.

Emergent Features of Social Organisation

Reverse salients of technological implementation

One of the features highlighted by this study is the phenomenon characterised by T.P. Hughes (1987) as the “reverse salient” of technological innovation. This Technical Documentation department displays what must be a common situation in many organisation, large and small, of a mixed environment of technologies which, in some sense, are intended to support work activities but which represent varying levels of technological sophistication. Hughes’ point is that one of the impulses for technological innovation occurs when an element in a system of technologies lags behind development in others, such that it becomes an impedance to the further development of the system as a whole. In this case, however, the problem seems to be not so much a lack of technological innovation but a “reverse salient” created by, presumably, lack of investment in appropriate technology by the company. However, the point is not so much to criticise the company, but to draw attention once again to the fact and consequences that work takes place within an organisational environment which itself becomes a feature of the day-to-day work. Having to “get the job done”, “making the best with what

we have", become mundane and routine aspects of the daily work, and aspects of the work which owe much to the policies of the company and, moreover, outside of the control of those who are largely responsible for getting the daily work done.

Although it may suit designers to think in terms of idealised environments, these are, we suggest, extremely difficult to find and, more often than not, what is likely to prevail is a piecemeal introduction of technologies into work settings such that, at any point in time, they become a mix of the old and the new. What impacts such a situation may have upon the work is an empirical question, but it is one, we suggest, which has to be a consideration of CSCW design.

"Making do with the technology-to-hand"

One of the features of the work which came out of the air traffic control study was the way in which controllers were able to 'make do' with less than adequate technology.[46] As a highly experienced, highly skilled group, controllers were able to follow a number of skilful strategies to get round some of the deficiencies of the technology, including refusing to use some of the facilities provided. More often, however, it was more a matter of developing a 'working sense' of the limitations of the technology, a sense of 'what to trust' and 'what to be careful about', 'which radar beacons were reliable', and so on, but within a very practical understanding of the interweave of the technology and the work. In other words, their trust in the technology was part and parcel of their trust, or the lack of it, in the competence of others, an intimate knowledge of the procedures, a trained sense of 'what looks right', and more.

Though the Technical Documentation department is not a safety critical environment, it displays a similar feature of 'making do with the technology-to-hand' in order to get the work done. In other words, although the technology is less than ideal, although there is often a great deal of pressure from elsewhere in the organisation to 'get things done', by and large ways are found to get round the problems, to solve them 'with whatever is at hand'. This ability, of course, is very dependent on the experience of those who work in the department and their knowledge of what can be done using 'whatever is at hand'.

Managing organisational pressures

The Technical Documentation department is, in effect, a service department for other departments, such as marketing and project management. Although proper technical documentation is essential to any product, from this study it is also clear that its production is subject to a number of pressures from other departments; pressures to do with 'getting the product out in time', 'getting the style right for the appropriate users', "making sure it conforms to the relevant legislation", and so on, which make it difficult for the department to plan its work with any strong expectation that the plan will hold good for very long. In such a situation the

---

46 We have not reported on this aspect of the research in the summary provided in this Deliverable, but see Hughes et al, 1988.

ability to "work with whatever is at hand" becomes vitally important for managing, to some degree at least, the varying pressures placed on the work group.

Technology Centre[47]

Date	Approx. Duration of Fieldwork	Relevant References
1994	4 weeks and continuing	O'Brien, Hughes, and Rodden, 1994

Objectives of the Study

This fieldwork is linked to a project concerned with the potential use of Virtual Reality techniques in facilitating communication between engineers in a Technology Centre and staff at the production sites of a multinational group of cable-making companies served by that Centre.

Such a remit entails the study of day-to-day work of staff at the Technology Centre, both technical and 'administrative'. Not only is this work observed as it occurs within the Centre, but in addition particular attention is afforded to the current practices involved in the establishment and maintenance of links with external production sites.

Duration and Pattern of Fieldwork

Approximately four weeks fieldwork has taken place over a period of two to three months, and is continuing. Typically a single fieldworker spends a working week with staff at the Centre and then returns to Lancaster for a number of days' debriefing and report writing, before returning for another week's fieldwork.

In many respects, this pattern conforms closely to 'concurrent ethnography'.

The Setting

The Technology Centre is situated in the north west of England, employing approximately forty staff on a site consisting of three 'elements': the Office, the Labs and the Pilot Plant. The latter, formerly used for the production and testing of fibres and cables under development, is no longer directly controlled by the Technology Centre since it has recently been turned over primarily to the production of fibres for the market by one of the group's production companies.

---

47 This study is part of an ongoing research project, VirtuOsi, funded by the DTI/SERC initiative on CSCW.

The Labs are made up of a suite of laboratories containing a range of optical and thermal testing equipment used in the development and testing of fibre optic cabling.

The Office is a large open plan office space on the first floor of the Technology Centre building. Arranged around a central aisle are a number of 'bays', each made up of four desk spaces, with connected drawer and cupboard space. The 'closed end' of the bay is typically filled by a filing cabinet, computer or, less frequently, empty wall space. The majority of members of staff have desk space in the office[48], although it is not uncommon for as many desks to be empty as occupied at any given point in a working day. It is also not uncommon to see staff sitting at desks which 'belong' to someone else.

The Work Process

Work in the Technology Centre is organised into broadly defined 'Projects' each overseen by a Project Manager; the term 'Projects' refers to a number of 'strategic goals' formulated two years ago by the Centre's Engineering Director in order to best serve the needs of the Centre's customers who are mainly the group's production companies. These Projects act as 'umbrellas' to a number of smaller sub-projects for each of which a 'Project Leader' is responsible.

Thus work in the Centre can be outlined, broadly speaking, in terms of, on the one hand, the management/administration of these projects and, on the other, the technical work such projects involve.

Both of these types of work are subject to the company's Quality Management System, which, in order to satisfy various national and international accreditation criteria, sets out standard, 'correct' procedures for the completion of work. Central to this system is the writing of standard format Technical and Test reports to cover all technical work carried out in the Centre, and the appropriate sets of paperwork to document the management and progress of each project. Thus, while much of the work is highly technical it also involves completing the formalities of documentation required not only for the reporting of the technical results of various tests so that they are widely available to relevant parties, but also for satisfying the Quality Management System.

Informal patterns of communication and support

However, it is clear that the formal system is by no means the major channel of communication and support. In contrast to the 'official' working regime, there is an 'informal' working culture within the Centre supported by the continuous flow of interaction between members of staff as they go about their work on a day-to-day basis. Technical expertise, updates on "who's busy doing what", and the

48 There are exceptions to this: two members of the Advanced Materials Project have desk space in the labs alone, as does a contract engineer. A number of members of the Pilot Plant staff also appear to have their desk space in the Pilot Plant alone. It is important to note, however, that there are occasions on which these members of staff carry out work at desks in the office.

results of important tests of fibres under development appear to be transferred quickly and effectively from engineer to engineer by such means, with the relevant official documentation often completed weeks after the informal communication of relevant information.

Staff at the Technology Centre also make use of their informal networks of contacts both within the Centre and at other sites in the company. Such networks are important not only for contacting persons with the "right expertise and experience" but what is also important is maintaining contacts of people "one can trust" even when these are geographically remote. Many of these contacts are an important resource for "getting things done".

Indeed the importance of this interaction is not to be underestimated, particularly since it is *the* crucial ingredient in the development of the 'local knowledge', the unwritten working knowledge not only of what the job entails but of the organisation itself, that makes up the expertise drawn on by the rest of the multinational group. As highlighted below, this type of knowledge is paramount in the successful completion of complex technical tasks.

Working knowledge, 'knowing the procedures' and documentation: the case of calibration testing

An important element in the technical work of the staff of the Centre is the testing and calibration of many of the products the company produces or is in the process of developing as well as many of the tools used in this process. The case detailed below involved the calibration of the 'alpha' tensile testing rig and its supporting software. Calibration of such machines is required by the Quality guidelines operating within the Technology Centre, in conjunction with those stipulated by the NAMAS accreditation scheme as well as providing the accurate measurements required for design engineers.

The 'alpha' Rig and its related software is a piece of equipment used in measuring the extension and point of 'failure' (the point at which the fibre or cable breaks) of fibres and cables when placed under varying degrees of strain. It produces graphs of fibre and cable extension over time, and is used in testing fibres under development to the required specification. This piece of equipment was no longer under calibration and thus had to be re-calibrated before it could be used as a tool to support the work of various projects.

As a result of a problem discovered in the manufacturer's recommended calibration system, staff at the Technology Centre had developed their own system for calibrating the rig, a draft proposal for which had been written up into a Technical Report in 1991, with related confirmatory tests written up into Test Report format in 1993. The diagram of the test rig in Fig 3.11 is taken from the Technical Report on the 'alpha' Calibration.

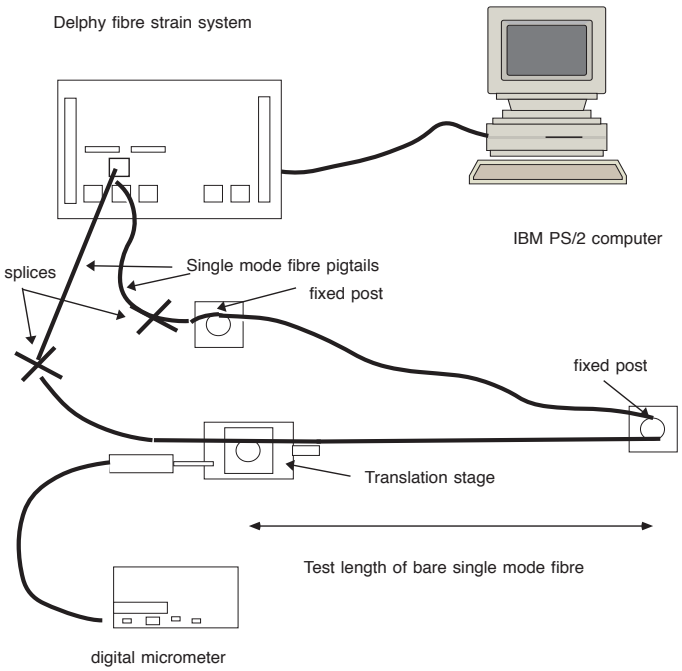

Fig 3.11 Schematic of 'alpha' Calibration Set-up

The experiment basically involves applying a known extension to a length of fibre, and measuring its elongation over a period of time, using the 'alpha' fibre strain measurement system. The test length of fibre is of the order of 10 metres, with the fibre wrapped ten times around each of two posts. Double-sided sticky tape is used to ensure that no slippage can occur. One post is mounted on what is known as a 'translation stage': this is the means by which a known extension can be applied to the fibre. A micrometer is used to measure extension of the fibre.

The experimental procedure is outlined in the following way in the Technical Report:

> i - Switch all the equipment on and allow to warm up for at least one hour. Apply tension to the test length of fibre. The fibre should be taught.
>
> ii - Set up the computer to log the data from the 'alpha'. Leave the equipment for two minutes to give a zero value baseline.
>
> iii - Apply an elongation to the test length of fibre of 3.0mm.
>
> iv - Wait for at least one minute.
>
> v - Repeat steps (iii) and (iv) until an elongation of 15.0mm is reached.
>
> vi - Remove the elongation in steps of 3.0mm waiting at each elongation for at least one minute.

> vii - Leave the equipment running after the elongation has been removed to give a final zero value baseline.

In this way a known extension, measured on the micrometer, is applied to the fibre and the graph produced by the 'alpha' can then be compared with what is known about the fibre, and the 'alpha''s accuracy and reliability ascertained.

This was the test being undertaken by GP, a Technology Centre engineer. The following description is taken from fieldnotes and tape recordings covering one afternoon and the following morning's observation.

As observation began, NH (the Section Manager, Measurements), GP's line manager, approached GP in the office and informed him that he was to run a calibration test on the 'alpha' Machinery. "*You know how to do it, yeah?*" rounded off his request of GP.

GP set about gathering together the equipment from the labs for carrying out the test immediately, without referring to any of the documentation outlined above: this clearly was a test that he was familiar with. At the same time GP was being shadowed by a B.Tec. student on a shorter placement in the Technology Centre, and he explained to this student what he was doing.

Having gathered together test fibre, posts and translation stage they moved from the labs into the Pilot Plant where the tensile testing rig and 'alpha' equipment are set up. GP switched all the equipment on and accessed the appropriate PC software first, explaining "*You've got to let it warm up and then settle down*"; he then fixed the posts and translation stage to the base of a larger testing rig and reeled out an adequate length of fibre for the test. Next he wrapped the fibre around the first post and secured it with tape, ensured that the fibre was taught to the second post and wrapped it around the latter, again securing it with tape, and looped the remaining fibre back towards the 'alpha'. "*It's quite easy to break when its under tension"* he explained *"so you have to be careful when doing this bit*". It is also important that the fibre is not twisted as it is reeled out, as there is an outside chance that this could distort the findings.

GP elaborated a little more on the context within which this test was taking place. They also wanted to be sure that the problem of inaccurate readings lay with the possible inconsistent quality of fibres being used to check the accuracy of the machine:

> In the past we've only ever calibrated it with this kind of fibre ... this time we want to calibrate it with two different types of fibre, to compare the three (sic) results, 'cause the last time I did it with this, we found out that it wasn't working accurately enough ... When it was checked it was found that it wasn't working within what the spec. said it should be ... I can't remember what that is off hand ... it's in one of those documents (Manufacturer's specification documents, kept with equipment). But whatever it is, whether it's five millimetres or five percent or whatever, it wasn't working accurately enough ... so we're going to try it with more accurate fibre (sic) ... there's a fibre called 'Dispersion Shifted', which is ...err, whenever anyone's using something that needs high reliability, they always use 'Dispersion Shifted' fibre ... so we'll try it for the test. So we're hoping that it (the inaccuracy) might be as much to do with using this particular fibre as anything else ... but then again it might work with this fibre this time!

In this test run, then, it has been decided that an attempt should be made to compare results across fibre samples, in an attempt to "get to the bottom" of the equipment's inaccuracy. Although GP's final comment highlights his expectancy that things might not be as straightforward as hoped.

He then spliced both ends of the fibre onto the 'pigtails' which connect the fibre to the 'alpha' equipment. To do this he first stripped and cleaned the sheathing of the end of both fibre and pigtail, before inserting them in a fibre splicer which gives a greater accuracy of connection, the process being observed under a microscope attached to the splicer. When satisfied that pigtail and fibre were flush against one another, GP switched the splicer's heating device on to fuse the two together.

GP explained his actions to the student with him, and on a number of occasions asked him to "have a go" and see if he could do it, giving him the opportunity to attempt splicing fibres to pigtails, for example. At no time, up to this point, however, had GP made any reference to documentation related to the equipment.

At this point NH came down to the Pilot Plant to check everything was alright with the test. He explained that the 'alpha' equipment had been moved to the Whiston site and had failed to operate accurately. This had, naturally, brought the machine's calibration into doubt, and as a consequence it had been transported back to the Technology Centre for calibration testing. NH expressed a certain degree of dismay that the equipment had been moved at all, expressing doubts as to whether any other piece of Centre equipment worth tens of thousands of pounds would be moved "*at the drop of a hat*" as this one had been. He also stated that he wasn't happy with the equipment and its associated cables being 'open', and that he was having them cased up in the near future in order to protect them. Having surveyed GP's progress, he left saying, "*As you're obviously getting on fine, I'll leave you to it, let me know how it goes*."

As mentioned earlier this type of informal 'updating' on progress is a characteristic feature of work in the Technology Centre, and whilst it is often cursory, almost casual in appearance to the outsider, it clearly holds a position of great importance in "checking how things are going".

Continuing with the test, GP moved to the PC and worked through the 'alpha' software, 'naming' the test run as prompted - such names usually consist of reference to the type of test - in this case 'calibration', the fibre being tested, and the date tested, so that they can be stored on and recalled from the hard disc as and when required. He seemed uncertain as to whether the software had accepted the 'name' he had entered, as the package appeared to be behaving "*slightly unusually*", although he could not be too sure in what way. He decided to press on once he felt he had entered the name successfully, and he continued to enter relevant information as prompted - of particular note was his choice of interval between readings: he knew that the 'alpha' could only make a maximum of 200 readings, and thus had to calculate appropriate intervals between these readings in order that it might generate data to produce a meaningful graph. This calculation was based upon the amount of time required between extensions of the fibre (as

stipulated for this particular test) multiplied by the amount of extension to be placed upon the fibre giving the total length of the test. This length was then divided by 200 to give the appropriate interval rate between the 'alpha''s measurements.

Again it is worth noting that GP made these calculations by calling upon his own knowledge of the relevant requirements of the test and the equipment without reference to any documentation related to the test he was undertaking. When asked if people often read Technical reports related to such tests, he replied

> No ...The tests you do you'd generally ... know the method already ... 'Cause you've seen it done ... the first time you did it, someone told you how to do it ... by word of mouth rather than giving you a great wadge of paper. It's (the Technical Report) more just for Quality Control ... you *know* the methods you're going to need to be using... once you've seen it done once then you'll be alright ...

GP continued to work through the 'alpha' software package, and reached a stage where a 'baseline' was appearing on a graph on the PC monitor - this 'baseline' consists of the 'alpha''s reading of a signal that it sends through the fibre on the test rig: at this point GP pointed out the marked socket on the 'alpha' equipment into which the fibre under test always had to be inserted:

> You see the connector? You've got a hole at the top - you've got a piece of copper there so that you know every time you put it (the test fibre) in, it's in the same place. It doesn't matter in this test, but on some tests, where you've got to be very accurate, even by say, instead of putting it there (the marked connector), you put it there (another connector on the 'alpha') it'll make a difference. You're only talking .02 db's, but in some tests that's the accuracy you need, whereas this (the current test) you're going to 80 db's, so it's not going to make any difference. I mean, .02 - it's not going to mean sod all!

He then turned his attention back to the PC monitor, and examined the readout of the 'baseline' as it emerged. He wasn't particularly happy with what he saw:

> This is what we had last time ... the very first time we did this, we were getting problems, we were getting a really small signal ... which is what I've got there (indicates graph on PC screen). So that's a really weak signal ... should be getting something that's about 60, 70, 80 ... so it might tell us that the signal's too low and it isn't enough to do the test.

However he continued through the package's prompts which basically outline a method for attaching fibres or cables to the 'alpha' in a similar manner to that he had already undertaken (laying out fibre on the rig for test, splicing on pigtails, connecting pig tails to 'alpha' etc.). He then reached the prompt displayed below.

Select Channel For Use
Fibre 1 - (N)
Fibre 2 - (Y)
Fibre 3 - (N)
Fibre 4 - (N)
Fibre 5 - (N)
Fibre 6 - (N)
Fibre 7 - (N)
Fibre 8 - (N)
Fibre 9 - (N)
Fibre 10 - (N)
Fibre 11 - (N)
Fibre 12 - (N)

At this point he hit return instantaneously, saying:

> Well normally, when this machine's sold they expect you to use 12, that's why there's that many connectors and points ... but we only actually use one. So for us all we ever do is type 'yes' that we're using Fibre 2 ... Number 1 is to check your reflection, but we don't use that, we never use that ... so Number 2 is the first one.

Having hit return the baseline graph reappeared, with the information that the 'alpha' was checking the strength of the signal. GP looked at the graph and felt that the signal was borderline strength.

> *It's just checking the signal level's okay. Normally anything that's about there, sometimes it can be a bit dodgy. Normally you'd really want something up there ... but it was like that last with time this fibre.*

He offered an opinion as to what might have been the cause of this:

> *The only thing we could think of was that ... while you're wrapping it around the thing* (post) *so that it doesn't slip, maybe that bend, although that bend should be okay, maybe it's just too small to get all the light through, maybe some of you're light's getting lost in that bend ... Every single cable and fibre has a minimum radius it can be bent to ... this* (fibre under test) *is* ***supposed*** *to be about 50 mil. which is what that* (post) *is, in fact just over, so it should be alright - but ... you just don't know ...*

The 'alpha' equipment did, however, accept the signal as strong enough, and GP hit return to indicate his acceptance that this was alright. At this point he also explained to me how care was needed when using this particular keyboard as certain keys had a habit of sticking, and thus running through the entire package, as the PC registered repeated keystrokes: this is clearly to be avoided if possible.

It was only at this point that GP consulted the Technical Report related to the 'alpha' equipment: he scanned through it quickly, nodding as he checked that he had followed the procedure correctly. He was doing it, he said "*just to follow the rules*". He then made the point that the pulley (translation stage) used in the method as written up in the Technical Report was different from the one he was having to use today:

> *See the pulley* (translation stage) *he's done it with had a hundred millimetres of extension, that one's been lost - someone's borrowed it and not given it back - so we're only using a fifty millimetre extension ... so the method's going to be slightly different in this case.*

This makes specific the contrast between the nature of the work as described in the appropriate documentation and the work as-it-happens: the documentation specifies the use of a translation stage that simply doesn't exist anymore in the Technology Centre, and as a consequence GP has to modify his method accordingly, in particular setting up the 'alpha' measurement intervals and in completing the subsequent calculations.

Such modifications are very much based upon his obvious knowledge of both the equipment and testing in general, knowledge which does not appear to be gleaned from 'official' documentation: nowhere in the Technical Report, for example, does it state that the 'alpha' equipment makes only 200 measurements in any test run and that the intervals between such measurements should be calculated accordingly; nowhere does it state that the choice of connector into which the pigtails are inserted can affect the accuracy of the test's results; nowhere does it state the appropriate level for the signal is around 60 to 80 dB's and that anything below this might cause problems and it certainly does not state anywhere that the PC keyboard is prone to jamming, thereby ruining one's progress through the software package.

Instead, as GP himself makes clear, this knowledge comes from his experience of using the machine - "*This is what we had last time ... the very first time we did this, we were getting problems,"* - and the experience of others as conveyed to him, apparently in 'informal' discussion, rather than in any 'official' documentation. Such 'local knowledges' are, then, built up as Technology Centre staff undertake their technical work over a number of projects.

It is also worth noting that one of the basic premises of this calibration test seems to have been determined by this 'local knowledge' of testing procedure. The need to assess the performance of another "*more accurate fibre*" is made clear by the engineer's knowledge that a certain fibre alone has been used to date, raising the possibility of certain idiosyncrasies in the behaviour of the fibre (perhaps the loss of light as it is wound around the post) that is skewing the 'alpha''s readings.

GP continued with the calibration test once he had glanced through the Technical Report. Taking his wristwatch off and using it to mark time throughout

the test, he extended the fibre by 3 mm every two minutes until the extension reached 15 mm, and then took the extension off in the same steps, with the same interval in between steps.

Throughout this process GP moved between the translation stage and the PC monitor, upon which the 'alpha' equipment was plotting a graph of the test fibre's extension. This emerging graph charted the strain within the fibre as it extended, and increased in steps that reflected each 3 mm extension as added - an idealised, schematic representation of this graph is presented in Fig. 3.12 below.

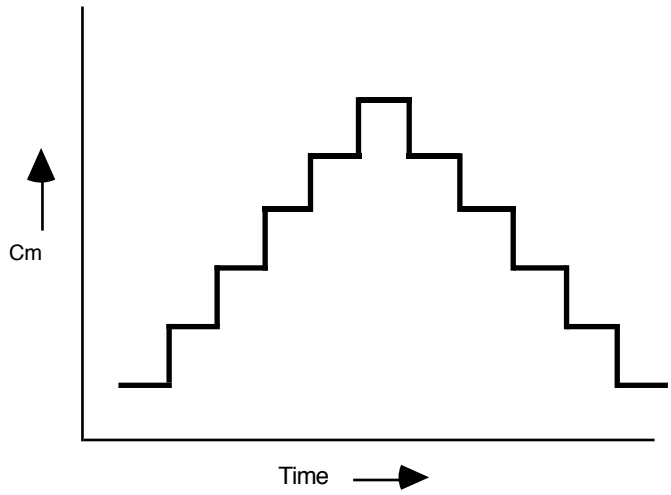

Fig. 3.12: Schematic Representation of the 'alpha' Graph of Fibre Extension

GP monitored the emerging graph and expressed satisfaction as it emerged on the PC screen. As the end of the test was reached he was happy that the results looked acceptable he entered the save test run option on the software package. It was here that problems emerged, however, for despite his attempts to save the test run, GP could see no evidence of the saved results on the PC's hard drive. Alarmed, he searched through all the test runs saved previously, in the hope that he might find the results he had lost. This proved fruitless, so he phoned AP, the Centre's Network Administrator and acknowledged 'trouble-shooter' on all computer-related matters, and asked him to come and have a look.

AP came down to the Pilot Plant as soon as he had finished the problem he was dealing with prior to GP's call. Whilst he looked the 'alpha' over, GP went upstairs to the Office to inform his Section Manager of the difficulty that he was experiencing with the PC. NH (the Section Manager in question) asked whether he had informed AP of the problem, as he was the "*best bet to sort out any computer nightmare*". GP made it clear that this was the first move he had made, and expressed his sincere hope that AP "*would come up with something*".

When he returned to the Testing Rig in the Pilot Plant, however, the news was not good. AP informed GP that the hard drive in the PC was full. He also said that

the software package did not have the facility to warn users of the fact, but instead it only became clear when a test run failed to save. Both AP and GP expressed a degree of disbelief that a software package could be so poorly designed, and the latter was yet further dismayed when AP said that he would have no choice but to 'dump' old test runs to create space on the PC and repeat and save the test once again.

GP returned to the office and informed NH of the situation once again. They decided that the test would have to be run the following day as it would not be completed by 4.45, the end of the working day. This decision was followed by rearrangement of some testing that GP was to have carried out the following day. NH then informed the Project Manager whose work the postponed tests related to of the situation whilst GP returned to reassemble the Test Rig to test another section of fibre on the following morning. AP also approached NH to give the full details on the situation with the PC's hard drive. He suggested that in future, test results would have to be more carefully stored on the hard drive, and, more importantly, removed when no longer of use. "*At least its not going to happen again*" concluded NH.

As the Calibration Test concluded and, ultimately, failed for the day, the immediate updating of other members of Centre staff became apparent. The interaction amongst those involved in the test carries the news to other appropriate individuals: GP's commitments are re-arranged for the following day, and staff affected by this are informed. Thus the contingency of a failed test is managed with rearrangement of the appropriate resources facilitated by the immediacy and flexibility of such interaction. NH's concluding comment appears to indicate that such an experience will be added to the stock of local knowledge, drawn upon in the future when carrying out tests on the 'alpha' equipment. The fact that such information was not available anywhere in the 'official' documentation related to the equipment brings to the foreground once again the interplay between 'informal' and 'official' working procedures in the Technology Centre.

In this example, the engineer's actions could not be gleaned from reading the 'official procedures'. In general, and GP was not alone in making this point, when encountering a new piece of technology or undertaking a new test, individuals are for the most part 'taken through' the technology/test by a more experienced member of staff, rather than being left to read the relevant Technical Reports. Technology centre staff seem to feel that is difficult to capture and convey adequately an intimate 'working knowledge' of, for example, a given piece of technology in such a written form, whereas talking someone through that technology whilst demonstrating it is far more effective.

GP also expressed a certain reservations with regard to the formal writing down of training requirements for each piece of testing equipment used in the Measurements Section - a move proposed by the then Section Manager - in order to comply with the requirements of certain of the Technology centre's major customers as well as NAMAS accreditation criteria:

> It would be okay for some of the more complicated stuff like this (the 'alpha' equipment), but for more straightforward stuff it would take you longer to actually write it down than it would to show ten people!

GP is also, in common with a number of other staff, asserting that such paperwork exists primarily for Quality Management and purposes and related audits by external institutions, rather than for use by the engineers and technologists themselves.[49] A similar contrast was also made through GP's comments on the role of Test Reports in the testing process. As quoted earlier, GP highlights the retrospective nature of such Technical and Test Reports, their completion is to be fitted in when more pressing demands upon the individual are satisfied:

> Test reports are the sort of things that, unless someone says to you, "I need that Test Report for tomorrow morning", you do when you've got time. So ... whenever you haven't got any testing work you'll sit down and write your report up ... It's quite often a fortnight or so after you've done the test

The fact that the completion of such documentation is not seen as integral part of the technical work *as-it-happens* was made clear at other times during the field work. On one occasion a member of staff was observed hurriedly completing work on a project in the Labs, the Technical Report for which had already been written and submitted to the project's funding body stating that the technology being developed *had* been completed, *was* operating to the required specification and *was* ready for demonstration at the end of the week. In this situation, then, the 'formal' working practices were, unusually, ahead of the 'informal' or, 'actual' work.

The retrospective nature of the vast majority of this report writing is clearly necessitated by the more pressing need of staff members to complete technical work required by the goals and 'milestones' of projects to which they are connected. This formal system of result reporting is, instead, replaced by 'informal' immediate reporting of results by word of mouth both within the Centre and to external production sites, in the latter case over the telephone.

Thus the role of the Test Report is not one of immediately informing interested parties of the results of important tests as GP makes clear:

> Generally before the Test Report is issued, you've already shown the person who's ... err ... got you to do the test ... generally before the Test Report's written you've already shown them the results, you know, you've shown them the graphs - so he knows that everything's right, he already said, "In that case we've to do this test instead". That's

---

49 It is important to note that this is not intended to 'debunk' the value of such documentation for the purposes for which they are intended. It is essential for the range of accreditation criteria which must be satisfied for the Centre's work to be recognised as reliable and trustworthy. There are, however, arguments about the nature and the amount of such documentation.

> (points to Test Report) just like formally finishing it off ... it's got to be done for Quality ... On the Test Report I've just finished off we knew the data the day we did the test, and we phoned them up and told them ... even though the Test Reports weren't done for three or four weeks.... I mean when you get the Test Report, he won't often read it, because he already knows the results ... he'll just maybe have a quick glance at the last page, just to check ...

To both author and recipient the importance of the Test Report is in "*formally finishing off*" that particular element of the work, satisfying the Quality programme and ensuring that the test results are formally 'acceptable' (reliable and reproducible). It is interesting to note GP's own contrast between the rate and nature of progress of the actual technical work informed by tests carried out and the paperwork which carries on at a somewhat slower pace.

Such a contrast resonates with the account offered by another member of staff of the nature of requests for information made by staff at production sites: he described the work done in ten minutes of 'informal' phone conversation as equivalent to that achieved with three days' wait for the 'formal' paperwork.

Contacts of his at production sites will often phone and enquire as to whether the Centre has data or expertise on a given technical matter, and in the majority of cases he is able to answer their question immediately, and set in train the Centre's response: forwarding the appropriate documentation, on occasion simply providing answers verbally or making contact with other members of Technology Centre staff whose expertise is more appropriate to the query. After these telephone conversations have taken place, Request for Information forms are despatched to the Technology Centre and the 'official' exchange of information is recorded in the required manner (this is the three days' wait for the paperwork' referred to). Once again it is clear that such documentation fulfils a different purpose for both parties to its transaction than the actual transfer of information, and once again it is important to stress that this should not be seen to devalue its role, but rather to note that its role *is* distinct.

Emergent Features of Social Organisation

The interweaving the formal and informal

The distinction between 'formal' and 'informal' organisation is deeply rooted in organisation theory.[50] What we wish to draw attention to from this study is how the analytic distinction is exhibited in various practices of the members of the Technology Centre not as so much as an 'analytic distinction' but as consisting of practical matters within courses of organisational work activities.

---

50 See, for example, Selznick (1948). The concept of formal organisation is traceable to Weber (1947). The discussion here owes a great deal to Bittner's (1974) seminal treatment.

From the point of view of an engineer in the Technology Centre the work is organised around Projects and, beneath this umbrella label, a series of practical tasks, such as "tests to do", "reports to get out", "experiments to set up", "results to collate", and so forth. These activities will also be subject to variable constraints of time and resources often imposed from without. The engineers are also often involved, at any one time, in a number of projects, all of which have to be interwoven into each other. It is rare, for example, for a large project to be carried out all the way through as a single stream of activities. Much of the work, in other words, has a piecemeal character to it which is not uncommon to many work settings. Much of the work, too, is under pressure to complete not least because the Centre is dependent on its resources from other companies within the corporation. This means that the engineers have, as a very practical matter, to "juggle with" and prioritise their commitments, almost on a daily basis. In other words, a constant feature of the work is "what needs to be done now", and organising commitments around the answer to this question.[51]

Accordingly, what have been referred to as 'informal' channels of communication are often seen as expeditious ways of communicating "what is necessary" to the relevant parties, and completing the formalities later when time permits. In other words, the 'informal' patterns of communication are not particularly subversive of the formal, but represent a means of realigning the overhead of meeting the letter of the rules in order to meet the spirit of what the rules intend. Thus, the Quality Management System is not derided as pointless - though it is often seen as a "chore" - but is more a case of ensuring 'quality' in the doing of the work, and then satisfying the auditing requirements of the system "when time allows". The Quality Management System along with the various standardised documents such as the Technical and Test Reports is a paperwork system intended to represent aspects of the work deemed relevant to "determinations" of its quality.[52] As indicated earlier, working distinctions are drawn among the engineers at the centre between "doing the work" and "satisfying the paperwork" although these are not always seen, by any means, as totally unconnected activities. For them, the important thing is to "get the work done" - and they are proud of the quality of their work - and much of the paperwork is seen as having a lower priority. In other words, the paperwork is embedded in what else they have to do and, inevitably, this means that questions of priority emerge and which have to be accommodated. In this case the distinction between the 'formalities' of the work and its 'informalities' is conceived in terms of the respective 'overheads' involved in "getting the work done" and "satisfying the paperwork". In Bittner's (1974) terms, what is being displayed is "organisational acumen", that is, complying with the requirements of

---

51 A similar pattern of working has been observed in the case of the police, though the circumstances which provoke the pattern are rather different. See Ackroyd et al, 1992.

52 There are issues here as to whether or not such representations do actually represent 'quality', but this is not an issue we are taking up here.

the organisational rules, and thereby giving them respect, but in practical ways that accommodate to the various exigencies of the daily work.

The achieved standardisation of procedures

Another place where we can see the interweaving of the 'formal' and the 'informal' as practical matters that enter into the courses of action which constitute the work is in the testing procedure described at some length in the report of the study. In particular, what is striking is the importance of the 'tacit knowledge' that has to be assumed in order to perform the test satisfactorily. Although there is a procedure laid down, what is clear is that following these procedures strictly is not sufficient to practically perform the test. They have to be understood in light of other particularities, such as the nature of the materials, what levels of accuracy "will do for this kind of test", what idiosyncrasies of the equipment can be safely ignored, what setting up parameters are needed, and so on, not all of which are written down anywhere and, even when they are, still need to be understood in terms of the kind of practicalities just mentioned.

It is also important to understand that much of this 'tacit knowledge' is the exercise of informed judgements about the materials, the equipment, the test procedures, and so forth. They are not, in other words, ways of 'getting around doing the test properly' but, rather, ways of 'getting the test procedures to work out' in the particular circumstances which obtain *at* the time and in *this* place. The procedures, as comes out in the report, cannot say all that is required. As Garfinkel (1967) points out in respect of rules, procedures carry around with them, as it were, 'etc.' and 'ad hoc' provisions which need to be 'filled in' in order to make the procedure usable in particular circumstances.

This observation is relevant to one of the distinctions that is drawn between a 'craft' and 'engineering', namely, that what is distinctive about the latter is the greater degree of standardisation in techniques, materials, and tools.[53] Without entering into this issue in any depth, what is being suggested is that what transformed the individual, singular and discrete skills of a craft into a profession in which knowledge can be passed on, written about, tabularised and calculated, formally taught, schematised, accredited, standardly visualised, and so on, were innumerable developments in the synoptic representation and regulation of techniques, processes, tools and materials. However, what is suggested by these materials, and it is consistent with other studies of scientific and technological work, is that much of engineering work consists of what can be described as 'craft-ship'; that is, the tacit knowledge of the practitioner and the skills which arise through experience (Clarke and Fujimura, 1992; Knorr-Cetina, 1981; Lynch, 1993). Thus, a great deal of the practical work of science and engineering, and the

53 See, for example, Böhne et al, (1978). Ferguson (1993) argues that one of the imortant innoavtions in the development of engineering was the gradual dissemination of standardised visualisation tools, engineering drawings, standard graphical symbols,. models, etc., which vastly increased the rate of dissemination of engineering knowledge. See also Latour (1990) and Lynch (1990) for interesting reflections on the standardisation of visualisation.

test reported earlier is a good example of this, consists of 'tinkering', that is, a way of working which makes use of "what is to hand", "making do", "patching things together", "getting it to work", and so on; a way of working which has to be understood in terms the situational contingencies in which it takes place.[54]

What this means is that the standardisation of activities such as the procedures for carrying out a test need to be understood in terms of their local and situated character. In other words, the argument is not to dismiss the idea of 'standard procedures' but to see them has having an 'achieved character' from within the circumstances of the work itself. As Jordan and Lynch (1992) note in connection with the 'plasmid prep', a standard recombinant used in DNA research, the standardisation of this 'ready made technique' is simultaneously associated with processes of dispersion. That is, although the technique is widely used and described as a "highly rationalised procedure" and widely disseminated through written instructions, exactly how it is done in particular circumstances is not effectively communicated through these. The technique is mastered through repeated, often solitary, practice. Moreover, the procedure tends to change its configuration whenever new reagents and items of equipment are introduced. In other words, while the technique is standardised there are serious practical questions about how 'deep' this standardisation penetrates each researcher's practise since it must be custom fitted to the task at hand. As Jordan and Lynch report, attempts to standardise 'by the book' sometimes result in cumbersome procedures which are later streamlined as practitioners develop 'personalised' variations.

As in the case reported here, the ultimate rationality of the technical procedure is not called into doubt despite sources of variation and uncertainty. What is also parallel to the case discussed by Jordan and Lynch is the use of 'ritual' and 'superstition' involved in 'following the procedure', where these are understood to mean practises or steps which reflect incomplete knowledge of the procedure, especially in light of local circumstances. It is often easier to repeat a procedure without extensive inquiry into a potential source of trouble. "It should be alright - but...you just don't know", becomes a ritualised step based on habit and 'following what others do', but also an operating rationale for the procedure since it is not known if leaving a step out will make a difference. It displays, again according to Jordan and Lynch, an attunement to the "demonic contingencies" of working life (Garfinkel, et al, 1989).

---

54 We would make the argument that this is a characteristic of all work. However, our focus here is on the work of engineering.

Software Production

Date	Approx. Duration of Fieldwork	Relevant References
1991	2 months	Hughes et al, 1994

Objectives of the Study

This study was part of an investigation into the software engineering design process in order to inform the requirements for a support tool for the design process. It sought to examine the software engineering aspect of the system design process as a 'real world' collaborative activity involving large number of people co-operating on a complex task. This aspect of the research took place in the plant of an aircraft manufacturing concern which, during the period of the study, was designing and developing the avionics system for a small but highly advanced fighter-trainer being developed as a private venture by the company.

A new management team had been brought in to oversee the project because of earlier cost and time overruns and the research team had been invited to study the 'occupational culture' of this site which, it was felt by one or two of the more senior managers, to be "less than committed" compared to other sites in the company.

Duration and Pattern of Fieldwork

The fieldwork took, in total, approximately two months during which the fieldworker observed and interviewed the software engineering team responsible for designing the avionic system. In particular, the fieldwork focused on the management of the software integration process of what was a large-scale project. The development was intended to take three years and the project was already well into the design and development phase by the time the fieldwork took place.

The Setting

The Avionics Unit is based on a medium sized site which used to be one of the company's airfields. The Avionics Unit was responsible for the project and was organised into several subsystem groups, including Systems and Computing, Detailed Design, Test and Integration, Quality Control, Architectures and Computer Facilities. The project group made use of other site facilities involved in the design and development of aircraft, including the Software Test Rig. A Quality Review Group was responsible for oversight of a number of projects throughout the company.

During the fieldwork the project employed about 160 engineers, including the executive engineers who were responsible for the management of the separate subsystem groups. However, the mix of skills as well as the number of personnel

varied according to the stage of the project. See Fig 3.13 for graph of approximate manning levels expected during the lifetime of the project. Thus, and typical of large scale projects such as this, the composition of the project team was expected to change over time both through natural wastage and as engineers were brought in from 'line departments' when their specialisms were required, and as others left when their work was completed.

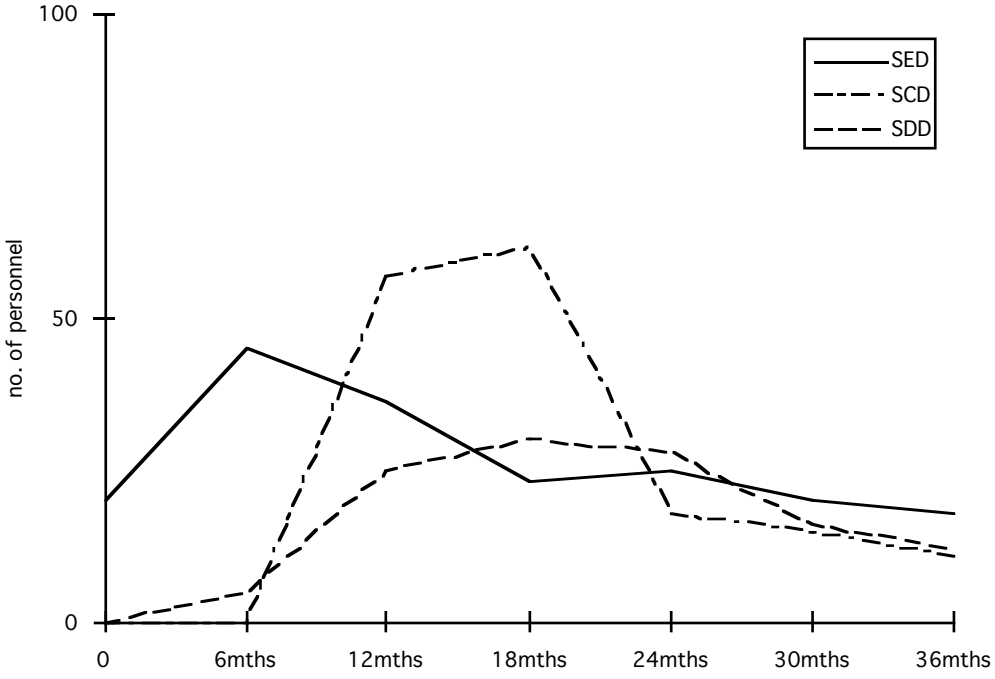

Fig 3.13: Expected Manning Levels by Project Months

The Work Process

Project life cycle management

In large scale projects such as this considerable co-ordination and organisation of interdependencies between specialist groups is involved. The project had three phases not untypical of life cycle models used in software development: systems design, hardware and software implementation and systems development testing, which included flight testing.

The system design phase had been done by the Top Level design team who had translated the conception of the aircraft down to a level at which code could be produced. This group was relatively small and had been disbanded after this phase of the design had been completed, though some members of the team became

leaders of the various subsystem groups noted earlier.[55] The hardware and software implementation phase consisted of the detailed design of the system components and integrating them into the fuller avionics system. This stage involved many more engineers than in either of the other two phases. At the time of the fieldwork they were reaching the end of the system integration phase, that is, the last part of the second stage of the development life cycle.

The final development phase, System Development Testing, consisted of ground tests, when final 'fixes' are done, flight testing, which includes 'fixes' on pilots' reports on functionality, then into production and, finally, into service. During this phase the project team 'winds down' to disbandonment. See Fig 3.14 for diagram of project phases.

At the time of the fieldwork the project was 23 months into its design life cycle of three years and was roughly on-target for completion as planned. There were no substantial problems with the basic design of the aircraft which had been well proven on previous versions. The main problems for the Avionics Team had to do with updating and innovating display configurations.

Organising such a large scale development project efficiently and to the necessary quality was the major responsibility of management. In practical terms, in this company this meant implementing the plan as set out in the company's Technical Directorate Exposition.

The Project Delivery Plan

This consisted of a large number of hierarchically organised documents which specified management responsibilities and procedures for quality control, review stages, planning and design verification, document control, inspection and test status, corrective action, training, techniques and administration. This plan constituted but part of an elaborate collection of documentation which formed the Exposition Hierarchy for the company's projects. This laid down a life cycle model for generating the 'product', in this case the avionics subsystem, for the customer, initially the Chief Engineer. Thus every phase and task of a project was specified in detail through this documentation. This focus on strict process management was justified in terms of the needs of safety critical systems.

In effect, what the plan laid down was a division of labour for the production of, in this case, the avionics system. It laid down not only what needed to be produced as a series of system components, but when and how they should be produced, using what resources and for how long. It also recorded, in fine detail, the progress of the plan through a set of procedures laid down in Configuration Control.

---

55 This created some problems for the Detailed Design group who complained that they did not understand why the Top Level design decisions had taken them down particular routes rather than others. Because of the use of object-oriented design techniques, this meant that the Detailed Design team were given objects which had been decomposed by the Top Level team in ways which were not always clear.

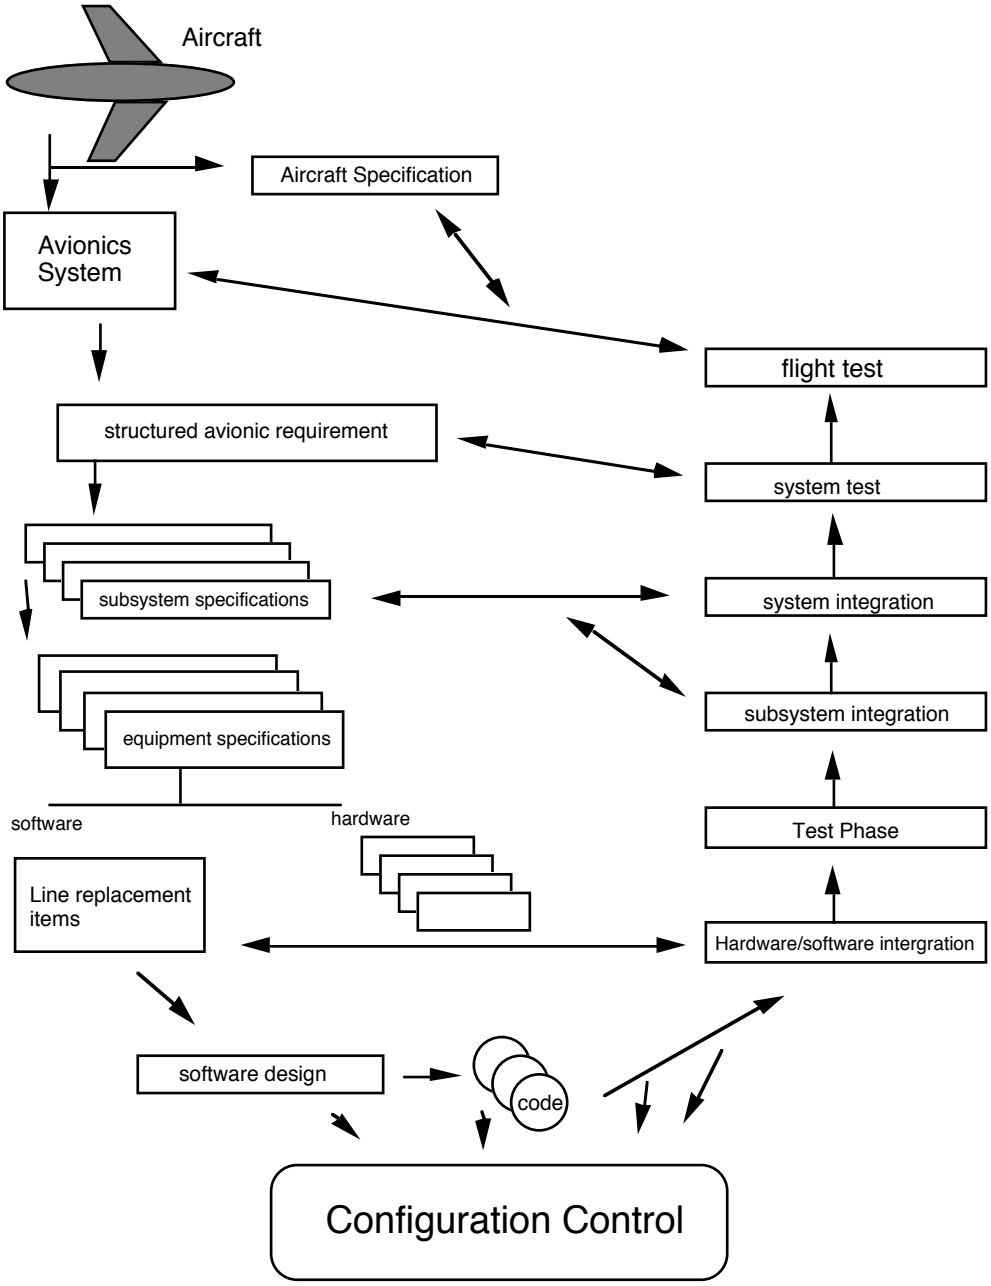

Fig 3.14: Scheme of Project Life Cycle and Its Management

Configuration Control

One of the major tasks in large scale software design is keeping track of things as the work proceeds. Another is obtaining an overview of the various stages. As in many large-scale projects the software development process in the Avionics Unit was ‘document driven’ and this was the responsibility of Configuration Control, a team that consisted of 5 people including the team leader.

Configuration Control Management System (CCMS) had the responsibility of managing any change to the work, documentation, integration, software and hardware, in the project by tracking documentation, chasing system problem reports (SPRs) for closure. Every configurable item is managed, after its issue, by the highly specified and documented change mechanism of CCMS. It is required to monitor and maintain a standard within the project and keep overall control. Although not part of the life cycle model itself, CCMS is involved at all levels and stages of the project life cycle by instantiating a set of procedures for organising and acting as a funnel into which all stages of the life cycle pass through (See Fig 3.14).

One of the more important functions of CCMS is to serve as a library keeping all master copies of documents, mostly electronically. If a document is required, then a request has to be made to the CCMS librarian who then makes a hard copy which is numbered. If the document is configurable, that is, still capable of being changed, then it is stamped 'CONFIGURABLE ITEM' to indicate that the document may be altered but only in the authorised manner. The requested documentation is given a 'fly sheet' signed by the recipient and returned to Configuration Control, and filed until the document is returned. If changes are required a System Design Change Report (SDCR) or a System Problem Report is raised - a one page sheet giving a brief description of the problem, who is raising it and who it is actioned to - and requires the signature of an executive engineer. This is given via a Change Control Meeting whose remit is to discuss requested changes and sent to the relevant group for action.

All documentation is numbered with a copy number, the project, site and department numbers. The system does not authorise configurable items to be changed or altered in any way; engineers are even discouraged from making personal notes on the documents.

Configuration Control acts as a 'tracking mechanism' and has responsibility for the 'whole picture', including integration, time keeping, audit and accountability. It is the formal means by which the project was tracked through its stages. All notifications for software changes, which could emerge from the test rig, implementation, or from attempts to integrate with other components, needed to be authorised by CCMS.

Design contingencies and fatalistic dependencies

CCMS is a strict procedure while acknowledging that changes to the design will be needed at various points during its development. Keeping track of such changes is what the procedures are for. However, the problems that are raised depend upon many factors not least upon the quality of the work being done as well as upon contingencies that arise during the course of the development of complex systems. What the procedure does not lay down, and could not lay down, is how many or what changes will be required. All it does, and this is by no means trivial, is specify how such changes should be authorised and recorded. It is by this means that it provides information for management as to how the project is

progressing and what work needs doing in order to carry the project forward within time and cost.

The procedure, in other words, fatalistically depends for its practical realisation on such matters as the competence of the engineers, their goodwill, the adequacy of the original design, not to mention the other exigencies which can afflict any project in any organisation.[56] In this case, although it seems that the basic design was sound, it had to be developed within a fixed budget and within firm time constraints. Any contingencies that arose had to be dealt with within these parameters.

Fixes

For the software engineers in the Avionics team, 'fixing' code, that is, making necessary changes as a result of test failure or requirements changes as well as fixing bugs, was a prevalent and routine feature of their daily work. However, because of CCMS procedures such 'fixes' had to be formally notified to the monitoring process. The Change Control Procedures laid down that code be formally released for testing and, if a result of the test, changes are required then a SPT be raised against the code. The Quality Group was responsible for testing and formal release who then raise SPRs on any problems that arise. The SPRs are logged and actioned via Configuration Management for work to be carried out to solve the problem whatever it might be.

For the engineers this procedure was not only cumbersome but, they felt, reflected badly on their work if there were too many of these reports. As a result ways and means were used to reduce the level of formal reporting. Rather than rely on the prescribed formal testing procedures, engineers often made use of 'engineers' testing' which allowed them to informally test code by reporting it as 'still under development'. This was a way of testing new or changed software on the dynamic test rig prior to formal build. This allowed some of the problems to 'show up' and be fixed before formal test and, by this means, improving the quality of the code prior to formal testing by the Quality Group. Engineers' testing was not acceptable as any formal proof of the quality of the code, and management were fully aware of the practise as commonplace.

However, one of the problems the team faced was that SPRs were "out of control". Management response, understandable in many ways, was to create more pressure for closing out SPRs quicker and improving quality by invoking more procedures; a management style which many of the engineers felt unnecessary, undeserved and compromising of their professional abilities. As was indicated, from a managerial point of view, the pressures arising from the contingencies of the development process when aligned against the project plan have to be dealt with which tended to take the form of exhortation, urgings, and

---

56 These are too numerous and indefinite to list but include such things as shortage of labour, the general economic cycle, cancellation of the project, changes in govenment policy, as well as more local contingencies.

"kick ass" responses which challenged the professional pride of the software engineers.

CCMS review

During the period of fieldwork, CCMS reviewed their procedures, partly as a result of one of the key personnel leaving and partly to improve the service aspects that CCMS was supposed to deliver to the project.

The first of these refers to the 'old hand' phenomenon. That is, despite the strict and detailed procedures which were laid down, much of how the system operated remained 'in the head' as 'tacit', taken-for-granted, 'know how' of those who operated the system's procedures.

The second of these tried to address the problems of delays in the system. For example, and among many others, there was a need to update and improve areas of Configuration Control work including ways of making the work more visible and less centralised. Another was to update the database to be current to half a day and, in this way, speed up access. A constant complaint of the engineers was the lack of an "overall picture" in terms of SPRs which, in turn, reflected on the quality of their work. On the other hand, CCMS was concerned to find ways of making the work more responsive to the needs of the project without the amount of paper passing across desks getting out of hand. Also, Quality Control required CCMS to keep them informed of the database changes in order that they could control fixes and chase SPRs.

Accordingly, CCMS were trying to juggle the various priorities of subsystem groups with their own priority of managing the work of the project.

Emergent Features of Social Organisation

The embeddedness of work activities in organisational structures

The work, as it is in project engineering generally, is of the problem solving variety. There were those which arose from the technical problems of the project and those which arose from the necessity to get the work done. The project was tightly resourced and, again as is typical of projects of this type, some tasks were in danger of exceeding the resources originally estimated for them even though the project as a whole was about "on time". Accordingly, many of the problems involved getting the work done given the resources, manpower, time and latitude, etc., factors which differentially afflicted different subgroups within the project. These are, of course, only the tip of the iceberg relative to the innumerable organisational constraints which entered into the resourcing and supporting of the project in the first place, and the extent of knowledge of the 'organisation's ways' in the negotiation of a project's course.

A prevailing concern in the organisation of projects is whether or not they are 'on track'. It is an omnirelevant concern in that checking the progress-so-far of the work against its schedule is a pervasive aspects of the tasks with the constant consultation of schedules, reviewing problem lists, their rates of clearance, and so on. Meeting such targets is a high priority, though not necessarily an inflexible restraint on projects. Many of the tasks, for example, were rescheduled in light of emergent problems. The question of whether a project is 'on track' determines how work tasks are prioritised within the teams. the 'outlook' for tasks in hand is a matter of perpetual inquiry in that how one person's task may be achieved may well hang upon the specific character of some else's achievement of their tasks. In the case of the software team studied here one of the organisational factors affected the achievement of the work was Configuration Control.

However, the determination of whether or not a project is 'on track' is not only about the pacing of the work. There are equally strong concerns with whether the work has been done in conformity to the specifications, whether the work done by different individuals is going to interface when it comes to the assembly of the final product, and so on. The prospects of unilateral decisions, overlooked tasks, unwitting deviations from schedule are all constant sources of anxiety within the project, and there are multifarious mechanisms, ranging from the complex, formal machineries such as Configuration Control to informal interpersonal consultations, to provide for 'double checking' against the schedule.

Branch Office and Specialised Services Banking

Date	Approx. Duration of Fieldwork	Relevant References
1994 and continuing	4 months to date	

Objectives of the Study

The study is linked to a DTI/SERC project, Systems Development and Co-operative Work: Methods and Techniques, in association with Digital, NatWest Bank and Syncho, a managerial cybernetics consultancy company. The project is concerned with examining the respective contributions of two methods of research and requirements elicitation in CSCW, that is, ethnography and managerial cybernetics using the bank as a research site and potential user of CSCW systems.

Duration and Pattern of Fieldwork

So far the fieldwork has lasted some four months. The first phase, of approximately two months, was spent studying the working of a 'typical' Branch Office, and the second phase, again some two months and continuing, involves

the detailed observation of the work of specialised centres, the Securities Centre, the Lending Centre, the Service Centre and the Business Centre, to which most of the 'back office' work of the local branches has been or is being transferred. Despite the comparatively long time spent in the field, because of the highly distributed nature of the work observed, distributed both in space and time, the fieldwork should more accurately be characterised as a series of 'quick and dirty' ethnographies with the fieldworker spending a number of days at a particular desk or with an individual worker becoming familiar with the work and its place within a complex division of labour. The completion of this phase of fieldwork will be followed by an 'evaluative' ethnography, concerned with providing an 'illuminative' evaluation of the organisational changes and team working suggested by managerial cybernetics and implemented by the Bank.

The Settings

The settings have been various, including two 'traditional' local branches which dealt with most of the work involved in local banking. The 'back office' operations of these branches is now in the process of being transferred to specialised centres leaving only cashiers and enquiries with the local branch. The specialised units tend to concentrate on one banking function and serve a number of branches whose main function is to serve as an initial point of contact with the customer. The Securities Centre, for example, which was concerned with taking, maintaining and releasing securities for loans, was located in a modernised former branch office and served 60 local branches. The Lending and Service Centres were located on a number of floors in a large modern building. The latter deals with a variety of 'back office' processes such as Customer Records, Entries Handling, Commercial and Foreign, Standing Orders, and so on, that were formerly the responsibility of the local branch. It also included a telephone room.

An important element in the rationale for this reorganisation of the infrastructure of the Bank's services was not only to improve efficiency through centralising specialised functions, but also to change the working culture of the bank from one which was essentially concerned with administering customer's money, to one which was also concerned to sell a range of services that the bank had to offer.

The Work Process

The two settings of the traditional local branch and the specialised centres provide an interesting contrast in the different ways in which the same functions can be realised as patterns of work. In the former, most of the functions were performed in the same place and although there was a degree of specialisation within a division of labour, in the specialised centres, as their name implies, this specialisation has much greater scope as well as specificity. Each specialised unit is dealing with, at most, just a few of the functions previously met within the

branch, involves more people performing just this function and servicing the requirements for that function of a number of local branches. Thus, a whole room of workers in the specialised centres would concerned with functions which were often the sole responsibility of a single worker in the traditional branch. For example, the accounts of deceased persons was typically only part of the work of a clerk in a branch, but in the centralised processing units it constitutes the entire work of one worker.

The Branch

In the main Branch studied the work was situated on three floors of a large town centre bank, and comprised a number of clearly identifiable units; for example, the Machine Room - Entries Handling Unit; the Lending Unit; Customer Services, the Enquiries Desk and the Cashiers.

The Entries Handling Unit is primarily concerned with the checking, keying in and administration of data, with most days following a pattern of the morning being spent in filing and checking (cheques from clearing, keyed entries, rejected credits, stopped cheques), and administration (computer printouts, statements) and the processing (transferring data from the paperwork into the computer, providing a balance of the day's work) taking place in the afternoon.

The Lending Team is concerned with the management of loans and overdrafts and the general running of accounts. Each day they work through the 'WE008', the 'out of order' list - a computer printout of people who have exceeded their credit limit and 'actioning' their accounts. Others are meanwhile concerned with any response, or failure to respond, to the letters the Bank has sent out to borrowers who are exceeding their credit limit. This takes the form of either writing to or phoning the customer in question. All decisions and correspondence are then entered into the customer's record card.

Customer Services comprised a number of personnel dealing with a wide range of work processes; records (opening, closing and transferring accounts); foreign business; share dealing; wills, and deceased accounts; credit cards; mortgages. Besides these relatively specialised activities all the staff were concerned with answering the telephone (within three rings) and acting on any requests made (usually a balance enquiry or money coming into an account). The working atmosphere here appeared to be the most intense in the Branch, most subjected to frequent and sometimes prolonged interruptions and most dependent on informal teamwork and the sharing of 'local knowledge'.

The layout of the branch main office is set out below in Figure 3.15.

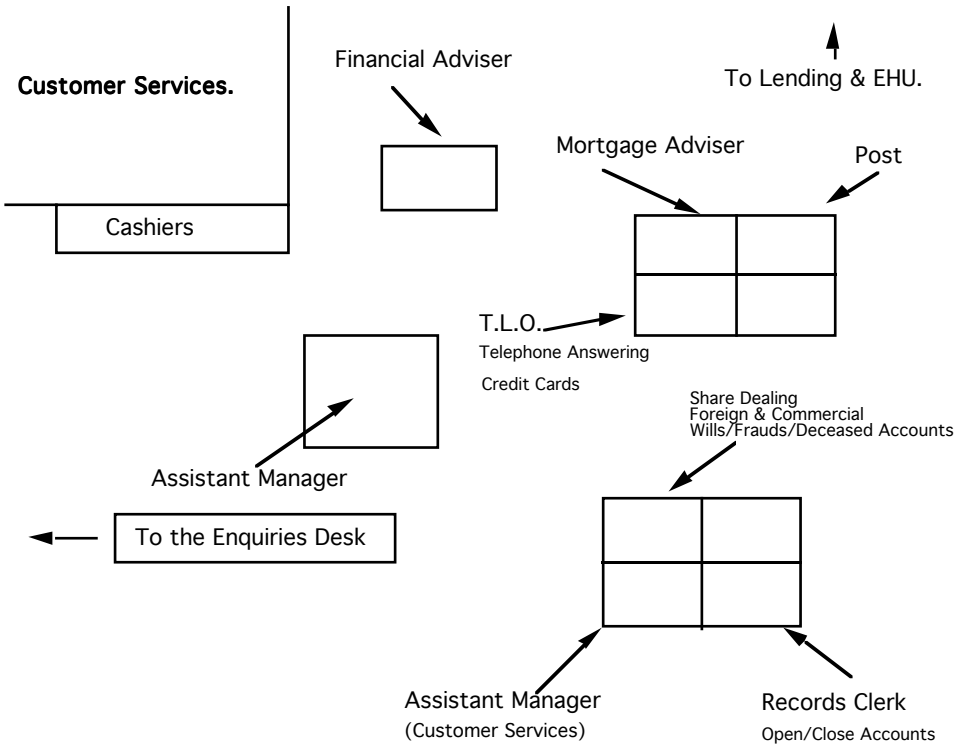

Figure 3.15: Office Layout of Branch Office

The Enquiries Desk is often the customers' initial point of contact when entering the branch. From here they would either avail themselves of the standard services, such as withdrawing cash, paying in, and so on, or be directed to enquiries for whatever other service they may require. Staffing at the Enquiry Desk varied with customer flow and although additional staff could be summoned when queues built up, typically it consisted of staff with considerable experience who were able to deal with the range of enquiries regularly encountered.

Most customer enquiries were routine although multiple enquiries were common. Although the most frequent was a balance enquiry this was often combined with requests about standing orders, debits, buying or selling shares, changes of address, opening accounts, forgotten PIN numbers, statement queries, and so on. Nevertheless, enquiries were reasonably routine.

The cashiers desk - depending on the flow it was staffed by between 2 to 5 clerks - dealt with a number of routine transactions, mainly the withdrawal and deposit of funds and the provision of balances. Less routine business was normally sent to the Enquiries Desk.

Centralised Processing Units

Although the work of the Branches is essentially replicated in the specialised processing units, the setting and the context in which the work takes place is very different. The Securities Centre, whose main task is concerned with taking, maintaining and releasing securities, is located in a modernised former Branch

office and serves some 60 branches. Each worker has a PC on their desk acting as a dumb terminal for using TecSec, a new software package resident on the server. Work within the Centre is performed by two teams each consisting of a hierarchy of checking responsibility, from the 'Doer1' at the bottom, who processes the requisition forms, to the Senior Securities Officer and Assistant Manager at the top, who deal with complex indemnities and debentures and who are responsible for checking the work of the teams. The layout of the Centre is sketch in Fig. 3.16.

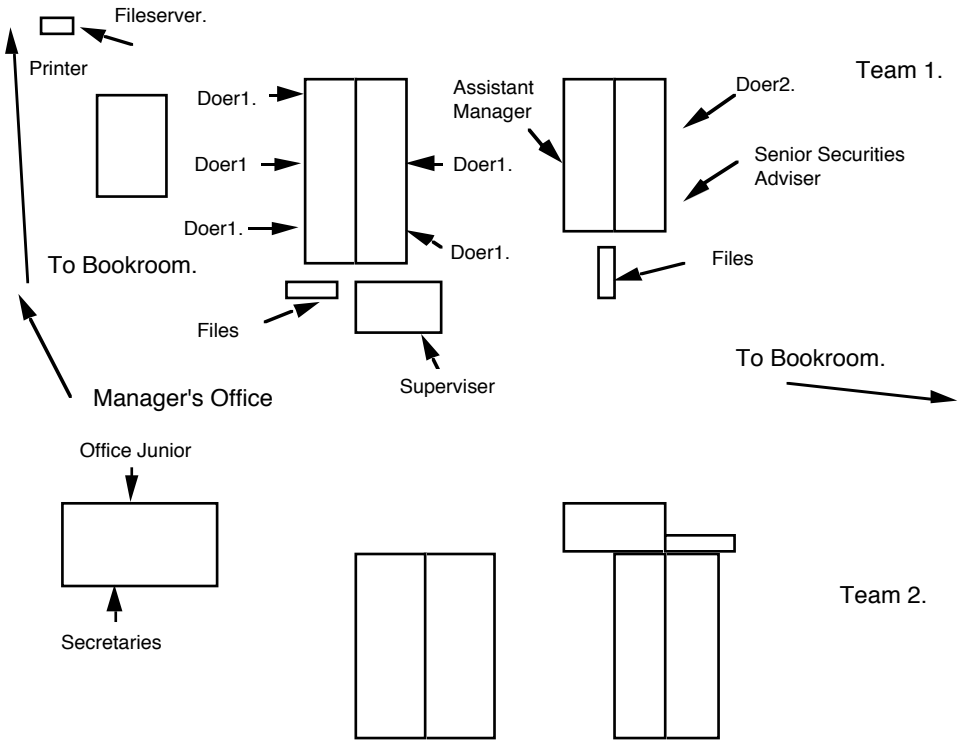

Fig 3.16: Securities Centre Layout

The Lending and Service Centres were located on a number of floors in a large modern building and perform the lending functions for the branches in the region as well as having a Small Business Lending section and Debt Recovery. Branch Lending Control is divided into a number of teams, each occupying a particular desk, that service particular branches, and generally consist of a Lending Controller and two Lending Assistants, overseen by an Assistant Manager and a Manager - a division of labour and responsibility that is instantiated in the lending limits attached to each position.

The layout of the Service Centre is set out below in Fig. 3.17.

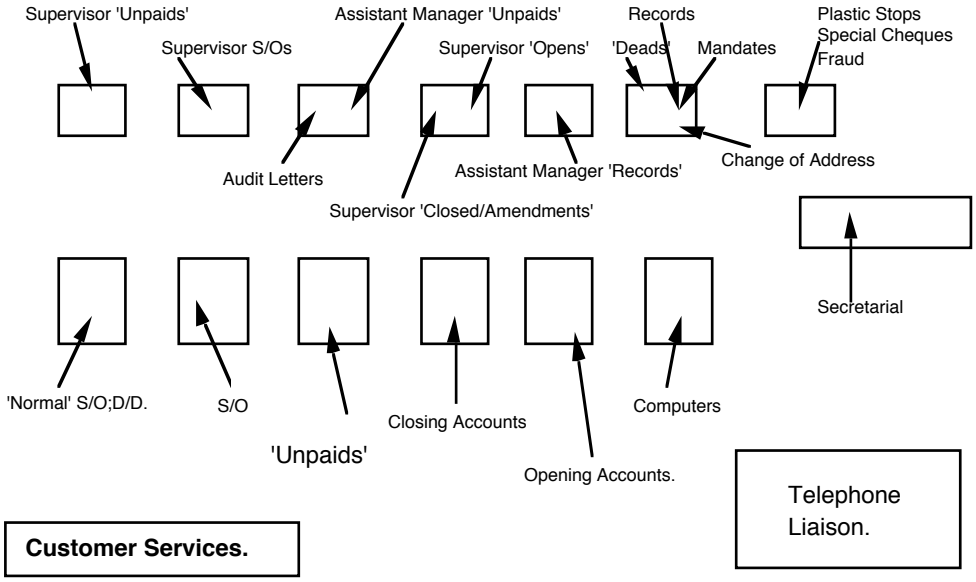

Fig 3.17: Service Centre Office

The Service Centre deals with a variety of ‘back office’ processes - Customer Records; Entries Handling; Commercial and Foreign; Standing Orders and so on - that were formerly the responsibility of the branches. A Telephone Room handles incoming calls, generally aiming to answer some 80% of enquiries but otherwise transferring the calls to the requisite desk.

The technology used is an ‘expert system’, known as TecSec, and many of the work activities are driven and controlled by this system. Other systems in use include BAF and ISS, though as yet there is little useful integration between them.

Emergent Features of the Social Organisation of Work

Informal team working and the working division of labour

As is typical of much office work, activities in both the branch and in the centralised units largely consisted of applying specific routines and procedures to tasks which they identified. In the case of the Bank, this is important not simply in order to do the work, but also to maintain an ‘audit trail’ in case of any problems, specifically the possibility of fraud and theft.

However, it is clear from the fieldwork that these routines are not slavishly followed or, better, in following them judgement is often required in order to apply them to particular transactions. This is as much the case in the Securities Centre where the work is driven by the system as it is in the branch offices. Whether or not routine is strictly adhered to or whether ‘short cuts’ or modifications are employed is a matter for ‘occasioned determination’. That is, as the result of a loose assemblage of contingencies that can arise during the course

of a working day. These include, for example, items that "do not quite fit the procedures", 'unusual' items, ignorance of arcane procedures, and so on. Both fieldwork sites showed that the speedy and efficient completion of work tasks depended a great deal on informal teamwork and 'constellations of assistance' as well as the use of 'local knowledge' both by individuals and as shared by members of the team. Thus, although much of the work is routine and procedurally driven, as it has been designed to be, dealing with the inevitable contingencies often meant departure from the procedures in order to get the work done 'as is required'. In many cases, this is a co-operative activity in that occasioned reliance is made on the knowledge and competencies of others in the team - a feature which the use of computer systems do not obviate.

The egological organisation of work

By this term we intend to point to the way in which the activities of the work, from the point of view of a person within a division of labour, are organised by a scheme of relevances. That is, the practical accomplishment of the work activities of a person within a division of labour requires learning, knowing about, knowing how to use artefacts of various kinds, etc., to perform the tasks but also knowing 'what is my responsibility' and 'what is not my responsibility', 'what are tasks that I have to do' and 'what are tasks that others have to do'. In other words, it is being aware of the interdependency of tasks but, at the same time, seeing these very much from the point of view of 'my tasks'. This is not simply a matter of implementing some mental finite list of relevancies or of necessarily seeing this as a constant issue. For much of the time, the work is routine and done without over much reflection. Each person operates within 'horizons of relevance' in which events, items, persons, information, incidents, etc., move in and out depending upon the eventualities that can arise during the course of the working day.

The "egological principle" is both illustrated by and yet is more than one worker's comment: "I just do things...I don't think about it", since it is precisely this ability to 'just do things' - that is, 'just do the things that I have to do' - that enables the work to get done. In the Lending teams, for example, the Controller 'knows' that he or she will be interrupted by the assistants with requests for advice, with information about accounts, and so on. He/she will know that "phone calls are somebody else's problem", that once a decision has been made 'actioning' that decision is for "someone else in the team to get on with", and so on. As we say, these 'horizons of relevance' are not known as itemised schemes of knowledge; rather they are learned and used as understandings about the work which only occasionally surface explicitly.

Although the "egological principle" is important for the skilful performance of tasks within a 'working division of labour' it is not always beneficial and can lead to a mentality which fails or refuses to acknowledge the implication of one's own task for others. Examples of these include failing to report theft from a till and

failing to seize selling opportunities, as the staff are increasingly required to do by the Bank, when they present themselves.[57]

Teamwork and skill

The issue of 'skill' becomes prominent in this ethnography because of the contrast between the 'turn to the social' in much of CSCW research with its emphasis on the social basis of skill and competences, and the widespread organisational emphasis on predominantly 'technical skills'. For example, for the Bank the development of the various specialised processing centres is clearly directed at reconfiguring the 'skills' required in a manner more commensurate with the technology required for much of modern banking processes; that is, through implementing a complex division of labour, familiarity with the technical processes and, hence, increasing processing speed. In practice, however, technical competence is (perhaps unwisely) simply assumed and the emphasis returns to predominantly social skills and competencies, such as "communication skills"; "team working"; "problem-solving" and the like (Bertrand & Noyelle 1988; Austrin 1988).

Nevertheless, observations suggest that the relationship between computer work and skill, especially team working skills, is complex, contingent and context dependent. In other words, 'skill' is a complex and contested concept that needs to be 'unpacked' and that while clearly work can be analysed in terms of skill and skill requirements such an approach tends to ignore the complex interweaving of skills in work; the variety of approaches (and skills) to task completion and the context in which skills may be exercised. Skill requirements are constructed through a social process and in a particular social context and in many ways it is this process and this context, revealed through ethnographic research, that it is important to describe and understand. Work context influences both the opportunity and the motivation workers have to exercise skills that may be regarded as important. It is, perhaps, for this very reason that the creation of a 'selling' culture within the bank has proved difficult; turning, "tellers into sellers" requires rather more than executive mission statements, 'visions' or slogans. It requires an organisational context in which selling skills become relevant to the work.

Even 'technical' skill, skill as a form of knowledge, involves the acquisition of the culture in which a particular work activity is accomplished (Collins 1987) and the subtle but essential competencies involved in making sense of, and thereby being able to make available to others, what is 'going on'; competencies required for 'mutual intelligibility' on the part of the members of a work team. For example, using the 'TecSec' system in the Securities Centre is not simply a mastery of the specific skills incorporated in it, it is also knowing and being able to rely upon the fact that others in the team know the relevance of such knowledge for the work being done. As has been shown in other contexts (see, for example,

---

57 Workers, in effect, become 'bystanders' to their own working lives (Latane and Darnley, 1970).

Harper and Hughes, 1992), this kind of competence is essential to a whole range of informalities involved in performing work activities, including 'knowing where others are in their work', 'getting round the inadequacies of the system', 'knowing who to rely upon to get things done', and more. The issue of 'skill' in CSCW is then complex and problematic and approaching the issue through a relatively narrow conception of skill is likely to produce limited answers. It is not that the technical skills involved, for example, in being able to use computer systems are unimportant - they are clearly essential - but so too are the skills or, better, competencies, involved in working within a team. knowing how and in what ways to use it as a resource and knowing how, and in what ways, to contribute to this resource.

Emotional teamwork

The final area of skill and of team working apparent in both fieldwork sites was that of 'demeanour work' and 'emotional teamwork'. While traditionally considerations of 'emotion' have been missing from accounts of organisational life, with its emphasis on the 'rational' - and by implication 'non-emotional' - actor; and while obviously banks are hardly emotional hothouses, 'emotion work' - the managing, control and occasioned display of appropriate emotion - was also clearly part of 'doing the work'. Individuals working in the various sections of the bank are expected, as members of a team, to embark on various forms of demeanour work or emotional control. This is not simply a feature of bank policy, reflected in "smile, your next customer may be the mystery shopper" or the comment "a complaint is a sales opportunity". Demeanour work - often translated or interpreted as "smiling down the telephone" - was a regularly observed feature of the customer contact. 'Emotional teamwork' was also manifested in group responses to and discussion of customers and phone calls, and seemed to serve both as a release of emotion and as some kind of indicator of group solidarity.

Local knowledge

The use of 'local knowledge', that is, knowledge of specific customers or particular processes or routines that facilitated work completion through the employment of modifications or shortcuts in the formal procedures, comprising what Bittner (1974) calls "gambits of compliance", was observed in both the fieldwork sites. And this despite the organisational belief that the move towards centralised processing units, and the accompanying increase in account volume, would (and should) reduce this phenomenon.

Within the Branch office it was evident that the accomplishment of work tasks (especially their speedy and efficient completion) was, to some considerable extent, dependent on informal teamwork, 'constellations of assistance' and the provision and employment of the 'local knowledge' of colleagues, thereby avoiding lengthy perusal of the PIF, the manual of job routines. The most striking example of customer knowledge concerned an overdrawn account. One of the

staff remembered the name and suggested that the case be referred to Lending saying, “he’s a God to this branch ... ask Lending what to do..” so preventing the branch unreasonably charging one of its major customers.

In the specialised processing centres despite the anticipation of the disappearance of local knowledge and the strict adherence to standardised procedures, local knowledge and ‘local logics’ continued to flourish. For example, workers were familiar with the specific requirements or circumstances of particular (usually the ‘awkward’ or the ‘bizarre’) customers, but there were also other instances when local knowledge was used to speed up the work process. Local knowledge was used for the customer’s benefit and local logics were deployed to meet the substance of the customers requests even though it involved negotiating or overcoming standard procedures. Thus from the fieldnotes - in the tale of the ‘Icelandic Fishgutter’ or the difficulty of transferring money to Morocco - the idea emerges that local knowledge is important in enabling them to quickly make sense of the work they are doing, and in orienting themselves to the requisite activity (which forms to complete; which forms to issue/check etc.). Thus local knowledge is still used, and used with judgement and skill, to expedite the subtleties of the work which ‘following procedure to the letter’ would ignore or with issues which the procedures or, more likely, the technology does not support.

Computer work and paper work

Financial institutions were among the first wave of business organisations to computerise many of their routine processing operations and, like other innovators, they have suffered the consequences in terms of overloaded, unreliable and outdated systems and software. Even in the more modern centralised processing units the machines regularly ‘went down’ or were slow, or impossible to access because of overloading the system. Two software packages - an accounting/bookkeeping package, developed in the 1960s with ‘bits bolted onto it’, and a modern relational database - were in more or less constant use throughout the organisation. But the fact that both were frequently used in conjunction and yet could not ‘talk to’ each other meant that quite routine and regular enquiries were unnecessarily time-consuming. TecSec, the new system used in the Securities Centre which effectively ‘drove’ the work of the unit through the completion and release of formalities, was obviously more impressive. Yet even here the system seemed markedly inflexible, slow to adapt to changes in procedure or legal precedent thereby pointing to a more general issue of ‘usability’ in system design and the need for systems to resonate with the actualities of use and respond to changes in usage. Organisational responses to change and deficiencies in the system have tended to be cumbersome; an indication, not necessarily of organisational inertia but of the difficulty of modifying systems in use or building flexibility and responsive monitoring into system design.

The sheer volume of paperwork and the varied ways in which the paperwork essential to banking, and the information in the bank's computer systems are compared, checked and used in everyday work was a notable feature of all the fieldwork sites. Even in the centralised processing units paperwork was still an important and inescapable feature of the work, though to what extent this was an inevitable product of centralisation without appropriate computer support or a consequence of the continuation of the 'administrative culture' is difficult to determine.

As modern banking moves toward increased electronic techniques and methods, the functionality of paper and the relationship between paperwork and computer work will become increasingly important to understand.

The centrality of paper in the everyday work of the various centres is demonstrated by the frequency with which the initial activity of the day centred around the perusal and 'actioning' of computer printouts. The paperwork frequently took centre stage in a series of 'modal transformations' and understanding such transformations in relationship to particular work activities may assist in thinking about the design of CSCW systems in respect of the regulation and management of types of information. However, paper and its very tangible existence, was also a part of the 'audit trail' - the need to attribute blame for mistakes or to be able to trace through errors - and as such observable in all the fieldwork sites being an especially important feature of various 'checking routines' that were implemented in the different sites and providing 'accounts' or 'mundane reasoning' behind work activities.

The 'Bank': the Organisation and Organisational Change.

The issue of the perception of the Bank as an organisation became relevant in this study because of the frequently articulated management desire - in staff meetings, 'house' magazines and so on - to move from an 'administrative' organisation to one which incorporated some kind of 'selling' culture - and the impact this appeared to be having, alongside other factors, on staff 'morale' and their sense of the work. The perception of the organisation - its 'image' - far from simply being a management problem, also raises the issue of the relationship between CSCW and organisational form, specifically whether computer supported collaborative work needs to be or can be 'tailored' to specific organisational forms; whether CSCW is simply to responsive to existing organisational cultures or can itself encourage or provoke organisational change.

Within the Branch there was a great deal of concern and uncertainty, as well as cynicism and scepticism, about the organisational changes combined with a sense of what Gabriel (1993) calls "organisational nostalgia" - reflecting not necessarily a longing to return to some mythical past but discontent with the present and the perceived direction of the future. Unlike that noted by Fearful (1992), this mood was not associated with the movement towards, or introduction of, computerisation but seemed to be linked more to the movement away from the

historical administrative function of the Bank towards a new emphasis on selling. In the processing centres morale appeared less of an issue partly because some attempt had been made, and was being made, to fit identifiable 'competencies' to work tasks, and partly because, their major function remained the administrative activities which they historically associated with bank work, and with which they felt comfortable.

The final aspect of organisational culture and organisational change to be raised here is the relationship between organisational change, technology and control - a relationship explored in some detail by Yates' (1989) historical study of the development of office technologies in late 19th century United States. What the study demonstrates, among other things, is the close relationship between technologies which support the formal and standardised forms of communication, filing of information and reports, retrieval systems, etc., and the need to control patterns of distributed working within organisations. To this extent, many of the organisational, procedural and technical changes introduced in the bank, are an extension of an evolution of work technologies which began with the growth of organisational size and complexity. Thus, the PIF, the use and storage of the computer printouts that are actioned by the staff, the KSAs, the in-house magazine, and PMP, the 'Mystery Shopper' campaign and Customer Service Index, may be viewed as attempts at 'control through communication'. Clearly, as Yates notes, the available technology has played some part in this process but essentially the technology seems to be responding to changing managerial needs and methods of management and is instantiated, most obviously, in the TecSec software in the Securities Centre since it incorporates all the essential features of systematic management - a careful definition of duties and responsibilities coupled with standardised ways of doing these duties and a specific way of gathering, handling, analysing and transmitting information.

There is, of course, a rationale for introducing and relying upon such procedures and the technologies which support them; rationales which are familiar aspects of the organisation of work within industrial societies. The issue, however, which we want to point to, is that although large organisations of all kinds depend upon 'control through communication', such mechanisms 'work' not only to the extent to which persons using them use them reliably, consistently and skilfully as part of their daily work activities - in this way turning them into resources for the accomplishment of the work - but also to the extent to which they are able to exercise judgement as to when, on what occasions, and with what result, they can 'circumvent' them in order to accomplish to 'spirit' of what the procedures intend. This issue is particularly relevant to the design of CSCW technologies.

Summation

This Section has reviewed a number of projects connected to the COMIC programme of research. They have been used to build up a corpus of fieldwork materials on the social organisation of work in order to:

- gain experience of using ethnographic fieldwork in a number of work settings and under a variety of conditions;
- increase our understanding of work within a variety of settings;
- provide materials which are capable of informing the development of a framework for the analysis of the social organisation of work.

What we have also tried to do is offer some analyses of the features of the social organisation of work which have emerged from the studies. These are not, of course, the only ones which could have been identified, and nor have we made any rigorously systematic attempt to develop a comparative analysis across the studies, even though many of the features are exhibited in a number of them. From our point of view, it is important when carrying out fieldwork for CSCW design that sufficient attention is given to the specifics of a domain rather than attempting to abstract and/or generalise too early. This is the importance of the case study approach to work analysis as a preliminary to CSCW design. However, while strongly affirming the importance of a case study approach and its attention to the specifics of work settings, we intend to use these materials to develop a framework which identifies some of the main properties of the social organisation of work, properties which are variably instantiated in particular domains.

To this end, in the following section we discuss more general themes which have emerged from the studies as a foray into the development of such a framework which will constitute much of the research work in the coming year. As a starting point to outlining a set of perspectives suitable for a CSCW design we look toward a set of perspectives which:

- encapsulate many of the emergent features highlighted across the set of field studies presented in this section;
- allow meaningful design implications to be drawn from them.

The structuring provided by perspectives serves the purpose of bridging between the observed nature of work and the needs of CSCW systems development. In particular, the perspectives should provide a means of sensitising designers to the actual nature of work while also allowing them to focus on the concerns essential to effective system design. To service this tension between representing the actual nature of work and the abstract structuring demanded by systems development we have selected a small set of exemplar perspectives.

Distributed Co-ordination

This perspective is focuses on the means by which groups co-ordinate their work activities and, in this respect, is relevant to the design CSCW systems to support the co-ordination of work.

Plans and Procedures

Plans and procedures are central to co-ordination work and, consequently, to the computer systems which are intended to support it. This theme intends to convey to designers the actual nature of plans in action.

Awareness of Work

The need to support and promote awareness is a key factor in organising co-operative work and developing systems to support co-operative work. This perspective considers the various ways in which awareness is manifest and supported in work settings.

The choice of these perspectives as our initial set is principally motivated by the needs of designers and developers who highlight these concerns as central to the development of CSCW systems. It is not intended as an exhaustive set of perspectives or, necessarily, a collection which, in the long run, will prove the most useful as the grounding for the development of co-operative systems. For us they are a starting point for the attempt to provide a suitable grounding for a social perspective to underpin CSCW design and development.

In addition to meeting the concerns of developers, and to repeat, the perspectives are also intended to encapsulate the emergent features of work highlighted by the studies reviewed earlier. Each perspective should allow different features of work to be represented to developers of CSCW applications. The different perspectives of work and the emergent features of work they reflect are summarised and presented in the following table.

Distributed Co-ordination	
Emerging Feature of Work	**Study**
Real Time Decision Making	Air Traffic Control
Co-ordination through paperwork	Small Office in Training Centre
Coordinating work in real time	Small Office in Training Centre
Representation and 'at a glance' availability of the work	Air Traffic Control
Organising complexity	Manufacturing Engineering
Teamwork and skill	Specialised Services Banking
Plans and Procedures	
Emerging Feature of Work	**Study**
Local Knowledge	Small Office in Training Centre, Specialised Services Banking
The locality of organisation	Building Society
Working Division of Labour and Team work	Air Traffic Control
Aggregate predictability and local contingency	Building Society
The local management of conflicting demands	Building Society
Demeanour work and 'getting round the technology'	Building Society
'Making do with the technology-to-hand'	Technical Documentation
Managing organisational pressures	Technical Documentation
The embeddedness of work activities in organisational structures	Software Production
The interweaving the formal and informal	Technology Centre
Representing the work	Technology Centre
Awareness of Work	
Emerging Feature of Work	**Study**
The egological organisation of work	Specialised Services Banking
Informal team working and the working division of labour	Specialised Services Banking
Computer work and paper work	Specialised Services Banking
The Organisation and Organisational Change	Specialised Services Banking
Ecology of co-ordination and awareness	Small Office in Training Centre
Providing for mutual awareness	Manufacturing Engineering

The following section of the deliverable investigates in more detail the consequences of focusing on the three perspectives identified earlier. To illustrate their general applicability we also consider illustrative examples drawn from the set of field studies. These examples are complemented by considering the different consequences for design which emerge from these perspectives.

Section 4
Perspectives on the Social Organisation of Work

Section 4
Perspectives on the Social Organisation of Work

Introductory considerations

The studies presented in the preceding Section represent a variety of work settings which we have used to identify and illustrate some salient aspects of the social organisation of work. As we indicated earlier, many of these aspects could well have been illustrated using a number of the studies, and we have made no formally systematic attempt to develop a more generic framework for the characterisation of the properties of the social organisation of work. There are a number of reasons why, in our view, the development of such a framework is less than straightforward which are summarised below:

- the heterogeneity of domains;
- the identification of co-operation;
- interdisciplinary communication.

Heterogeneity of domains

As we discussed in Deliverable D2.1, one of the major problems of requirements elicitation, especially as far as the development of CSCW systems is concerned, is the variety of work domains; a variety which is too often obscured by the use of the large-grained characterisation of work and work domains which tend to predominate within the social sciences and, from a different direction, by the abstracted decomposition of work in terms of tasks which is characteristic of many software engineering requirements methods. Accordingly, and if this diagnosis is correct, studies of the social organisation of work will need to proceed in a manner which recognises this heterogeneity of domains and develops analytic tools which are capable of exhibiting the relevant scope of this variety.

In significant respects, it is this objective which underpins the use of ethnographic fieldwork and case studies within CSCW as a promising means of designing systems which are more responsive to the needs and the skills of users. As we discussed in Section 2, this involves giving detailed attention to the subtleties and the situatedness of the work activities as day-to-day phenomena. Such a focus, however, means that developing an analytic framework of some generality needs to be developed ‘from the ground up’ as it were, and capable of retaining a sensitivity to the details and the variety of work domains.

The identification of co-operation

Bound up with the above considerations are problems involved in identifying co-operative activities and, relatedly, the interdependencies of activities. Many of the early writings in CSCW attempted to identify 'co-operation' as a distinct, discrete type of activity whereas, and as many studies of the social organisation of work show, matters are much more subtle, and more complicated, than this assumes. Moreover, from a sociological point of view, with its focus on interaction, the collaborative and co-operative nature of work is assumed from the beginning as its topic of inquiry - which is not to say that sociologists necessarily agree on how this inquiry is to be pursued. The association, for example, of co-operation with synchronously, co-located persons working in a team, tends to ignore the pervasiveness of a variety of interdependencies within work settings which are immensely relevant to CSCW design. In other words, the relevant properties of the social organisation of work do not appear as 'readily packaged' within work domains but need to be brought out by an analysis of the ethnographic materials.

Interdisciplinary communication

One of the main objectives of a framework for the analysis of the social organisation of work for CSCW, is to develop a vocabulary to enable design to take cognisance of the social organisation of work. Such a vocabulary must serve the needs of design while, at the same time, being adequately grounded sociologically. Quite what this might mean is not yet altogether clear, though the research reported in Section 5 on using DNP is one way in which we are exploring this issue. It is also one reason why the notion of viewpoints is attractive. As we reviewed earlier, the 'gap', in this case, between ethnography and software engineering design is not necessarily one which is unbridgable as a matter of incommensurate mentalities; rather, we see it is a difference in the kind of problems that each addresses, but a difference which needs to be complementary rather than antagonistic. In other words, and for a simple example, one of the ways in which a system designer might want to represent a set of work activities is as 'flow of work' using standard graphical tools. For the ethnographer this mode of representation might appear to conceal the very richness of the work that the fieldwork has uncovered. However, if we regard this as a viewpoint on the fieldwork materials, that is, as a vantage point intended to highlight features which are of interest to the system designer and, as important, used in conjunction with the rich fieldwork materials, then we have, at least, a promising way of supporting communication between fieldworkers and system developers.

It is for the above reasons, and more could be adduced, that we take the view that the development of a framework for the analysis of the social organisation of work is less than straightforward. These are not, of course, arguments against the idea of a framework.

What follows in the remainder of this Section is a discussion which is intended as a the basis for the development of a framework the formulation of which will

constitute much of the research in the coming year.[58] What we have done is try to stand back from the details of the studies and identify what we see as promising perspectives from the point of view of CSCW design from which to view the social organisation of work. These perspectives are, we stress, vantage points from which the social organisation of work may be examined rather than discrete 'things'. In other words, they are not self-encapsulated themes but represent varying starting points from which to examine features of the social organisation of work.

The selection of the perspectives has been influenced by two considerations: first, the features emerging from the case studies and, second, a sensitivity to CSCW design issues. Some of the latter are discussed in relation to the emergent features of the social organisation of work. Given the emergent status of the research, the design consequences which are highlighted should be read as illustrative examples rather than fixed principles of design practice.

First of all, however, we preface the discussion of the perspectives on the social organisation of work by some remarks setting out the basis of the sociological approach we are using and which will be developed further as the framework evolves during the coming year. This will help to point up some of the more sociologically-oriented principles which inform the ethnographic fieldwork approach which is the focus of this Strand.

Locality of social organisation

Earlier we identified the problem of scale as a collection of issues, theoretical and methodological. We also suggested that one of the problems was the conceptualisation of organisation since this was also strongly related to some of the methodological issues surrounding the use of ethnography in work analysis. More specifically, we cast doubt on the idea of 'macro' conceptions of organisation as some 'transcendental', more inclusive real entity independent of the actions, the activities of its members which, in some sense, governs, dictates, causes these actions and activities. This conception is bound up with the assumption that social phenomena can be categorised as 'large scale' and 'small scale', as macro and micro phenomena, and that for effective sociological understanding the 'small scale' must be situated relative to the 'large scale'.

In much of the literature this is encompassed within the 'structure' versus 'agency' dispute; that is, the question of whether the actions of individuals in situations are beyond their own control or not. The idea is that the individual's actions must be explained in terms of characteristics of the structure. In other words, their actions are governed, dictated, caused by the positions they occupy in some larger structure(s). They act in the way that they do because the structure

---

58 Of course, such a framework alone will not bridge the 'gap' we referred to earlier, since what is pictured by the metaphor are a host of issues and problems which encompass training and education, the organisation of design work, the economics of system design, and much more.

requires it of them. Any sense that an individual might have that they are doing something because they want to do so, is effectively an illusion. They are doing what they are doing because forces of which they are unaware are compelling them to carry out these actions.

The issue is also conceived as one which has to do with the reality of the large structures; whether the structure has a reality of its own so that it can be a causal factor in the determination of someone's actions. This is the position associated with Durkheim (1964) and, more recently, with the 'realist' view (see, for example, Bhaskar, 1975). The alternative point of view, whose most distinguished representative is Weber (1949), claims that social structures are merely the composites of the actions of individuals and, thus, are the product of agency.

However, if 'social structure' is conceived as an 'underlying pattern' which organises and orders the everyday affairs of society, then the activities of persons in the world of daily life must be the 'particulars' in which that pattern is manifest. The pattern is seeable in the particulars and, reciprocally, the particulars themselves are not recognisable unless they are identified as parts of the pattern. The relationship between the pattern and the particulars is not a causal one but, for want of a better word, hermeneutic. The relationship between the pattern and the particular is mutually elaborative, such that the putative pattern and its constituent particulars are inextricably intertwined. The pattern can only be found through the particulars and the particulars only adequately identifiable as part of the pattern. The issue is neither one of causation nor one of eliminating either 'social structure' or the 'individual' and the 'structure' versus 'agency' debate is set aside. Instead, the problem becomes one of seeing how the activities of persons embody, implement or carry out, the affairs of the social order.

Thus, our own point of view emphasises the way in which persons 'make sense' of the social order from within that very order. This is one of the main rationales for ethnography as a strategy for work analysis and the notion of social structure which it employs is the conception of a pattern which, in some way, is findable within the affairs of daily life. Thus, the ethnographer's task is to find an intelligible social order from within that very social order, and the problem is concerned with the adequate characterisation of the activities of persons in social settings; that is, providing descriptions which are perspicuous with respect to the socially organised character of those actions. This is not a matter of methodology in the usual sense, namely, using methods which, according to some standards of scientific adequacy, are judged as valid means of producing generalised observations from particular instances, since the very identification of particular instances is inextricable from that of their constitutive pattern. The recognition of instances is as instances-of-a-general pattern.

Organisation from within

What should come across strongly in the field study reports are the ways in which parties to the work orient to the 'ordinariness' and 'routine' character of the work and its setting. The daily world of work is experienced as an everyday world

in which the 'usual' happens, in which things happen 'over and over again', in which occurrences and instances happen which are 'things seen everyday'. To put it more generally, these occurrences of working life are typically encountered and experienced as events-in-a-pattern, as activities-which-are-part-of-an-order, as 'regularly things which happen around here/within this organisation/this town', etc.. They are recognisable to an individual as routine, as a usual, as things which happen to this individual in just the same way that they happen to others. The person participates in social situations, in significant respects, as 'anyone', as someone who is entitled to behave in the ways that he/she does so, and to expect certain kinds of treatment merely by virtue of the fact that this is the way 'anyone' is entitled to be treated.[59]

This is bound up with the capacity of an investigator to recognise and come to understand the events and activities within a setting by making use of his/her own common-sense understandings. That is, the investigator can say what is going on in a social transaction by virtue of being another 'anyone' entitled to say because anyone can see that, for example, "the car stopped at the traffic lights" or "the couple were having a fight" without very much by way of hesitation, prolonged inquiry or close scrutiny.[60] Moreover, the investigator's ability to invoke such understandings is itself a function of the way in which social settings are organised. That is, settings are organised to provide for their recognition as the events-in-the-order that they are. Thus, the connection between the organisation and the observability of actions is an integral not a contingent one; that is, activities are not only organised to be what they are, but are organised to be recognised as, identified as, what they are. The visibility of actions is not some added ingredient in addition to the organisation of the actions themselves. The very organisation of the actions involves providing for their visibility.

This, of course, is a statement of what it means to characterise activities as social, that is, as actions directed toward others, and their performance as part of, in the course of, collective ventures means that they are directed toward occasioning responsive actions from others. This means that parties must be able to identify and recognise what the actions are. Accordingly, the way in which actions are done is in ways which ensure their recognisability as the actions-that-they-are.[61] The very organisation of actions involves provision for their 'visibility' such that actions are done to be recognised for what they are, organised

---

59 The notion of 'anyone' here comes with a subscript so we are not naively supposing that *just* anyone can adequately identify particular events. We use 'anyone' in the sense of 'wide awake adult' or 'anyone who spends a lot of time around here'.

60 We are not, of course, suggesting that such descriptions of 'what is going on' will inevitably turn out to be correct descriptions. The point we are drawing attention to is the inescapable reliance on commonsense knowledge for the description of social actions. Such descriptions, just as they are for ordinary members of society, are revisable in light of further information about the event, the persons and the circumstances.

61 This does not, of course, rule out the possibility of deception of some kind. Such occasions, in fact, support the argument here in that deceptions are carried out so as to be identified by those against whom they are directed as actions-different-from-those-they-truly-are; that is, integrally organised in order to secure a kind of recognition from others.

to be what they are and to be recognised as such are one and the same thing. The organisation of actions is self-explicating. Thus, the problem of finding 'social structure' within settings is resolved organisationally. It is the organisation of the social setting which reveals itself in the conduct of its affairs.

Thus, from the perspective just outlined, the features of an organisation are not transcendent in the way that organisational theory often presumes, but are very much part of the locality of the work's settings and its self-explicating character. Air traffic controllers, for example, are aware that they are working a system and that their work is organised as working-within-and-working-a-system. The staff of the Training Centre Office similarly are aware that they are working within a complex of activities which are connected to the hotel, the university, the organisations that use the facilities, and more. The workers within the various offices and centres of the Bank are aware of their work as part of some larger institution and that their work activities are the work of that organisation. The sense in which controllers, to develop this example a little more, are aware of the fact that they are working within a system that transcends their here-and-now, is also reflected in the organisation of the control room and in their relationships to each other. They are aware of the fact that the control room is but a location within a system of air traffic control which extends beyond LATCC; aware of the fact that they work within a bureaucratic structure of work organisation, subject to a managerial supervision and situated within a complex division of labour. They are aware of the fact that their working is dependent upon, for example, radar and computing support. Aware of the fact that, as a condition of their own work, they are dependent upon a number of radar stations distributed across the country, managed and supported by other specialist groups within the traffic management system, just as their computational support is provided and managed by specialists in computing and in software. These, and more, are known facts of their lives and their organisation-and-formulation of their own work activities involve acting and talking within and as part of that system. To see what they are doing requires being attuned to the fact that they are participants in that system.

Thus in referring to the locality of organisation we are not denying the fact that, for example, the workers in the Bank's branches or the controllers in the control room at LATCC, or the software engineers in the avionics team, or the staff in the technical drawing department, and so on, are part of larger organisations. The very characterisation of them as 'bank workers', as 'air traffic controllers', as 'software engineers', calls into them into play as people who are organisationally positioned and people who will have at least some organisationally prescribed responsibilities and tasks. What it does mean is that one can learn a great deal about an organisation from within its local elements.

For example, and to return to air traffic control, a great deal can be learned about the organisation of the whole system from within the control room of LATCC. Indeed, one must learn from *within* the control room if one is going to understand, in an organisationally relevant sense, what is going on in that site. The fact that controllers are situated within at least these various ways within a

system is not exigent to their activities because these are pervaded by their sense as operations in and of the system. The very sense of the activities is of them as 'air traffic control operations'. The controller's work is done in relation to other participants in the system, in relation to them *as* participants in the system. The controller is pervasively attentive to the affairs of the system, familiar with its intricacies and exigencies. The controller is not controlling a sector but controlling a sector which adjoins other sectors and is controlling a sector in collaboration with other controllers. Controllers do not just dispatch aircraft through the sky but seek to control aircraft so that when they arrive at the boundaries of a sector they will do so in such a way as to ease the work of other controllers.

Those who are doing 'the work of controlling' are, accordingly, oriented to themselves as situated within the order of air traffic controlling, and conduct their work as the practical administration of that order. Their activities are ones which carry out the business of the air traffic control system and are ones carried out in response to the unfolding state of the system.

Familiarity of the organisation from within

What we are drawing attention to here is the manifold ways in which, conventionally, what are treated within much of sociology and organisational theory as 'transcendent' properties of organisations, have their matter of course relevances for the local settings of organisations.[62] There is a matter of course awareness on the part of participants within organisations that their situation within the setting is within and part of a larger system and that their activities are, in ways of which they are varyingly aware, integrated with and dependent on the doings of others. They are aware that the 'system' has a history which, from their point of view, has evolved and is continuing to change. They can relate how things have changed for them, both within the work site and for other aspects of their lives. They have an awareness, albeit sometimes diffuse and certainly variable, of the technicalities of the system as determinants of their own working lives. Aware, that is, that the shaping and reshaping of their current situation is in the hands of persons other than themselves, such as the policies of their own organisations, government legislation and economic policies, their own organisational superiors, other staff, and more. All of these are just some of the ways in which the locality of social organisation manifests itself in day-to-day working lives.

The egological approach and the 'working division of labour'

The approach we have described above begins with individual actors and how they orient to the organisation of activities and to the people around them. That is,

---

62 We would be prepared to state this more strongly in that there are no transcendent properties of organisations as these are typically conceived by sociology and organisation theory. This is not, however, an argument we want to elucidate at this stage.

what is of interest is looking at social organisation from the 'actors point of view'; hence the term 'egological' (Anderson et al, 1989). This also means that in examining a division of labour within an organisation, we need to start from how that division of labour is encountered and experienced by those working within it. How, that is, the tasks and activities which constitute the work gain their 'real world, real time' character from within the social organisation of the division of labour itself? This is to focus on describing what the organisation of activities consists in as an environment within which tasks and activities are located, co-ordinated and implemented.

As an encountered phenomenon, the division of labour becomes not so much a smooth, highly integrated and planned framework but a fragmentation of activities, tasks and performances. That is, immersion in a division of labour on a day-to-day basis is experienced as stream of differentiated and discrete tasks-to-be-done. The accomplishment of work activities appears as an impersonalised stream of tasks in hand, tasks completed, and tasks to be done. Individual performance is bounded by 'structures of relevance' which formulate concerns relevant to the tasks to be accomplished.[63] Thus, and for example, controllers are not only attuned to what is happening in their sector of air space, they are also sensitive to what is happening on adjacent sectors, but less 'geared into' what is happening on sectors at some remove from their own, while other matters are simply taken for granted. Of course, for the accomplishment of any individuals task within the division of labour, a whole set of other practical activities will have to be done by others. The staff in the Training Centre, for example, do not have to worry about the postal service, the electricity supply industry, or the innumerable other activities and facilities which are necessary to their work. What they do have to concern themselves about are the specific tasks which are the responsibility of the office and, even in this case, there are tasks to be done which are then passed on to others.

'Gearing into' the division of labour[64]

Because a division of labour specifies which tasks have to be embedded within others, the co-ordination of tasks is possible. For any actor within a division of labour, competent task performance is the achievement of the routine embedding of that task into others. Some of this will involve invoking institutional resources, that is, anticipations of how institutional resources work. Memos can be put in people's pigeon holes in the anticipation that they will be read and acted upon, a Reservation Form can be filled in with the expectation that this will be acted on by the Hotel staff, and so on. Such anticipations of the routine operations of the division of labour provides organisational specific ways in which those who work within it can call it up, gear into it, and make it work for them. 'Gearing in' is a

63 The centrality of structures of relevance for the organisation of social action is discussed at length by Schutz (1970) and Gurswitsch (1964).

64 'Gearing into' is examined by Gurswitsch (1979)

means of ensuring the smooth performance of the flow of activities by ensuring the fit of one's task performance into that of others and by being able to carry out 'running repairs' when these are needed. Co-ordination becomes routinised, taken care of as a day's work and made 'invisible' except when 'troubles' occur and it becomes itself an object of attention for organisational members.

Thus, the central feature of the sociological approach used to inform the analysis of the ethnographic fieldwork materials is an egological approach to understanding the social organisation of work activities and, through this, examining and describing their 'real world, real time' character. Further development of the approach, along with an associated framework, will be, as indicated earlier, the focus of the coming year's work.

In the following discussion, the approach we have just been outlining has been directed at generating perspectives from which to view the social organisation of work for CSCW design. We have selected two as promising beginnings:

- Distributed Co-ordination
- Awareness of Work

Both of these vantage points enable us to highlight some of the rich and varied features of the social organisation of work using some of the fieldwork materials as illustration. As we stress, these are best seen as vantage points from which to gain a more perspicuous gaze at the complexities of the sociality of work. Another consideration is that they also have close parallels with a number of interests and concerns which are prominent in CSCW.

Distributed Co-ordination

By this term we refer to the manifold ways in which the co-ordination of people and tasks are accomplished as a routine feature of 'real world, real time' work.

One of the distinctive even if quite ordinary feature of all work settings is that tasks, activities and, of course, persons are embedded within an organised ensemble of some kind. By this we mean nothing more nor less than the fact that there is an irreducible sense in which we see activities and persons as not only 'belonging together' but interconnected in some fashion. Such interconnections can be relatively 'loose' as, for example, in the case of an audience at a concert, or more 'tightly' interconnected as in most work settings through the interdependencies of a division of labour. In the latter case what we have is activities and persons distributed in such a way that they have to be treated not as isolated activities and persons but as part of some *organisation* of activities and persons.

These remarks are obvious, but important reminders nonetheless, about one of the main features of work, namely, that it is organised for the interdependency of persons and activities which, in its turn, becomes part of the work itself since interdependency implies the need for the co-ordination of the separate activities and persons. Job descriptions and responsibilities, procedures and regulations,

hierarchical structures, plans, timetables, skills, and more, are all examples of the way in which the co-ordination of distributed tasks, activities and persons is accomplished. However, the point we want to emphasise here is that co-ordination does not consist in any one feature of the work but is deeply, and inseparably, implicated in the kind of things just mentioned by way of example. In other words, co-ordination is not simply, say, the task of a manager or a supervisor - although this, indeed, may well be a major part of their duties - but it is also integrated in various ways in the details of work activities. For example, much of the co-ordination of tasks in ATC is achieved by 'following procedures' where this is understood in the details of the work of organising the strips, communication with other suites and with pilots, marking the strips, and so on: so much so that it would not be unreasonable to characterise ATC work as primarily devoted to co-ordination as long as this is understood to refer also to the myriad details of the activities of the work itself. It is in order to achieve the co-ordination of aircraft that co-ordination within a division of labour around the suite and in the control room is necessary. One is done through the other.

Again, and for further example, much of the work in the Training Centre Office is explicitly about effecting the co-ordination of whole collections of activities throughout both the Training Centre itself and in the hotel. In this case, and again this is not untypical, much of this co-ordination work consists of distributing relevant information to relevant parties and keeping this flow of information going as a routine state of affairs. Similarly, Configuration Control in the avionic software development group performed much the same kind of function though in a different form, that is, keeping track of the work in order to co-ordinate other work. Many aspects of work in the bank, whether in the traditional branch or in the specialised centres, is explicitly concerned with co-ordinating interdependencies of various kinds in order to 'get the work done'. For one example: 'unpaids' involves co-ordinating the work of Lending, the local branches, other Banks, customers by working through the relevant paperwork which requires input and information from these varied sources.

The important point we are trying to stress using these examples, and one which has already been alluded to, is that co-ordination is not always satisfactorily treated as a separate, distinct activity, though it can be. In many settings, for example, there are meetings explicitly scheduled to achieve co-ordination among relevant parties. The Training Centre Office, again for example, has weekly meetings to review the work schedules for the coming week. In the Technical Documentation Department 'Scrutiny Meetings' are a salient feature of the work process. In other words, such meetings are not mere adjuncts to the work but integral to it.

Co-ordination through artefacts

Co-ordination can also be a feature of specifically designed or evolved artefacts. As Yates (1989) describes, memos, files, duplicators, standard forms, etc., evolved during the nineteenth century to solve problems of what we are calling

here, 'distributed co-ordination', as organisations became larger and the problems of management and control correspondingly increased.[65]

An example of artefacts as co-ordinating 'mechanisms of interaction' is the ABOM developed by the S4000 project team at Foss Electric. The form lays out the components required for a unit along with other identifying details as well as who in the project team is responsible for which tasks when producing components at different phases of the development cycle. As the work of the various parties progresses the ABOM accumulates relevant information about the distribution of tasks so that those who 'need to know' are informed about the progress of the work and where further co-ordination may be required. In other words, as an evolving representation of relevant aspects of the work's activities, the ABOM is able to function within the flow of those work activities and, through this, achieve the co-ordination of those self-same activities. Other artefacts achieve co-ordination through the modal transformation of information from one artefact to another. The 'TecSec' system used in the Securities Centre, for example, leads the user from one activity to another, from one modal transformation of information to another form as does much of the paperwork in the Training Centre Office.

Such artefacts facilitate the co-ordination of tasks by embedding descriptions of the task, along with other relevant information, within the format of a document as 'instructions', as 'items required', as 'number of residents', as 'specifications', as 'for information items', and so on, which makes available to those who know how to use the artefact, what the implications are for the actions of others. These artefacts serve as sets of instructions for a set of institutionally identified persons - in the case of the ABOM - the S4000 team - to perform particular tasks and, in addition, serve as a check on whether or not these tasks have been performed.

Of course, the use of such artefacts depends upon users knowing how to use them as a tool for the work that is done and, through this, investing them with sufficient legitimacy in order for them to operate as co-ordinating mechanisms. In other words, co-ordination artefacts presuppose a moral order of some kind so that 'completing the form correctly', 'entering the relevant details', 'passing them on through the channels', and so on, become sanctionable matters.

Related Consequences for Design

Distributed co-ordination is a central interest of CSCW. The communication facilities provided by network computer systems, for example, are appealing because of the potentialities they offer in this respect even when they build upon existing technologies which have been developed for just this purpose, such as the telephone, the FAX, 'snail mail', and so on. However, it worth reflecting on some

---

65 See, for example, Weber (1947) and the development of the office, filing systems, records, and so on. Possibly among the first examples of artifacts such as these are those used in the management and deployment of large armies and navies.

of the design consequences that emerge from the fieldwork corpus from the perspective of distributed co-ordination.

Identifying co-ordination features

Given the central role of co-ordination within co-operative work the identification of the features work which promote co-ordination is important for the purposes of CSCW. Equally the provision of electronic equivalencies to existing manual techniques is a central issue in the design of co-operative systems. Notable co-ordination features include:

- The ability to monitor the activities of others
- Access to shared and readily available information.

Both these features of co-ordination represent significant research activities across the COMIC project and the investigation of these issues is one of the ways in which we wish to uncover general design principles which have been informed by the nature of work. In particular, access to shared information underpins the development of the COMIC Shared Object Service (Deliverable D4.2), as well as the extensive investigation of event models, reflects the need to support the monitoring of activity as a means of supporting the co-ordination of activities.

Co-ordination features may also take the form of dedicated artefacts, as we have seen, within the work setting which can be translated into corresponding electronic equivalents. This approach to the supporting co-ordination underpins the investigation of mechanisms of co-ordination within COMIC (Deliverables D3.2, D3.3). A significant design challenge in the development of these mechanisms is the outlining of appropriate notations to represent their properties.

Minimising the overhead of co-ordination

As we have seen, co-ordination is a part of the generally fluid nature of co-operative work. Much of this co-ordination is implicit in the course of the work as people monitor the activities of others, through the public character of many of the activities and the artefacts as well as through the more explicit mechanisms such as meetings and conferences. Among the design consequences that flow from the recognition that co-ordination is intimately bound up with the work itself are the following:

- Co-ordination mechanisms need to be sensitive to the work taking place. It is important to avoid any impedance between the work and the co-ordination of the work.
- The cost of co-ordination with others needs to be kept low. If the cost of co-ordination is ‘too high’ it will either become an end in itself at the expense of the work, or become a such an ‘inconvenience’ that it has low priority

Appropriate lightweight mechanisms for co-ordination are a central feature of the work of the COMIC project. These mechanisms are principally being developed as part of the outlining of the shared interface service (Deliverable D4.2). For example, within the SOL prototype, tool kit users and designers can

specify the propagation of interaction to promote co-ordination. The co-ordination mechanism is a side effect of the work taking place rather than a separate endeavour. A similar approach to co-ordination is reflected in the work on adapters within the Shared Object Service (Deliverable D4.2).

Plans and procedures

As we indicated earlier, co-ordination is a prominent feature of work tasks themselves insofar as they are tasks organised into a division of labour. The preparation of flight strips by the assistant controller, for example, is directly related to serving the needs of the radar controller and is a task which has its point within a division of labour around the controlling suite. It is a task which is completed and then passed onto to someone else. In this case, the task is one which directly supports the work of the radar controller in that the latter's tasks are dependent on the successful completion of those of the assistant controller. Such relationships of interdependence can be much more complex than this example, of course. The development of large-scale software systems involves the interdependence of work tasks which are in themselves very complex and involve many more people.

One of the more prominent ways in which distributed co-ordination is achieved is through plans and procedures. Here we are referring to the wide ranging phenomena which includes manuals, Project Delivery Plans, project schedules, the Manual of Air Traffic Services, job descriptions, instructions for how to do a test on a fibre optic cable, office procedures, and so on. Although these examples point to a variegated collection of phenomena - ranging from the overall strategic plan of a large organisation to the filing procedures within a small office - what we are interested in is how they are used as co-ordinating mechanisms within socially organised 'real world, real time' work activities.[66] As we have already said, their explicit point is to co-ordinate the work of numbers of people in order that separate work activities and tasks come to have a coherence and, typically, through this meet other goals such as efficiency, time constraints, and so on. Plans and procedures are quintessentially about co-ordination. However, within this general characterisation there are some important issues as well as some misunderstandings which need consideration[67]

Critique of the Planning Model

Much of the critique of the planning model in CSCW, namely, that plans are ineffective arises from misinterpretations of Suchman's (1987) seminal work.[68] Most of these interpretations have overlooked the fact that the critique is directed

66 It is important to point out here that we are making no assumptions that all the items in the 'variegated collection of phenomena' are used in the same way.

67 Some of the following issues have been extensively discussed in Deliverable 2.1. Here we simply reiterate the main points of the debate. See also Pycock et al (1994).

68 See, for example, Vera and Simon (1993).

at the 'planning model of action' as it appears in cognitive science. It has, however, been interpreted as saying that people do not follow plans because, in actuality, they are *post-hoc* rationalisations of courses of actions. Suchman's critique is directed, however, at the notion that there are mental plans which operate as causal determinants of subsequent courses of action. In other words, the criticism is directed at the idea that people have 'plans in their heads' in the ways in which some cognitive scientists suppose, not upon the idea that people never follow plans in the sense in which we ordinarily observe them doing.[69] The essence of Suchman's critique of the planning model derived from cognitive science is that it requires that activities be thoroughly determined and causally directed in advance. Such a conception cannot be adequate to the phenomena it purports to describe because it does not allow room for the fact that the plan is an abstract construction which will, at the very least, require articulation with, and application to, the specifics of the circumstances in which it is to be followed. 'Following the plan' will always involve more than can be specified within it and, therefore, the relationship between the plan and the action it directs cannot be causal (Pycock et al, 1994).[70]

Thus, the critique of the planning model as it appears in cognitive science cannot be an objection to the utility of plans in ordinary life nor carry the implication that plans can never work.

Plans in Organisational Life

The construction and use of plans in 'real world, real time' activities do not typically involve the supposition that literally everything must be spelled so that nothing whatsoever is left out. This is not to say, of course, that plans and procedures are beyond criticism or that, within organisational life, they may be judged to work well or badly. The point is that in practical life plans are characteristically 'recipient designed', that is, spelled out to an extent to which those who are to follow them are, for example, familiar with the circumstances in which they are to follow them, sufficiently trained in the tasks involved, and a host of other possible considerations. There is always a trade-off between the detail incorporated in the plan and other ways of organising. Nor does the making of plans indicate any expectation that the course of actions which they specify will, of necessity, follow through. Indeed, the point of plans is often to direct courses of action to maximise the chances that these courses of action will ensue despite the contingencies which can arise. But this is not the same as seeing the plan as somehow necessarily resulting in the courses of action concerned. Plans often include 'fail safe' devices to cope with situations where things are 'not

---

69 Suchman's own formulations cannot be exempted from some responsibility for these misinterpretations. For example, she suggests, following Mead, that there is a "part of us...that actually acts" and which engages in an essentially situated and ad hoc improvisation, while another represents those actions by making "an object of them". Accordingly, "our descriptions of our actions come always before or after the fact, in the form of imagined projections and recollected reconstructions".

70 This is much of the burden of Ethnomethodology's treatment of rules and procedures. See, for example, Garfinkel (1967) and Sharrock and Anderson (1986).

going to plan' by specifying arrangements for adaptation of the plan to exceptions, unforeseen circumstances, even extensive revision, as well as mechanisms to oversee the implementation of the plan and enforce its requirements.

The planners of projects are not naïve to the fact that courses of activity will not necessarily or in fine detail follow as they are specified in the plan. In formulating and using plans they depend upon their typified conceptions of how courses of action are carried out, or on the presumed grasp that others on the project will have of how things are done. They are also aware that the efficacy of methods is variable, from the 'sure fire' techniques to those which are less than reliable, and to those which require some skill and experience in their application, and so on. In working out the sequence of operations that a project is to follow, it is often appreciated that the project may have to deviate from the plan due to unplanned for contingencies, etc.. It is routinely accepted that there are detectable elements of risk, and every prospect of problems arising which had not been foreseen, including the very operation which might be supposed to give the least trouble turns out to be the one that gives the most.

Planning, to reiterate, involves specifying what will have to be done, by whom and how long it will take to do whatever needs to be done, and using what resources are available. Many of these factors, such as how long a project is to take, are matters of guesswork and, sometimes, such guesswork can become enshrined as a condition which 'needs to be faced' or 'got around somehow'. Plans, we might say, unfold in real time which means that what the plan specified, what the plan agreed, what interdependencies there are, only become clear as the courses of action specified in the plan unfold. The Training Centre, for example, despite regular weekly planning meetings routinely had to deal with late changes in customer requirements, with sudden cancellations, with delays in customers actually specifying what there requirements might be, and so on; contingencies which were capable of reverberating through the whole operation and which had to be dealt with, often 'on the fly'. Similarly, in air traffic control, following procedure is always subject to contingencies of various kinds, sudden emergencies being the fortunately rare but architypical example. Less dramatic are the 'awkward flights', military aircraft crossing civilian airways, flight delays, and so on, which must all be dealt with as part of applying the plan. Similarly, and as noted before, the procedures of the building society office constitute a series of plans as to the way the various tasks that the office deals with are to be done, procedures formulated on a conception of what typical interactions with a customer are but, of course, any one customer may well arrive with 'untypical' requests, or requests which reveal other problems, but which, nonetheless, have to be dealt with.

Typically within organisations a given project or plan is one among a number. In software engineering, for example, individuals may be involved in more than one project and each project, over its course, may well involve a varying number of individuals, the kind of work to be done, and the mix of skills required. In other

settings, there may be a sense of an overall plan which is intended to enable the co-ordination of different activities or events. An example of this is the Training Centre once again, in which an overall set of procedures is designed to manage the flow of customers and clients, each with their specific requirements and needs which often emerge in the course of their stay. Thus, one of the purposes of plans is to represent the nature of the work in order to determine the kind of team required, the mix of skills and to ensure their availability at the periods required. In engineering, for example, projects have to make specific demands for time from particular engineers and, at the same time, make sure that overall they are optimally employed.

Making plans work

Of course, and again though this is perhaps an obvious point it is an important one, plans and procedures do not apply themselves but have to be used by persons in the course of 'real time, real world' settings. We have already more than alluded to this throughout the discussion, but there are aspects to this relationship of plans to situated activities which, once again, are worth drawing out a little more.

When earlier we described the ways in which artefacts can be seen as devices to co-ordinate complex collections of activities, we made the point, albeit in passing, that their use depends upon practically 'knowing how' and practically 'knowing that' in what ways they represent actual activities. By themselves, even if such an idea makes any sense, they could do nothing. They depend irreducibly upon understandings about what they stand for and how they can be interwoven into the work activities in order that they can operate as co-ordinating instruments. They are co-ordinating instruments because they are integrated, often in subtle and 'invisible' ways, into the work activities they help co-ordinate. They becomes constitutive of the work of co-ordinating activities.[71]

Plans, too, depend upon the practical understandings about what the plan specifies in *these* circumstances, using *these* resources, *these* people, and so on. In other words, although plans may be presented as abstractions, as manuals, as statements of procedures, and so forth, the 'just what' it takes to realise them is, as we have been saying, a practical matter for those using, or subject to, the plans. As we have mentioned before, implementing the procedures in the Building Society or in the Training Centre, for example, is very often routinely a matter of 'making the plan work' through all the various and inevitable contingencies that can arise. It is such activities which maintain the plan and the procedures by dealing with 'those things which arise', 'the things not planned for', the 'things which suddenly come up'.

In engineering such features can become endemic. Much of engineering work is of a problem-solving kind and, at the outset, it is not always well-understood how a project is going to carry through its objectives. The S4000, for example,

71 This is so whether or not they work well or badly.

was the first 'system instrument' for the company and which introduced new measurement parameters and an enhanced measurement speed; all of which increased the complexity of the design task. From the outset it was realised that the basic design would never be sufficient but would require a number of iterations through mock-ups and prototypes which would inevitably give rise to problems which would need to be solved; problems and solutions, moreover, which could well reverberate through the project in complex ways. In other words, it is the project's task to find a way of doing the work assigned. How the work is to be done is worked out over the course of the project itself.[72]

What this means is that plans often serve more as a specification of requirements, an enumeration of things that must be achieved, how they need to be scheduled, what the critical dependencies might be, and so on, rather than a detailing of how things are, in fact, to be achieved. There are many problems which cannot be dealt with at the level of planning. For example, there can be a tensions between the cost and the quality of the product which are present in the plan, tensions which are handed on to the project itself to resolve during the course of its implementation. Although there are many processes within organisations which are planned as step-by-step procedures - many of the tasks within the bank are close to this - in other contexts such stepwise solution of problems may involve passing problems 'down the line'.

An example of this is from the avionics project. One of the principal conditions laid down on the project was the strict allocation of 'person hours' in order to keep costs down. A number of hours were allocated to a task and these could not be exceeded if this added to the project's total hours. Time could be 'borrowed' from one part of the project but this would have to be repaid at a later date. Of course, the hours allocated for the completion of tasks or, for that matter, for the whole project, could not be precisely calculated. Such figures were 'guesstimates' closely governed by the need to keep costs down, and the problem of finding out how to do the work within the hours available was handed down to the workforce of the project itself. In this case, they resolved the problem by working overtime, subsiding it through work in their own time, working at weekends, putting in time which would not count against the official allocation of hours, and so on; motivated in significant part because their jobs were at risk due to a downturn in defence requirements and in aircraft manufacture more generally.

Another example from the same study illustrates not only the way in which 'what a plan means' has to be worked out in the course of the plan in order to 'make the plan work', but also the ways in which 'deviations' from the plan sustain, so to speak, its spirit. As many sociological studies of work and

---

72 The extent of this is variable. Some projects involve carrying out tasks which have been done before while others are exploratory. Also, it may be that those involved in a project have never done anything like it before. In other words, relevant here is the extent to which an organisational setting is favourable, the extent to which, for example, the project is more or less exempt from budgetary pressures, work overload, skill shortages, and so on. Even where a project involves work which has been done before it may not be within the organisation's power to staff the project with those who are familiar with such work (Pycock et al, 1994).

organisations have shown, real time real world work often involves 'cutting corners', 'bending the rules', and so on. Though sometimes this is for subversive reasons, in many other cases such activities are intended to support the objectives of the rules or the procedures. In the avionics case only when software was formally released and entered the formal process of testing could problems be officially logged and so become part of the monitoring process which kept track of the project. Problems became problems only when they were recorded as such and resulted in changes that had to be monitored and propagated according to the development process model. The 'fixes' of code, using engineer's testing, was routinely avoided not only because it was felt that it could be a reflection of the quality of the work, but also because it got around the cumbersomeness of the official procedure, saved time and got the work done more expeditiously. Similar patterns were exhibited in the Technical Document department where, and for a variety of reasons, various 'shortcuts' were employed in order to move the work forward.

From this social point of view, plans and procedures come to be seen as elements which enable workers to make sense of their own work and that of others. Of course, what sense is actually made of particular plans and procedures is a matter for investigation. What the perspective brings out is the importance of seeing how and in what ways plans and procedures are interwoven into a highly variegated set of phenomena that make up the social organisation of work. How and in what ways, for example, plans and procedures are related to the sequentiality of work; how and in what ways they formulate the work's interdependencies; how and in what ways they identify 'troubles' of various kinds; how and in what ways they make work accountable, 'open to criticism', 'subject of failure', and so on.

There are, of course, many other issues here, some of which will be discussed from a slightly different focus in what follows. To conclude this part, however, it is worth raising one issue of some direct relevance to CSCW design, namely, the ways in which plans and procedures may be incorporated into a computer system.

Plans in the machine

One of the ways in which plans and procedures appear in work activities, as we have already described and illustrated, is to formulate the work's activities as step-by-step stages. This feature is, of course, a standard one in bureaucratic organisations and one which has been explicitly designed as a means of making work subject to a system (Yates, 1989; Weber, 1947).[73] There are many examples throughout the case studies of this very general phenomenon. Of course, plans do not always have this degree of specificity nor is it always the case that highly detailed procedures, air traffic control is a good example of this, determine real time, step-by-step routines. In these cases, the procedures are more like the

[73] Much of the concern of Weber's work is with the extent to which the rise of rationality in the West is marked by an extension of planning and calculation to more and more areas of social life.

instructions for a game; that is, providing a means through which various tasks can be done but which do not determine when, at what stage, in the work when a particular procedure will be used or invoked. In other cases, however, detailed procedures readily lend themselves to incorporation into a machine system which does determine much more strongly when and in what ways the procedure is called into play. The static screens used in the Building Society, the TecSec system used in the securities centre, are both good examples of this. Such systems enable the work in that they automatically take the operator through the various steps involved in the completion of a task and, at the same time, effect relevant modal transformations of the inputted information.

The supposed benefits of such systems is that they automate much of the co-ordination of interdependencies within the work, involve lower training costs, require less experienced and knowledgeable operators and reduce the error rate. However, as we have seen in a number of the case studies, while such systems - and perhaps one ought to be hesitant about generalising here - do enable the work to meet at least some of the goals just mentioned they can also introduce rigidities into the work which often generate strategies of 'getting round the system'. There were, for example, many occasions in which workers in the securities centre 'lied' to the TecSec system in order to get the work done more speedily.

Related consequences for design

As mechanisms for effective distributed co-ordination, the role, representation and nature of plans is obviously of central concern for CSCW. Plans and procedures are a traditional technique for organising co-operative endeavours and have been used successfully in a wide range of contexts and for a considerable length of time. The advent of computer systems offers new possibilities in using plans to support co-ordinating activities and it is important such possibilities are designed in ways which are sensitive to the social context of the use of plans and procedures.

Relationship of People to the Plan

Existing co-operative systems which exploit the notions of plan assume the existence of an agreed plan which is followed by all. The general view is that the plan is closely related to the work taking place and that those subject to the plan interpret it in a similar way. However, even a cursory examination of the accounts of work to be found in the corpus of field studies, suggests that their is a massive variety evident in how plans are used and made relevant to the work involved. As a consequence, limited attention has been paid to representing the manner in which people relate to planned actions. We would suggest that a consideration of plans within co-operative work should:

- identify the different actors and their potential relationships to plans and procedures;

- understand and accept the existence of diverse viewpoints on plans and procedures;
- record and be sensitive to different notions of 'following a plan'.

The acceptance of diversity and the development of appropriate mechanisms for representing different perspectives is a significant design challenge for CSCW. Uncovering the design principles associated with this is a significant area of work for the COMIC project. As an initial step in addressing these issues the project is developing a framework which allows the representation of multiple perspectives on organisations (Deliverable D1.2). Similarly, and as we discuss in the following Section of this Deliverable, we are developing the notion of viewpoints as a presentation device for fieldwork within the CSCW systems development process.

Modelling and representing co-operative work

The development of models of activity and the representation of co-operative work has been dominated by a direct consideration of plans. Many of these have attempted to model work at a very detailed level. The variability evident in the use of plans outlined by existing studies suggest that overly detailed plans are problematic. This is particularly the case when the representation of activities embeds too much semantics related to the work taking place. We would suggest that a representation of plans needs to:

- allow users to alter and circumvent the planned activity;
- present a minimum overhead for those realising the co-operative activity;
- provide a representation of the status of an activity;
- allow users the freedom to interpret the plan and its current status in suitable ways.

A balance needs to be struck between sufficiently representing the work to allow users to monitor and co-ordinate their work, and presenting an overly detailed account of the work which is not flexible enough to allow the variability evident in the interpretation of plans. The outlining of flexible and malleable representations and notations is a significant theme in the work on Mechanisms of Interaction (Deliverable D3.2, D 3.3). An initial exploration of a lightweight model of activity based on shared information which builds upon the results of our study is also underway in the development of the shared object service (Deliverable D4.1, D4.2).

Awareness of Work

Another perspective on the social organisation of work, and one which is closely involved with that of Distributed Co-ordination is Awareness of Work. What we have in mind here is the way in which work tasks are made available to others and the important role that this plays in the 'real world real time' social organisation of work. Once again, this is a theme which involves interactional subtleties and,

once again, is not any single element within work organisation. Nonetheless, it is a theme which is relevant to CSCW design in a number of ways including, in particular, the means through which co-ordination of work tasks is achieved as a practical matter. However, our reason for emphasising this as a separate issue is to change the focus slightly to bring out the significance of spatial arrangements as well as other matters to do with the way in which collaboration is achieved, again as a practical matter. The aim is to bring to the fore some of the manifold ways in which work is made public and available to others as an essential ingredient in 'doing the work' as part of a socially distributed division of labour.

One of the features we drew attention to from the Training Centre Study was the 'ecology of awareness' afforded by the layout of the office and how this was used in the performance of the work activities. Similar affordances were also apparent in the Building Society office as well as in the Technology Centre and around the ATC suite, but the phenomenon is widespread. Many of these affordances arise simply because of the co-location of a collection of people doing 'much the same things' or who are performing interdependent activities. Being able to ask for advice just when it is needed, updating colleagues as one is passing, dealing with an emergency when 'all hands are needed', and more, are all informal, often intermittent, certainly unplanned occasional events which are facilitated by appropriate spatial layouts of work areas. The point we want to make about such affordances is that they are not necessarily designed or planned, though they can be, but represent arrangements which can be used in the day-to-day 'doing' of work.[74]

However, it is also important to point out that the affordances offered by spatial arrangements are only part of what we mean by the notion of awareness of others' work. Others are to do with the ways in which other's work is made available and visible, often 'at a glance'. For example, for those familiar with the control room at LATCC, what a controller is doing is visible by looking at the strips, the screen, listening to the talk around the suite, and so on. The incoming controller, in 'getting the picture', does not routinely have to interrogate the outgoing controller to determine what is happening but simply sits and observes the strips, the radar and listens to the talk and, using these resources, is able for the vast majority of shift changes to determine what the radar controller is doing and why. Similarly, though in a very different setting, in the offices of the Training Centre and in the Building Society, the state of desks, where someone is within the spatial setting, looking in a particular filing cabinet, working through a pile of paperwork, and more, furnishes information about what the person is doing, where that person is in the stream of work, how busy they are, and so on; information which, in subtle ways, is available so that others can update

74 In some respects the social use of space is a neglected area of sociological inquiry though, of course, it is the stuff of architectural design, town planning, ergonomics to mention but a few. The direction we have in mind here is with the ways in which spatial arrangements can be used to facilitate, or indeed hinder, collaboration and the coordination of work tasks. The importance of CSCW is the need to face up to the fact that computer mediated communication can abolish spatial affordances though it may need to emulate some of these. This is one of the thrusts of the research in Strand 4.

themselves on the state of the work, how it is going, whether 'we are behind', 'on top of it', etc.

Many of these affordances of making features of the work visible and available to others 'naturally' arise out of characteristics of the physical setting and out of the experience and knowledge which is accumulated through participation in the work setting and the activities which go on there. Others, however, as more specifically features of the work activities themselves and explicitly designed to represent aspects of the work *for* others. As part of the research involved in Strand 4, studies of document use (King and Hughes, 1994; Hughes and King, 1992; Hughes et al, 1993b) drew attention to the ways in which documents are massively public in the obvious sense that they are produced for specified or specifiable others. One of the features of their public character was the use of formatting; that is, conventions, many of which are heavily standardised - such as forms, memos, spreadsheets, reports, etc., - and which represent the features of the work in which they are embedded.

Representing Work within the Work

The issue of representation is receiving some attention in CSCW and is one which runs through a number of aspects of the COMIC programme.[75] Strand 3, for example, is concerned with the development of notations which can represent key aspects of co-operative work for CSCW design (Deliverable D 3.2). On parallel but slightly different lines, some of the research reported in this Deliverable is examining the ways in which ethnographic materials can be represented through the Designer's Note Pad. Our focus at this juncture is with the ways in which various kinds of representations operate within work domains to support the awareness of work.

We start from the presupposition that from the practical point of view of persons involved within a division of labour, where co-ordination is a prevalent and inescapable issue, one of the key tasks is representing the work as part of the work itself. Thus, while abstraction and selectivity are features of representation, our preference is to treat these as very practical matters for those who use them in the course of their work. In other words, the emphasis here is upon representing as itself an activity which can be done through artefacts of various kinds within work settings. For example, as a flight progresses through airspace the corresponding flight strip accumulates notations which represent the decisions and instructions given to an aircraft. That is, the strip is a means of representing elements of the work done by the radar controller in the course of the work. The point is that these representations are not isolated and detached marks or doodles which the controller happens to write on the strips but are embedded in the work itself reflecting the details of that work so that it is available to the suite team. In this

---

75 See, for example, Bowers (1992a, 1992b), Cooper (1992a, 1992b), Suchman (1993). See Pycock (1994) for an examination of the notion of technologies of representation in connection with field study material drawn from a project developing the hardware and software for photocopying and printing equipment.

way the work, its details, its progress, its state of affairs is made available to others involved in the work and, thus, becomes an instrument used to perform the teamwork around the suite.

In other words, what these representations do, among other things, is make the work 'visible' so that it can be 'taken note of', 'reviewed', 'queried', and so on, by others involved. They put the work on display so that others may be aware of it.

There are other aspects of this which need discussion which centre around the issues of what aspects of the work are to be represented, how they are to be represented and what roles they may play within the work setting. There is bound up with these the additional issue of who the representing is for.

Selection for representation

The overriding practical issue for representations is what to represent about the work, for whom, and how to represent it. What is it that people who are to use the representation need to know about the work and in what format is this to be conveyed? In other words, there are principles of selection involved. For example, the radar controller in writing on the strips does not write 'just anything' down. He/she uses a standard format - such as i14 to indicate an instruction given to an aircraft to ascend to flight level 14, or 14000 ft. - to provide relevant specific information regarding that aircraft. The notation used is extensive, and has a degree of flexibility to it, but it is, like all notations, a constrained system in that it is a means of succinctly describing actions that can be taken with respect to aircraft in controlled airspace. Similarly, the ABOM is a way of representing selective aspects of components relating them to phases of the development cycle, work activities and organisational personnel for whom they are relevant. In both these cases, the 'technologies of representation', to use this phrase, are highly specific, the flight strips especially, to the work settings in which they are used.

What seems to be at work here is a practically oriented principle of economy with respect to the information that is represented through the artefact.[76] We emphasise this as a practical principle for three reasons.

First, not all the possible information that might be gathered or represented is so. For simple example, the weight of the pilot of the aircraft is not recorded on the flight strip, nor is marital status, passenger names, and so on. What is represented seems to be governed by considerations of 'what is worth' representing or 'what is necessary'.

Second, the question of what is represented is not solved 'once and for all' but can be, and often is, a constant preoccupation in and through the course of the work.

---

76 Hutchins (1990) suggests that representation technologies be aimed at making activities easy and involve less effort by bringing the distant near, making the immediate durable and, therefore, delayable, by abbreviating complexity, and so on. See also discussion in Pycock (1994).

Third, understanding what the representation represents, so to speak, relies upon knowledge of the work itself and how the representation is embedded within the work as an instrument of the work.

The first point highlights the practical task of determining what features of the work are to be represented and in what form. That is, choosing from among the potentially infinite number of possibilities that are available, those that adequately, economically, represent the relevant features of the work. Representations are very often abbreviations and condensations of what it is they represent. This is not, of course, a problem which can be answered for all interested parties or, for that matter, for all possible uses. Managers, for simple example, might well have an interest in representing the work in a particular way but which is seen by workers as inimical to their interests. More subtly, the issue of representing work is also bound up with the kind of work that is done and what the representation is to be used for. As Garfinkel (1967) argues in respect of record-keeping, this is not some activity detached from the work itself but is "an element of institutionalised practices".

For example, one way of representing the work done at the Technology Centre is through the writing of Technical and Test Reports. These follow a standard format as set out below:

1. INTRODUCTION

A general outline of the context and aims of the technical issue at hand

2. EXPERIMENT

A basic summary of the experiment

2.1 Apparatus set up

A detailed description of the apparatus used and the manner in which it was set up: the rationale is that the experiment can be reproduced by someone familiar with the apparatus

2.2 Procedure

A detailed step-by-step account of the method followed to satisfy the aims of the experiment as outline earlier: the rationale is for 2.1

3. RESULTS

Graphs/tables of results are commonly appended as representations of the results

3.1 Comments

These are added as 'explanations' of the raw data presented in 3.

4. DISCUSSION

The results and comments are contextualised within the experiment's role in addressing the given technical issue

5. CONCLUSION

A summary of the contribution that the experiment has made to tackling the technical issue at hand.

Technical Reports are usually written by the most senior member of the staff involved in the project, although it is not uncommon for this task to be delegated to a more junior member. The reports are circulated to relevant parties and stored electronically with a hard copy archived.

The above format is a fairly standard one for technical report writing and is designed to set out as economically as possible what relevant interested parties need to know about the results of the project and how they may be assessed. It constitutes a record of the work done, and represents this in a form which highlights and presents the relevant information. This can be contrasted with the data presented in connection with the calibration test in the summary of the Technology Centre fieldwork. Much more occurs in a project than is mentioned in the report. What the report does is summarise by selecting the relevant features 'for those who need to know them'. It does represent the work done but does so 'institutionally'; that is, in ways relevant to the work practices of the technical staff.

Arising from this is a more general property of representations which has been identified in the literature (Pycock, 1994), namely, the mobility of representations. The Technical Report is just that: a representation of the work done in a form which presents the work in relevant ways for its audience, an audience which does not need to have been present at the test or during the project.[77] The report also constitutes an enduring record which is not only distributed when complete but stored so that it can be distributed as and when required in the future. It can be reproduced, combined and recombined, commented on, and so on.

The technical staff know, of course, that a report does not contain all that might be relevant. It is a document which needs to be read with some understanding of the work that it represents. Nor is a report treated as necessarily self-contained and sufficient 'unto itself' but is relatable to other reports and combined with them. Typically, if a test or a project report is of especial interest to a member of staff, he or she can talk to the people responsible for the report to be 'filled in' on any matters that the report itself does not contain. Observations suggest that this is a frequent practise among the staff.

Representations as modal transformations[78]

One of the features we have noted before in connection with the co-ordination of activities is the way in which information goes through modal transformations in which it is re-represented in a different form. A good illustration of this is the transformations of the information contained on the Function Sheet used in the Training Centre is fed into the staff roster, the Diary, memos, the 'move forward file', door signs, timing sheets, hotel reservations, among others. What the various transformations do in this case is facilitate co-ordination through appropriate representations of the 'basic information'. For example, the number of people on a two-day course translates into information about hotel rooms required, numbers for coffee, meals, as well as information about overheads projectors, seats and

77 Law (1986) describes some of the representations used in the 15th century, such as navigational charts and sail design technologies, which underpinned, he argues, the 'long distance influence of the Portugese over the East Indies. See Pycock (1994) for reservations about the direction of some of the literature on technologies of representation.

78 For a discussion of 'modal transformations' see Anderson et al (1989).

desks and other facilities in the training centre. It is not that all this information is contained on the Function Sheet; it is some of this information which, in effect, triggers the action implications for others. Similarly, in the bank and in the building society much of the work can be described as the step-by-step modal transformations of information.

Although this is a process which is deeply involved in co-ordination, it does so by making the work available to others in a form on which they work; that is, representing the work within the work so that others can do theirs. In these cases, they serve to routinise the work by using standard procedures and standard formats for representing the tasks done and the tasks to do.

Formatting as procedures

We have already made a number of points relating to the formatting of representations and it is worth identifying a number of other features relating to formatting in connection with awareness and co-ordination. As has been pointed out elsewhere (Hughes and King, 1992a), formatting is a ubiquitous property of documents and thus of documentary based representations. What we also want to stress is that this property is a social property in a number of senses. First, it imparts an element of standardisation so that the information it presents can be shared. Second, many formats are widely used as conventions within many activities and, thus, widely known. Third, a format is a set of procedures for creating and using representations as, for example, the format of the Project Report discussed earlier and, as another example, the ABOM. 'Knowing the format' is tantamount to knowing how to use it to create and use it as a means of representation. Knowing how to use a spreadsheet, for example, means knowing how its data is arrayed in a vertical and horizontal manner, knowing how to move between the rows and the columns, and so on.[79]

So, one of the properties of formatting is the way in which it can proceduralise representation and, through this, represent the work to others for particular purposes. That is, using a format means using its rules to construct the artefact of representation. The format functions as a set of instructions in both its creation and its use. Such rules can be simple or complex, strict or loose, and many of them technical, even mathematical and calculative. Among the latter, they include such constructs as lists, statistically based indices, timetables and more. Organising items and displaying them as a list, for example, or as the timetable of a project plan, serves not only to represent the work activities for others, but it is can also be a mechanism for summarising and, through this, organising other work activities. Once again the ABOM is a good example of this as is the System Problem Report used in the avionic software team.

---

79 These ideas are closely related to 'mechanisms of interaction' which is the concern of Strand 3 research in COMIC. Pycock (1994) explicitly draws parallels with this work and with some of the work which is part of this Strand. See also Suchman and Trigg (1992) on the format of the 'complex sheet'.

Some representations lend themselves to calculation, for example, as in the case of SPRs which were seen 'as getting out of control' because their number was increasing. Such a calculation can also be related to the plan and used as a measure of the extent to which the work is seen as 'falling behind' the plan, as the case may be. Within large organisations, of course, there are "centres of calculation", as Latour (1987) calls them, in which representations of various kinds are subject to statistical and other mathematical operations.

Formatting and relevance

Formatting is one of the ways of practically settling the issue of relevance; that is, answering the question what should be represented and how. The ABOM, for example, lays out for its users what is relevant to report and filling in the appropriate categories is a means of representing the work in a form considered suitable for the kind of activities it is intended to represent, similarly with the flight strip, the static screens in the building society, and the TecSec system in the Securities Centre. The transformation of the information on the Enquiry Form in the Training Centre Office is effectively a transformation of information relevant to the responsibilities of others in the complex, such as hotel staff dealing with booking accommodation, kitchen staff planning meals, and so on.

Of course, and as indicated earlier, the issue of relevance is bound up with that of selectivity - the question of what should be represented - and both intimately connected with ways of organising the work within the work activities themselves. In the avionics team, for example, particular tasks were the responsibility of a group of engineers which meant that an SPR raised against some code was also to raise further work tasks for those who produced the code. Pycock (1994) refers to this as the "alignment of representations" using the case of the linking, within a photocopier development team, of machine parts with specialist engineers. Thus, the various functional parts of the machine are associated with sub-groups of the project team, such that a part which is the source of a problem is also the responsibility of the relevant work group. In this way, the allocation of a problem to a part is the simultaneous distribution of work tasks to the group responsible.

Routine troubles

Although formatting and the construction of representations is an answer to the practical problem of relevance within a domain, as we said earlier it is not an answer 'once and for all', but can be a source of 'troubles'.

As we have indicated, formatting has different roles and for the purposes here we can identify at least two: first, that it provides for perspicuousness such that a person can know where to look for a particular item of information, where it will be located in a pattern, which is useful for 'at a glance' summations of what is going on; second, as a standardisation and disciplinary device. The ABOM, the SPR Form, the forms used in the Training Centre Office, etc., are all designed to

collect standard information, to make the information comparable and to control the information that is provided. There is a presumption, in other words, that the information the format provides for is information that is routinely required. In this respect they are solutions to the problem of the assembly of information in organisations identified by Garfinkel (1967). There are two aspects to this problem. When assembling information it is not always the case that what information is needed is known. How is information to be classified when it is not certain what the information is needed for?[80] The other aspect has to do with the worth of collecting the information. Will the use it has be worth the effort involved in collecting it?

One of the problems formatting encounters is that there are limits to it and it cannot always eliminate the 'how much is enough?' problem. Moreover, it is important to note that this is not simply a problem for the organisation, but a very practical problem for those who have to comply with the formats. The attempt to get people to comply with formats is often, with regard to those who must apply them, a disciplinary matter; that is, an attempt to ensure, by laying out a series of procedures to be followed, that the persons who need to comply actually do so. The issue for those who must comply is, very often and to put it bluntly, what is in it for them? Will complying with the procedures for providing information actually be of any use to anyone, not just for them? After all, much information that is collected is collected for other, perhaps unknown, participants in the process.[81]

Another routine trouble, and one which has already been mentioned, is the extent to which priority is accorded to the reporting format such that complying with it displaces the actual requirements of the work itself. This is, we suggest, particularly apparent when the reporting format is treated not merely as a record of the work it reports on, but also as a measure of the reporter's performance and activities. The SPRs in the avionics study is a good example of this as is the PIT in the banking study.

Thus, and as a more general issue, there is always the problem of articulating the formatted document with the flow of work; that is, the extent to which the information included is entered as a natural and integrated part of the work, and the extent to which it is regarded as an unwelcome overhead. The controllers' use of the strip fits in with the working, but in the case of the software engineers it becomes a structural problem of the software engineering process to build formats which can be similarly viewed. In the case of representing work there is always the problem of assessing the value of the information gathered relative to the effort involved in its collection and assembly. However, the values of this equation, to put it this way, are not the same for everyone in an organisation.

---

80 This is a similar problem to that faced when assembling the corpus of case studies in Section 3.

81 Relevant to this issue, especially in industrial environments, is whether entering the information will prove worthwhile since there are often strong possibilities that the project will be cancelled before anyone will use the information. This is endemically the case with CASE tools where much of the information is required for the purposes of future users.

There are, we might suggest, a diversity of rationalities within organisations. The value of particular information varies among different groups and what they want the information for. As a simple example, the value of information which permits control of the workforce and the location of responsibility for problems is attractive to management, but those being managed are providing information that may be used against them. For some in an organisation, some information 'might be nice to have', and there may be many things that could be done with it once it has been collected but, typically, for those who have to collect the information this becomes one more task on top of probably many others that they already have.

Related consequences for design

The reliance of co-operative endeavours on a shared awareness of work ensures that the support and promotion of awareness is a central design criteria for the construction of co-operative systems. Much of this awareness will result from the propagation of the effects of application interaction across a community of users. However, the sense by which we use awareness within this perspective is considerably richer than the mechanistic propagation of effect; rather it is closely tied to the features of the social organisation of work just discussed.

The setting of work

The setting within which work takes place is a prominent feature in understanding the social organisation of the work. This effect is particularly evident in the case of control rooms where the physical layout of the room is designed with a view to making common information shared and allowing controllers access to each other. Similar physical effects are evident within the layout of offices, factories and workshops. The physical affordances offered in the development of an 'ecology of awareness' within the work setting are important to the development of CSCW systems.

Computer systems are placed in real world settings and the manner in which they are used co-operatively is influenced both by the nature of the settings the properties of the computer system. Two important consequences are evident for the design of co-operative systems:

- An understanding of the physical nature of the work setting is central in understanding the nature of the work and informing the development of co-operative settings.
- Designers and developers of co-operative systems need to take account of the settings within which co-operative systems are to be placed and to plan their introduction appropriately.

In addition, computer systems have electronic effects that result from the work taking place. In the same way that the physical effects of work are central in supporting the co-operation taking place these electronic effects are essential to

the development of computer support for co-operative work. Understanding and supporting these equivalencies is a major design challenge for CSCW.

One approach to exploring these effects is to investigate the nature of space and how the space is used to support co-operative work. This understanding can be used to developed electronic equivalents which are applicable to work taking place within electronic rather than physical settings. This approach underpins the use of virtual reality within COMIC to develop a spatial model (Deliverable D4.2) of co-operative interaction and the construction of equivalent mappings applicable to co-operative systems in general (Deliverable D4.2).

Standardisation and Tailoring

The discussing of formatting and representation of work activities to promote awareness points to mechanisms of making information regarding the work available to others. The readily accessible nature of information is dependant on the use of agreed 'standards' and 'formats' of representation.

The design of these common formats and, as important, understanding the forms of co-operation they are intended to support is a significant challenge for the effective design of co-operative systems. Many of these 'standard' formats are closely tied to abstract representations of the work taking place which need to be agreed across the organisation within which the work is set. Consequently understanding the relationship between different representations of organisations, the use of organisational standards and procedures and the pluralism evident in the use of common organisational features is essential for design. The development of framework for reasoning about these different representations within COMIC (Deliverable D1.2) is one way in which we can explore these design issues as is, from a complementary direction, the research being undertaken in this Strand.

The use of commonly formatted artefacts is essential to the co-ordination of work. Users exploit their shared understanding of the common format to orient their actions in relation to the actions of others and exploit the ability to make public action as a means of articulating the work taking place. Many of these features are evident within the notion of mechanisms of interaction as they have been identified in the COMIC project (Deliverable D3.2, D3.3). Common formats thus provide us with a focus for analysis and the means of developing new mechanisms to support co-operative work. One design constraint on such new mechanisms is avoid presenting themselves as a costly overhead to those who have to provide the relevant information.

The reliance on shared common formats also highlights a tension in the design of co-operative systems. It is virtually impossible to predict in advance the detailed nature of the application interfaces without understanding the context within which it is used. Tailoring of the application properties is normally promoted in traditional approaches to interface design as a means of postponing design commitment until sufficient understanding is gained. Tailoring of this form

also allows user to match the properties of applications to their particular needs and preferences.

Traditional applications are essentially individual in nature and decisions about their properties lie within the remit of single users. However, co-operative applications are shared rather than individual in nature and alterations in the properties of applications have effects out with a single user. Understanding the role of tailoring is central to the design of co-operative application. In particular, a number of essential questions are posed:

- How do we represent and control the extent of tailoring across a community of users?
- How is the tailoring of common application properties controlled and managed?
- How do we detect and manage the divergence of representation which results from tailoring within co-operative applications?

These issues of tailoring and representation are explored in the development of user interface tool kits in co-operative systems. Within COMIC SOL (Deliverable D4.2) is a user model being developed to specify the extent to which different properties of interaction objects are common across a number of users rather than individual in nature.

Summation

This section of the Deliverable has examined in some detail the use of perspectives as a means of drawing out important features of work settings for designers. We have examined how the corpus of fieldwork presented in Section 3 can be used to illustrate the three perspectives we have focused on in this Section. Where applicable we have shown how the perspectives are manifest in the various settings covered by our field studies. We have also illustrated the importance of the different perspectives for design by highlighting, albeit in a general way, some of the salient design consequences which can be drawn out, as well as how they connect with other research in the COMIC programme.

We stress that the perspectives we have highlighted in this Section are but an initial set. The development of a framework for analysing work with a view to supporting CSCW systems development will require the confirmation of these perspectives as well as uncovering additional perspectives which illustrate additional significant features of work for CSCW applications. This is a significant methodological challenge for CSCW and for the work of this Strand for the coming year.

We have chosen to adopt the notion of viewpoints as a means of undertaking our investigation of applicability of multiple perspective on the fieldwork. Viewpoints are notational devices which allow multiple perspectives on systems to be recorded and reasoned about. The choice of viewpoints embraces and extends an accepted notion within systems development. This allows us to exploit

existing practices and to promote our results more generally within the systems development community. We have adopted a pragmatic approach to this work by recording and presenting the work of field studies.

Using a series of initial experiments as a basis (Sommerville, 1993) we have augmented a tool aimed at supporting design rational to act as means of presenting ethnographic information. The tool allows the rapid construction and amendment of information formulated as directed arcs. Given the prominence of this form of representation as a means of structuring information in systems development we have chosen to add facilities that allow the representation of viewpoints to this tool. It has been developed using a rapid prototyping approach based on the observation of its use in a series of early stage design and requirements specification sessions (Twidale, 1993). During all of the initial prototyping sessions the system was used as a structuring tool by more the one user. Similarly, a significant proportion of the presentation of the ethnographic information was developed jointly by the observer and the requirements engineer. It is to a report of this work that we now turn in the next Section.

Section 5
Developing a Framework for the Presentation of Field Studies

Section 5
Developing a Framework for the Presentation of Field Studies

Introductory considerations

Central to using ethnography in systems development is the effective communication of the results of fieldwork studies to inform the development of requirements. While the small scale nature of many of the endeavours outlined in Section 2 allowed the extensive use of informal debriefing meetings as a means of communicating these results, it is unlikely that this approach will scale up to industrial contexts. The demands of industrial development require the communication of the results of studies to be undertaken in a manner that is scaleable to large scale teams and to be sufficiently systematic to allow them be used to formulate effective design decisions.

Our chosen response to the challenge of larger scale development has been to focus on a systematic framework for presenting the results of ethnographic studies based on the outlining of multiple perspectives as suggested in the previous section. The development of this approach has built upon an existing design rationale tool (called the DNP) that supports the rapid construction of the directed graphs widely used in structured methods. We have complemented the use of this tool with an exploitation of viewpoints as a means of presenting the ethnographic material.

Viewpoints in Requirements Engineering

An extensive review of requirements methods and techniques was undertaken in the first year of the COMIC project (Deliverable D2.1). During this review we drew attention to the notion of viewpoints that has been incorporated into more recent methods of requirements engineering (Kotonya, 1992; Finkelstein, 1992) where it is explicitly recognised that requirements for a system derive from different sources which may have quite different perceptions of, and needs for, the system. These ‘viewpoint-based’ approaches are principally concerned with looking at viewpoints in isolation rather than as co-operating entities.

The basic premise of these approaches is underscored by an acceptance of the fact that different users of a system will hold different views of its operation, and argues that these views must be identified in order that inconsistencies and conflicts between them can be avoided. Such an approach clearly lends itself to the requirements specification process: viewpoint analysis of an application

domain (to use the language of viewpoint analysis) or workplace (to use the language of social science) can identify actors' relationships to the work process and to each other within the setting, and thus offers an account of the potential requirements of a system under development.

The temptation for social scientists might be to debate the validity of the notion of 'the viewpoint' as somehow an essential element of 'a reality', the essence of which is 'captured' with the identification of a sufficient number of viewpoints. This position is, however, a misrepresentation of the proposed contribution of viewpoint oriented analysis, and in practice does little to inform successful system design. Instead, our emphasis is on the practical requirements of system engineers as they set about the software development process. This approach represents a more practical and, we would argue, more constructive engagement with the problem of relating the contribution of ethnography to the design process

Nevertheless, it is equally obvious that the notion of 'viewpoint', in a variety of forms, has its sociological supporters and precursors and appears, implicitly or explicitly, in a number of sociological accounts of the research process. Thus, for example, 'strong viewpoint' approaches might be represented by Denzin's (1978) notion of "triangulation" or in feminist perspectives (McCarl Nielsen, 1990); and "weak" approaches, perhaps much closer to our own position, feature, again for example, in Law's (1994) outline of "modes of ordering" as well as in a variety of perspectives on evaluation, such as Parlett and Hamilton's (1976) notion of "illuminative evaluation". Our adoption of what might be termed a 'weak viewpoints' approach reflects a concern that viewpoint identification is fundamentally an empirical exercise without necessarily carrying any ontological presuppositions. As we have pointed out previously, our objective is to make use of the idea to effect a practical bridge between ethnography and system design.

Viewpoint oriented analysis

Our approach to investigating the applicability of the viewpoint approach is as a means of presenting ethnographic information in a form which can be used by system designers while, at the same time, preserving the richness of the fieldwork materials. We frame our account of the emerging use of the DNP in presenting ethnographic data in terms of a particular interpretation of the viewpoint approach.[82] Our aim is to examine the viability of the approach, its costs and its benefits, its advantages and disadvantages, by exploiting an existing technique.

In reviewing a number of viewpoint oriented approaches Kotonya and Sommerville (1992) offer a critique, based on dissatisfaction with the definition, inherent within these approaches, of precisely what it is that constitutes 'a viewpoint'. They point out that too much reliance is placed on the use of intuitive definitions of a viewpoint which, they argue, results in the situation whereby:

---

82 It is, at the present time, not altogether clear that a *definitive* viewpoint approach or method exists.

> two analysts specifying the same situation are likely to end up with completely different viewpoints.

Kotonya and Sommerville contend that a more rigorous approach to the definition of viewpoints in such analysis can facilitate an extension and enrichment of these approaches. To this end a "pragmatic definition" of 'viewpoint' is offered "that is both unambiguous and complete":

> **viewpoint**: an external entity that interacts with the system being analysed, but one that can exist without the presence of the system

Further to this the authors highlight four key elements of in viewpoint oriented analysis:

- viewpoint identification;
- viewpoint structuring and decomposition;
- information collection;
- reconciliation of information across viewpoints.

We have chosen to adopt these four elements as an essentially pragmatic approach to the organisation and presentation of ethnographic material in system design. However, we do not wish to suggest that such a viewpoint-oriented approach enjoys an entirely unproblematic relationship with an ethnographic study of a given work-setting. Rather, these elements are used to frame our account of the use of the DNP to present ethnographic information.

Viewpoints as a means of presenting ethnographic material

Rather than use the notion of viewpoints simply as an analytical tool for software engineering, we wish to deploy viewpoints as 'alternative perspectives' on a corpus of common ethnographic information that may be used to inform the developer or development team. The adoption of a technique based on viewpoints allows us to present information in a form that makes explicit the different but complementary interests involved in the design process, particularly the interests of providing an adequate characterisation of the sociality of the setting and those of software engineering.

The presentation of different viewpoints also allows alternative views and perspectives to be set alongside each other as a resource for designers and requirements engineers in formulating the more abstract models essential to the requirements definition process. In summary, a number of particular reasons motivate our choice of viewpoints as a means of involving ethnography in the development of requirements for interactive systems:

- They highlight the multiple orientations people may have to a supporting system, not all of which are conflicting.
- They are naturally sympathetic to the heterogeneity evident within the development of requirements.

- They provide a means of setting the multiplicity of user needs along side each other to inform the construction of requirements.
- They are generally understood within the development community and dissemination of information in this community will be enhanced.

The last of these points is central to the need to support multi-disciplinary working when formulating requirements from a human centred focus. The problems of differing vocabulary, disciplinary traditions and approach are significant when we consider the incorporation of any human centred method into the engineering tradition of systems development (Bentley et al, 1992). By adopting a common vocabulary that is understood within the target domain of requirements engineering we seek to minimise these. It is unlikely, given the nature and scale of the development process within real organisations, that requirements approaches will move significantly from the existing representation techniques, particularly given that they represent a serious investment of resources. Thus, rather than offer a critique of these methods, or try to develop a notational vocabulary capable of encompassing both ethnographic materials and traditional requirements techniques - and it is difficult to see clearly what such a creature would involve - we seek to develop ways in which the two means of representation can be more strongly coupled than they are at present

Our starting point is to see the problem of relating the findings of fieldwork studies to engineering design as one of communication, particularly in representing the fieldwork findings in a form usable by engineers; that is, a means of transforming the primarily textual form of ethnographic reportage to a graphical means of representation. The focus of our investigation is on the initial examination of fieldwork materials within a design context. Our set of presentation techniques aim to relate the current discursive textual information of ethnographic data to the graphical network based notations exploited by engineers.

Presentation viewpoints

It is important at the initial stage of formulating requirements that the techniques adopted are not too restrictive in approach but instead promote some flexibility of interpretation. We focus on the use of viewpoints as a simple means of representing particularly pertinent aspects of the study rather than as a means of analysis and requirements specification. One consequence of this methodological decision is that we have avoided attributing the specific meanings or interpretation to viewpoints that is characteristic of much of the literature, and offer them instead as a means for supporting specification judgements.

Currently, ethnographic records are closely associated with the fieldworker undertaking the study of work. There are strong practical and ethical considerations for ensuring that this remains the case. Fieldworkers construct the ethnographic record essentially as an *aide-memoir* to allow them to recall later features of the work. The record is a heterogeneous collection of information

involving a wide range of media. It includes examples of existing paper information, audio and video records, and sketches and notes made by the observer. The collecting of this information is essentially a personal endeavour that usually results in a "contingent corpus of materials" which often only makes very much sense to the fieldworker (Hughes et al, 1993a).[83] In addition, the ethnographic record is normally considered to be the property of the fieldworker who has control over its access and presentation. This arrangement make the ethical position of the observer explicit and unambiguous in acting as a gateway between the domain of study and that of development.

Altering the relationship between the fieldworker and the ethnographic record in any substantial manner would have significant methodological implications for the ethnographic study. Consequently, we have chosen to focus on the provision of facilities that allow a fieldworker to deliver information from a study to developers involved in the construction of a requirements specification. These facilities are not intended to replace the diverse collection of materials that currently constitute the ethnographic record; rather they offer an additional means of recording and representing the information which is geared to the requirements of CSCW design. In other words, this is not offered as a analytical tool to use in the field by an ethnographer. Its use is very much to do with presenting fieldwork materials to developers once they have been gathered and written up. In other words, and to repeat, what we are not developing is a fieldwork tool.

In structuring the fieldwork materials we have specified a small set of viewpoints, each of which represents a particular focus on the social organisation of work activities within the domain. The viewpoint categories have been chosen in order to bring out key aspects of the sociality of work in a way which facilitates systems development based on the fieldwork materials. Consequently, the starting point for these viewpoints is the set of perspectives outlined in the previous section.

Two principal perspectives were identified in Section 4, the awareness of work and distributed co-ordination. These two perspectives, we argue, not only bring out important features of the social organisation of work but also highlight central concerns to CSCW design.

Each of these perspectives motivate particular presentation viewpoints.

- A viewpoint that presents the *ecology of work* draws directly on the concerns expressed by the awareness perspective explored in Section 4. This perspective highlights the important role that the ecological setting of work has to play promoting awareness.
- A viewpoint that presents in some manner the *abstract flow of work* reflects the central concerns expressed in the distributed co-ordination perspective. Of particular note is the need to relate work to abstract plans and representations.

---

83 This closely parallels the problems we highlighted earlier in connection with the assembly and the presentation of the corpus of fieldwork studies set out in Section 3 of this Deliverable.

It would be misleading to give the impression that the perspectives outlined in Section 4 are complete or that they are the only ones which might be worth developing. Rather they represent a starting point in uncovering a framework for presenting ethnographic material as well as an opportunity to begin the exploration of the technique. Nevertheless, even as a starting point a further viewpoint is needed to allow for linking between the ecology of the work and the flow of work viewpoints with information about other features of the social organisation of work in the setting. The three viewpoints that emerge from the perspectives presented in Section 4 are briefly set out below.

Ecology of work

This viewpoint seeks to represent the spatial distribution of the work place in terms of its participants, the work they do and the local resources that they use. The purpose of this is to provide a sense of 'where the work takes place' and the socially constructed affordances that this offers as an arena of various kinds of interactions that take place. In this respect, it is a view upon the 'workaday' character of the world within its setting. It is also relevant to understanding some of the ways in which the awareness of work activities are made available as a property of the physical affordances of the work site.

Work Flow

This viewpoint, and it is probably best seen as a collection of potential viewpoints, focuses more directly on sequences of work activities, information flows, and so on. In this respect it emphasises and exhibits the division of labour within the work along with its various interdependencies; interdependencies, it is important to stress, which are not always organisationally formally specified. The kind of fieldwork materials germane to this viewpoint include 'tracking work' through its sequences and transformations, such as a particular piece of software through error testing, the flight of particular aircraft through UK airspace, invoice processing, etc.. Once again, such materials will consist of reports of activities, the relationships among parties to the work, how the interdependencies are achieved as 'real world, real time' phenomena, the contingencies that can arise, how they are dealt with, and so on. This material is normally textual though use is often made of simple diagramming techniques.

Perspectives on work

This viewpoint is broadly formulated in response to the relatively loosely structured character of ethnographic studies of work which, typically, do not begin by using a strictly formulated set of theoretical categories. Rather, their emphasis is on providing materials on the 'real world, real time' nature of work which can be used for sociological analysis. These materials furnish portraits of the practical nature of work, often presented as illustrative vignettes within a larger report. The analysis, once again typically, tries to bring out the day-to-day

experience of the work from the point of view of various actors within the setting; points of view which can, themselves, be used to generate viewpoints for design depending on the interests of the designer and the system being developed.

This viewpoint aims to collate summaries of this body of observational information in such a manner that they are accessible by developers as a resource for requirements specification, open to flexible use and, importantly, capable of adding substantive information about the social organisation to the other viewpoints. Given the relatively discursive presentation of this material as well as its focus on the subtle, often 'invisible', often tacit, features of workaday activities, this viewpoint is really a collection of potential viewpoints from which such materials can be examined depending, as said earlier, upon the interests of the designers. Observers and developers are free to add additional perspectives relevant to the study as a means presenting this information. For example, it may be presented from the point of view of a particular actor within the work setting, from the point of view of a sequence of tasks which involves a number of actors, a collaborative endeavour, the policy of the organisation, and so on.

These presentation viewpoints are intended to support the uncovering of additional perspective useful in the design and development of co-operative systems. They suggest areas of particular concern or interest to designers. The eventual presentation of field studies using these perspectives is intended to support developers. However the process by which this presentation is constructed is of equal interest in outlining a framework for incorporating studies of work into the development process.

An Overview of the DNP

Originally created for software engineers, the Designer's Note Pad (DNP) is a flexible system providing support in the early stages of design and allowing for the rapid change and expression of ideas. As Monk (1994) characterises it, DNP is a prototype flexible information management and browsing system providing a range of facilities for representing, changing, rearranging and referencing information. Haddley and Sommerville (l990) suggest that the DNP represents a first step towards the next generation of CASE tools which provide effective support for the whole of design process and not simply support for detailed design activities. The emphasis is on lightweight support for the designer with the DNP being used to express design ideas which are generated. In a similar fashion, Sommerville and Rodden (1992), in stressing the need to recognise and understand the social aspects of systems design, emphasise the lightweight character of the DNP and its role in providing support for informal design activities.

For Haddley and Sommerville (l990) the potential of the DNP is therefore substantive in that it is not restricted to computer-based systems, while Monk (1994) argues that its flexibility has led to DNP being used in a variety of

contexts; as a basis for computer-supported learning, the development of safety cases for software systems; as a generic software design tool; and, in this instance, in the organisation of data produced during an ethnographic study of work processes.

The remainder of this section is concerned with outlining some of the features encountered in using the DNP to represent ethnographic information. The intention here is relatively restricted to an examination of a small number of the facilities of the DNP - designs, text notes and cross references - and is certainly not to produce a definitive list of features, problems or possibilities since our usage of the DNP is both at an early stage as well as an ongoing, evolving concern.

Representing Designs in the DNP

The DNP is designed to allow information stored in it to be reorganised flexibly and easily. This supports the idea that for such a tool to be useful, the user must be freed from as much 'cognitive overhead' as possible in creating the design. Thus, commonly entities may be moved from design to design across window boundaries by direct manipulation, groups of entities converted in a sub-design and the text of annotations converted into entities, and vice-versa.

Designs are seen by the user as a window on which objects can be placed and moved around. The various objects (cross-references, text-notes and sub-designs) are structured annotations to entities or links. The DNP supports hierarchical decomposition by allowing an entity to represent a sub-design. This sub-design behaves in the same way as its parent design so that entities, links and sub-designs can be created in it.

An extensive symbol library is available so that users can define the graphical representation of the types required in their method. The ability to manage multiple families of types means that the DNP can support a comprehensive yet simple system of information hiding. Therein, the DNP has been commonly used as an outliner/ideas organiser. The ability to create entities and links quickly and to rearrange them spatially makes this a common use of the DNP. The free-format is very fruitful as ideas are put down as entities and then rearranged, grouped and deleted as they become more coherent.

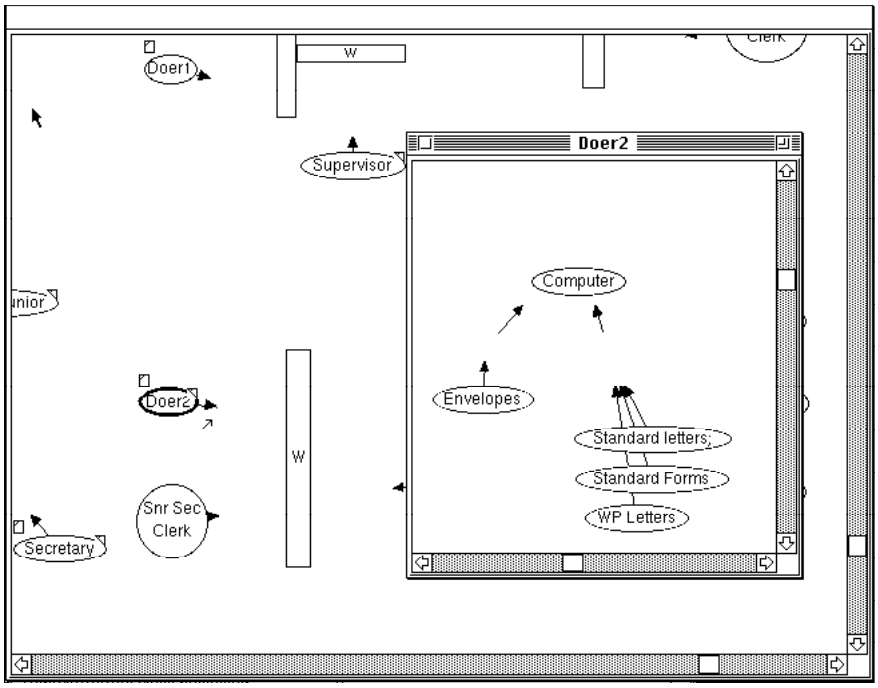

Figure 5.1: Using a design in the DNP to represent a work setting

In the case of ethnographic fieldwork notes the work site can be represented as a simple plan which utilises entities, links, text-notes and sub-designs to valuable effect. Various different colours and symbols can be used to signify social actors, their positions and spatial relationships (fig 5.1). A sub-design can be used on an individual worker in order to display a information concerning their work flow. Sub-designs may also contain be complex entities with text notes attached.

Text-notes

The manipulation of text-notes are a very important, indeed key, feature in the DNP for ethnographers in the management and organisation of their fieldnotes. Text-notes are annotations that can be attached to any type of entity or link. They are analogous to sticky ‘post-it’ notes and can be simple unformatted text areas or complex, lengthy user-defined terms. Text is usually typed into the note pad and then attached to the entity or link. Any number of text-notes can be attached to an entity and they can be browsed, edited and printed.

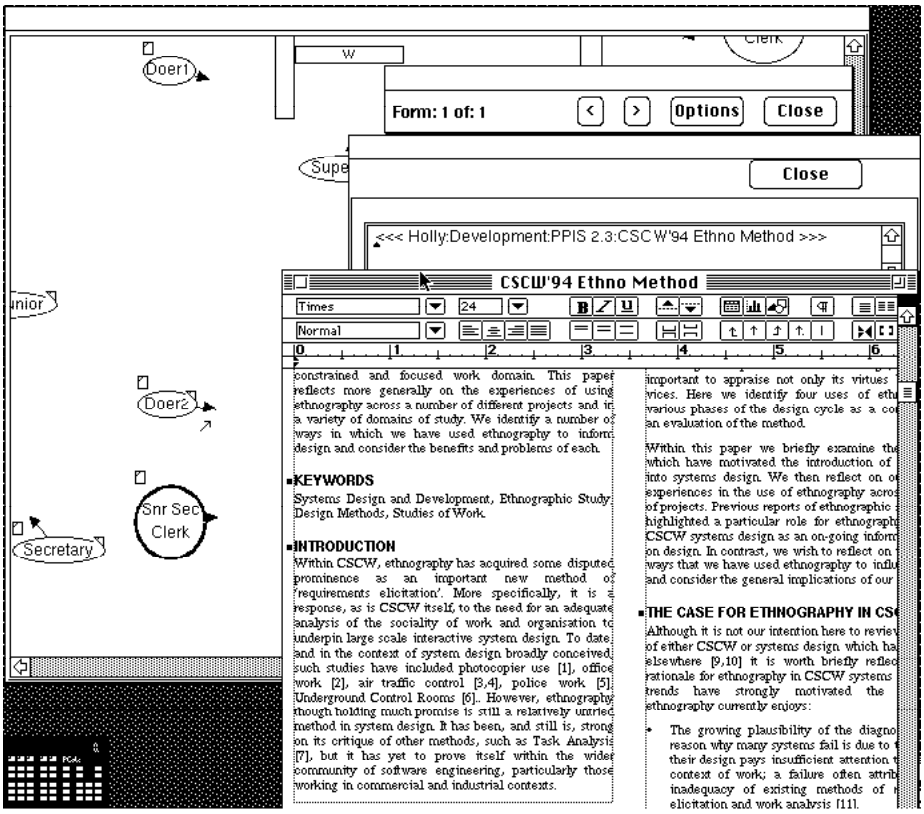

Figure 5.2: Linking from the DNP to external applications

A large number of documents are created during a study which are of different types. These can be transcriptions of conversations, interviews, statistics and observations etc. This information is often held in files created across a wide variety of applications. The information in text-notes need not be necessarily held within the DNP. A text-note may be linked to a file so that its contents are derived from that file. Changes to the file may come about by the application of another tool. This facility allows material generated in existing tools to be directly linked to material within the DNP. For example, in figure 5.2 fieldwork material stored in a word processor file on the Macintosh is linked to an entity within the DNP.

Cross-references

The DNP allows the construction of a hypertext network between different entities and designs in the system. Links can be established between any entity and any other entity which can be in different designs. Such cross-references, which are designated by a small arrow next to that entity, can be bi-directional. If more than one cross-reference exists, when asked to search these, the DNP lists them in an options menu (fig 5.3).

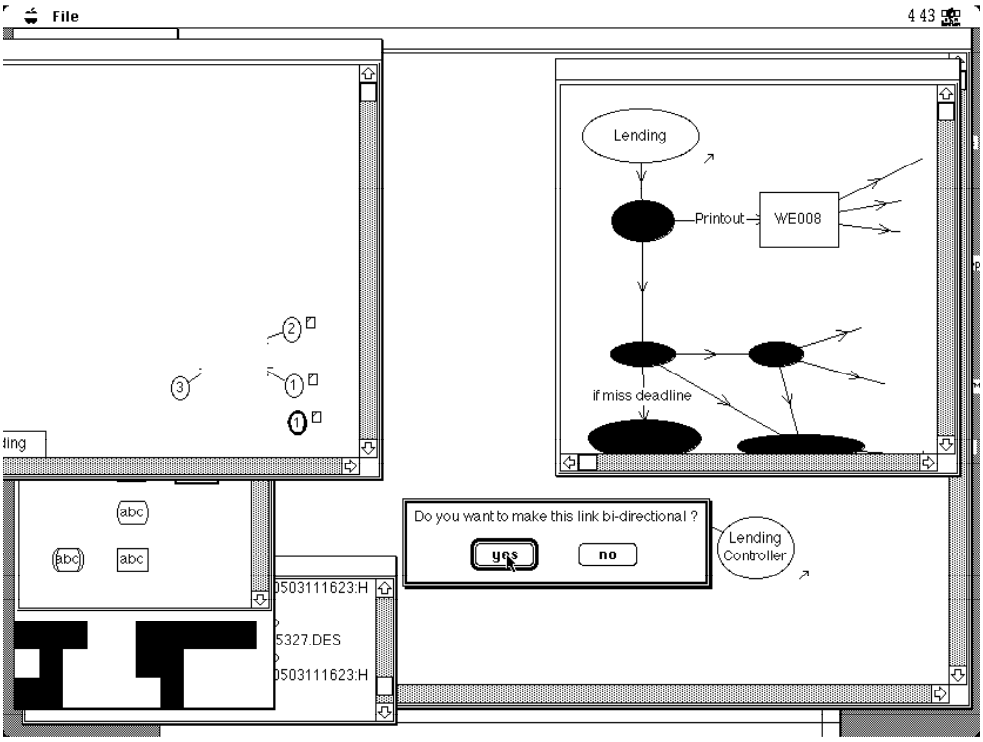

Figure 5.3: Developing links using cross referencing

Viewpoints in practice

The strength we identify in our initial experiences of using DNP, and again we stress that these are merely observations of our experiences of an application under development, is that it provides an opportunity to maintain and develop the viewpoint-oriented analysis in the early stages of the requirements elicitation and system design processes, whilst also remaining consistent with the complexities of the social characteristics of the application domain. To illustrate and elaborate upon this assertion we present an account of using the DNP, both in organising a number of sets of ethnographic data and in the initial stages of requirements elicitation for the visualisation of a document scanning and storage system. This account will be 'framed' by and Kotonya and Sommerville's 4 stages of the VOA process, and will offer a consideration of the ways in which we feel the DNP has the potential to contribute to the enrichment of this process by facilitating complementary involvement of a systematic representation of socially generated viewpoints.

Viewpoint Identification

Much of the current literature on viewpoints in effect glosses the act of viewpoint identification. Given Sommerville and Kotonya's (1992) emphasis on the need for clarity and practicality in 'viewpoint' definition we feel there is an equivalent

need for the act of viewpoint identification to be similarly rigorous in nature. By this we mean that the socially-generated viewpoints are identified within the ethnographic record in a practical, consistent manner. This does not mean, to repeat a point made earlier, that analysis of the ethnographic data should be governed by, or moulded into, a limited set of *a priori* representational viewpoints, but rather that these representational viewpoints should, first of all, come out of a full consideration of the ethnographic material. As we have stressed earlier, DNP is not intended as a substitute for the analysis of ethnographic fieldwork materials. Nevertheless, as a means of representing these analyses in a design context, various views on such materials can be explored in a systematic manner.

In this case we have found that a number of viewpoints have emerged in the use of DNP to organise and represent fieldnotes. These presentation viewpoints are consistent with the broader ethnographic findings at the field sites in question and, in this respect, are findings from the fieldwork rather than brought to and imposed upon the data. A particular focus here is the ecology of the work place and the egological organisation of the work

Ecology

Detailed descriptions of the work setting or ecology are a general feature of most ethnographic accounts of work; in this instance the DNP permits a graphical representation of that ecology to overlay and further explicate the textual representation. Such a graphical representation marks the boundaries of the application domain, and offers an overview of the physical relationships between the actors.

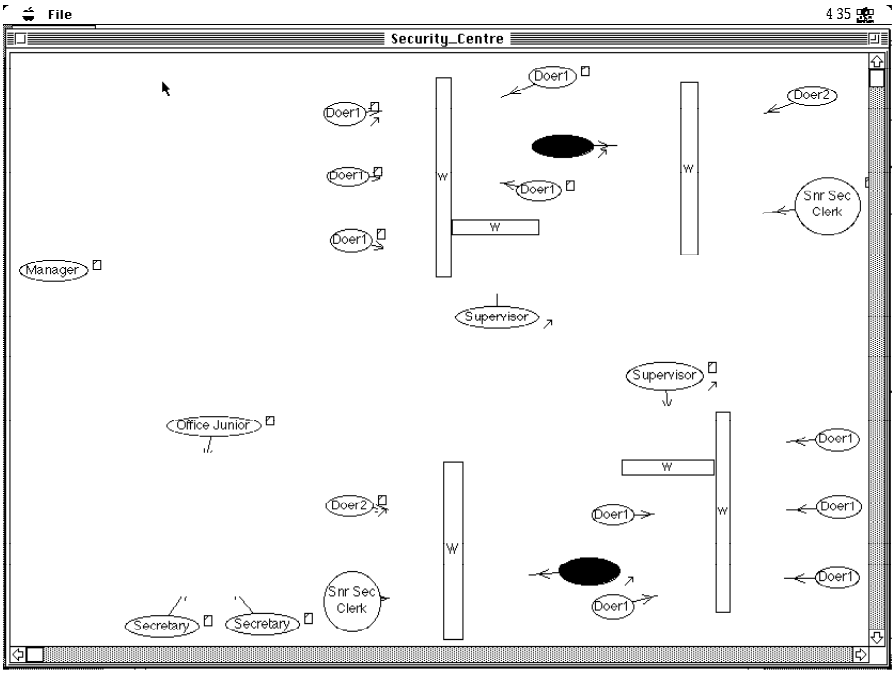

Figure 5.4 - Office Ecology

In the example given (Figure 5.4) it is clear that the ecology of the workplace is organised around desk spaces, each belonging to an individual actor - in this case the physical proximity of workers is important to consider in offering an account of the accomplishment of work through teamwork, informal interaction and the sharing of local knowledge and local logics. Put simply, team members, by their proximity are able to share, by design and serendipitously information about customers, knowledge of proper procedure, passwords and the like.

It is also worth noting in this example that the physical relationship between the team work setting and other important work areas, such as the book room, where all the paper records are kept and the managers office, where actions are sanctioned, is also made visible. Whilst this may seem a trivial point, this sense of 'meaningful physical organisation' can often be obscured in conventional accounts of work activities. We take the view that such ecological details are meaningful in that they offer a framework for understanding some types of work behaviour. In this case, and for example, workers would typically wait until they had a number of files to collect, or would ask other team members whether they needed any files, before proceeding to the book room because of the time taken in searching records.

A relatively static ecology, such as is depicted in Figure 5.4, clearly lends itself to the diagrammatic form of representation, and it is perhaps worth sounding a note of caution to the effect that we have not yet attempted to represent in this way an ecology of work of a more open and flexible kind although the use of cross-referencing, sub-designs and other hyper-text links may prove advantageous in this context.

Clearly, such graphical representations are not intended to replace or even subordinate the conventional ethnographic record, instead we suggest that the two representations are complementary. In particular, by laying out the ecology of work in this way, along with the supporting text notes, is a means of exploring the ethnographic record to see, for example, the extent to which it covers all the relevant roles and processes, and where additional investigation might prove profitable.

A further significant issue is uncovering and representing the relationships between different viewpoints. For example, Kotonya and Sommerville (1992) describe a viewpoint as "a hierarchy with differing levels of abstraction" . The DNP facilitates the exploration and representation of the ethnographic record in terms of relationships such as these. This is most pertinent in the case of more abstract representations of work. The work flow viewpoint is the principle example of this within the research reported here.

Work Flow

A schematic representation of work flow documents just that, the flow of work, as inputs are transformed into outputs. Needless to say, the very premise of the ethnographic enterprise is one which foregrounds the fact that such a process is

not a simple or straightforward one but, instead, highlights the importance of local knowledges and collaborative skills that are integral to the production and achievement of such a flow.

We have used DNP representations of work flow viewpoints in order that the complexities of such interdependencies between the various stages of work are formally specified. This is clearly of importance to system design, in that a given system under development must support such interdependencies, and, indeed will have to operate effectively within such a flow.

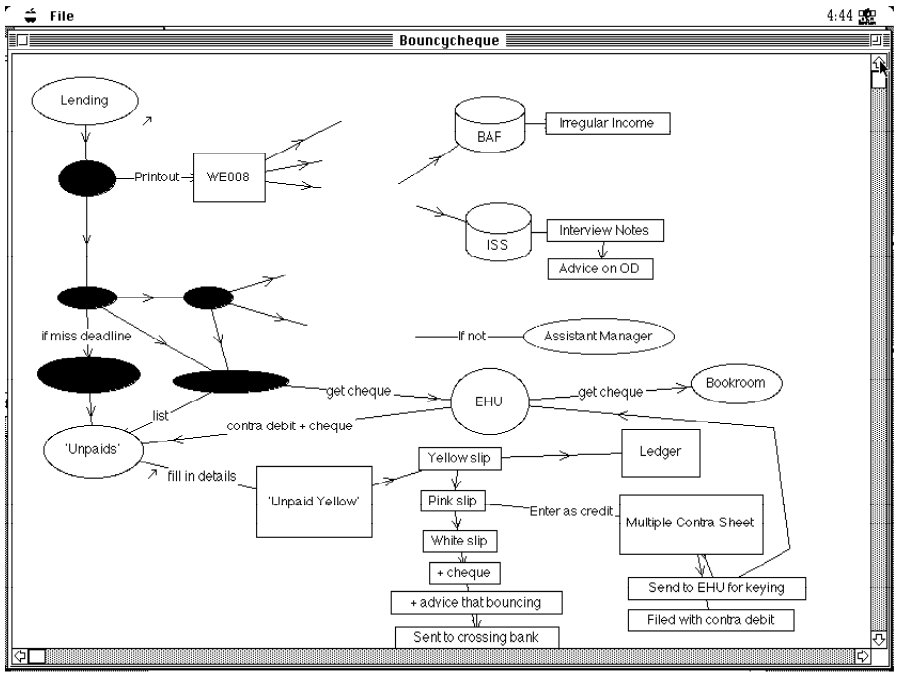

Figure 5.5 Work flow schematic.

When representing work flow within a DNP design the ability to follow hypertext links (in the form of cross-references, sub-designs and post-it notes) to other representations of the ethnographic record facilitates the assessment of constitutive levels of specialisation ‘beneath’ such an abstraction.

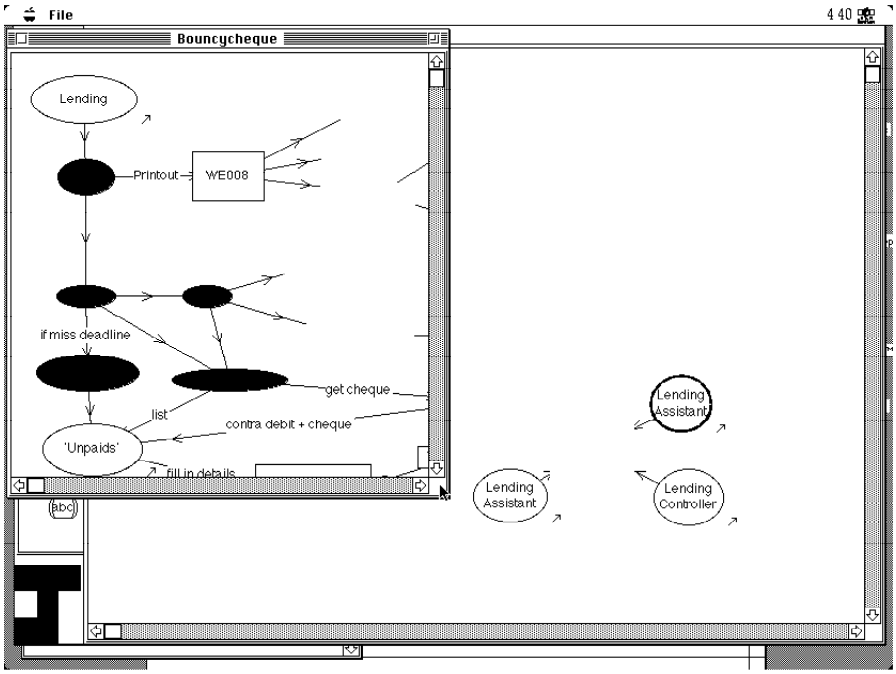

Figure 5.6 Work flow and Office Ecology.

The requirements engineer is therefore able to assess the ethnographic material which has been gathered in the construction of such an abstraction or representational viewpoint, and the myriad complexities of the work flow which it covers can be explored and assessed by the system designer.

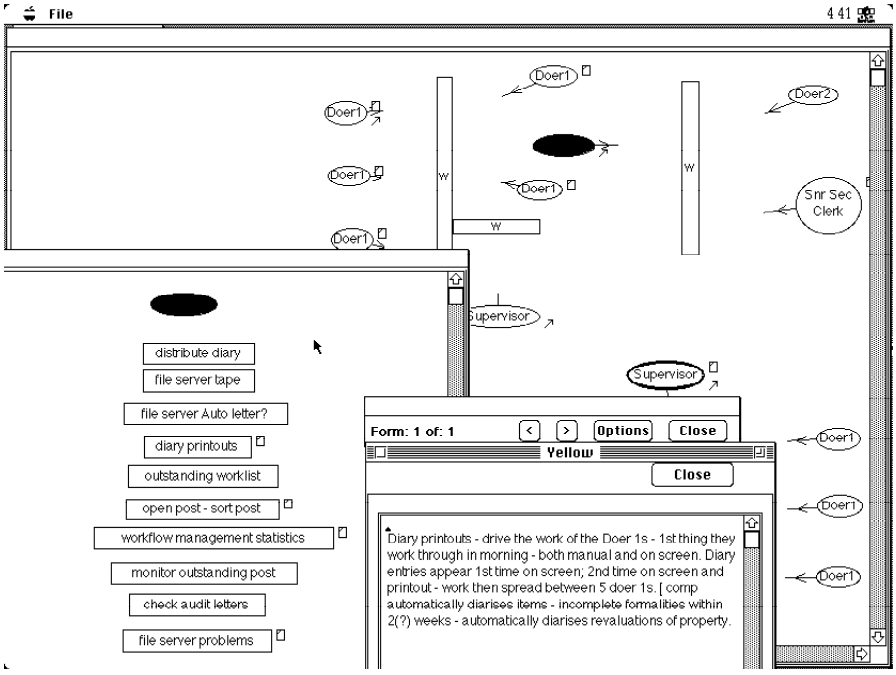

Figure 5.7 Work flow Representation and Text note

In summary then we feel that the DNP facilitates effective, systematic presentation of the work flow viewpoint, and furthermore enables the 'reader' of the DNP design to assess rapidly the ethnographic data that has been drawn upon to construct such an abstraction.

Egological organisation of labour

The term 'egological' refers to the manner in which the individual actor fits into a complex division of labour involving a number of collaborating individuals. It encompasses what the individual needs to accomplish and needs to know, for example in terms of a practical understanding of processes, the location of files or information, knowledge of computer packages, and the myriad of other details that are part of the work's organisation. The actual accomplishment of work may bear little relationship to idealised models of the work process since it is likely to include aspects of local knowledge or local logics that lie outside, indeed might be regarded as subversive to, more formalised specifications of activities, roles and tasks. As an admittedly extreme, and unusual example, in one location where the work flow was rigidly driven and generated by the computer, - in that certain 'formalities' had to be completed before the computer would 'allow' any other work to proceed - the simple expedient of 'lying to the computer' - telling the computer that a particular procedure had been completed when in fact it had not - enabled workers to maximise their efforts and accomplish tasks more speedily.

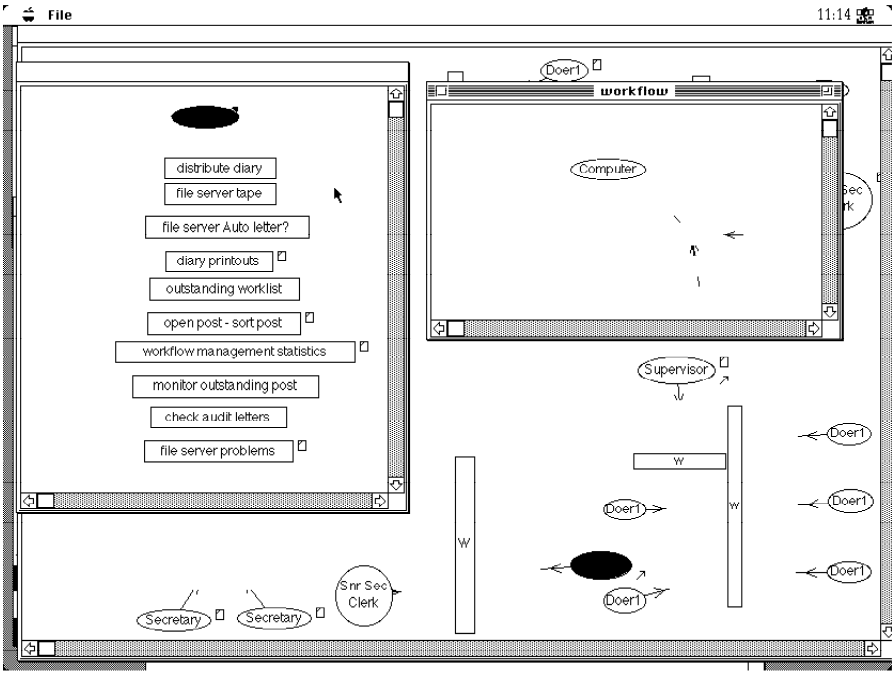

Figure 5.8 Egological Organisation - Office Ecology and Sub-designs of Work flow and Desk Organisation.

Egological organisation can be represented on the DNP by a combination of sub-designs and 'post-it' notes giving graphical representations of 'typical' work

flow; the organisation of deskspace and access to relevant files; and accounts of how the work is practically accomplished, including the myriad of interruptions and disruptions that punctuate the working day. In figure 5.8 above the work flow and desk organisation of the Supervisor of Team 2 is displayed using sub-designs alongside his overall position in the office ecology. The sub-sub-design of the desk organisation details the contact lists and the paperwork requirements that are required on a regular basis; which any attempt to redesign the work would need to take into account. The work flow sub-design gives an outline of a 'typical' working day and the accessible 'post-it' notes on the sub-design provide more detail on aspects of the work flow (Figure 5.9 below).

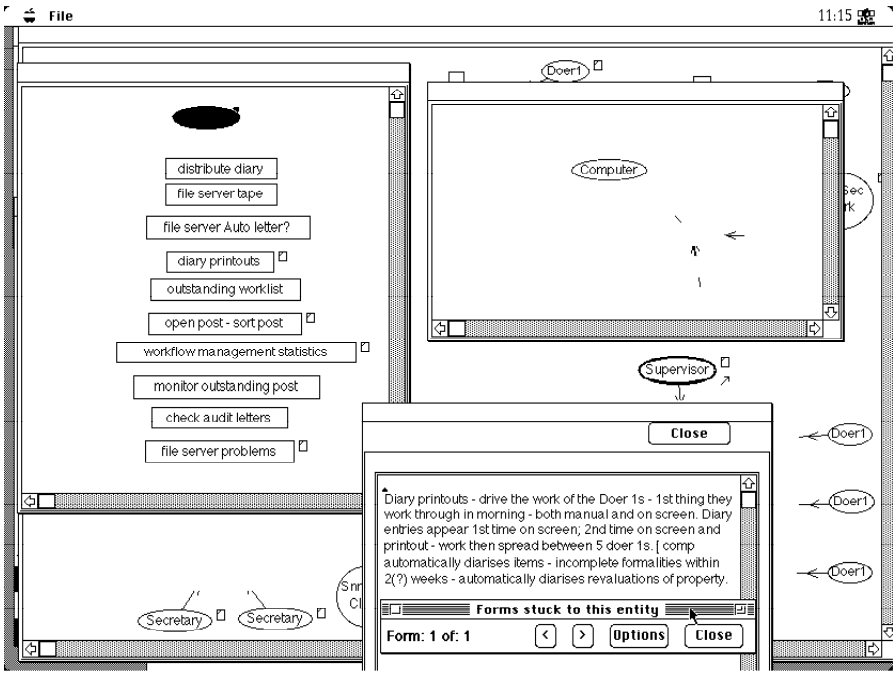

Figure 5.9 Egological Organisation, Sub-designs and 'Post-it' notes.

In figure 5.10 below the raw observations of the Supervisor at work, recorded in 'post-it' notes, have been accessed, providing, with the other sub-designs and post-it notes and alongside the 'conventional' ethnographic account, a reasonably detailed description, a 'viewpoint', of the position of the Supervisor in this particular system.

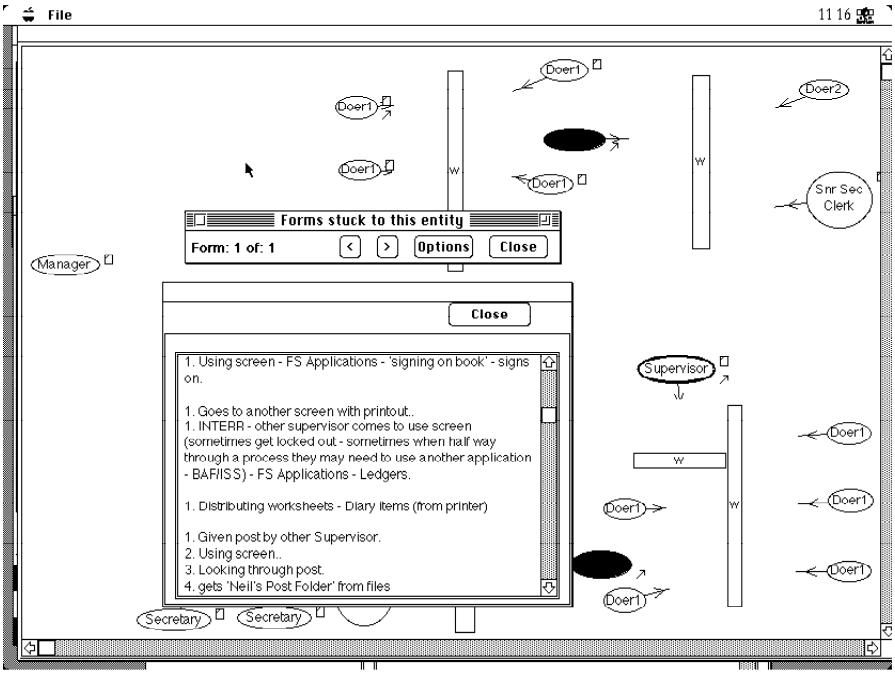

Figure 5. 10 Egological Organisation, Office Ecology and 'Post-it' notes.

Information collection and the role of viewpoints

In terms of a viewpoint oriented analysis, this stage in the process is characterised as the gathering of information to provide detail relevant to the viewpoints identified earlier. In Sommerville and Kotonya's summary, such an enterprise is presented, more as a necessary precursor to the final, and crucial, stage of reconciliation of this information across viewpoints, than as the opportunity for the identification of co-operation and collaboration between actors in the work-place. This in effect implies that certain viewpoints have been defined *a priori* and applied to a study of the workplace. It is in such a process, we would argue, that the danger of over-schematisation is run: the assumption that a generic set of viewpoints exist that can be brought like a 'palette' to be drawn upon in representing any and all work settings.

In contrast, our structuring approach exploits the notion of viewpoints to highlight and record the diversity of perspectives evident in a work setting. Offering a set of viewpoints suggests only a more systematic approach to presenting ethnographic studies within the development process. Within our example, field studies viewpoint identification has emerged in the process of information collection, and its subsequent organisation and presentation in DNP designs. We feel that representation of this kind benefits from the absence of *a priori* over-schematisation of the work settings covered.

Reconciliation of viewpoints

The final phase of the viewpoints process outlined by Kotonya and Sommerville focuses on the reconciliation of information across viewpoints. The general hypertext facilities provided by the DNP enables the linking of information central to this process. As an example we shall consider the development of an abstract representation of work drawn from a field site which provides a framework for reconciling the different viewpoints.

The focus for our study was a technology centre within a large multi-national manufacturing company. The centre serves various manufacturing businesses that are part of one of the divisions of the multinational company, and is one of three such centres, all established in the mid-1980s. Links with the various businesses are through a combination of regular formal meetings, informal personal contacts and 'fire fighting' problems as they arise. Until recently, the company's policy was to locate technical expertise in particular centres which would provide consultancy on specific problems and technical issues as they arose in various manufacturing sites as well as undertaking product related research and development. The company is currently actively seeking to provide more direct access to this expertise using electronic communication facilities. A central part of this initiative has been the migration of information from existing paper based technical reports to an electronic representation. This involved the development of an electronic data system into which existing reports were scanned. The company had invested in electronic storage facilities and scanning software to undertaken initial prototyping trials.

The centre comprises three organisational elements: the office, the laboratories, and the Pilot Plant. The office is a large open plan on the first floor of the technical centre building. The Pilot Plant is a large factory building adjacent to the office and labs which contains a variety of testing and production line equipment. It is occasionally used for the post-development production of highly specialised products. The nature of the work of the technology centre requires a sophisticated technical infrastructure including not only the equipment concerned with research and development but also workstations for each of the technical staff.

The Technology Centre

Given the nature of their role in the larger organisation as a service provider, the technology centre places considerable emphasis on the importance of quality control and has invested significant resources is gaining accreditation for its work processes from a range of professional bodies including ISO-9000 certification. This feature of the technology centre is often crucial in securing contracts for the manufacturing division and the organisation wishes to preserve it.

One result of the focus on the work process is that a number of representations and flowcharts currently exist as part of the work of the technology centre. These representations provide a useful starting point for structuring the results of our study given their general acceptance across the organisation. The screen shot in

figure 4 shows the representation of the development and accreditation of a technical report and the archiving of technical reports within the centre. These representations are based on documents used across the organisation which were transposed into our system.

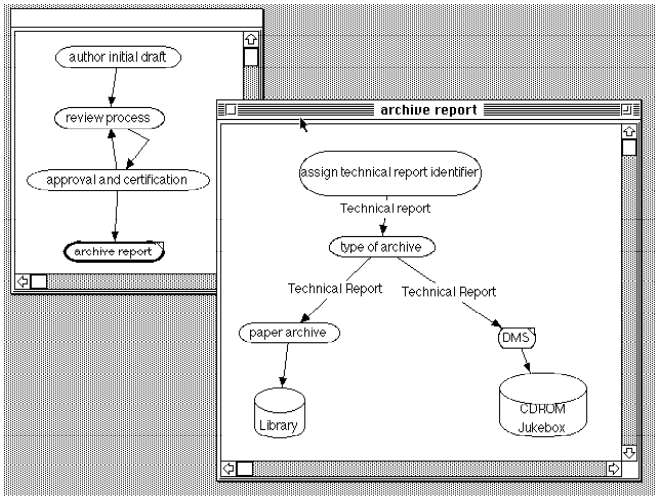

Figure 5.11 : The technical report process.

Our particular focus in reporting the results of the study is on the document management centre. This is the point in the process where documents are converted to an electronic form and stored in a CD-ROM jukebox for later recall. This part of the process was unspecified and our work required us to develop an abstract representation of the work involved. This provided us with a natural focus for this part of the study and allowed us to exploit the simple viewpoint structuring facility shown in figure 5.11. This feature of the structuring tool collects together the information associated with three viewpoints outlined in the previous section and the abstract representation of the process resulting from an examination of this part of the process. In this case we see that we have recorded a selection of different views of work associated with different participants and two major ecology of work viewpoints, one of the DMS (Document Management System), and one of a worker called Sherryl.

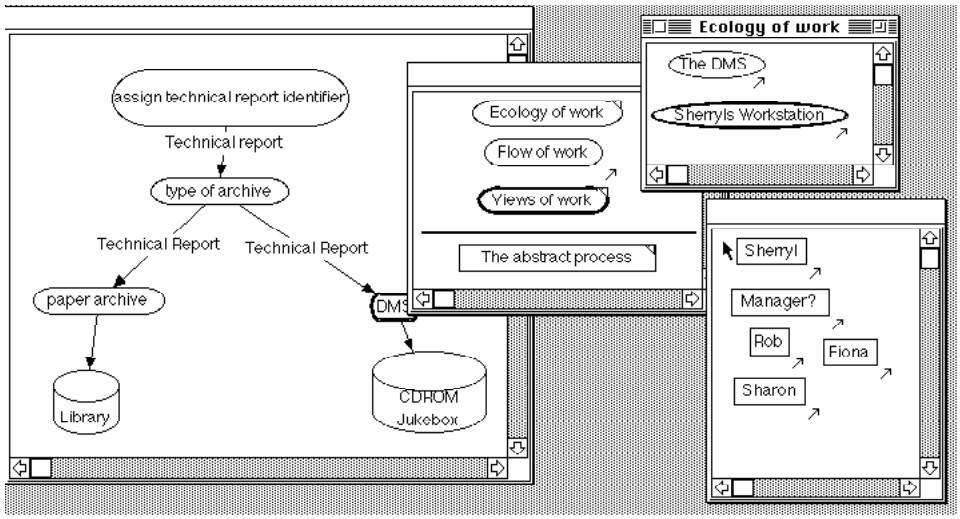

Figure 5.12 : Representing the viewpoints.

The ecology of work is principally represented as a plan diagram of the work setting. This plan representation shows the principal participants and resources involved in the work setting and their physical relationship to each other. The plan is also annotated with a series of notes that are placed in the diagram using 'post it' facilities inherent in the tool. Figure 5.13 shows an example of the ecology of work associated with the document management system and the form of annotations placed on it. The open annotation contains a summary of an interview with one of the participants represent in the ecology of work.

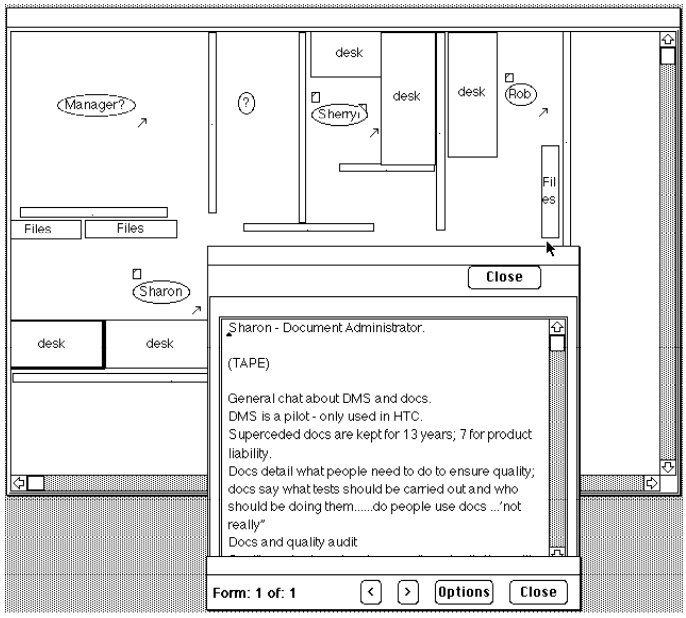

Figure 5.13 : An ecology of work viewpoint.

Each of the views in the viewpoint windows exploits cross reference facilities provided by the tool to provide direct access to comments and interview information that the observer has decided to convey to developers and, in this way, not only allows for the comparison of viewpoints, but also preserves some of

the richness of the fieldwork materials. In the case of figure 5.14, the view of work associated with rob shows the portion of the ecology of work associated with Rob and a part of the ethnographic record which includes a summary of an interview.

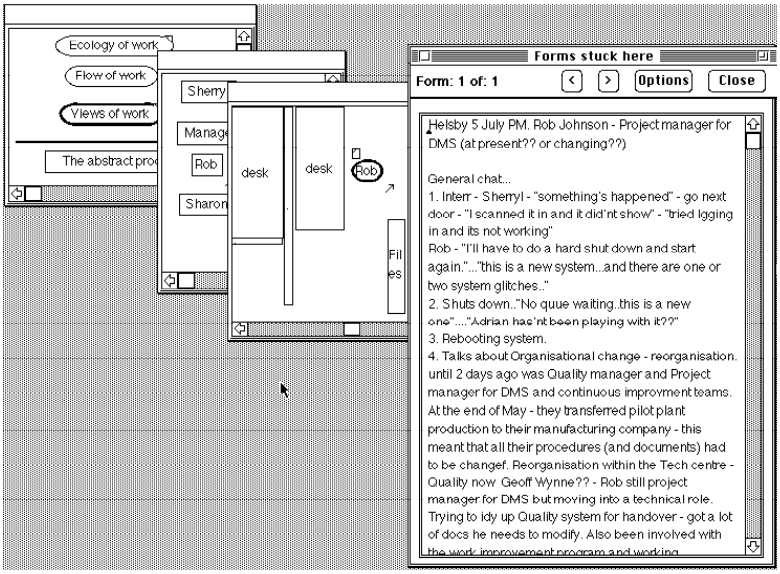

Figure 5.14 : The view of work associated with rob.

Our final viewpoint is more closely connected with the abstract representation of work developed by the requirements engineer and recorded along with the three distinct viewpoints. In our case one of the investigators involved with the research project acted as the requirements engineer in formulating an abstract representation of the work taking place. The flow of work viewpoint (fig 5.15) represents the work taking place in archiving documents at the DMS, this is presented graphically to show the general sequencing of work and the resources used. This viewpoint often provides a starting point for more standardised abstract representations (fig. 5.16). Significant use is made of cross referencing facilities to link elements of both these representations and to associate them with items in the other viewpoints. This is important in order to maintain the idea that this abstract representation is just that; that is, a representation done for the purposes of engineering and which can be cross-checked, validated, and assessed against other viewpoints stored in the system.

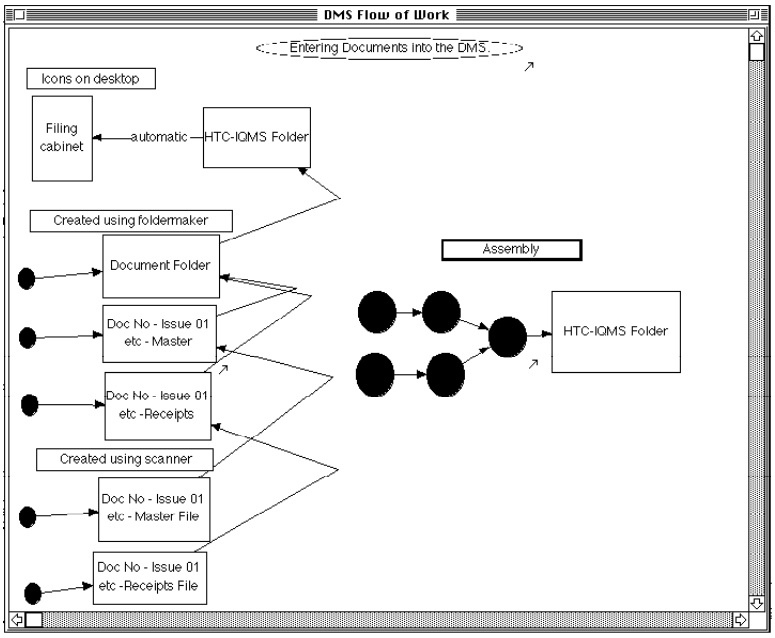

Figure 5.15 : The flow of work viewpoint for the document management centre.

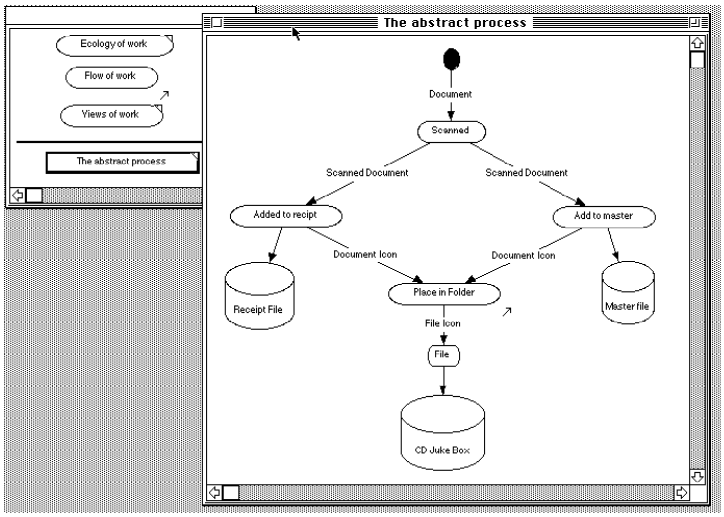

Figure 5.16 : The associated abstract representation of work

Summary

Ethnographic study has a long tradition as a means of understanding the everyday aspects of work settings. However, its use in the development of computer systems is recent and relatively unproven. All of the suggested models of involving ethnography highlight the importance of finding appropriate mechanisms to present the results of ethnographic studies to the developers of requirements. It is essential that systematic techniques are found for conveying this information which are resilient to the demands of the industrial context in

which systems are developed. In this section we have focused on the use of viewpoint techniques as a means of structuring the presentation of the results of study to enable the construction of abstract representations of work.

The traditional role of the private set of fieldworker notes makes information from an ethnographic study difficult to access. We have tackled this by developing more structured techniques for presenting ethnographic information to developers. This allows traditional models of work to draw upon a larger corpus of informal material extracted from an ethnographic study. This informal material would not normally be recorded or remembered in traditional approaches to developing these models.

The use of viewpoints provides us with the communication between observer and developer that is needed for an investigation of different models for using ethnography in requirements engineering. In addition, we feel that meeting the demands of a defined set of viewpoints provides a way of developing a more structured approach to ethnographic studies within the context of requirements engineering. Our intent is to exploit the construction and use of support tools as a means of developing these research themes. Consequently, it is important that we pause and reflect on the nature of the DNP.

In particular, we briefly turn to the problems and limitations encountered in our use of the tool. The logic of such reflection, we stress, is part of the evolution of the tool, highlighting opportunities for its potential development, rather than indicating its redundancy. These can be briefly listed as:

- The DNP cannot display multiple text-notes and sub-designs simultaneously on the screen if from the same entity.
- The DNP cannot search on more than one entity at a time (no multiple searches) nor give simultaneous display of results (one text-notes/sub-design at a time) which limits comparative work.
- Cross-referencing search, after option menu, takes the reader to entities/text-notes (particularly large ones)/designs and sub-designs but not to a specific location within those features: in complex designs this might lead to difficulty in locating the linked entity in a new design.
- The DNP cannot adequately deal with large amounts of textual data in terms of displaying a large amount of text on the screen at any one time.

These points are of relevance to the technical development of the DNP's role as a tool for presenting data to the most flexible of extents, and as such is worthy of note, in terms of, for example, its use in comparing a significant number of viewpoints. In addition, prominence is given to the graphical expression of information. This bias reflects the tool's history and may be seen as problematic if the DNP was to be used solely as a tool to support 'conventional' ethnography. However, we do not perceive DNP as a 'stand-alone' technology replacing the typical practices and tools used by fieldworkers, but rather suggest that it offers significant potential to operate as a tool that complements the ethnographic exercise and supports the endeavour of systems design. Indeed, given that the

target domain for our field studies is systems development, this bias is potentially useful in that it offers a presentation of textual ethnographic data complementary both to existing fieldnotes and the diagrammatic representations underpinning accepted design methods.

Also of importance, is the development of the DNP in use: it is clear, for example that many of the assertions made herein presuppose a DNP 'reader' with a significant understanding of software notation systems. It seems logical to suggest that the development of a more systematic form of representation of viewpoints upon *ethnographic* data would assist in the utilisation of this tool.

We feel that the use of the DNP to organise and present ethnographic data facilitates the discussion across a broad development team of the representational viewpoints emerging from such data. The complex sociality of the work setting - which the system under development must support - as represented in the ethnographic record is available for assessment and interpretation by users of the DNP. The identification and decomposition of viewpoints for the requirements elicitation process can thus be supported by such a representation. Consequently we would argue that the use of the DNP can provide a systematic means of viewpoint identification and representation. This represents an emerging development of a bridge between the complexities of an ethnographic account and the requirements of a system designer, and consequently a first step towards the development of an analytical framework.

Section 6
Concluding Remarks

Section 6
Concluding Remarks

In this Deliverable we have reviewed the research work of Strand 2 during the past year. This has had two general and related concerns:

- the development of a corpus of fieldwork materials and the identification of the emergent features of the social organisation of work;
- the investigation of techniques to support the involvement of ethnographic studies of work in the CSCW systems development process.

Corpus of materials and the emergent properties of the social organisation of work

This element of the research has been primarily concerned with developing a series of case studies of work settings in a variety of field sites and, as a parallel endeavour, building up a corpus of fieldwork materials which are capable of serving as the basis of the development of a framework for the analysis of the social organisation of work.

We have discussed some aspects of the complementary relationship between the research reported here and that being undertaken in other areas of the project particularly in Strands 1 and 3. The research in this Strand is more inductively inclined in developing a framework for the analysis of the social organisation of work using a corpus of field studies taken from COMIC and related CSCW research projects.

Techniques to support CSCW system development

In addition, we have begun particular refinement of the DNP, a tool for representing field study materials in a form which is able to support design work. We have used the tool to exploit the notion of viewpoints as a series of vantage points from which to present field study materials.

These concerns have been focused around the longer term objective, based on the ground work reported in Deliverable D2.1, of integrating ethnography and field studies into the CSCW design process.

The practical use of the DNP has also outlined a series of shortcomings with the DNP. These shortcomings form the basis of a development agenda for the forthcoming year's work. The intent is to incorporate particular facilities to represent and reason about viewpoints on field study materials.

Integrating ethnography in the CSCW design process

The work of this Strand is primarily methodological and, in particular, concerned with exploring and calibrating ethnography as a method for informing the design and development of CSCW systems. As we discussed in Deliverable D2.1, our attitude to many of the issues surrounding the role of sociology and ethnography in design is to approach these in a pragmatic manner in order to situate ethnographic field studies of the social organisation of work within a design process which will necessarily involve the concerns of software engineering as well as giving more attention to issues which go beyond the technology. This we see as one of the distinctive features of the CSCW thrust in system design.

As part of this objective a distinctive feature of the research in this Strand is to come to a better understanding of the ways in which ethnographic field studies can contribute to CSCW design and development outside the research laboratory. That is, coming to understand how and in what ways the method can be incorporated in the world of commercial and industrial design. This has involved an attempt to specify, based on our own experience of using the approach, the variety of roles that ethnographic fieldwork is capable of playing in the design process and, as part of this, the kind of information it is capable of providing during the various stages of that process.

It is no part of our argument that ethnography is a panacea to the problems of system design. We do, however, claim that it is capable of performing an important role in CSCW, given that one of the essential motivations of CSCW is the need to examine the 'real world' social context into which systems will be placed. There are three related aspects to the current work of this Strand that we would particularly like to emphasise as contributions to the above objectives and to CSCW more generally:

- understanding the nature of the social organisation of work;
- coming to terms with problems of representation and communication;
- exploiting the viewpoints notion.

The social organisation of work

As we have pointed out previously, our approach to understanding the social organisation of work is predicated on the presupposition that this needs to begin from a close involvement with work as a 'real world, real time' phenomenon; hence the argument for ethnographic fieldwork and case studies. However, we also acknowledge the need to derive a framework for such investigations which is capable of sensitising design to the social character of work and, importantly, the design issues that are likely to need attending to in order to realise the aspirations of CSCW.

In many respects, the research reported in this Deliverable is but the beginning of the steps toward developing such a framework. The assembly of the corpus of materials is intended to provide a resource for such a development as well as

constituting the basis for a broader understanding of the features of the social organisation of work.

Representation and communication

One of the main objectives of the corpus of field studies is a resource to sensitise design to the social character of work. However, what is also clear is that within the interdisciplinary process of design, such understandings need to be communicated in an appropriate form; that is, by being able to relate the sociological understandings to the concerns of design while, at the same time, preserving the verity of the sociological accounts. We see this as essentially a practical issue though not one which is to be dismissed because of this. The DNP is the means by which we are exploring, using field study materials, the problems of presenting such materials in a form which allows designers, along with fieldworkers, to represent aspects of the work and its settings as a way of setting out design options alongside the fieldwork materials. By supporting both a textual and graphical means of representation, DNP is intended to facilitate communication between fieldworkers and system designers. There are, of course, possibilities in such a tool for supporting larger, and even distributed, design teams and, in this way, make a contribution to 'moving ethnography out of the control room'.

Viewpoints

We exploit this idea in two main ways: first, as a means of acknowledging the different but complementary interests of fieldworkers and system designers; second, as a means of displaying different vantage points from which fieldwork materials can be examined. In other words, our conception of the viewpoints notion is less than highly structured but, nonetheless, it does provide a framework which can be flexibly developed so as to avoid, on the one hand, imposing rigidities on the fieldwork materials yet, on the other, allowing for a disciplined exploration of design ideas.

Future Work

As we have already indicated, this Deliverable is more by way of an interim report ongoing research. A number of the field studies, for example, are continuing and will provide more materials to support both a better understanding of the social organisation of work and how this may be related to CSCW concerns. These will also feed into not only the further work of this Strand but also that of Strands 1 and 3. With regard to Strand 1, the fieldwork materials will enable us to develop a sociological understanding of working life within a variety of organisational settings ranging from clerical work in banks to engineering work within different settings. Such materials, as is shown in Deliverable D3.2, will

also illuminate the notions of mechanisms of interaction and provide a sociological underpinning to this framework.

More particularly, in the coming year we intend to focus on the following:

Handbook of Field Studies for CSCW

As we have said earlier in this Deliverable, the corpus of materials is seen as a resource for the development of a framework for the analysis and presentation of the social organisation of work. We have already sketched out some preliminary formulations of this framework and intend to develop these further in the coming year.

In order that the framework and the corpus of materials is accessible to the design community, we are thinking of producing a handbook of fieldwork materials along with the framework. This will make use not only of the studies reported here but also other ongoing studies related to the COMIC project that we have not been able to use in this Deliverable. These will extend the corpus considerably.

We anticipate this Handbook to be an important means of sensitising CSCW design to the social characteristics of work and its organisation.

Handbook of Fieldwork Method for CSCW

In some respects, this will be the culmination of the methodological research of Strand 2. This will represent a distillation of the relevant research within COMIC and related projects using ethnographic fieldwork methods and an effort to formulate 'best practise' within CSCW design. The objective is not to try to turn system designers into ethnographers, but to provide an accessible way in which they can come to understand the purpose of fieldwork and its findings can be made relevant to design issues.

Once again, this will draw on the corpus of fieldwork studies for illustration as well as the experience and expertise of the fieldworkers and designers involved in other Strands of the COMIC programme.

Also incorporated in the Handbook will be the experiences gained through developing DNP in an effort to ensure the relevance of fieldwork to CSCW design.

Assessment and Evaluation

The work of this strand has a strongly methodological emphasis with a particular focus on incorporating ethnographic studies of work into the development process. The emergence and use of co-operative applications over the last few years has increased the prominence of evaluation in the CSCW community. As the applications developed in COMIC mature over the next year we shall investigate the issues surrounding CSCW evaluation and assessment in practice. This will have two principle themes:

- The nature and role of evaluation and assessment in the development of interactive applications;

- The features of evaluation and assessment to which ethnographic fieldwork methods can contribute.

Work in both these areas is already underway in the project. This is reflected in the two papers included as the appendices to this Deliverable. This will provide the basis for the work of the project as further CSCW systems become mature enough to be assessed within the project.

In the Longer Term

One of the recurrent features of the field studies, and it is manifested in a variety of ways, is the drift within many organisations toward dispersed groups connected by means which are principally electronic in nature. While much of this phenomena is an extension of well known and familiar technologies such as the telephone, the fax and other networked services, it is also clear that much more sophisticated technologies, such as those being examined in Strand 4, are capable of tremendously extending the capabilities of these technologies in ways that are likely to have significant impacts on work and its activities. This move poses a number of challenges for CSCW and for the fieldwork orientation which supports CSCW design.

First, as is shown in a number of the studies, electronic communications are not only capable of enhancing the flexibilities of work but also can involve a great deal of deskilling and the routinisation of activities. The design of appropriate CSCW systems which are capable of exploiting the new technologies of distributed computing will still need to be informed by studies of the social organisation of the activities they are intended to support.

Second, dispersed and electronically connected groups pose a challenge for observational studies which will need to be faced. In particular, it might well require respecifying the local-global distinction as well as trying to understand the differences such systems might make for a number of the features of the social organisation of work we have already identified, such as ecologies of awareness, the use of local knowledge, and many of the other informalities of interaction which constitute such vital ingredients of ‘real time, real world’ work settings.

Finally, the exploitation and use of viewpoints as a means of presenting field studies shows considerable promise. Viewpoints are currently gaining considerable prominence within the software engineering community as a methodological approach to the capture and structuring of requirements. The relation between the structuring and presentation of ethnographic information and the methodological approach suggested by these techniques requires considerable investigation. A balance needs to be struck between a systematic approach to ethnographic studies within requirements and the freedom central to studies of this form. The increased visibility offered by viewpoints allow us to make ethnographic studies more visible within the development process however it is important that this is not achieved at the cost of the methodological rigour which we has proven problematic in the past.

Bibliography

Ackroyd, S. and Hughes, J.A. (1992), *Data Collection in Context,* London, Longmans.

Ackroyd, S., Harper, R., Hughes, J.A., Shapiro, D., and Soothill, K. (1992), *New Technology and Practical Police Work*, Milton Keynes, UK: Open University Press.

Andersen, H.H.K., (1993a), Production of Technical Documentation in a Large Scale International Company, Report Cognitive Systems Group, Risø National Laboratory.

Andersen, H.H.K., (1993b), The Distribution List, Report Cognitive Systems Group, Risø National Laboratory.

Anderson, R.J., Button, G., and Sharrock. W.W. (1993), Supporting The Design Process Within An Organisational Context, *ECSCW '93. Proceedings of the Third European Conference on Computer-Supported Cooperative Work, Milan*, edited by G. d. Michelis, C. Simone and K. Schmidt, Kluwer Academic Publishers, pp. 47–59.

Anderson, R.J., Hughes, J.A., and Sharrock, W.W. (1989), *Working for Profit: The Social Organisation of Calculation in an Entrepreneurial Firm,* Aldershot, UK: Gower.

Anerson, R.J. and Sharrock, W.W. (1993), Can Organisations Afford Knowledge?, *CSCW,* 1, (3), pp. 143-162.

Bannon, L. (1991), 'From human factors to human actors: the role of psychology and human-computer interaction studies in systems design', in J. Greenbaum, M. Kyng (eds), *Design at Work: Cooperative Design of Computer Systems,* Hillsdale, NJ, Lawrence Erlbaum.

Bannon, L. and Schmidt, K. (1991), CSCW: Four Characters in Search of a Context, in *Studies in Computer Suppoerted Cooperative Work: Theory, Practice and Design*, ed. J. Bowers and S. Benford, Amsterdam: North Holland.

Bansler, J.P. and Bødker, K. (1993), A reappraisal of structured analysis: design in an organizational context. *ACM Transactions on Information Systems*, 11(2): 165-193.

Benson, D. and Hughes, J.A. (1991), Evidence and Inference, in Button, G. (ed), *Ethnomethodology and the Human Sciences,* Cambridge, Cambridge University Press.

Bentley, R., Hughes, J., Randall, D., Rodden, T., Sawyer, P. Shapiro, D and Sommerville, I (1992), Ethnographically Informed System Design for Air Traffic Control, in *Proceedings of CSCSW '92*, ed. J. Turner and R. Kraut, 123-129, Oct.31-Nov.4, Toronto, Canada: ACM Press

Bertrand P. and Noyelle, D. (1988) 'Employment and Skills in Financial Services: A Comparison of Banks and Insurance Companies in Five OECD Countries' The Service Industries Journal, Vol 8., No. 1. January 1988. pp 7-18.

Bhaskar, R. (1975), *A Realist Theory of Science,* Leeds, Leeds Books.

Bittner, E. (1974), The concept of organization, in R. Turner (ed), *Ethnomethodology,* Harmondsworth, Penguin.

Böhne, G., Van Den Daele, W., and Krohn, W. (1978), The "Scientification" of Technology, in Krohn, W., Layton, E.T. and Weingart, P. (eds), *The Dynamics of Science and Technology: Social Values, Technical Norms and Scientific Criteria in the Development of Knowledge,* Dordrecht: D. Reidel, 1978, pp. 219-30.

Borstrøm, H., Carstensen, P. and Sørrensen, C. (1994), Two is Fine, Four is a Mess: Reducing Complexity of Articulation Work in Manufacturing, Working Paper, COMIC-Risø-3-9.

Borstrøm, Henrik and Carsten Sørensen (1994b): CAD Models are not Mechanisms of Interaction — Cooperative Aspects of Design For Manufacture. In *Advances in Agile Manufacturing. Integrating Technology, Organization and People. — Fourth International Conference on Human Aspects of Advanced Manufacturing and Hybrid Automation, Manchester, England, July 6–8, 1994*, ed. Paul Kidd and Waldemar Karwowski. Amsterdam: IOS Press, pp. 113-117.

Borstrøm, Henrik and Carsten Sørensen (1994c): The More The Merrier — Reducing Complexity of Articulation Work in Design For Production. In *Issues of Integration in Manufacturing. Proceedings of the ninth IPS Research Seminar, Fuglsø, Denmark*, ed. Jens O. Riis. Institute for Production at Aalborg University, pp. 157—172.

Borstrøm, Henrik, Peter Carstensen, and Carsten Sørensen (1994a): "*Two is Fine, Four is a Mess"— Reducing Complexity in Manufacturing*. Technical Report Risø National Laboratory.

Bowers, J. (1992a), The politics of formalism, in M. Lea (ed), *Contexts of Computer Mediated Communication,* Hassocks, Harvester/Wheatsheaf.

Bowers, J. (1992b), Get Cyborganised! Fragments for a Techno-Ethnography of Human-Machine Associations, Paper presented at conference on Technologies of Representation, University of Warwick, December.

Bowers, J. (1994), Understanding Organisation Performatively, COMIC Working Paper, Man 1-4.

Bowker, G., Gasser, L., Star, S.L., and Turner, W. (1993), Presentation at Workshop on Social Science Research, Technical Systems and Cooperative Work, Apr 8-10, CRNS, Paris.

Brooks, F.P. (1975), *The Mythical Man-Month,* New York, Addison-Wesley.

Burrell, G. and Morgan, G. (1979), *Sociological Paradigms and Organizational Analysis,* London, Heinemann.

Button, G. (ed), (1991), *Ethnomethodology and the Human Sciences,* Cambridge, Cambridge University Press.

Carstensen, Peter (1994): The Bug Form as a Social Mechanism of Interaction. In *Social Mechanisms of Interaction*, ed. Kjeld Schmidt. Lancaster, England: Esprit BRA 6225 COMIC (Deliverable D3.2).

Carstensen, Peter and Carsten Sørensen (1994): The Foss Cases. In *Social Mechanisms of Interaction*, ed. Kjeld Schmidt. Lancaster, England: Esprit BRA 6225 COMIC (Deliverable D3.2).

Clarke, A. and Fujimura, J.H. (1992), Introduction, in A. Clarke and J.H. Fujimura (eds), *The Right Tools for the Job: At Work in Twentieth-Century Life Sciences,* Princeton, Princeton University Press.

Collins, H. M. (1987), Expert systems and the science of knowledge in W. Bjiker, T. Hughes, and T. Pinch (eds), *New Developments in the Social Studies of Technology,* Boston, MIT Press.

Comic Deliverable D2.1, Informing CSCW System Requirements, Department of Computing, Lancaster University. ISBN 0-901800-29-5

Cooper, G. (1992a), Context and its representation, *Interacting with Computers,* 3, pp. 243-252.

Cooper, R. (1992b), Formal Organization as Representation: Remote Control, Displacement and Abbreviation, in M. Reed and M. Hughes (eds), *Rethinking Organization: New Directions in Organization Theory and Analysis,* London, Sage.

Coulter, J. (1982), Remarks on the conceptualisation of social structure, *Philosophy of the Social Sciences,* 12, pp. 33-46.

Curtis, B., Krasner, H., and Iscoe, N. (1988), A field study of the software design process for large systems, *Communications of the ACM,* 31: 1268 — 1289.

Denzin, N. (1978), *The Research Act in Sociology,* London, Butterworth.

Diaper, D. (ed.), *Task Analysis in Human Computer Interaction*, Ellis Harwood.

Durkheim, E. (1964), *The Rules of Sociological Method,* Glencoe, The Free Press.

Fearful, A. (1992), The introduction of information and office technologies: the great divide, Presented at 10th International Aston/UMIST Labour Process Conference, Aston University.

Ferguson, E.S. (1993), *Engineering and the Mind's Eye,* Cambridge, Mass.: MIT Press.

Fielding, N. and Lee, R.M. (1991), *Using Computers in Qualitative Research,* London, Sage.

Finkelstein, A. and Fuks, H. (1989), Multi-party specification, in *5th International Workshop on Software Specification and Design,* Pittsburgh, IEEE Computer Society Press.

Gabriel, Y. (1993), Organisational Nostalgia - Reflections on the Golden Age, in S. Fineman (ed), *Emotion in Organisations,* London, Sage.

Garfinkel, H., Livingstone, E., Lynch, M., Robillard, A., Taka, P. (1989), Respecifying the Natural Sciences as Discovering Sciences of Practical Action: (I & II) Doing So Ethnographically By Administering a Schedule of Contingencies in Discussions with Laboratory Scientists and Hanging Around Their Laboratories, Unpublished manuscript, Department of Sociology, University of California, Los Angeles.

Garfinkel,. H. (1967), *Studies in Ethnomethodology,* Englewood Cliffs, Prentice Hall.

Gibson, J.J. (1979), *The Ecological Approach to Visual Perception,* New York, Houghton Mifflin.

Goguen, J. (1993), Social issues in requirements engineering, in *Proceedings of RE 93: International Symposium on Requirements Engineering*, Jan 4-6, San Diego: IEEE.

Grudin, J. (1990), The Computer Reaches Out: The Historical Continuity of Interface Design. *Proceedings of ACM CHI'90 Conference on Human Factors in Computing Systems. Evolution and Practice in User Interface Engineering*. pp. 261-268.

Grudin, J. (1991), CSCW: The convergence of two development contexts, In *Proceedings of CHI'91*, New Orleans: ACM Press.

Gurswitsch, A. (1964), *The Field of Consciousness,* Pittsburgh, Duquesne University Press.

Gurswitsch, A. (1979), *Human Encounters in the Social World,* Pittsburgh, Duquesne University Press.

Harper, R. and Hughes, J.A. (1992), "What a F-ing System! Send em all to the same place and then expect us to stop em hitting": Making technology work in air traffic control, in *Technology and Working Order: Studies of Work, Interaction and Technology*, ed. G. Button, 127-144, London: Routledge.

Harper, R., Hughes, J.A. and Shapiro, D. (1991), Harmonious Working and CSCW: Computer Technology and Air Traffic Control, in *Studies in Computer Supported Cooperative Work: Theory, Practice and Design*, ed. J. Bowers and S. Benford, Amsterdam, North-Holland.

Harper, R., Hughes, J.A. and Shapiro, D.Z. (1991), Working in harmony: an examination of computer technology technology in air traffic control, in J.M. Bowers and S.D. Benford (eds), *Studies in Computer Supported Coopertaive Work: Theory, Practice and Design, Amsterdam, Elsevier*.

Harper, R., Hughes, J.A., and Shapiro, D.Z. (1989), *The Functionalities of Flight Data Strips,* London: CAA.

Harrington, J. (1984), *Understanding the Manufacturing Process: Key to Successful CAD/CAM Implementation,* New York, Marcel Dekker.

Heath C, Jirotka M., Luff P., Hindmarsh J. (1993) Unpacking Collaboration: The Interactional Organisation of Trading in a City Dealing Room, in de Michelis G., Simone C., Schmidt K. (eds): *Proceedings of ECSCW'93*, Kluwer Academic Publishers, pp 155-171.

Heath, C. and Luff, P. (1992), Collaboration and control: Crisis management and multimedia technology in London Underground line control rooms, *Computer Supported Cooperative Work*, 1(1-2): 69-94.

Helander, M. and Nagamachi, M. (eds) (1992), *Design for Manufacturability: A Systems Approach to Concurrent Engineering and Ergonomics,* London, Taylor and Francis.

Hirsch, P., Michaels, S., and Friedman, R. (1987), "Dirty hands" versus "clean models", *Theory and Society,* 16, pp. 317-336.

Hughes J., King V., Mariani J., Rodden T., M. Twidale, (1993b), Paperwork and its lessons for database systems, *Proceedings of the 12th Schaerding International Workshop on Design Of Computer Supported Cooperative Work And Groupware Systems*, June 1-3 1993

Hughes, J.A. and King, V. (1992a), Paperwork, Working Paper, COMIC-LANCS-4-1.

Hughes, J.A. and King, V. (1992b). Sociology for Large Scale System Design, Paper delivered at CRICT Conference, Software and Systems Practuce: Social Science Perspectives, Reading, Nov. 30th-Dec. 1st.

Hughes, J.A., King, V., Randall, D. and Sharrock, W. (1993a), Ethnography for System Design: A Guide, Working paper, COMIC-LANCS-2-4.

Hughes, J.A., King, V., Rodden, T., and Andersen, H. (1994), Moving out from the control room: Ethnography in system design, Paper presented at CSCW'94, Chapel Hill, North Carolina.

Hughes, J.A., Randall, D. and Shapiro, D. (1992), Faltering from ethnography to design, In *Proceedings of CSCW'92*, Toronto, Canada: ACM

Hughes, J.A., Randall, D. and Shapiro, D. (1993), From Ethnographic Record to System Design: Some Experiences from the field, *CSCW,: An International Journal* , vol. 1 no 2

Hughes, J.A., Shapiro, D.Z., Harper, R., Sharrock, W.W., Anderson, R.J., and Gibbons, S. (1988), *The Automation of Air Traffic Control,* Final Report, SERC/ESRC.

Hughes, J.A., Somerville, I., Bentley, R. and Randall, D. (1993), Designing with ethnography: making work visible, *Interacting with Computers,* 5 (2), pp. 239-53.

Hughes, T.P. (1987), The evolution of large scale technological systems in W. Bijker, T.P. Hughes and T. Pinch (eds), *The Social Construction of Technological Systems,* Cambridge, Ma., MIT Press, pp. 51-82.

Hutchins, E. (1990), The Technology of Team Navigation, in J. Galagher, R.E. Kraut and C. Egido (eds), *Intellectual Teamwork: The Social and Technological Foundations of Cooperative Work,* Hillsdale, Lawrence Erlbaum Associates.

Jirotka, M. and Goguen, J. (eds) (1994), *Requirements Engineering: Social and Technical Issues,* London, Academic Press.

Jordan, K and Lynch, M (1992) The Sociology of a Genetic Engineering Technique: Ritual and Rationality in the Performance of the Plasmid Prep in A.E. Clarke, and J. Fujimura, (eds) *The Right Tools for the Right Job,* Princeton, Princeton University Press.

King, V. and Hughes, J.A. (1994), If All the World's a Stage Are Documents the Script?, COMIC Working Paper, Lancs-4-

Kotonya, G. and Sommerville, I. (1992), Viewpoints for requirements definition, *Software Engineering Journal,* 7(6), pp. 375-87.

Kuutti, K. (1993), Information systems and the development of work activity, COMIC Working Paper, OULU-2-6.

Latané, B. and Darley, J. (1970),. *The Unresponsive Bystander: Why Doesn't He Help?,* New York, Appleton-Century-Croft.

Latour, B. (1987), *Science in Action,* Milton Keynes, Open University Press.

Latour, B. (1990), Drawing things together, in M. Lynch and S. Woolgar (eds), *Representation in Scientific Practice,* Cambridge, Mass.: MIT Press, pp. 19-68.

Law, J. (1986), On the methods of long-distance control: Vessels, navigation and the Portugese route to India, in J. Law (ed), *Power, Action and Belief: A New Sociology of Knowledge,* London, Routledge and Kegan Paul.

Lynch, M. (1990), The externalised retina: Selection and mathematization in the visual documentation of objects in the life sciences, in M. Lynch and S. Woolgar (eds), *Representation in Scientific Practice,* Cambridge, Mass.: MIT Press.

Lynch, M. (1993), *Scientific Practice and Ordinary Action: Ethnomethodology and Social Studies of Science,* Cambridge, Cambridge University Press.

McCarl Nielsen, J. (ed) (1990), *Feminist Research Methods: Exemplary Readings in the Social Sciences,* London, West View Press.

McGrath, J. (1990), Time matters in groups, in J. Gallagher, R.E. Kraut, and C. Edigo (eds), *Intellectual Teamwork: Social and Technological Foundations of Cooperative Work,* Hillsdale, NJ: Erlbaum, pp. 23-62.

Morgan, F. (1986), *Images of Organization,* London, Sage.

Moyse, R. (1991), Multiple Viewpoints Imply Knowledge Negotiation, *Interactive Learning International,* 7, 21-37.

O'Brien, J. and Hughes, J.A., Rodden, T. (1994), Interim Fieldwork Report, VirtuOsi Workpackage 1, Department of Sociology, Lancaster University.

Parlett, M. and Hamilton, D. (1976), *Beyond the Numbers Game,* London, Macmillan

Potts C. (1993) Software Engineering Research Revisited *IEEE Software* Vol. 10 No 5, pp 19-28.

Pycock, J. (1994), Mechanisms of Interaction and Technologies of Representation: Examining a Case Study, COMIC Working Paper, MAN-3-3.

Pycock, J., Calvey, D. Sharrock, W.W., King, V., Hughes, J.A. (1994), Present in the Plan: Process Models and Ethnography, COMIC Working Paper, MAN-1-7.

Quintas. P. (1993) ed. *Social Dimensions of Systems Engineering: People, Processes, Policies and Software Development,* New York, Ellis Harwood.

Randall, D. and Hughes, J.A. (1994), Sociology, CSCW and Working with Customers', in Thomas, P. ed. *Social and Interaction Dimensions of System Design,* Cambridge, Cambridge University Press.

Rittel. H. and Webber, M. (1973), Dilemmas in a general theory of planning, *Policy Sciences*, 4(2): 155-169.

Rodden T., King V, Hughes J., Sommerville I. Supporting the software process as a social process, *Proceedings of EWSPT 94, 3rd European Conference on Software Process Technology,* 7-8 February 1994, Springer Verlag.

Rouncefield, M., Hughes, J.A., Villers, S. and Rodden, T. (1994), Working with Constant Interruption: CSCW and the Small Office, Paper presented at CSCW'94, Chapel Hill, North Carolina.

Sacks, H. (1984), Notes on Methodology, in Atkinson, J.M. and Heritage, J. (eds), *Structures of Social Action: Studies in Conversation Analysis,* Cambridge, Cambridge University Press.

Schmidt, K. (1990) *Analysis of Cooperative Work. A Conceptual Framework*, Risø National Laboratory, DK-4000 Roskilde, Denmark.

Schmidt, K. (1991a), Cooperative work: a conceptual framework, in *Distributed Decision Making: Cognitive Models for Cooperative Work,* ed. J. Rasmussen, B. Brehmer, and J. Leplat, 75-109, Chichester: Wiley

Schmidt, K. (1993) The Organization of Cooperative Work — Beyond the Leviathan' Conception of Organization, in *Issues of Supporting Organizational Context in CSCW Systems*, ed. by L. Bannon and K. Schmidt, COMIC, Esprit Basic Research Project 6225, Computing Department, Lancaster University, Lancaster, U.K., 1993b, pp. 119-139.

Schmidt, K. (1993), The Sociological Bonanza?, Working Paper, Comic Working paper, COMIC-RISØ-2-3, Department of Computing Lancaster University, Lancaster LA1 4YR, UK.

Schmidt, K. and Carstensen, P. (1993), Bridging the Gap: Requirements Analysis for System Design, Working Paper, COMIC-RISØ-2-2.

Schutz, A. (1970), *Reflections on the Problem of Relevance,* The Hague, Martinus Nijhoff, trans. M. Natanson.

Selznick, P. (1948), Foundations of the theory of organization, *American Sociological Review,* 13, pp. 25-35.

Shapiro, D. (1993), Ferrets in a sack? Ethnographic studies and task analysis in CSCW, presented at the *12th Schaerding International Workshop on Design Of Computer Supported Cooperative Work And Groupware Systems*, June 1-3, Elsevier Press

Shapiro, D., Hughes, J.A., Randall, D., Harper, R. (1991), Visual rerepresentation of database information: the flight strip in air traffic control, *Proceedings, 10th Interdisciplinary Workshop on Infomatics and Psychology: Cognitive Aspects of Visual Language and Visual Interfaces,* Schärding, Austria.

Sharrock, W.W. and Anderson, R.J. (1986), *The Ethnomethodologists,* London, Tavistock.

Sharrock, W.W. and Button, G. (1991), The social actor: social action in real time, in Button, G. (ed), *Ethnomethodology and the Human Sciences,* Cambridge, Cambridge University Press.

Sharrock, W.W. and Watson, R. (1988), Autonomy among social theories, in N. Fielding (ed), *Actions and Structures,* London, Sage.

Sørensen, Carsten (1994a): The Augmented Bill of Materials. In *Social Mechanisms of Interaction*, ed. Kjeld Schmidt. Lancaster, England: Esprit BRA 6225 COMIC (Deliverable D3.2).

Sørensen, Carsten (1994b): The Product Classification Scheme. In *Social Mechanisms of Interaction*, ed. Kjeld Schmidt. Lancaster, England: Esprit BRA 6225 COMIC (Deliverable D3.2).

Sørensen, Carsten, Peter Carstensen, and Henrik Borstrøm (1994): We Can't Go On Meeting Like This! Artifacts Making it Easier to Work Together in Manufacturing. In *OZCHI 1994, Melbourne, Australia*, ed. Steve Howard.

Strauss, A. (1985), Work and the Division of Labor, *The Sociological Quarterly*, . 26, (1), pp. 1-19.

Strauss, A. (1988), The Articulation of Project Work: An Organizational Process, *The Sociological Quarterly*, vol. 29 (2), pp. 163-178.

Suchman, L. (1983), Office procedures as practical action: models of work and system design, *ACM Transactions on Office Information Systems*, 1(4): 320-328.

Suchman, L. (1987), *Plans and Situated Actions:The Problem of Human-Machine Communication,* Cambridge, Cambridge University Press.

Suchman, L. and Trigg, R. (1992), Understanding practice: video as a medium for reflection and design, in J. Greenbaum and M. Kyng (eds), *Design at Work,* Hillsdale, Lawrence Erlbaum Associates.

Suchman, L. Office procedures as practical action. *ACM Transactions on Office Information Systems*. 1, 1983, 320-328

Suchman, L., and Wynn, E., (1984), Procedures and problems in the office, *Office Technology and People*, 2: 133-154.

Sudnow, D. (1972) ed. *Studies in Social interaction*, New York, The Free Press.

Vera, A. and Simon, H. (1993), Situated action: a symbolic interpretation, *Cognitive Science,* 17, pp. 7-49.

Wattam, C. (1992), *Making a Case in Child Protection,* London, Longmans.

Weber, M. (1947), *The Theory of Social and Economic Organisation,* New York, The Free Press, trans. A. Henderson and T. Parsons.

Weber, M. (1949), *The M<ethodology of the Social Sciences,* Glencoe, The Free Press.

Weber, M.. (1947), *The Theory of Social and Economic Organisation,* New York, The Free Press.

Yates, J. (1989), *Control through Communication: The Rise of System in American Management,* Baltimore, The Johns Hopkins University Press.

Zimmerman, D. (1970), The practicalities of rule use, in J.D. Douglas (ed), *Understanding Everyday Life,* London, Routledge and Kegan Paul, pp. 221-38.

Appendices

Appendix A
Use, Design, and Evaluation : Steps towards an Integration

Liam J. Bannon

Dept. of Computer Science & Information Systems
University of Limerick, Ireland

This paper argues for a shift in perspective away from thinking of the activities of design, use and evaluation as quite distinct activities, but as activities that are necessarily interleaved and mutually constitutive. Adopting this view has implications for the organisation of design teams, and for the role of 'evaluation' in the larger design process. After some discussion of these points, the paper reviews some recent evaluation studies in the area of CSCW and notes some methodological issues that need to be addressed in evaluation work.

1. Introduction

The argument of this paper, expressed in a nutshell, is that too often the work done under the rubric of design and that done under the label of evaluation is carried out as two completely independent activities, if indeed one finds cases where there has been any evaluation done at all. While this problem is endemic to much of design in general, the focus in this paper is on design and evaluation in the context of software development, focusing even more narrowly on CSCW experiences in this area. The paper is organised as follows: Section 1 provides a perspective on design that emphasises the iterative nature of design and the role that evaluation through use can play in this iteration. Support for such a view is drawn from a variety of sources. Section 2 introduces and critiques the concept of evaluation, arguing for a demythologising of the topic, and for the development of simple but useful tools for evaluation. After this general discussion, we then document the situation with respect to certain CSCW applications in recent years, giving examples of evaluations and their strengths and weaknesses.

2. A perspective on the design process

While empirical studies of the design process - both design generally and software design in particular - are still relatively few and far between, the last few years has seen an increased interest in understanding the nature of this process more fully, in order that it may be supported better through a variety of means. The simplistic idea of designs materialising 'out of the blue' more or less completely formed in the head of a lone designer has given way to understanding the work of a design team who engage in a fine-grained analysis of the work needs of specific people

in a specific context, and the slow and laborious process of developing a design from a tentative first sketch through to a fleshed out design model that can be prototyped, tested and refined in an iterative fashion, before the process becomes fully industrialised. My intent in this paper is not to attempt to study the details of this whole process, but rather to highlight certain aspects of the process which I believe have not been given the attention they are due to date. These aspects are all to do with the areas of use, testing and evaluation, not as terminal stages in some linear design model, but as necessary and interleaved aspects of iterative design. Since I wish to focus in this paper on these aspects of use and evaluation, I do not intend to spend much time arguing for my particular view of the design process over other frameworks, as my view has been developed and builds on a large body of work on information systems development, much of it Scandinavian in origin, which has evolved and been reported on extensively over the past several years, and to which readers are referred for further information, e.g. Ehn (1989), Floyd (1987), Greenbaum & Kyng (1991), Bødker & Grønbæk (1991), Grudin (1991), Bannon & Bødker(1991), Henderson (1991), Henderson & Kyng (1991). Common to these accounts is a concern for the evolution of a design practice that encompasses an understanding of the use situation, of the needs and concerns of the actors involved, and of the importance of an ongoing dialogue with the involved parties through concrete instantiations of design ideas (prototypes) that can be worked with by 'end-users'. Furthermore, there is an acceptance of the fact that user needs are difficult to articulate as system requirements, and that such requirements are at best temporary and local, for as time unfolds changes occur in the use situation and surrounding context that will inevitably affect system requirements. Part of the inherent nature of the design process can indeed be viewed as managing such contingencies.

Thus we should see design, use and evaluation as interleaved and inter penetrating practices, not as distinct steps in a linear development process that moves from analysis through design to implementation then use and, ultimately, evaluation. Indeed use can be seen as the basis for design. This need to reframe our accounts of the design process has become more accepted in recent times - Henderson & Kyng (1991) discuss "continuing design in use", and the notion of a cycle or wheel of design - moving from use -> observation -> analysis -> design -> implementation back to use - has been well articulated by Henderson, (1991)[84]. This alternative conception of the design process allows one to see how design can emerge from an understanding of current (mis)-use of systems, as a spur to re-design. The iterative nature of the process thus focuses attention on the interleaving of design and use as fundamental to design. Observations of use can be conducted in a variety of ways, from the designer's own experiences with a

[84] One can argue for at least two distinct views on this wheel of design. The micro-level focuses on the shorter term design process during which the initial system is built. The macro-level involves a much more long term view, where evaluation is conducted over a number of months if not years, in several settings, and leads to ideas for re-design or new designs. This latter perspective can be viewed as normal application maturation. This distinction, although worthy of further investigation, will not be developed here.

simple model, to experiences with a more fleshed-out prototype, through to more rigorous evaluation of the system prototype in a field study with potential users. The latter activity has often been labelled Evaluation in older flow models of design and occurs at the end of the development process. Changes made as a result of such studies are relatively minor, as the system has been more or less frozen at this stage. Changing perspective to the wheel of design does not imply that formal evaluation studies per se are required at the early stages of design, but that some form of use and evaluation, be it simply with mock-ups or simple prototypes, or even simple tests with storyboards, get carried out from the very beginning of the design process, and are allowed for, nay built into, the whole design process. The depth and extent of these studies of use will of course vary as the design concept develops, but it is the shift in perspective that sees such studies as an integral part of design that I wish to stress, as a perspective for what follows in this paper.

The relevance of some of the points in this paper to the area of CSCW is especially important for a number of reasons. For one thing, as noted by a number of people (e.g. Grudin, 1989) the intuition of designers about useful software for groups is likely to be poor, and thus understanding of the use situation even more important than usual in design, thus leading to the need for early prototyping and feedback from users. For another, the early days of CSCW were characterised by a large number of papers that described design models for 'supporting' various aspects of group activities, embedding a model of group communication or group co-ordination activity that seemed open to question (see Robinson & Bannon, 1991). There was little or no evidence that such models were of any practical relevance to the task at hand, and further, no discussion of such issues as possible testing or evaluation of the designed systems took place. Worse still, when refinements to the initial model were presented at a later date, often these refinements were also based on abstractions rather than on any clear empirical evidence for the relevance of these new features in actual work situations. There appeared to be a total neglect of the need to understand the use situation, and of the need for iterative design, to take the experiences of use of systems seriously.

The question is not simply whether such abstractions were of use or not, but to actually test them and be able to say, based on some empirical evidence, no matter how meagre, that certain features were found to be of use and others not. Even more frustrating to the longer term development of the CSCW area, in certain cases we have seen the abandonment of one form of model for yet another model, without any clear rationale for the switch, other than the designer's whim or fancy. If such changes in direction have come about through some set of experiences in using a prototype, such information would be of great value to our community, as currently the number of reports of use of systems, be they successes or failures, from early mock-ups to full-blown systems, is quite meagre.

3. Another look at Evaluation

While much time and effort has been spent on formally developing evaluation concepts and methodologies, that is not our interest here. We rather focus more pragmatically on how we can improve the quality of our systems. Thus our focus is on evaluations in practice, rather than the nature of Evaluation as a concept - with a capital 'E'. On a simple level, there is an obvious need to know whether or not some designed system does actually perform its intended function. Evaluations are supposed to examine this aspect among others. There are a variety of forms of evaluation: expert investigation, questionnaire surveys, verbal reports, controlled experiments, design reviews, informal observations and formal analysis (See Karat, 1988 for an overview). Evaluations can be formative or summative (Scriven, 1967). Summative evaluations focus on the results that can be achieved with the designed artefact. Formative evaluations are concerned with improving aspects of the design during the design process itself.

Recently, there has been a surge of interest in the use of ethnographic methods as another form of data collection and analysis that might be useful in the systems design process (see Jordan (1995) for some further information on ethnography). While the notion that ethnographic methods could be seen as evaluation tools might seem strange to some, given their qualitative, interpretative nature, there is no doubt but that such accounts can be useful, both in helping to inform the requirements process, and in understanding how people work with or around designed systems. The reason why the term evaluation might seem out of place here is because one normally assumes some metric or set of metrics against which a system is being 'evaluated', and while this is relatively easy in the assessment of hardware or software features of systems (speed and size of memory etc.) the whole issue of exactly what we are measuring against, what our criteria for evaluation are, is much more problematic in the case of measuring usability or utility of systems, from the point of view of end users at different levels in an organisation. Attempts have been made to firm up such metrics, under the rubric of "usability engineering", but such attempts engender their own problems, as noted by some of the pioneers in this particular field (Whiteside, J. Bennett, J., & Holtzblatt, K., 1988).

Given this uncertainty, a careful systematic account of what happens in particular settings when a prototype or system is installed, and how the system is viewed by the people on the ground, can provide useful information for 'evaluating' the system and its fitness for the purpose it was designed. The next issue is how such information can influence the design process, as traditionally the focus of ethnographers has been on understanding a particular setting, and not on re-designing artefacts for that setting. While many ethnographers eschew this issue, but prefer to present their account and leave it to others to try to glean some design implications from the material collected, there are a number of attempts to try to make the results of ethnographic work more usable to the design

community (see Blomberg, J. , J. Giacomi, A. Mosher, & P. Swenton-Wall, 1993).

Much useful insight can be gained about the success or failure of systems simply from the subjective assessments of designers and users, whether noted and collected informally, e.g., in diary reports, or from informal discussions and observations. User-based evaluations can be of a number of forms, and can provide a variety of kinds of useful information, depending on who the users are. Often, users are randomly selected people, perhaps with a level of education expected of intended users of the new interface. Simple empirical studies such as "Talking -Aloud" can provide much useful information on problems with specific features of the interface. Indeed, we can even perform evaluations of hypothetical systems through 'Wizard-of-Oz' techniques, where the potential functionality of a system is simulated by a person rather than actually being built into the software. Such studies can give insight into whether the intended functionality is actually deemed useful by users, in advance of any attempt to actually build software to produce such functionality. Evaluating whole systems becomes more difficult, as first of all it is important that the 'users' (a term that is fraught with problems - see Bannon, 1991) tested with the system are indeed drawn from the community of people who will be operating the new system. Failure to test appropriate users can lead to systematic problems that may not be uncovered until the system is operational (see Whiteside and Wixon, 1987 for an example). Also the time taken for empirical study is often far too brief for it to be possible to answer larger questions of how people may adapt and use the system over longer periods of time in real work situations (For more on the problems of evaluation in HCI, see Bannon, 1991, Bannon & Bødker, 1991, Thomas & Kellogg, 1988).

Given that our designs will inevitably be flawed, the important point is that the results of these designs are tested and the findings used to help in the re-design process. Indeed, in recent years, this realisation of the inevitable need for re-design has become a commonplace - as noted in the oft-quoted dictum : "Design to throw one away - you will anyway". Iterative design, moving from initial design to a prototype, followed by evaluation, and back to re-design has become a buzzword in systems development. Also within HCI, we have seen an increasing concern with the need for appropriate evaluation of systems under the rubric of 'usability'. Despite the existence of the 'human factors' area, which supposedly evaluated the usability and utility of designs, it appears that our methods and techniques were not able to prevent a large number of inappropriate designs coming to the market. This has also pushed the need for methods of early evaluation of prototype systems so that modifications can be fed back into the design process rapidly. Within the HCI community, Landauer (1991) stands out as someone who has consistently argued for the need to move away from narrow, detailed and laborious experimental studies which are of little use to the hard pressed design team keen to make improvements to the prototype design as rapidly as possible, towards more 'lightweight' evaluation techniques, as has Nielsen (1989).

The key point here is to impress on designers the need for early evaluation of systems, whether under the guise of user involvement, rapid prototyping or user testing, or all three combined. Failure to do this can be attributed to a number of factors, some hinging on the personal history of design team members, others that may be structural. For example, it may be the case that the design team has not in the past worked much with users, and thus the necessary skills (often really quite simple) are not available, leaving no person who feels responsible or capable of assuming responsibility for such activities. Such activities do of course cause some disturbance in the life of the design team, (the real world is a bit 'messy' after all) but the rewards can more than justify the disturbance. My view is not necessarily to attempt to handle such problems by adding another design team member who is solely responsible for 'human factors' or for 'doing ethnographies', as this often just increases variance in the design team, but rather to make all the team members aware of the need for studies of use, formal and informal.

So, having argued for the desirability of performing some form of evaluative activity at every step in the design process, I now wish to show, through discussion of a particular case, some of the problems that can occur in the conduct of an evaluative study.

4. Critiquing a CSCW Evaluation Study of THE COORDINATOR

While the argument above has been to the effect that ongoing evaluation can significantly improve the design of an application, it must also be noted that a number of the existing evaluations of prototypes or systems can themselves be critiqued from a variety of standpoints, conceptual, methodological, empirical. In such cases, the import of the published results must be carefully weighed. The old question of who evaluates the evaluators needs to be taken quite seriously. While addressing this thorny problem fully would require a separate paper, I base this section of the paper around a discussion of one particular study, by Carasik and Grantham (1988), which has been cited frequently in the CSCW evaluation literature. The rationale for this is that, in my opinion, this study exemplifies a number of important issues that need to be taken seriously in the interpretation of any evaluation study.

The Coordinator is a commercially available software system that has become one of the most talked about CSCW applications in recent years due to its articulation of a well-developed theory of "language as action" that has exerted considerable influence in the research community (Winograd, 1986) as well as having received enthusiastic commercial endorsement, at least within certain companies (Johnson, B., Weaver, G., Olson, M., & Dunham, R., 1988) who claim that use of the system has increased their productivity enormously. It has attracted both praise and criticism at both a conceptual and pragmatic level. The actual system can be simply described as a fancy electronic-mail-cum-project-

management system[85]. The system is built on the belief that human action is based on conversations, primarily conversations of a particular form, for action. Thus people using the system do not simply send mail, but make requests, or promises, or offer or decline to perform certain activities. (The system does allow for 'free-form' responses, but this choice indicates abdication of the underlying framework on which it is built). Within this framework, the system then keeps track of the commitments made by individuals.

Support for the benefits that are provided by the system comes from a number of sources. In one survey and observation study on groupware systems in organisations, Bullen & Bennett (1990) note how the *ability to link messages* in electronic messaging systems (e.g. in Higgins, The Coordinator, All-in-One) was found to be a very useful facility across the board. Through its concept of a 'conversation', The Coordinator supports such linking quite explicitly. At the same time, there is an at times quite bitter, almost ideological, dispute between different groups as to its explicit design goal, which is to change the way people in organisations think and act. Whether speech act theory itself is an adequate theoretical framework on which to erect any computer-mediated communication system is open to question (Bowers and Churcher, 1988, Levinson, 1983) but the main complaint against the system in use to date has been that it seems to exclude negotiation. Considering the volume and vehemence of rhetoric on the system, there is a relative paucity of case studies, and the evaluations that do exist are rather partial, and of varying quality. Existing evaluations of the system have generated conflicting results (Johnson et al, 1986, Bikson et al., 1988, Bair & Gale, 1988, Carasik & Grantham, 1988, Bullen & Bennett, 1990, Schäl, 1995). We will investigate one of these (Carasik and Grantham, 1988) in a little more detail here, in an effort to point out some of the problems that can occur in the design and execution of an evaluation study[86].

In any evaluation study, it is important to be aware of the original context into which one is 'parachuting in'[87] or introducing the new application, and how it will impact on the selected 'users'. Of course a crucial question is what do we hope to learn from such an evaluation? Is it concerned with testing some highly specific feature of the interface, e.g. a set of menu commands, to see if they are usable, or is it attempting to test the utility of a whole new application for a particular group of people working in a specific environment? In the former case, the selection of 'subjects' for the study may be a lot easier than the latter, where it may be much more difficult to find a suitable test environment. Performing

---

85 For further details of The COORDINATOR, and another interesting case study of its use, see the Chapter by Schäl (this volume).

86 It should be noted here that while the text that follows is quite critical of the aforementioned study in places, the author is well aware of the difficulties encountered by the authors of the study, and indeed has written about some of them from his own experience in a related HCI context some years ago (Bannon & O'Malley, 1984).

87 This term refers to the literal "dropping" of technology onto an unsuspecting populace, and then watching what happens. While such an approach can on occasion be fruitful and provide evidence for completely unexpected uses devised by the populace, in general, the lack of relation between user needs and the technology provided usually results in the total abandonment of the technology.

'evaluations' of systems in inappropriate environments means that we cannot put any weight on the obtained results, even if other more local aspects of the 'evaluation' were conducted carefully. Due to the difficulty of finding appropriate test sites for field study evaluations of systems, a common problem with such kinds of studies are that other environments are found, less appropriate in terms of scientifically testing the case, but infinitely more tractable. If the evaluators are open about the situation, then such studies, although they must be interpreted with caution can still provide information of use, but care must be made in generalising any findings. The Carasik and Grantham (1988) case study is problematic on a number of counts in this regard, and since it has been cited quite widely as an example of how people 'reject' the Coordinator, these problems need to be noted.

An initial problem with the study is in the selection of the test group, colleagues of the authors. This seems somewhat unfortunate, even if convenient for the researchers. In the published case study, it appears that one of the authors took on the role of championing the product, and suggesting to the 'group' in which he was a member that they should all try out the system for a period. The group consisted of 15 professionals who provide consulting in information technology, are located in 2 different locations 35 miles apart, and are frequently travelling, or working at home. The authors openly admit that the people do not work extensively with each other, so to what extent they can be seen as a 'group' other than that they are organised administratively into the same unit is questionable. It is also noted that many of these people currently have access to other email systems, and use a variety of different hardware and software (PROFS, UNIX, PCs, Macs). For the purposes of the trial PCs were made available to all (in order for people to be able to interact with the Coordinator), but it is not noted how many *already* had them on their desks, and how many were encumbered with an additional machine simply for this software trial. The intent of the study was to see if people would find the Coordinator system useful for intragroup communication. Given that, as the authors admit, the organisational division where the trial was carried out does not really constitute a work group of any coherent form, the likelihood of being able to adequately test this particular hypothesis under such conditions does not seem very high.

The actual implementation of the study raises further serious issues. One of the experimenters apparently took on the role of arranging to provide people with the software, and the additional PCs required by some people in order to be able to use it. Within the organisation, the manager of the group expressed support for the study, but did not use system himself, rather having his secretary use it for him. The software vendor gave a 4-hour training class for the whole group on use of the software, but it appears that many users were unfamiliar with PC's and the authors note that many users never really made an effort to use the system. As they note: "some group members used the software extensively, while others simply avoided it". Such a result is not surprising, given the conditions, however, there is no teasing out of the interdependencies between problems with PCs in general and problems with the particular software, so one has no clear way of

separating out the relative contributions of each of these factors to (non-)use. The authors do however note that the new system had no interface to other mail systems, especially a problem where people already use PROFS or UNIX mail, as the Coordinator users' conversation base was restricted to the other people with PCs. But amazingly, no figures are given as to how use or non-use related to previous and current use of other email systems, which would seem to be a crucial factor.

It might be argued that my critique is overly harsh on a study which did at least attempt some form of evaluation of a system in a use context, and this is an area where, as noted earlier, there is a paucity of studies in existence. However, part of my reason for critiquing it in such strong terms is that the paper does not read as an informal anecdotal study of an attempt to experiment with the use of the system in-house, which is what, in my view it actually is, but purports to be more than it in fact is. So, for example, the authors discuss their evaluation in more formal terms, to the extent of discussing their methodology, and the use of 3 "instruments" in order to assess the effects of the "experiment". Let us look at the instruments used in the study and their appropriateness. The first was the use of a semantic differential scale which measures changes in subject's cognition of test items, the second was a network analysis questionnaire, which noted the attitudes of people to other members of the work group, and the communication patterns observed, and the third was a matrix questionnaire looking at the potential substitutability of different forms of communication within the group. While each of these "instruments" may in certain settings, under certain conditions, produce useful information, their relevance in this situation, given the site set-up described earlier , is in my view deeply problematic[88]. Given the informality of the "experiment" and the conflating of a number of issues outlined above, at best the study could report on anecdotal reports of people, and subjective impressions, even still laced with caution due to the extreme limitations of the study. But it is difficult to imagine how such instruments could provide much new insight, given the informal nature of the "experiment".

The authors usefully note the anecdotal reports of their users that are very negative on the "new language" which the system introduces to label communicative acts, on the poor interface etc. It is then noted that only at management urging would people continue to use the system, but there is no detailed logging of use of the system displayed, so one has no idea of what level of use was ever achieved with the system, and by whom. It is noted that after 6 weeks the group decided to discontinue use of the Coordinator, and adopt PROFS instead, citing the following difficulties: lack of language clarity, difficulty in learning the tool, and cross-system compatibility. These appear to be quite sensible reasons for stopping use of the system in this particular setting. To put it from a 'users' perspective, they were asked to use a new piece of software that

88 Indeed, I could go further and claim that they give a spurious conception of scientific rigour to what was by all accounts a very informal study, and should be reported as such. In other words, instruments should be used in appropriate settings, and not applied simply because they are available.

was foisted on them at the behest of one member of staff who was curious about the potential of the new tool, and was backed up, at least initially, by management. It is difficult to see how the system meets any real identified need of this group. The trial also involved for some people having to use a completely different computer system to what they currently used, to communicate with co-workers that they did not need to communicate with extensively in the first place, so is it any wonder they were reluctant to use it?

In sum, what this study is able to say about the general utility or usability of the tested product is very debatable. The lack of quantitative data published from the study makes it difficult to check out the details of the case - in terms of total time for different users on the system, number of messages sent, correlation's between use of the new software and availability of alternatives, etc. The one table produced, which shows the change in the semantic differential before and after the test simply shows that they rated the concept 'Collaborative Work' as more boring after their experiences. The likelihood of such a measure providing any useful information in this kind of informal trial is very low. Indeed the one significant before-after measure could simply be seen as a response by those involved to the hype surrounding the study itself. To be fair to the authors, in the final part of the paper, they do raise some important points about the Coordinator, but these points become attenuated due to the low quality of the empirical evidence that is adduced.

To conclude this Section, my purpose in discussing this study in some detail is twofold. On the one hand, some pitfalls in the conduct of a user software trial have been noted. Concomitantly, I wish to bring to the reader's attention the fact that this study has been cited numerous times in the CSCW literature as empirical evidence for the general 'failure' of the Coordinator in actual work settings. While the paper does have some interesting commentary on the Coordinator, much of which I personally agree with, and provides some informal empirical information, it should be clear by now that it certainly should not be understood as in any way clearly demonstrating that the Coordinator is a failure. Evaluations are important yes, but it is also important to be aware of the quality of the evaluation, and of what can legitimately be learned from any particular study. In the present case, I have tried to show that, based on the evidence presented in the published paper both about the set-up of the study and the results proffered, very little of substance can be concluded about the utility of the Coordinator, even for the particular setting studied.

In the next Section, I wish to briefly outline some other evaluation studies that have been done on certain CSCW systems and note some of their features, positive and negative.

5. A Further Sampling of CSCW Evaluations

The MIT Information Lens Project

The Information Lens system (Malone et al., 1987) has been the subject of a number of research reports and the ideas behind it have now been incorporated into several commercial products. The system is designed to support people in managing their electronic mail. It has at times been referred to as an 'intelligent' information sharing system. The filtering available in Information Lens is designed to screen users from 'junk' mail and cull other messages of interest from a larger set, even if not directly addressed to specific users, thus extending the information sources available to individuals. It provides capabilities for organising mail based on various aspects of the incoming message. It allows users to make message templates of various forms and have rules (of an IF-THEN-ELSE variety) that act selectively on these 'semi-structured' messages. If the sender has selected a *colloquium* form for the mail message, and a message form of type: *colloquium* has been defined by the group, then the sender can be provided with support for composing the message through a partially filled ('semi-structured') message template, and the receivers can make rules that utilise the information that a message is a colloquium announcement to file it appropriately. One can see how this could be quite useful to help put some structure on the myriad of different forms of email communication which at present are insufficiently disambiguated. It helps the sender to structure messages appropriately, and can serve a reminder function for what information is necessary for certain announcements (e.g. to remember to specify the location of a meeting[89]) as well as helping the receiver to sort incoming mail appropriately, rather than have all kinds of messages mixed together in the incoming mail file, as has been the case in most email systems up until the past few years.

In an empirical investigation of the use of the Information Lens system, Mackay et al (1989) summarised their findings as follows: people without significant computer experience can create and use rules; useful rules can be created based on the fields present in all messages *without* special message templates; people use rules both to *prioritise* messages *before* reading them and to *sor*t messages into folders *after* reading; and people use *delete* rules primarily to filter out messages from low-priority distribution lists and not to delete personal messages to themselves. Mackay (1990) also shows the wide variability in patterns of use of the system, though overall there seems little doubt that prototype systems (after some iteration) are being used effectively in work situations.

Undoubtedly some of the ideas embodied (over time, evolving through use) in Information Lens have proved useful in practice. So it seems that people can makeup rules that are useful, but it this does not imply that they can be

89 This might seem a rather trivial example, but the number of email messages that I have received over the years concerning meetings without any mention of a location is not insignificant.

encapsulated into an 'agent' and allowed to be triggered automatically. For instance, one key point noted by Mackay (1990) was how people tended to make up rulesets but then run them manually, i.e. the people themselves determined when to run the ruleset, in particular occasions of use, rather than have it done automatically, according to some pre-specified formula. This supports the notion that it is very difficult for people ahead of time to specify clearly the conditions under which certain rules should be run. Luckily, the technology allowed for the user to manually 'trigger' the rulesets, although this was not part of the initial idea of how the system would be used.

In the field study, Mackay shows the importance of the social environment in affecting the use and development of Lens, with a local expert exerting considerable influence, as well as information sharing going on among the participants - sharing of rules developed by one person and then picked up by others, even without any explicit support for such sharing in the system itself. Mackay shows how quickly people settle into a routine, with changes to rules, addition, deletion, modification, often prompted by outside forces, such as a new version of the system, or a break from routine work, or a visit from the Lens researcher[90]. For many (though not all) users, even in this fairly brief evaluation period, months would go by without further changes to the rulesets. Mackay shows how both individual differences and task differences can have a big effect on the use of the system. People in control of their mail do not see the need to invest in learning a new system, while people swamped already feel taking time out to learn the new system will only make matters worse. An interesting fact is that the initial idea of use was that Lens rules would be automatically run on the incoming mail, yet it turned out people found it useful, and it was possible (through a debugging feature), to run rules on a particular folder. So one user found that this debugging feature allowed him to run rules after he had read his mail. Others liked this idea and began to do the same. This finding subsequently was reflected in the later releases of the technology where there were possibilities of creating multiple rulesets that could be triggered by different events.

As Mackay notes: "Software does not remain static when it is introduced into an organisation. People in the organisation evolve their individual patterns of use, share them with each other, react to external changes, both technical and non-technical, and sometimes pro-actively modify the system to produce significant innovations." In the current discussion, these observations are important, as they show how ongoing evaluation of prototype Lens systems led to new ideas about the very conception of Lens as a tool, which could be factored into the next iteration of the system. In other words, the very conception of what Lens was, and how it was to be used, changed for the design team as they witnessed the way people actually used their prototype system. This provides a very powerful

---

90 This latter case points out how the "experimenter" affects the experiment by his or her very presence. Such effects are inevitable in these kinds of studies, but the important thing is for the experimenter / observer to be aware of this effect, and to discuss how their presence affected their data and interpretations.

example of how important ongoing evaluation studies are as the original design idea becomes articulated and reified in a particular piece of software. On another level, it changes how we think of the design process, and the actors which it encompasses, as we see original ideas about what the tool is and how it can be used coming from the 'users' themselves. It is for this reason that the term 'user testing' is not one I favour, as it places unnecessary limits on the conception of design, and its separation from use, which I have decried earlier in the paper. Users as designers, or at least co-designers, becomes more than simple rhetoric if we shift our perspective in this way.

Xerox PARC CoLab

This project involved building a computerised meeting environment to support small (2 to 6 people) face-to-face meetings. A special room was constructed containing several workstations connected on a local area network. A number of software tools were developed to allow users to jointly work on documents and share the same views on these documents (WYSIWIS - What You See Is What I See). The project has now ended. Stefik et al. (1987) describe the design goals of the project and some of the software tools. This system has been the subject of an interesting evaluation project (Tatar, Foster, & Bobrow 1991), which presents a very thorough analysis of problems in use of the original system, and their possible cause, together with recommendations for improving the design of one of the tools.

Many people in CSCW are of the opinion that CoLab was used and tested extensively by a variety of people over its lifetime. However, while extensive use of the system was made by the designers during the design process, it was only quite late in the project that any more thorough evaluation of the system with outside groups was performed. It is important to note here that both of the groups that were studied were familiar with the underlying interface technology, and also had worked as groups previously - overcoming 2 major methodological problems encountered in several other attempted evaluations of systems. These 2 experiments done by Tatar et al. showed up some serious problems, in that neither of the groups ended up using the shared computational workspace that was the core idea of the project[91]. In one case, the group stopped using the system altogether and resorted to working together with a pad of paper. In the other, they "managed to find a successful way of using the tool by using the video network to look at the screen of whoever was typing, thus employing the shared video workspace instead of the shared computational workspace" (Tatar, Foster, & Bobrow 1991)

What is significant in this evaluation study however, is the persistence of the investigators to provide a well grounded explanation, based on a conversational model of interaction, for why the 2 groups had problems. The major problems of

[91] To be fair, it should be noted that these experiments were of relatively short duration, thus not allowing much time for users to adjust their work practices or evolve new ones to fit the new system.

users had to do with the visibility of certain operations and with problems in reference. Studies showed that people at times had problems in interpreting others comments when their views of the shared world did not match. So, for example, if people re-sized or moved the shared window on the system, peoples' references to spatial locations might not always be appropriate for the other participants. This caused considerable disruption to the work of the group. The focus of work on this project would appear to have been more on examining the technical issues involved in developing software for the real-time computer support of groups than on an understanding of how people could or would use such a system in everyday work activities. To their credit, after enumerating these difficulties they show how they used these findings as a basis for redesign of a number of key aspects of the system. Results of studies on the new system are noted as positive though there is not much detail on them. One large question that arises from this work is whether the original conception of providing separate screens was a good one as separate screens allow for loss of gaze and gesture information that turned out to create problems. Perhaps, tongue-in-cheek, it should be 'Back to the Chalkboard' - in the sense of a common reference - rather than 'Beyond the Chalkboard' ?

DOMINO Office Procedure System

This prototype procedure system is interesting as the underlying model behind it has been extensively described in the literature and more recently evaluated informally by some of the design team (Kreifelts, T., Hinrichs, E., Klein, K-h., Seuffert, P & Woetzel, G., 1991). The system makes a number of assumptions about the nature of office work, and provides 'support' for a number of work activities. A working prototype has been developed and in use in a research organisation, where initial studies of its use have been performed. The initial system model had been the subject of some criticism concerning its view on work activities, but what is interesting is to see what actually happens when in use. While in certain respects, having designers themselves perform the evaluation is open to critique, what we are arguing for here is exactly that members of the design team do try to understand the use of their system as early in the design process as possible, in order to evaluate the effectiveness of otherwise of their proposed model. From the viewpoint of a formal evaluation methodology, their procedure is indeed problematic, but it is precisely such informal studies that I would encourage design teams to engage in as early and as frequently as possible in the design cycle, as but one aspect of ongoing evaluation.

The small internal study of Kreifelts and colleagues shows that, indeed, the system was seen as problematic on the grounds of not allowing for sufficient flexibility, for example allowing necessary informal communication, for lack of integration with other tools - for example electronic mail and spreadsheets. While some of these difficulties could have been predicted without having to empirically test the system, such kinds of informal evaluations of experiences of use can have a powerful effect for the design team itself who see with their own eyes some of

the difficulties experienced by their users. The point is not that such systems have no future, but that we must take seriously the findings that people do not simply "follow procedures" in an office (Suchman, 1983), and thus office support must be very tailorable and flexible if it is to be of practical use to the people doing the work. Subsequent work on the DOMINO system is explicitly taking account of the difficulties experienced by the users of the prototype, and has lead the designers to understand the need for a richer conceptual framework for understanding office activities. Supporting cooperative work with technology requires that we understand a lot about the details of how people together *achieve* a shared understanding, it is not something that passively occurs. The technology must support, or at least not hinder the subtle activities that people engage in to accomplish the apparent orderliness of their work .

In sum, what is of particular interest here is that, while conceptual arguments about the veracity of the initial underlying model of DOMINO had been in existence for some time, it was as a result of a small simple empirical 'evaluation' of the prototype system in a work setting that the design team, on reflection, began to re-conceptualise the system and re-design it.

6. Conclusion

The intent of this appendix was twofold: to re-think the role of observation / evaluation studies in system design, providing a perspective that integrates use, design, and evaluation in the design process; and to explore briefly one or two examples of evaluation studies to note some methodological issues that affect their interpretation. We have attempted to demystify the methods and purposes of evaluations, and have emphasised the need for quick and dirty methods for informal evaluation, quite distinct from more formal studies that may be conducted well after the design has been frozen. Conceiving of design as being part of a larger and inevitable cycle of observing use, developing requirements (formal or informal), designing, building and again observing, allows one to plan from the outset for various forms of evaluation in the design process, from very informal to more formal studies. After each kind of evaluation, care should be taken in interpreting the results of such studies, as we have seen how certain empirical studies have been conducted in a manner which makes them very difficult to interpret. Finally, just as in other aspects of the design process, methods alone will never suffice, as there is always a place for common-sense in their application, and in the interpretation of the results.

Acknowledgments

Thanks to a number of people for commenting on a draft of this paper: Marina Jirotka, Yvonne Rogers, Jonathan Grudin, David Jennings, and anonymous referees. Not all of their suggestions are reflected in the present text, for a variety

of reasons, but their comments and critique are being taken on board for future work.

References

Bair, J. & Gale, S. (1988) An Investigation of the Coordinator as an example of computer supported cooperative work. (unpublished ms.)

Bannon, L. (1991) From Human Factors to Human Actors: The role of psychology and human-computer interaction studies in systems design. In Greenbaum, J. & Kyng, M. (Eds.) (1991) Design at Work.: Cooperative Design of Computer Systems. Hillsdale: Lawrence Erlbaum Associates, pp. 25-44.

Bannon, L. & Bødker, S. (1991) Beyond the Interface: Encountering Artifacts in Use. In J. Carroll (ed.) Designing Interaction: Psychology at the human-computer interface. (Cambridge: Cambridge University Press). 227-253.

Bannon, L. and O'Malley, C. (1984) Problems in Evaluation of Human-Computer Interfaces: A Case Study. In B. Shackel (Ed.) INTERACT 84 - Proceedings of IFIP Conference on Human-Computer Interaction, London, United Kingdom, September, 1984.

Bikson, T. (1988) Panel Discussion: Evaluations of Coordinator. Proceedings of CSCW'88, Portland, Oregon, Sept. 1988.

Blomberg, J. , J. Giacomi, A. Mosher, & P. Swenton-Wall. (1993) Ethnographic Field Methods and Their Relation to Design. In D. Schuler & A. Namioka (Eds.) (1993) Participatory Design: Principles and Practices. Hillsdale, New Jersey: Lawrence Erlbaum Associates. pp123-155.

Bowers, J. & Churcher, J. (1988) Local and global structuring of computer mediated communication: Developing linguistic perspectives on CSCW in COSMOS. In Proceedings CSCW '88, Portland, Oregon, pp 125-139.

Bullen, C. & Bennett, J. (1990) Learning from user experience with groupware. In Proceedings, CSCW '90, October, Los Angeles, CA, pp291-302.

Bødker, S. & Grønbæk, K. (1991) Cooperative Prototyping: Users and designers in mutual activity. In S. Greenberg (ed.) Computer-supported Cooperative Work and Groupware. London: Academic Press, pp. 331-356.

Carasik, R. & Grantham, C. (1988) A case study of CSCW in a dispersed organisation. In Proc. ACM CHI '88, 61-66.

Ehn, P. (1988). Work-oriented design of computer artifacts. Falköping, Sweden: Arbetslivscentrum /Almqvist & Wiksell International.

Floyd, C. (1987) Outline of a paradigm change in software engineering. In Bjerknes, G., Ehn, P. & Kyng, M. (1987) Computers and Democracy: A Scandinavian Challenge. Aldershot, UK: Avebury.

Greenbaum, J. & Kyng, M. (Eds.) (1991) Design at Work.: Cooperative Design of Computer Systems. (Hillsdale,NJ: Lawrence Erlbaum Associates).

Greenberg, S. (Ed.) (1991) Computer-supported Cooperative Work and Groupware. London: Academic Press.

Grudin, Jonathan (1989) Why groupware applications fail: problems in design and evaluation, Office: Technology and People, vol. 4, no. 3, 1989, pp. 245-264.

Grudin, J. (1991) Obstacles to user involvement in software product development, with implications for CSCW. Int. Journal of Man-Machine Studies, 34, 3, pp.435-452.

Henderson, A. & Kyng, M. (1991) There's no place like home: Continuing design in use. In J. Greenbaum & M. Kyng (Eds.) Design at Work: Cooperative Design of Computer Systems (New Jersey: Lawrence Erlbaum Assoc.), 219-240.

Henderson, A. (1991) A development perspective on interface, design and theory. In J. Carroll (ed.) Designing Interaction: Psychology at the human-computer interface. (Cambridge: Cambridge University Press). 254-268.

Johnson, B., Weaver, G., Olson, M., & Dunham, R. (1988) Using a computer-based tool to support collaboration: A field experiment. In Proceedings CSCW '86, Austin, Texas, pp.343-352.

Jordan, B. (1995) Ethnographic Workplace Studies and CSCW. In D. Shapiro, M. Tauber & R. Traunmller (eds), The Design of Computer-Supported Cooperative Work and Groupware Systems, Amsterdam: Elsevier Science, (forthcoming).

Karat, J. Software Evaluation Methodologies. In M. Helander (ed.) Handbook of Human-Computer Interaction. Amsterdam: North- Holland, pp.891-903.

Kreifelts, T., Hinrichs, E., Klein, K-h., Seuffert, P & Woetzel, G. (1991) Experiences with the DOMINO Office Procedure System. In Bannon, L., Robinson, M. & Schmidt, K.(Eds.) Proceedings of the Second European Conference on CSCW - ECSCW'91 (Dordrecht: Kluwer).117-130.

Levinson, S. (1983). Speech Act. Ch. 5 in Pragmatics. Cambridge, UK: Cambridge University Press.

Landauer, T. (1991) Let's get real: A position paper on the role of cognitive psychology in the design of humanly useful and usable systems. Book Chapter in J.M. Carroll (Ed.) (1991) Designing Interaction: Psychology at the Human-Computer Interface, pp.60-73. (New York: Cambridge University Press)

Mackay, W. (1990) Users and Customizable Software: A Co-Adaptive Phenomenon. Doctoral dissertation, Sloan School of Management, MIT.

Mackay, W.E., Malone, T., Crowston, K., Rao, R., Rosenblitt, D. & Card, S. (1989) How do experienced Information Lens users use rules? ACM CHI '89 Proceedings, Austin, Texas, 211-216.

Malone, T. Grant, K., Turbak, F., Brobst, S., & Cohen, M. Intelligent information-sharing systems. Communications of the ACM, Vol.30, No. 5, May, 1987, pp. 390-402.

Nielsen, J. (1989) Usability engineering at a discount. In G. Salvendy & M. Smith (eds.) Designing and Using Human- Computer Interfaces and Knowledge-Based Systems. Amsterdam: North-Holland, pp. 394-401.

Robinson, M. & Bannon, L. (1991) Questioning Representations. In Bannon, L., Robinson, M. & Schmidt, K.(Eds.) Proceedings of the Second European Conference on CSCW - ECSCW'91 (Dordrecht: Kluwer). pp. 219-233.

Schäl, T. (1995) System Design for Cooperative Work in the Language/Action Perspective: A Case Study of The Coordinator in an Italian Company. In D. Shapiro, M. Tauber & R. Traunmller (eds), The Design of Computer-Supported Cooperative Work and Groupware Systems, Amsterdam: Elsevier Science, (forthcoming).

Scriven, M. (1967) The Methodology of Evaluation. In R. Tyler, R. Gagne, & M. Scriven. Perspectives of Curriculum Evaluation. Chicago: Rand McNally, 1967, pp.39-83.

Stefik, M., Foster, G., Bobrow, D., Kahn, K., Lanning, S. & Suchman, L. Beyond the Chalkboard: Computer Support for Collaboration and Problem Solving in Meetings. Communications of the ACM, 30, 1, 32-47, 1987.

Suchman, Lucy A.: Office Procedures as Practical Action: Models of Work and System Design, ACM Transactions on Office Information Systems, vol. 1, no. 4, October 1983, pp. 320-328.

Tatar, D., Foster, G. & Bobrow, D. (1991) Design for Conversation: Lessons from Cognoter. In S. Greenberg (Ed.) Computer-supported Cooperative Work & Groupware. pp 55-79.

Thomas, J. & Kellogg, W. (1989) Minimizing ecological gaps in user interface design. IEEE Software, Jan. 1989, 78-86.

Whiteside, J. Bennett, J., & Holtzblatt, K. (1988). Usability Engineering: Our experience and evolution. In M. Helander, (Ed.), *Handbook of Human-Computer Interaction* (pp. 791-818). Amsterdam: North-Holland.

Whiteside, J. & Wixon, D. (1987) Improving human-computer interaction - a quest for cognitive science (Discussion). In Carroll, J.M. (Ed.) Interfacing Thought: Cognitive Aspects of Human-Computer Interaction. Cambridge: Bradford Press, 1987, pp. 337-352.

Winograd, T. (1986) A Language/Action Approach to the Design of Cooperative Work. In Proceedings CSCW'86, Austin, Texas. Reprinted in Greif (Ed.) (1988)

Appendix B
Situated evaluation for cooperative systems

Michael Twidale
Computing Department
Lancaster University
Lancaster LA1 4YR
UK

David Randall
Department of Interdisciplinary Studies
Manchester Metropolitan University
Manchester M15 6BG
UK

Richard Bentley
GMD FIT
Schloss Birlinghoven
53754 Sankt Augustin
Germany

This paper discusses an evaluation of the MEAD prototype, a multi-user interface generator tool particularly for use in the context of Air Traffic Control (ATC). The procedures we adopted took the form of opportunistic and informal evaluation sessions with small user groups, including Air Traffic Controllers (ATCOs). We argue that informal procedures are a powerful and cost effective method for dealing with specific evaluation issues in the context of CSCW but that wider issues are more problematic. Most notably, identifying the 'validity' or otherwise of CSCW systems requires that the context of use be taken seriously, necessitating a fundamental re-appraisal of the concept of evaluation.

Few practitioners in CSCW would wish to contest the importance of evaluation work. In principle, evaluation should be a significant check of a system's capacity to deliver what is required of it. As Grudin [12] has noted, the evaluation of CSCW systems is especially difficult for a variety of reasons, including the effect on performance of the behaviour and personalities of other group members, the effect of social, motivational, economic and political dynamics and the importance of time, in that group interactions may unfold over days or weeks.

Grudin cites the difficulty of evaluation as just one contributory factor in why CSCW systems fail to deliver the benefits intended. Indeed, we may regard all his case studies as examples of a failure to adequately determine what is being evaluated, when it is appropriate to evaluate, and what methods are likely to prove suitable when the focus of evaluation moves from system functionality to system use [1, 21]. These problems became evident to us through our experience of evaluating the MEAD system [3, 5], a rapid prototyping tool developed as part of a project investigating multi-user interface design in the context of Air Traffic Control [4]. Although, as we recount below, we believe our evaluation studies contributed useful information concerning the effectiveness of MEAD, in the course of the studies we became increasingly conscious of their limitations. Our problems in finding an acceptable basis for 'validating' the system led us to suggest there is a pressing need for the reappraisal of evaluation philosophies and techniques. In particular, the view that evaluation should be regarded principally as a summative process which takes place at a given stage in the software life cycle and which yields 'objective' results is, we believe, deeply problematised by

CSCW's interest in the 'real world' context of use. Rather, we argue there is a need for practitioners to identify common experience of problems in the evaluation of CSCW systems.

Our aim in this paper is to begin this reappraisal process by recounting our own practical experience of evaluation work, primarily in the context of MEAD, but drawing also on other research work in which we have been involved. We aim to identify approaches that resulted in useful information about the performance of MEAD and its relevance to the concerns of those who might use it, and perhaps most importantly, to specify the limits of these approaches. We do so in the hope that we will prompt others to engage in a similar exercise, with the ultimate aim of uncovering evaluation measures which relate directly to CSCW problems.

What is Evaluation?

The issue of evaluation was raised fairly late in our Air Traffic Control (ATC) project. This happened largely because we had presumed, somewhat naively as it turned out, that as such exercises tend historically to occur late in the cycle, it was advisable to allocate specific time towards the end of the project for our own activities. This assumption was profoundly challenged, as we shall see, in the course of the work. Our initial concern was to identify the different ways in which the primary question of 'what is evaluation?' can be answered. Our review suggested rather more competing viewpoints than we had expected, reflecting the differing intellectual backgrounds and histories of those who had been involved in this kind of activity. Some of these are listed in Table 1.

1.	An assessment of the overall effectiveness of a piece of software, ideally yielding a numeric measure by which informed cost-benefit analysis of purchasing decisions can be made.
2.	An assessment of the degree to which the software fulfils its specification in terms of functionality, speed, size or whatever measures were pre-specified.
3.	An assessment of whether the software fulfils the purpose for which it was intended.
4.	An assessment of whether the ideas embodied in the software have been proved to be superior to an alternative, where that alternative is frequently the traditional solution to the problem addressed.
5.	An assessment of whether the money allocated to a research project has been productively used, yielding useful generalisable results.
6.	An assessment of whether the software proves acceptable to the intended end-users.
7.	An assessment of whether end-users *continue* to use it in their normal work.
8.	An assessment of where the software fails to perform as desired or as is now seen to be desirable.
9.	An assessment of the relative importance of the inadequacies of the software.

Table 1: Definitions of evaluation

The problem of competing definitions is of course compounded by the fact that evaluation activities may well be directed at more than one of the above definitions. Adding to the complexity of the evaluation task is the multitude of techniques that can be used. These can be classified in a number of orthogonal dimensions including:

- summative <-> formative
- quantitative <-> qualitative
- controlled experiments <-> ethnographic observations
- formal and rigorous <-> informal and opportunistic

Although it is quite possible to discover evaluations in the literature that fit into each of the possible dimension combinations, in practice we may conflate these dimensions to a single dimension with the two most common styles of summative-quantitative-controlled experimental versus formative-qualitative-opportunistic having an oppositional relationship. For convenience we shall refer to these as formal and informal techniques.

The influence of software engineering can be discerned in at least three of the definitions mentioned, in that definition (2) broadly corresponds to the software engineering concept of verification, definition (3) to validation, and definition (6) to acceptance testing. Amongst the numerous evaluation techniques, it is clear that a scientific paradigm influenced many of them. The influence of the paradigm is apparent in the tendency to regard formal techniques as somehow more 'proper' because the 'proving' of the system renders it 'complete'. Certainly, the summative evaluations characteristic of psychology, education and HCI research belong to the scientific paradigm. Part of computing also belongs to that paradigm in that it draws on these resources, but it equally draws on the paradigm of engineering. That these are different is illustrated by the fact that within the engineering paradigm, and unlike in science, proof by construction is permissible. Thus the historical relationship between engineering approaches and 'scientific' evaluation is not a necessary one.

In our view, both the scientific origins of formal evaluation and the more pragmatic combinations of formal and informal approaches associated with commercial systems development may be inappropriate in a context where a new paradigm is struggling to be recognised [16, 22]. In this respect, the surfacing of doubts concerning the relevance of historically-evolved and discipline-specific procedures for evaluation parallels the concern over requirements analysis which is characteristic of the CSCW community.

Formal techniques tend to be the most frequently advocated in the academic research literature, arguably because some 'fit' between science and engineering is commonly assumed. It would appear that controlled experiments are generally perceived to be more objective and to offer the potential of reproducibility. (Strangely however, unlike in the natural sciences, it seems that evaluation experiments are rarely reproduced.) Associated with their 'scientific' status, experimental techniques can additionally be attractive to developers and their

clients as they offer the promise of risk reduction precisely because they proffer measurements of productivity or effectiveness. On the other hand, there are a number of disadvantages associated with formal techniques, not least their cost in terms of materials, equipment usage, acquisition of subjects, and the labour of the experimenters in planning, administering and analysing the results. A related problem is that the whole activity takes a great deal of time which may not be available within the deadlines of a project. Despite these problems, the immense literature on evaluation in the interlinked fields of education research, psychology and HCI provide only a minority of cases in which more informal techniques have been advocated, mainly for rapid formative evaluations [2, 19, 20].

Our interest in the first instance lay in the possible value of using formative, incomplete, and 'subjective' evaluation as an alternative. We were and are satisfied that they are appropriate, given that our concern was with obtaining information that would allow us to change the system, rather than with definitive information about the system as it stood. Such informal evaluations have been used in the engineering approach of rapid prototyping or iterative design for the development of both general computer systems and CSCW systems [7]. We were conscious that the presumed benefits of objective results do not seem to accrue universally, even when the systems under consideration have not been historically regarded as CSCW systems (see for instance Vincent [31] for a discussion of productivity and office automation), and were interested in what value might be derived from a more informal stance. The fact that scientific testing does not, on the face of it, universally confer the benefits that one might expect was reassuring given our own opportunistic approach, though at the time we were not in a position to understand quite why this might be so.

However, and we stress that at the outset we had little purchase on these issues, radical critiques of evaluation have recently become available. Bannon [1], for instance, argues that design, use and evaluation should be viewed not as distinct activities, but as being necessarily interwoven. Evaluation can be understood as a process which should saturate and be constitutive of the design process precisely because the 'context of use' is central to the analysis of CSCW systems. If one accepts such an argument there are immense implications for the issues we have mentioned above and we draw on this and on the similar arguments in [21] to argue that while informal evaluation procedures have a value for specific purposes, they constitute only modest revisions of the status quo. Identifying the 'completeness' or 'validity' of MEAD turned out to be quite beyond the scope of our chosen methods, largely because we had not at that stage taken the step of imagining what such a 'situated' evaluation might look like, and the analytic purchase that shifting the focus of evaluation from the computer system to the socio-technical whole might bring. Our discussion concludes with an examination of this issue.

Evaluating MEAD

We re-iterate here that these issues were only vaguely glimpsed when we began the process of deciding how to undertake an evaluation. At the outset our problem was to decide which, if any, of the purposes cited in Table 1 were relevant to our concerns, and which of the associated methods, if any, should be adopted. Our stance throughout the exercise was a pragmatic one, largely because it was unclear how the ongoing collaboration between sociologists and computer scientists [4, 17] would affect our evaluation procedures. Nevertheless, we shall argue in the light of our experiences that the 'objectivist' stance associated with validity and acceptance testing in particular, and the attendant mathematisation of technique, is misconceived.

The problems we encountered led us to view informal evaluation techniques as useful given an orientation towards 'incompleteness' rather than 'completeness', although no objective conclusions about system validity or acceptance are possible. This implies that summative strategies, aimed at providing such objective conclusions, are likely to prove inappropriate in domains where the context of use may vary in significant ways, and these domains are likely to be those in which CSCW has an interest. We return to these issues in the latter part of this paper.

At the outset our problem was to decide on appropriate purposes and techniques. We drew substantially on the literature on evaluation in education research, particularly on research into computer based learning systems and Intelligent Tutoring Systems. Many of our concerns had already been debated in this context and a number of approaches proposed [24]. An important example is the distinction between evaluating the power of the software, the effect it has on learners in experimental conditions and the degree to which the software continues to be acceptable in the intended environment. The latter may be characterised as, "do the teachers go on using it after the evaluators have left and it is no longer a novelty for the students?"

As a matter of practical purpose, and given that our project had been from the start concerned with the potential value of ethnographic insights for systems design, we also decided that our methods should at least in part utilise the domain knowledge of the ethnographer. Given the impressionistic and informal nature of ethnographic data, we felt that a collection of informal methods would be suitable, although we had no firm convictions at that point as to the relative merits of the various techniques available. Our reasoning was in part also due to the difficulty of deciding what criteria might be used to determine 'completeness' in a tool used to support rapid prototyping, as our aim was not to prove the completeness of the system, but to identify incompleteness. In other words, we wanted to be able to say something about what the system failed to do, and perhaps what should be done about it. We opted to undertake a formative evaluation with a view to identifying problems with the system, their relative importance, and discovering how such a prototyping tool might be used.

We were conscious that much of our argument elsewhere, and that of others in the CSCW community [26, 15] has been for a naturalistic approach to the study of work, on the grounds that such approaches have much to tell us about the 'situated' nature of the collaboration. By implication, a similar approach to evaluation work would confer many benefits about the actual use of new systems. However, for us, the prospect was infeasible, since it would have been difficult to set up, costly, and time consuming to do, particularly for continuing formative evaluation of an unfinished prototype. Fortunately a characteristic of formative evaluation means that less 'authentic' studies may yield valuable preliminary results. That characteristic may be phrased as, "all problems with a system scale up and out; any success may not", and is particularly true of a system's user interface (Figure 1). So, to take the extreme case, if even a system developer finds a feature of the system clumsy or confusing to use, we can be reasonably confident that the same will be true for a novice end-user.

The converse of this characteristic is obviously untrue and has been noted by HCI evaluators and used as evidence for the importance of evaluation with end-users [9]. Although what such evaluators say is quite true, they miss the point that as part of ongoing formative evaluation, less authentic testing can be a valuable way of eliminating grosser system errors in a more economical manner. If those errors remain until later, not only will the delay mean that they are more difficult (and costly) to rectify, but the gross errors will completely swamp the interaction making it unlikely that more subtle but still significant system errors will be observed at all, as shown in Figure 1. In this regard, our approach to evaluating MEAD was informed by our earlier work on the design of the user interface for the Designers' Notepad - a system to support the early stages of design [29, 30]. Here, initial interface testing was performed by the system developers who undertook authentic tasks such as preparing talks or papers. This revealed some problems, including lack of functionality, which were rectified. Subsequent testing involved end-users increasingly remote from the development group including project members not involved in systems development, postgraduate computing students, visiting academics, colleagues from Sociology, undergraduate computing students, and undergraduates with minimal computer experience. The advantage of this approach is that it does not take large numbers of subjects to reveal problems with a developing system, and so one can discover what is wrong quickly and cheaply.

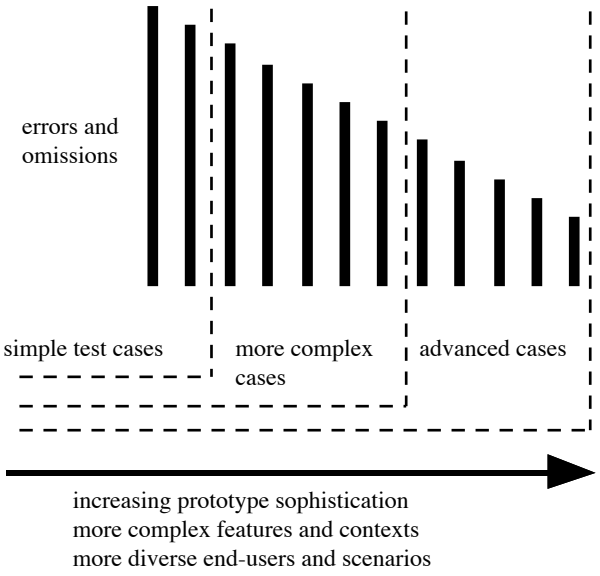

Subtle errors/omissions are unlikely to be detected until less subtle errors are rectified. Simple test cases do not involve some activities and so cannot detect problems concerning them.

Figure 1: Scaling up

Although it was not possible to locate the MEAD prototype in the actual work situation, we made some attempt to maintain authenticity by providing tasks which subjects would regard as authentic tasks to perform. In this manner we hoped to avoid the inherent danger of MEAD's designer embodying his possibly incorrect beliefs about the activities to be supported by the tool - if the subjects are asked to perform tasks which relate to their particular concerns, the tool will be directly challenged as to whether it can support those tasks. Equally, we were persuaded that in depth observation of a small number of individuals would enable us to discover something of the kind of problems they would encounter in using the system, with no suggestion that we might discover all relevant problems.

Our interest lay in revealing the kinds of misconceptions that end-users may have. In the context of a summative evaluation, a possible objection to this would be that the small number of end-users might be unrepresentative. However, in one sense no end-user is unrepresentative in that all end-users' viewpoints and requirements reflect a context in which the system may have to function and, in the first instance, we were concerned to identify some and not all ways in which the system might under-perform or cause problems. In this respect we were influenced by a study reporting problems that arose with Cognoter [27]. This study was very much in the style we chose to adopt, involving authentic tasks and in depth small-scale study (two groups of three each having two 2-hour sessions). Despite the limited size of this study, major problems with the system were discovered and generalised.

A further problem we had to contend with concerned the relationship between MEAD's user interface and underlying features of the software. A poor user interface can distort the performance of any system by confusing or distracting end-users. Likewise, if a system has a particularly good interface, much of the improvement in performance may be due to how that interface supports an end-user, freeing resources to concentrate on the task in hand and acting as a supportive environment, thereby reducing the number of issues to be maintained in working memory [28]. One must therefore be careful to correctly ascribe both advantages and disadvantages of a system, as performance variations may result more from changes to the interface than from underlying software features that purport to support the collaborative nature of a task more effectively. In an ideal world one would have the time and resources to undertake work to determine these relative effects, but this is a luxury available to few issues and projects.

The development of the MEAD multi-user interface prototyping system was informed by a series of ethnographic studies of ATC [4, 14, 25]. These studies fed into the functionality provided to support rapid prototyping of novel ATC interfaces. Hence the evaluation of MEAD was intended to assess the extent to which the system supported the process of such prototyping, and the system's effectiveness was to be determined by the speed and ease with which different interface designs could be implemented and iteratively evaluated and refined. (Note the recursiveness of the evaluation problem in this context: partly we are evaluating how our system supports evaluation!) In the rest of this section we describe the preliminary evaluation of a completed prototype of MEAD.

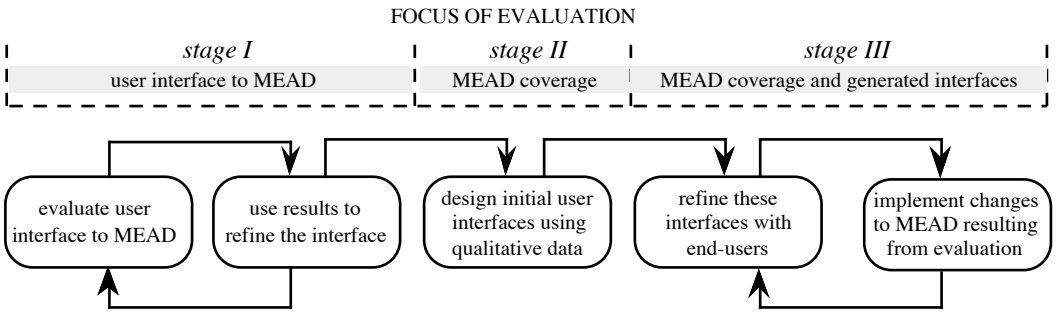

Figure 2: Stages of the evaluation process

What to evaluate

The nature of the MEAD system meant that it could be evaluated in three distinct but interconnected ways:

- *The user interface to MEAD*
  (The ease of learning and use of the tools provided for interface developers)
- *The coverage of MEAD*
  (The ability to generate the kinds of interfaces required by end-users)

- *The user interfaces produced by MEAD*
  (The usability of the interfaces that can be generated)

All three aspects contribute to the system's overall effectiveness and are closely related. However, we first needed to ensure the system did not fail to support end-users' requirements due to the inflexibility of its user interface. Other aspects contributing to effectiveness could only be investigated once it had been established that the tools MEAD provides could be used flexibly. This removal of potential 'swamping effects' of errors in the MEAD user interface was the focus of stage I of our evaluation process (Figure 2). Stage II involved the design of initial user interfaces for ATCOs and in stage III these interfaces were refined with end-user participation.

The evaluation was intended to be formative and opportunistic, merely to discover some of the most visible issues for future work. Even for these purposes, it would be worthwhile to undertake simulations of realistic controlling situations and the environment in which they take place. Such simulations were however beyond the budget and available time of the project. Even for informal ethnographic studies we had to bear in mind the cost and scarcity of end-users (ATCOs). For this reason, stage I of the study involved undergraduate computing students. The scaling up and out argument presented above means that any problems identified and rectified by this study could be found at a cheaper cost than by using an authentic ATC activity.

Stage I: Evaluation of the MEAD user interface

For the purposes of this study, aimed at discovering problems with MEAD's interface, two undergraduate computer science students acted as interface developers. The students sat side by side working at separate machines and were encouraged to discuss problems they were having; it was hoped that this arrangement would establish a "constructive interaction protocol" [11] where the emerging conversation between the students would give information about what they are trying to achieve and problems they are having. To focus attention on MEAD's interface, the session was based on an old course work exercise very familiar to the students - the design of a user interface for a simple Theatre Booking system. Help was only provided when the students could not see a way forward.

Discussion

The students completed the exercise comfortably. The majority of the problems encountered were either with unfamiliar concepts or with aspects such as the wording of button labels and menu entries. The students rarely asked for detailed help about what to do, but rather wanted to know how the various features of the tools could be accessed and manipulated. The study did reveal usability problems with the system, causing it to crash with certain action sequences. This problem arose due to a lack of error checking caused by assumptions made by MEAD's designer about how the system would be used; an example of system errors due to

the designer's implicit assumptions. This simple study shows the relative ease with which many of these errors can be detected by informal testing. Following the study additional error checking was added to MEAD and some menu options and button labels which had caused confusion were re-worded.

Stage II: Design of the initial ATC interfaces

Using the results of social studies of ATC work, interfaces were prototyped with MEAD that offered slightly different facilities to those already in use by ATCOs. Currently ATCOs use two main information displays. Firstly a radar shows current aircraft position as *blips* with *datablocks* next to each blip giving identification numbers and current height. Extra information can be displayed; for example, a *track* of three dots for each aircraft can be displayed indicating its past course, speed and heading. The second display is a set of paper *flight strips* held in racks in front of the controller. Much of ATCOs' work involves positioning and writing on these strips. All controlling decisions, such as instructing aircraft to climb, descend or change heading, are recorded on the strips in ink. This information is then available to different members of a team of controllers. The public visibility of information is vital in controlling work; even during periods of heavy traffic density when flight strips often overflow the bays, they are still positioned where they can easily be seen by all team members.

Electronic flight strip display

While replacing the paper strips with electronic strips 'under glass' would have many advantages [25], it would also have problems such as limited screen space. Some previous attempts to produce electronic strip systems have provided a scrolling mechanism to allow access to strips which cannot fit on the screen (e.g. [23]). Social studies of ATC work have stressed the importance of strip visibility, and scrolling is therefore not suitable for an electronic strip display. Our studies of ATC did however show that most of the information on the strips is provided for use by one member of the controlling team - the ATCO in direct contact with the aircraft or the *radar controller*. This information is not required at all times, so we used MEAD to prototype an initial electronic strip display that allowed the radar controller to 'fold' strips to smaller versions that displayed only the information potentially needed by other team members.

Augmented radar display

The data block on the radar screen shows only a small amount of information about each aircraft, and ATCOs must often relate a radar blip with the aircraft's corresponding strips. As the strips are currently paper, there is no link between them and the electronic radar display, so the task of locating strips involves visually scanning the strip racks. Our studies suggested a need to display information on the radar screen that is currently only presented on the flight strips. Using MEAD an augmented radar was designed that allows blips to be expanded to alternative representations displaying more information.

Stage III: Refinement with ATCOs

Both the electronic strips and the augmented radar were based on existing ATC displays, but were modified in the light of information from social studies of ATC. The next stage of the evaluation was to refine these initial designs with the participation of real end-users. Our approach was similar to that sometimes described as cooperative prototyping [6], where feedback from end-users informs immediate prototype refinement. To support this, the ATCOs were not given information about what MEAD could or could not do so they would not feel constrained in discussing the features they wanted in the evolving interface (and thereby allowing identification of features missing from and inflexibly realised in the existing version of MEAD). For the sessions MEAD's designer acted as the interface developer with access to the system's development environment. (Again the scaling up argument makes this appropriate for finding problems - the next evaluation stage would require someone other than the designer to take this role).

The first ATCO session

The session began with a demonstration of the initial display designs developed in Stage II. The ATCO chose the augmented radar as a starting point based on the flexibility with which information could be added and removed. To explore a novel arrangement of information, the ATCO suggested we expand the augmented radar to present all the dynamic, constantly changing information about aircraft, such as heading, position, coordinated height and so on, while restricting the strips to static, reference information such as aircraft type and expected route. This division would allow investigation of different patterns of sharing, as team members could each have personal displays configured to their own tasks, but could all share the static, reference information display.

The ATCO wanted to keep the track facility from existing radar displays that super-imposes three dots on the screen to show previous aircraft position. The system could not provide this as it requires displaying a function of the position history for each aircraft; there was no way to specify that the position of graphical components on the screen should be based on functions, as in graphical constraint systems such as Rockit [18]. Instead it was decided to try and modify the radar blips to show only the dynamic information required by ATCOs. This revealed more system limitations; an early design decision meant that in order to modify the radar blips a shutdown and re-start process was required. This process took about a minute and if needed several times in a session, was far too long to maintain the interest of end-users as well as being tedious for the interface developer.

During one shutdown/re-start pause, the ATCO discussed a facility provided by some radar displays that automatically indicate if an aircraft is equipped with an on-board Traffic Collision Avoidance System (TCAS). ATCOs find this useful as TCAS can cause aircraft to deviate from their course without ATCO instruction. This facility was added to our display so that a 'T' was shown if a represented aircraft was TCAS-equipped.

Discussion

The shutdown/restart problem arose from a difference in approach to interface development taken in this session to that used by MEAD's designer during system development. During this session a large part of the interface refinement required experimentation with different layouts, sizes, colours and so on - the ATCO was keen to try many different configurations to find the most acceptable, requiring many small changes to the interface definition. This was not the approach used by MEAD's designer for the initial displays, where there was an understanding of what was to be developed from the start, and the definitions were completed before the results were displayed on the screen. A further distinction was the designer's ability to see how a display would appear by examining its definition; the ATCO did not want to look at the definition and was only concerned with how it looked as part of the display on the screen.

It is perhaps surprising that this problem was not discovered before, during system development or when the initial displays were developed. What is more surprising is that the problem was not caught in stage I of the evaluation, which focused on just this sort of problem arising from incorrect assumptions on the designer's part regarding how MEAD would be used. With hindsight, it was obvious that the reason for this lay in the designer's formulation of the theatre booking task performed by the students. This problem is summarised by Twidale [28]:

> There are inevitably implicit assumptions about the nature and style of use (and about the user, task etc.) in the design of the tool. In making up test problems, developers are in danger of incorporating the same assumptions. Thus the study will fail to reveal them.

By trying to simplify and structure the task to be performed by the students, MEAD's designer imposed the same process of interface design as had been assumed during the system's development.

Although the results of this session seemed overly negative, it was still possible to prototype alternative information representations and, in particular, incorporate the TCAS facility described by the ATCO. In the light of these results MEAD was modified to resolve the shutdown/restart problem and it was hoped that the system could now be used in another session to support the rapid prototyping process more effectively and build a more realistic user interface.

The second ATCO session

A different ATCO participated in the second session, which commenced with another demonstration of the initial display designs. Based on the idea of strip 'folding', it was decided to prototype an electronic strip interface similar to one the ATCO had seen during a secondment to the UK ATC development centre. The interface resulting from this session, which closely reflected the ATCO's requirements, is illustrated in Figure 3. These requirements were based on a series of rough sketches quickly made by the ATCO, which were then translated into preliminary strip and display designs and iteratively refined.

A detailed description of the session which resulted in the interface shown in Figure 3 is not appropriate for this paper and interested readers are referred to [3]. It is however worth briefly describing the development of one feature of the interface that is characteristic of the process of cooperative prototyping MEAD supports. A facility was added to the new strip and radar displays that allowed aircraft to be *hooked* by an ATCO, causing their flight strips and radar blips to highlight and removing the need for visual scanning of strip racks to find all an aircraft's strips. Using MEAD it was possible to experiment rapidly with different methods of hooking and highlighting before the final strategy was selected.

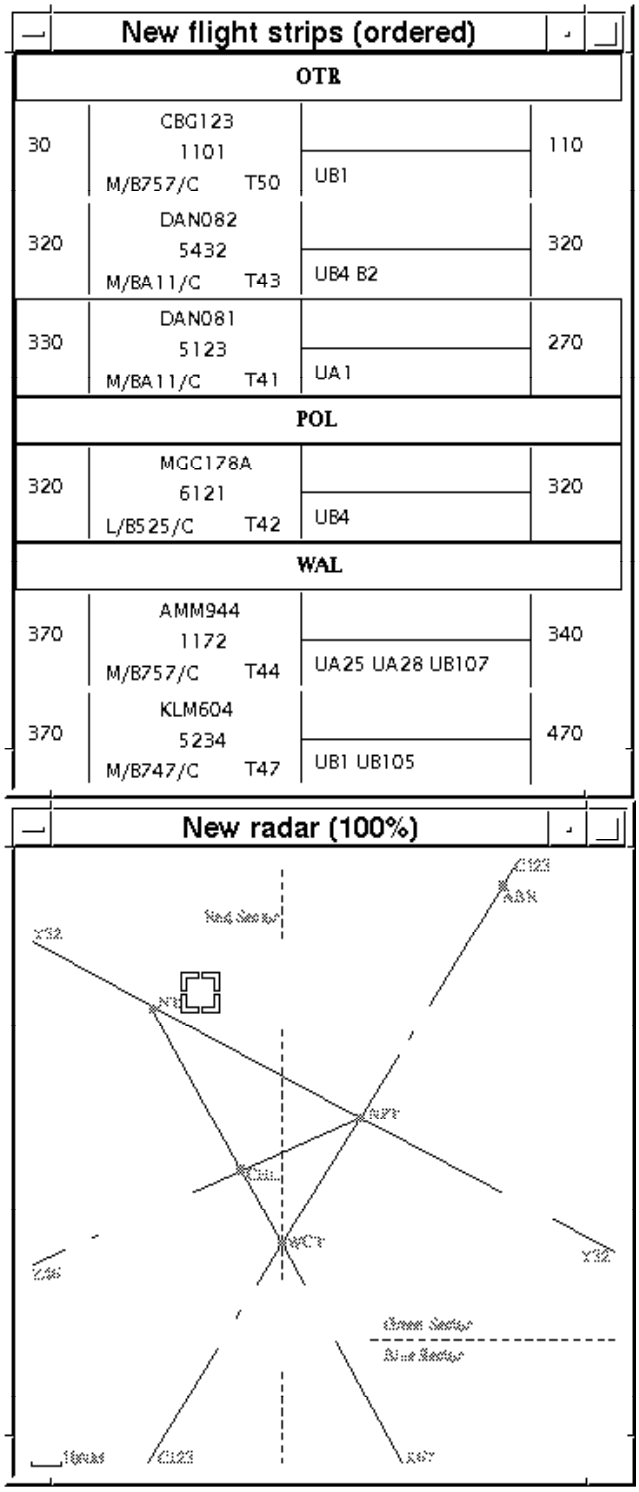

Figure 3: The interface developed with the ATCO.
The aircraft 'DAN082' has been *hooked* and its flight strip and radar blip are highlighted

Discussion

The interface developed during this session closely reflected the ATCO's requirements. It was possible to experiment with different styles of presentation

and interaction until the ATCO was satisfied, and correction of the shutdown/re-start problem encountered in the previous session heightened the sense of end-user participation in the development process. By removing this problem a number of subtler and less severe problems were discovered. For example, the tool MEAD provides for designing information representations such as flight strips and radar blips was lacking in functionality that would have further accelerated the process of display refinement. These smaller errors would not have been discovered if the much grosser effects of the shutdown/re-start problem had not first been eliminated.

Limits of the Evaluation

The evaluation of MEAD was based on a number of small, formative studies of the system in use in limited contexts. In certain respects this decision was justified by the results gleaned from the studies. In particular, the highlighting of specific system features that needed attention is unlikely to have come out of a formal analysis, where overall effectiveness tends to be the focus of the evaluation. Our informal approach revealed the grosser system errors which could be rectified before more time-consuming, costly evaluations were performed. Thus, the initial evaluation with undergraduate computer scientists had little relevance to the issue of whether such a system could be used to generate usable interfaces for ATCOs, but did reveal some significant problems that could be rectified. If these problems had been carried into the ATCO studies they would have made it difficult to assess the ability of MEAD to produce interfaces for ATC. There is little doubt that these exercises fulfilled the function of evaluation definition 8 in Table 1.

During the second session with an ATCO, MEAD was used to generate a realistic prototype of a new user interface. This was iteratively refined using MEAD's development tools with the expert end-user critically assessing each new version. The final version of the user interface prototype incorporated both the required presentation of the information and some behaviour details to allow selection of flight strips and radar blips. By this stage we felt that we had gone some way towards meeting definitions 3 and 6 from Table 1. If nothing else, the study illustrates how informal evaluation techniques can yield a substantial amount of information at relatively little cost and in a short space of time. Further, the overall impression of the system given by end-users was a positive one.

Nevertheless, it was on further reflection and particularly on analysing the transcripts of our discussions with ATCOs that we began to realise that, while our evaluations had provided useful and usable results, our procedures had been very limited in their scope. We were satisfied that we had identified ways in which MEAD was incomplete, and that in these circumscribed settings it was capable of doing useful work. Indeed, discussions with ATCOs who were actively involved in the development of new ATC interfaces gave further support to the idea that MEAD had potential value for ATCOs as a rapid prototyping tool. Thus:

ETHNOGRAPHER: I was quite interested in what you said, about how that would have been extremely useful to you right at the start, and you termed this two phase prototyping, when you play and try out all the things you might want to do ...

ATCO: Yeah ... 'cause when you've not seen anything before and you have people with lots of imagination and you come up with hundreds of ideas, you've got to reject a lot of those ideas unless you can look ... you can't afford to put them onto a large scale simulator 'cause of time and cost .... you lose ideas before they ever get off the drawing board ... something like this would have been ideal for us because it would have helped the deadline as well because we have always been up against that ... and we never managed to get ahead ... it was always 'go firm' too quickly ... behind you is coming the people who are going to write the reqs ... they've got to know what you want ... you've got to make the decisions before you're ready to ... we didn't have anything like this ... a system was bought which took hours to change a track data block from green say to yellow ... to me as someone who's been involved in simulations for some time now you want to be able to say to the programmer 'well, we'll take a quarter of an hour break now' and just change the colour and we'll start again ... you can move along fairly quickly .... what we were into was, you write out a form and it would be changed days or even weeks later ... totally unacceptable at that stage ... that comes later when you're trying to stabilise ... you don't need to write it down 'cause you're not affecting the software ... you don't need to trace what you've done ... you're using a tool which is there to be changed ....

However, the nature of our evaluation method necessarily focused attention on the problems with MEAD rather than conclusively demonstrating its effectiveness. It was, for instance, clear that the exercises in no way enabled us to assess continued acceptance of the system, or whether in fact the interfaces that the tool was capable of generating would prove generally acceptable to ATCOs in the course of their work. This has caused us to consider the problems evaluating systems-in-use might present, and aroused our interest in the idea of 'situated' evaluation.

Our discussions with the ATCOs provided many oblique references to the kinds of problems we were interested in. For example, it became increasingly clear that although there had been a series of trials of proposed interfaces for actual use in the ATC environment, little confidence could be expressed about their acceptability to the ATCOs. One of the most significant aspects of this was that the proposed interfaces were designed for use in a context where the work itself was likely to change. Hence:

ATCO: we can see crew chiefs and assistants being pushed aside and ultimately controllers as well ... and you haven't got the assistants to do the writing and say there's something to work on when there's no paper strips ... I suspect they'll just say well you close the system down but what happens while you're actually doing it I'm not so sure ... that's the message I get anyway ... and aircraft can't park ... there's something else I want to say about teamwork actually ... maybe I've given the impression that teamwork is dead ... one must bear in mind that *[ATC]* does involve teamwork ... the tactical and planning controllers ... but they each have their own jobs to do and we've been alluding to something else ...

In other words, the 'validity' of the proposed interface might depend at least in part on the success or failure of new working practices. Thus 'acceptability'

seems to be intricately bound up not only with the functionalities of the system but also with the potential re-design of work practice. This implies that a 'situated' evaluation would need to address not only the capacity of the 'system', but the flow of work around it.

On the face of it, such a focus seems to have much in common with currently deployed methods of organisational change, including Workflow management, Business Process Re-engineering (BPR) [8], and Total Quality Management (TQM) [10]. However, leaving aside the actual or potential value of such methods, we believe that 'situated' evaluations which deal with work and systems are nevertheless quite distinct from them. Space precludes close examination of the foundations of these managerial philosophies, but we believe they share concepts of 'process', and implicit or explicit mechanisms for monitoring 'process', which ignore the very contingencies, interruptions, and problems which arise in use, and which a 'situated' evaluation might identify. In this vein, the emphasis on teams in the TQM literature apparently parallels CSCW interest in teamwork. Our work on ATC has already highlighted teamwork as an important component of effectiveness, and the following extract re-affirms that importance:

> ATCO: teamwork shifts aeroplanes basically ... it all flows because of people who know when to point the finger at the right time and make the odd comment ... and it all flows very well ... but you cannot teach that ... it relies on something which you can't write down ... you can't design an Air Traffic Control system around it ... or at least to do it ... but you can write it out so its not necessary ... except when you've got that big black cloud coming across, or whatever ... well is it a human problem or a system problem when you take the flexibility ... the flair .. out? You still need to train ATCOs for when things go wrong ...

Several points can be made about this remark. Firstly it suggests the elusiveness of teamwork is such that the specification of teams associated with TQM might be improved by a better understanding of teamwork and the use of technology. Secondly, and crucially for evaluation processes, a continuing uncertainty about the relationship between prospective teamwork and system elements is also evident, despite the fact that the design and implementation of new ATC systems are well under way and due for introduction in 1996. This has to do with a whole gamut of imponderables including training, transition time, responses to new working arrangements, and individual and group capacity to learn the new system. Thus, where we were confident that MEAD was a usable system for generating interfaces that were envisaged by ATCOs on the prototyping team, there was no means to determine whether in the end a 'workable' system would result in the absence of detailed knowledge about such factors. If, as we believe, ultimate decisions about the appropriate interface for use in the real working environment depend for their validity on a complex interrelationship between system functionality and the social organisation of work, evaluation work will need to contend with the intricacies of change and change management. The 'situated' evaluation of CSCW systems should, we feel, be oriented towards the reduction of uncertainty through examination of these areas.

The total elimination of uncertainty, however, is not possible. A striking feature of our conversations with ATCOs was the considerable uncertainty about the outcome of the introduction of new technology and practices. This was felt even by those people who knew the workings of the current system from the inside. For example, during our discussions the issue of the accuracy of Flight Plans (filed by the airlines prior to aircraft take-off) and the flexibility of the current system was raised:

> ETHNOGRAPHER: but presumably you'll still get input errors ...
>
> ATCO: well they'll be corrected centrally and from that point on they'll be correct for the area they're going to fly through ... in theory all flight plans should be correct ...
>
> ETHNOGRAPHER: is that plausible ...
>
> gap ...
>
> ATCO: it is ... but what it does is take the flexibility out the system ... the flow controller will do it all .... at the moment we can say Bristol looks busy, we can say right let's reroute ... I don't know if we'll retain that .... if we lose that we're in deep trouble ... it is possible I suppose that every flight plan will be correct .... there aren't actually that many major errors in flight plans, not as many as there used to be ... the wings are less useful than they used to be but its exactly because they're losing their skills ... I mean it'll happen with the ATCOs down the line somewhere and it comes back to the human or system problem ... whatever the problem is in your less than perfect system, a big black cloud or whatever, do you train them specifically for those situations, whereas now you don't have to 'cause with the expertise we've got we can cope ...

That is, technological and procedural changes have already considerably reduced the level of error in Flight Plans with a concomitant reduction in the usefulness of the assistants or 'wings'. What cannot be known with any degree of certainty is what the consequences of errors, when and if they arise, will be in a situation where experience of dealing with error has been attenuated. It is not as if the evaluation methods we are advocating unequivocally resolve problems of this nature, for we do not believe they can. They do, however, draw attention to their existence as problems and as such constitute resources for decision making which are not available from process-driven models.

Conclusions

Of course systems have to be evaluated, and we are not trying to suggest otherwise. We are confident that our informal procedures were a powerful, cost effective means of evaluation, given a constrained view of what that process should involve. Anecdotal and impressionistic evidence of specific cooperation activities and breakdowns in the use of specific systems in specific situations can contribute to knowledge about the general issues of collaborative working and the potential role of new systems. The procedures were cheap, relative both to more formal evaluations and to the costs of software development and were capable of informing design intuitions. The number, complexity and varying importance of the design decisions which are made during development of a computer system

suggest there is a case for specific evaluation procedures targeted at particular kinds of problem, and in effect this is what we did with our small-scale evaluation studies.

However, we are aware that these procedures alone did not begin to address the issue of the situated character of work. We came slowly to the view that our discussions with ATCOs raised issues of very considerable importance for evaluation. Not least the uncertainties discussed above - articulated in the course of discussion between ATCOs and an ethnographer with considerable experience of their work - which concerned working with systems rather than systems in isolation. The salient issue, both in principle and in practice, should be how we determine whether systems can be said to 'work' or not.

The issue of 'validation' is very much the focus of our argument here. The fundamental problem lies in the combination of assumptions concerning purpose, timeliness, and method that typically (and our evaluation of MEAD is no exception) surround the evaluation process. More specifically, the idea that evaluation should occur late in the development process, should be concerned with machine or software functionality, and should concern itself with 'objective' results, sits strangely with the concern for the social organisation of work that characterises CSCW enquiry. We were led to question whether systems for use in cooperative work environments can indeed be evaluated for validity in isolation from the work. Significant doubt must be cast on the notion that we can 'validate' a system at a given point in the project if we accept that the use of systems is not completely determined by the functionalities designed into them. There may in principle be a vast range of reasons why usage may vary even within a single organisation. Systems put in place may initially fail because they do not resonate with existing practices. Training failures, the prevalence of 'fear and loathing', the breakdown of new organisational processes and so on may all impinge on the speed with which systems become 'usable'. Equally, tried and trusted systems may begin to fail as changes in the environment begin to impact upon them.

To some extent, and we have many reservations, the relationship between the design of technologies and the design of work is addressed in the BPR literature. As Davenport [8] puts it:

> The term process innovation encompasses the envisioning of new work strategies, the actual process design activity, and the implementation of change in all its complex technological, human, and organisational dimensions .... it implies a strong emphasis on how work is done within an organisation, in contrast to a product focus's emphasis on what.

That is, Information Technology design and organisational change are inextricably linked to one another. Systems innovation must associate with a set of assumptions about how 'improving work' is constituted. This shift in emphasis away from the system as technical artifact towards the system at work has a number of important implications for the evaluation process, implications which place ethnographic insights of the kind that the ATC project has utilised at its centre, at least if system 'validity' or 'acceptance' is the problem being addressed. In order to gain a purchase on 'validation', evaluation work will need to focus

increasingly on the examination of the relationship between the system, existing work and organisational practices, and the re-design of both.

In CSCW, this places evaluation at the core of all the design activities which normally precede it, but which we argue is better conceived of as surrounding it. This is no easy matter. As Davenport points out, changes in work activity may take years to manifest, and the impact may not, even if apparent, be straightforwardly measurable. If true, it indicates that evaluation must be extended not only into the whole of the conventional design process, but also well into the system's useful life. That is to say, evaluation work will have to be conceived of not as something separate from other stages in the design process but as a necessary feature of all design work. Further, substantial re-conceptualisation of the notion of the 'system' and its boundaries will be necessary if we are to be serious in our attempts to evaluate use.

With the admitted benefit of hindsight, we came to feel that all of the ethnographic work undertaken during the course of the project can and should be regarded as ongoing evaluation, proving useful in various ways at different stages of the design of MEAD, and in principle in systems development at large. In other words, there are good reasons for regarding evaluation as extending well beyond the point of closure associated with both formal and informal techniques. Indeed we would say that it should be undertaken throughout the entire design process and well beyond. However, the relationship between these activities and more specifically located evaluations needs clarification. Hence there is a need to devise mechanisms which both identify the specific evaluative purposes of ethnographic enquiry, its timeliness, and mechanisms to enable us to cope as well as possible with parallel activities.

We began our evaluation procedures with a considerable cynicism concerning the use of controlled experiments which, we felt, could not take the situated nature of work into account. Such experiments take place in a laboratory, involve ingeniously designed, but consequently artificial and de-contextualised tasks, and due to the usual constraints on obtaining subjects, are generally short term and specifically located [12]. We felt at that time that informal and opportunistic evaluation work might yield useful results without the need for the baggage carried by formal procedures. We came to the realisation that whilst our methods presented us with considerable information concerning 'incomplete' aspects of the system, they gave us no purchase on what 'completeness' would look like. The highly dynamic and variable nature of much co-operative work and the organisational context in which it takes place means that both formal and informal methods are subject to limitations imposed by the assumptions that inform them, and led us to conclude that there is an urgent need for triangulation [13] of the results, scope and limitations of evaluation activity in the context of CSCW enquiry. We would encourage other researchers to report their findings in different situations, so that a growing set of case histories can provide us with a means to assess the lessons, value and purpose of evaluation techniques which relate the system to its use.

Acknowledgements

Our thanks are due to our partners in the ATC project and colleagues at Lancaster, notably John Hughes, Dan Shapiro, Pete Sawyer, Ian Sommerville and Tom Rodden. We would also like to record our gratitude to Jacqui Forsyth for helping us turn the wildly idiosyncratic styles of three individuals into a single, grammatical and readable text.

References

1. Bannon, L., Use, design and evaluation: steps towards an integration, in Proceedings of the International Workshop on the Design of CSCW and Groupware systems, to appear in D. Shapiro et al. (eds), The Design of Computer Supported Cooperative Work and Groupware Systems, Elsevier Science, Amsterdam, 1994.
2. Barnard, P., Applied cognitive psychology: research for human-computer interaction, in Engineering the Human-Computer Interface, A. Downton (ed), McGraw-Hill, Maidenhead, 1991, pp. 28-61.
3. Bentley, R., Supporting multi-user interface development for cooperative systems, unpublished PhD thesis, January 1994, available as a technical report, Computing Department, Lancaster University.
4. Bentley, R., Hughes, J. A., Randall, D., Rodden, T., Sawyer, P., Shapiro, D. & Sommerville, I., Ethnographically-informed system design for air traffic control, in Proceedings of CSCW'92, Toronto, ACM Press, Nov. 1992, pp. 123-129.
5. Bentley, R., Rodden, T., Sawyer, P. & Sommerville, I., An architecture for tailoring cooperative multi-user displays, in Proceedings of CSCW'92, Toronto, ACM Press, Nov. 1992, pp. 187-194.
6. Bodker, S. & Gronbaek, K., Cooperative prototyping: users and designers in mutual activity, in Int. J. of Man-Machine Studies, 34, 1991, pp. 452-478.
7. Cool, C., Fish, R. S., Kraut, R. E., & Lowery, C. M., Iterative design of video communication systems. in Proceedings of CSCW'92, Toronto, ACM Press, Nov. 1992, pp. 25-32.
8. Davenport, T., Process Innovation, Ernst and Young, London, 1993.
9. Downton, A., Evaluation techniques for human-computer systems design, in Engineering the Human-Computer Interface, A. Downton (ed), McGraw-Hill, Maidenhead, 1991, pp. 325-355.
10. Feigenbaum, A. V., Total Quality Control (3rd edition), McGraw-Hill, New York, 1991.
11. Draper, S. & Wately, D., Practical methods for measuring the performance of human-computer interfaces, notes accompanying talk at the JCI HCI Summer School, Queen Mary College, Aug. 1991.
12. Grudin, J., Why CSCW applications fail: problems in the design and evaluation of organisational interfaces, in Proceedings of CSCW'88, Portland, Sept. 1988, pp 85-93.
13. Hammersley, M. & Atkinson, P., Ethnography Principles in Practice, Routledge, London, 1983.
14. Harper, R., Hughes, J. A. & Shapiro, D., Harmonious working and CSCW: computer technology and air traffic control, in Studies in Computer Supported Cooperative Work: Theory, Practice and Design, J. Bowers and S. Benford (eds), North Holland, Amsterdam, 1991, pp 225-234.
15. Heath, C. & Luff, P., Collaborative activity and technological design: task coordination in London underground control rooms, in Proceedings of ECSCW'91, Amsterdam, 1991, pp. 65-80.

16. Hughes, J. A., Shapiro, D. & Randall, D., CSCW: discipline or paradigm?, in Proceedings of ECSCW'91, Amsterdam, 1991, pp. 309-323.
17. Hughes, J. A., Randall, D. & Shapiro, D., Faltering from ethnography to design, in Proceedings of CSCW'92, Toronto, ACM Press, Nov. 1992, pp. 115-122.
18. Karsenty, S., Landay, J. & Weikart, C., Inferring graphical constraints with Rockit, in Proceedings of HCI'92, A. Monk et al. (eds), York, 1992, pp. 137-153.
19. Landauer, T., Relations between cognitive psychology and computer systems design, in Interfacing Thought, J. M. Carroll (ed), Cambridge, MIT Press, 1987.
20. Nielsen, J., Usability engineering at a discount, in Designing and Using Human Computer Interfaces and Knowledge Based Systems, G. Smith & M. Salvendy (eds), Amsterdam, North-Holland, 1989, pp. 394-401.
21. Randall, D., Perspectives on evaluation: putting your money where your mouth is, paper to the UK CSCW SIG, DTI, London, September, 1993 (available from the author).
22. Randall, D., Hughes, J. A., & Shapiro, D., Systems development - the fourth dimension: perspectives on the social organisation of work, in The Social Dimensions of Software Engineering, P. Quintas (ed), Ellis Horwood, London and New York, 1993.
23. RSRE, Prototype Electronic Flight Strips Video, Malvern, 1991.
24. Self, J., Special Issue on Evaluation, Journal of Artificial Intelligence and Education, 4 (2/3), 1993.
25. Shapiro, D., Hughes, J. A., Randall, D. & Harper, R., Visual re-representation of database information: the flight strip in air traffic control, in M. J. Tauber et al. (eds), Cognitive Aspects of Visual Languages and Visual Interfaces, Elsevier Science, Amsterdam, 1994, pp 349-376.
26. Suchman, L., Plans and Situated Actions, Cambridge University Press, Cambridge, 1987.
27. Tatar, D. G., Foster, G. & Bobrow, D. G., Design for conversation: Lessons from Cognoter, in Int. J. of Man Machine Studies, 34(2), 1991, pp. 185-210.
28. Twidale, M. B., Redressing the balance: the advantages of informal evaluation techniques for Intelligent Learning Environments, in J. of AI and Ed., 4 (2/3), 1993, pp. 155-178.
29. Twidale, M. B., Rodden, T. & Sommerville, I., The Designers' Notepad: supporting and understanding cooperative design, in Proceedings of ECSCW'93, Milan, 1993, pp. 93-108.
30. Twidale, M. B., Rodden, T. & Sommerville, I., Developing a tool to support collaborative dialogues and graphical representation of ideas, in Collaborative Dialogue Technologies in Distance Learning, F. Verdejo (ed), in press.
31. Vincent, D., The Information Based Corporation, Dow Jones Irwin, 1990.