

Workplace Studies

*Recovering Work Practice and Informing System
Design*

Edited by

Paul Luff

Jon Hindmarsh

Christian Heath



CAMBRIDGE
UNIVERSITY PRESS

PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE
The Pitt Building, Trumpington Street, Cambridge CB2 1RP, United Kingdom

CAMBRIDGE UNIVERSITY PRESS

The Edinburgh Building, Cambridge CB2 2RU, UK <http://www.cup.cam.ac.uk>
40 West 20th Street, New York, NY 10011-4211, USA <http://www.cup.org>
10 Stamford Road, Oakleigh, Melbourne 3166, Australia

© Paul Luff, Jon Hindmarsh, Christian Heath 2000

This book is in copyright. Subject to statutory exception and to the provisions of relevant collective licensing agreements, no reproduction of any part may take place without the written permission of Cambridge University Press.

First published 2000

Printed in the United Kingdom at the University Press, Cambridge

Typeset in Plantin 10/12 pt in QuarkXPress™ [SE]

A catalogue record for this book is available from the British Library

Library of Congress cataloguing in publication data

Luff, Paul.

Workplace studies: recovering work practice and informing system design/Paul Luff, Jon Hindmarsh, Christian Heath.

p. cm.

Includes bibliographical references and index.

ISBN 0 521 59186 4 (hardback) – ISBN 0 521 59821 4 (paperback)

1. Organisational change – Management. 2. Technological innovations – Management. 3. Technological innovations – Employee participation. 4. Communication in design. 5. Communication and technology. I. Hindmarsh, Jon. II. Heath, Christian, 1952– . III. Title.

HD58.8.L84 2000 99-30983 CIP

ISBN 0 521 591864 hardback

ISBN 0 521 598214 paperback

Contents

<i>List of contributors</i>	vii
<i>List of figures</i>	x
<i>Preface</i>	xii
1 Introduction	1
PAUL LUFF, JON HINDMARSH AND CHRISTIAN HEATH	
Part 1: Exploring the workplace	
2 Making a case: 'knowledge' and 'routine' work in document production	29
LUCY SUCHMAN	
3 Design by problem-solving	46
GRAHAM BUTTON AND WES SHARROCK	
4 Analysing cooperative work in an urban traffic control room for the design of a coordination support system	68
JACQUES THEUREAU AND GENEVIEVE FILIPPI	
5 Expert systems in (inter)action: diagnosing document machine problems over the telephone	92
JACK WHALEN AND ERIK VINKHUYZEN	
6 The critical role of workplace studies in CSCW	141
KJELD SCHMIDT	
7 From individual action to collective activity and back: developmental work research as an interventionist methodology	150
YRJÖ ENGSTRÖM	
Part 2: The interface between research and design	
8 Analysing work practice and the potential role of new technology at the International Monetary Fund: some remarks on the role of ethnomethodology	169
RICHARD H. R. HARPER	

9	Ethnography, communication and support for design JOHN HUGHES, JON O'BRIEN, TOM RODDEN AND MARK ROUNCEFIELD	187
10	Where the rubber hits the road: notes on the deployment problem in workplace studies BOB ANDERSON	215
11	Situating workplace studies within the human-computer interaction field LIAM J. BANNON	230
12	Analysing the workplace and user requirements: challenges for the development of methods for requirements engineering MARINA JIROTKA AND LINCOLN WALLEN	242
13	Supporting interdisciplinary design: towards pattern languages for workplaces THOMAS ERICKSON	252
	<i>Bibliography</i>	262
	<i>Index</i>	281

1 Introduction

Paul Luff, Jon Hindmarsh and Christian Heath

In the past few years we have witnessed extraordinary pronouncements concerning the ways in which new technologies will transform the ways we work together. In both the popular press and in academic debate, an interest principally focused on extensions to existing computer networks, new forms of telecommunications and the potential of faster and cheaper systems, all have suggested that we are soon to be faced with a very different workplace. Workers will be more mobile when all the technological support they need can be provided wherever they are located and it may even be no longer necessary for individuals to travel to a particular site when they can work from home. The actual 'organisation' for which they work will become fragmented, geographically dispersed and possibly 'virtual', being transformed into a business with no physical location and little organisational structure.

Such pronouncements may seem curiously reminiscent to those familiar with the predictions associated with the microchip in the 1970s, or the motor car in the 1940s, or even earlier with the potential afforded in the nineteenth century by the telegraph, telephone and electricity (cf. Evans, 1979; Hall, 1988; Marvin, 1988). It is certainly the case that in the last few years the personal computer (PC) and electronic mail (email) have greatly transformed the way that work is accomplished in a large number of organisations. However, despite the grand intentions of proponents of novel technologies it is frequently the case that their impact is more modest. Indeed, it is not unusual for new systems once they have been introduced to be ignored, used to only a small degree of their capabilities or worse to be the cause of some great disaster. It appears that rather than radically transforming current work practices it is difficult even to achieve the less ambitious hope of supporting workplace activities, whether these are accomplished in a particular location or geographically dispersed. It appears that we need not only further technological developments to mobile devices, telecommunications and distributed computer systems but also a better understanding of the nature of workplace activities that are being intended to support, transform or replace.

It may seem remarkable, given the great body of work undertaken within the social and cognitive sciences concerning the use of information systems, that these do not seem to provide the resources relevant to developers of new technologies. Neither the multifarious studies of the processes surrounding the introduction of different technologies into organisations nor the detailed examination of individual activities carrying out pre-specified tasks appear to offer an account of naturally occurring workplace activities that is relevant or sufficient for developers of new technologies. Of course many of these studies have been undertaken for quite different purposes, with respect to debates within psychology, sociology and to theories of work and management. Nevertheless, even those fields with an expressed orientation to informing the design of systems through the detailed examination of individual activities with computers, like that of human-computer interaction, have had a surprising lack of influence on the development of new technologies (Carroll, 1991). Hence, more applied fields have emerged, such as requirements engineering, with a direct concern for providing practical advice and methods with respect to the needs of users.

In the light of these difficulties and a growing interest in developing technologies to support collaboration and group work, a corpus of studies has emerged that has been concerned with revealing the details of how activities are accomplished in real-world workplaces. Although many of these workplace studies have not been directly concerned with the development of any specific technology, they have begun to influence designers of novel systems, particularly of technologies to support collaborative work. They have suggested not only broad issues and topics which should be of concern to designers, but also ways of conceiving collaborative activities which can shape the development of novel technologies to support activities in the workplace. Indeed, a field has emerged which has acted as a forum of debate between developers of new technologies to support collaborative work and researchers of workplace activities: Computer Supported Cooperative Work.

In this chapter we discuss some of the background to workplace studies, both with respect to recent technological developments and to current debates within the social sciences. We begin by briefly outlining some developments in technologies aimed at supporting groups and collaborative activities. Despite their novelty, certain difficulties emerged with these systems and it became apparent that designers required a better understanding of the contexts in which these technologies were to be placed, particularly collaborative activities and social interaction in workplaces.

We then review some of the recent workplace studies that have been undertaken. Although most of this work has been related to the interests of CSCW, it is not the case that the principal motivations behind it have

been towards the design of new technologies. Certainly, there is no method which transforms a study of a workplace into a set of design guidelines. Indeed, many researchers would question whether the development of such a method would be the most appropriate way for workplace studies to be relevant for design. Moreover, there are also several orientations that can be taken towards the analysis of workplace activities. Despite having a common focus on naturally occurring workplace activities, these orientations are themselves developments from a range of earlier work in the cognitive and social sciences and therefore can utilise quite different conceptions to their particular domain of study. In this chapter we outline some of the principal analytic orientations that have informed the study of collaborative activities in the workplace.

Although the implications of workplace studies are frequently considered in terms of their potential for informing the design of a new computer system, this may not be their principal contribution. Workplace studies may not only suggest requirements for specific or generic technologies, but also provide for a respecification of the conceptions that underpin various of the applied and academic fields that take technology as their focus. We review some of the many directions in which researchers have developed the outcomes of their studies of workplaces. These include not only particular exercises that seek to shape new technologies, but also those that could inform the practices of designers and software engineers. Workplace studies also appear to offer a contribution to disciplines as diverse as the study of human–computer interfaces, the social study of technology and organisational behaviour. So, although workplace studies can contribute to the design of new technologies, even suggest some radical alternative ways for computers to support collaborative work, their more significant contribution will be in reshaping the ways in which we conceive of everyday social actions and interactions in the workplace.

This chapter provides some preliminary background, nevertheless it is hoped that the contributions in this volume, through illustrative case studies, discussion of relevant conceptual issues and debates concerning the relationship of these studies to the development of new technologies, will provide a critical resource for both those interested in the analysis of social activities in the workplace and those aiming to relate this analysis to design. These two concerns are reflected in the structure of this volume, an outline of which concludes this chapter.

Background

The prevailing deployment and use of computer systems like personal workstations linked together on networks and through applications like

email have led researchers to consider more advanced ways of providing support for workplace activities (e.g. Winograd, 1988; Sharples, 1993). In particular, designers have been especially concerned with extending the technology's capabilities for supporting individual activities so that computer systems can support collaborative work either when individuals are co-present or when they are remote (e.g. Stefik *et al.*, 1987; J. Olson *et al.*, 1990). Some of these innovations have directly built upon existing email capabilities, for example providing asynchronous support for individuals who are working on a common project, commenting on one another's work or writing a document together. Others have sought to provide synchronous support for several individuals working together at the same time. Some of the systems have been developed into products like Lotus Notes, but the success of more advanced designs has been harder to ascertain. For example, users appear to be ambivalent, at least, towards the kinds of technological support offered by shared drawing tools and desktop conferencing systems. Many other technologies using projection techniques, locator technologies and video and audio infrastructures have remained as prototypes (Ishii, 1990; Harper, 1992; Bly *et al.*, 1992), it being unclear whether and how they would be deployed within workplaces. Even more straightforward developments of systems for managing collaborative tasks, though requiring only a simple technical infrastructure, have met with little enthusiasm and even hostility from users. Although there appear to be a wide range of possibilities for developing technologies to support collaborative work, and a great number of suggestions have been proposed and prototyped, it appears to be hard to actually develop and deploy such systems in real-world settings (cf. Grudin, 1988). Those systems which have met with some success appear to be more due to happenstance than design. It may be that the difficulties associated with collaborative technologies may not be so much associated with poor design but more related to the general objectives underlying the systems, particularly with respect to how designers are considering the activities they are aiming to support.

Hence, it may be worth exploring a few of these developments in a little more detail, not only to provide an insight into work which has been undertaken within CSCW, but also to reveal how system designers have characterised the collaborative workplace activities that they aim to support. The heterogeneous range of technologies, systems and devices considered by researchers in CSCW makes it meaningless to select a 'representative' set of cases. Instead, we briefly examine three developments within CSCW that have been the focus of some debate within the field: Group Decision Support Systems (GDSS), workflow technologies and media spaces. Each of these aims to provide quite different kinds of

support (both synchronous and asynchronous) to various workplace activities (both co-present and distributed).

Group Decision Support Systems are typically comprised of a range of devices within one local domain which are configured to facilitate meetings, particularly, as their name suggests, to assist the individuals to formulate 'decisions'. So in the case of an early example, COLAB, private workstations were provided for each individual in the meeting and these were linked together and also connected to a public display visible to all at the meeting (Stefik *et al.*, 1987).

Various applications aimed to facilitate the generation of suggestions within the meeting, the distribution of these to colleagues and the collaborative formulation of arguments to support the decisions that were being made. In later developments of GDSS such tools have been refined so that quite sophisticated techniques have been provided to allow members to comment on the suggestions of colleagues, to categorise 'ideas' and to rank and analyse alternatives (Vogel and Nunamaker, 1990). Although each of these tools could be used separately, their use is considered with respect to an overarching serial process through which a problem is identified, vague solutions are proposed, then clarified and analysed and finally options are ranked, voted upon and decisions are made.

Experiences with the early use of COLAB revealed some problems due to the fragmented nature of technology (Tatar *et al.*, 1991). The public and private screens and the various windows on each made it difficult to recognise which participant was making which contribution, typically entered as typewritten statements. An underlying objective of the system, common in GDSS, to ensure anonymity of the participants did not help, making it even harder to ascertain whether different contributions were being made by the same participant. It was also hard to make sense of individual contributions, particularly when references were made to other statements through the system. Although there was an intention to support decision-making by providing for natural 'conversations' through the system, the technology did not support the interactional resources participants utilise to make sense of one another's contributions, that is the sequential nature of the conduct. Even the efforts to preserve the anonymity of users, an idea that was meant to provide for greater participation, that made decisions less biased and perhaps more rational was not necessarily an advantage. It seemed to undermine the practices that participants utilise to make sense of the contributions of others and assess those contributions.

Indeed, these drawbacks with COLAB echo more general concerns with the conception of decision-making embodied within GDSS. So, for example, March (1991) contrasts the implicit assumptions underpinning

such systems with observations concerning how decisions are made in organisations. He describes how ‘decisions’ as such, rather than prefiguring an action, are often *post-hoc* justifications for outcomes which have already arisen. March’s rather ironic analysis of the work within organisations does offer an alternative to the rational and formal characterisation of decision-making by developers of GDSS systems. It also, as he suggests, leads to the possibility of considering an entirely different kind of support for ‘decision-makers’ – tools which focus on the presentation of decisions rather than the processes through which they are made. Hence, March (1991), even by utilising general observations of how activities are accomplished in organisational settings, provides not only a radically different conception of workplace activities, but also an entirely different direction in which to proceed for technologies being developed to support them. March’s analysis suggests the ways in which even general observations of an activity may have some practical implications for system design. More importantly, it reveals how initial presumptions concerning an activity, for example that decisions are the outcomes of prior reasoning performed by groups of individuals through largely rationalistic argumentation, can be set in stark contrast with the everyday accomplishments of participants in organisational settings. Nor is the case that such observations merely present the deficiencies of everyday conduct against some ideal process, rather they reveal the ‘good reasons’ for such *ad-hoc*, situated and contingent practices.

Rather than supporting a real-time synchronous activity accomplished by co-present individuals, workflow technologies aim to support asynchronous collaboration between physically dispersed individuals. Moreover, they do not rely on being located in a dedicated predesigned setting, being based on more straightforward technological foundations they can be typically used on conventional personal workstations. As their name suggests these systems are designed to support the representation, dissemination and presentation of workflows – sequential relationships between activities (Winograd, 1988). Tasks which are to be accomplished by several individuals, like the preparation of a document, can be laid out, usually graphically, using the system. Then, as the workflow is accessible to all the individuals through a computer network, the workflow can be invoked. Careful preparation can allow for some flexibility to the ways in which the workflow is accomplished, nevertheless the system aims to ensure that the appropriate individuals participate in the activity at the relevant time.

Even though users could produce their own workflows with optional paths, early experiences with the technology revealed that users still found the systems too constraining (e.g. Carasik and Grantham, 1988). It being

impossible to predefine all possible contingencies, there were necessarily going to be occasions when the workflows would have to be transformed once they were underway. Of course making such changes could cause problems for other users and may undermine the very reasons for using such a system. Hence it was perhaps not surprising that users would circumvent the system and use other means to collaborate and communicate between colleagues. More recent developments in workflow technologies have sought to address these problems, but their apparent inflexibility may not be so much due to the ways in which tasks may be ordered and changed but in the very ways that tasks are specified and categorised (Suchman, 1993a). The explicit definition of tasks may itself be problematic for users. It may not be straightforward to circumscribe the tasks which are relevant to users. Not only may their specification gloss critical features of the work, particularly with respect to their collaborative accomplishment, but also actually making the tasks explicit, and each individual's contribution to them, may interfere with the smooth performance of workplace activities. The pre-specification of tasks actually accomplishes quite a different activity from outlining a flow of future actions, with respect to the ongoing concerns of participants, at that moment. So, despite the good intentions of designers in making work activities more visible and manipulatable by those who undertake them, the technology might actually undermine their accomplishment. The work of the participants may thus be augmented with efforts to get the technology to work, to make the pre-specified tasks fit the moment-to-moment demands of the setting. Although aiming to develop a flexible technology that is open to redefinition by its users, designers of workflow systems may still be neglecting the ways in which workplace activities are situated and contingent. The very conception of tasks embodied within such a system appears to have been misconstrued; a stipulative and circumscribed characterisation of task actually makes it more problematic to accomplish activities through the technologies designed to support them.

Innovative communicative technologies, either in the form of desktop conferencing systems or more novel media spaces, offer the potential not only for supporting collaboration between physically dispersed individuals but also providing this in real-time (Bly *et al.*, 1992; Gaver *et al.*, 1992; Mantei *et al.*, 1991). Although such systems typically offer common access to an electronic workspace through specially designed 'shared applications', their novelty lies in the capabilities afforded by continuous access to a remote domain through both audio and video links. Through the combination and configuration of conventional audio-visual technologies, proponents of video-mediated technologies can offer systems to support collaboration that should be straightforward to operate by their

users. Indeed, the more optimistic hopes for such technologies are to provide new spaces for collaboration, where informal, typically face-to-face, communication and the peripheral awareness of a remote colleague's activities can take place. Rather than refining a pre-specified task or aiming to transform the way work is accomplished, the technology should provide a resource through which collaborative activities can be seamlessly interwoven within the everyday work of the participants. However, in the new space the participants are now physically separated, in distinct offices or even in geographically dispersed sites (Abel, 1990; M. Olson and Bly, 1991). This would appear to be a straightforward deployment of audio-visual technologies to support collaborative work. Nevertheless it does not appear to have the impact its proponents would have hoped for.

Studies of the extensive use of such audio-visual infrastructures reveal that accomplishing everyday interactional activities, such as the production of gestures and their coordination with talk, may not be that straightforward through the technology (Heath and Luff, 1992b). The symmetries that underpin the accomplishment of visual conduct within interaction, that are relied on by participants in more conventional settings, are transformed through the technology. Participants are not able, in the same way, to rely on the resources they typically utilise in the production and receipt of visual conduct. This would seem to be a difficulty for proponents of such systems, particularly those who have characterised their advantages in terms of the technology's ability to provide for gaze direction, gestures and other features of 'face-to-face' interaction. These very features have been typically considered critical in providing better support for 'informal' interaction, not offered in other media, like the telephone or through the computer network.

The detailed analysis of interaction through media spaces coupled with the ambivalent results of more quantitative measures of conduct through video have led designers to rethink the focus on supporting 'informal' work. Hence, several researchers have proposed that the design focus of such technologies be redirected towards supporting the more mundane collaborative accomplishment of workplace activities, more focused on the objects used within an interaction than the remote colleague (Nardi *et al.*, 1993). Moreover, other studies of video-mediated communication have noted how participants appear to rely more on the views offered by document and object centred views, rather than those of the other. However, despite these proposals and observations, the resources offered by these technologies to support the actual accomplishment of collaborative activities are relatively undeveloped. Typically users are given some shared workspace or a document view from a separate camera, but these

are provided in distinct and fragmented domains; there is little support for tying these resources to the ongoing conduct of their co-participant. It appears that by focusing on supporting informal interaction, designers have overlooked how to support more focused collaborative activities. Even advanced systems like media spaces remain largely unused in the organisations in which they have been deployed and desktop conferencing systems remain a novel, but under-utilised application. To refocus these developments, however, requires a greater understanding of the resources that individuals utilise when accomplishing work activities within interaction.

Despite the obvious differences between GDSS, workflow technologies and media spaces, they each have been designed in different ways to support collaborative activities, and in each case there appear to be profound problems integrating these technologies into the ways individuals accomplish their everyday work activities. Of course, developments in the design of CSCW systems are in their preliminary phases. Examples of these three cases are still largely prototypes or early implementations. However, it may seem surprising that technologies explicitly aiming to support collaboration, often with considerable attention being paid to how they will be used, appear to be so ill suited to the contingent, emergent and collaborative aspects of the work they aim to support. So, COLAB fragments the resources that individuals make use of in accomplishing interactions, workflow systems make explicit activities that are usually implicit and media spaces transform the conduct they are meant to support. In each of these, what appears to be a straightforward conception of a collaborative activity, a 'decision', a 'workflow' or 'informal interaction', which seems to be in need of technological support, turns out to be problematic. Activities which appear distinct, indivisible and possible to circumscribe are revealed to be emergent, complex and interwoven with others features of conduct when their accomplishment is examined in everyday organisational settings. When the complexity of collaborative work is considered it is perhaps not so surprising that examples of 'successful' CSCW products are so rare. Email is a noteworthy, and perhaps questionable, example of a CSCW product, and even groupware technologies like Lotus Notes do not appear to be used as designed or fail when introduced to support inappropriate organisational activities (Grudin, 1988; Orlikowski, 1992).

Although this may appear to be a failing of the emerging field of CSCW, it may not seem so unusual when the problems associated with the more general introduction of new technologies are taken into account. The newsworthy examples of computer failures coupled with more numerous mundane examples of unused or underused systems

point to a lack in our understanding of the everyday work activities they are designed to support. This is despite considerable effort being devoted to developing methods and approaches for the design of technologies that are appropriate and easy to use. Within the field of human–computer interaction, for example, not only has there been an interest with developing an understanding of how technologies are used by individuals, but also there has been a concern for developing applied findings of relevance for design (Barnard, 1991; Card *et al.*, 1983). Researchers in this field have thus paid considerable attention to the ways in which their findings could be applied to the design of new technologies and to developing methods for supporting a more ‘user centred’ approach to design (Norman and Draper, 1986). These approaches have drawn on an analytic framework developed within cognitive science, accounting for the behaviour with computer systems in terms of ‘mental models’, ‘task grammars’, cognitive schemata and rules (Norman, 1983; Payne and Green, 1986). Such conceptions have informed a range of methods, typically utilising an experimental paradigm, that not only seek to provide an account of human–computer interaction, but also offer ways of evaluating and even suggesting guidelines for the design of computer systems.

Recently, HCI’s orientation developed from cognitive psychology and cognitive science, focusing on the individual user, and often utilising an experimental paradigm has been called into question. Too constrained a conception of human–computer interaction appears to overlook the collaborative, social and organisational nature of how conventional technologies are used in everyday settings. Too much emphasis on the use of computers to perform circumscribed experimental tasks neglects the contingent ways in which activities are accomplished. They may also unnecessarily constrain the ways of informing the design of technologies for real-world domains. Hence, consideration has begun to focus on methods and approaches that explore the achievements of participants in naturalistic settings and in developing the ways in which computer use is conceived, particularly with respect to the social and situated nature of this conduct.

With respect to the more practical concerns of designing computer systems, a field has recently emerged that has concentrated on exploring ways of eliciting, describing and specifying user requirements for new technologies. Motivated by the practical problems associated with discovering and defining what users might need from a computer system, requirements engineering has sought to develop techniques for requirements capture, modelling and specification. Within requirements engineering there have been shifts, similar to those in HCI, towards the social. However, despite these initiatives it appears that this approach to a more

systematic development of computer technologies also has its shortcomings.

Although there have been a number of interesting studies within requirements engineering, the pragmatic concerns of the field have led to a rather curious amalgam of conceptions, models and approaches drawn from a range of disciplines. For example, proposed requirements methods have utilised tools and techniques that break work activities into the distinct tasks performed by individuals, that require naturalistic observation of work practices and which involve facilitators to assist with design meetings. These developments have drawn on a disparate range of work within the social sciences and elsewhere, including task analysis, cybernetics, socio-technical systems, participative approaches towards design and the analysis of group processes within social psychology. Incorporated within a method for design, the derivation and provenance of any technique can become unclear and the associated underlying conceptual assumptions can be masked. Hence, methods proposed to reveal users' requirements often draw on frameworks that are stipulative and tightly circumscribe how the activities under scrutiny are analysed. Although there may be a worthy aim to elicit the requirements of actual users often through some analysis of their current activities, the methods utilised and conceptions adopted may constrain how work activities are examined and what possibilities can be explored to support them.

In recent years, therefore, initiatives in fields associated with the development of new technologies in CSCW, HCI and requirements engineering have all involved a turn towards the social. These have been motivated by quite different concerns, whether these be the demands implied by developments in new technologies, the constraints of existing analytic frameworks or the requirements for novel approaches to the design process. Although a range of methodological orientations could be seen to be relevant to these requirements, there has been particular interest in those that are naturalistic and not stipulative, and account for the contingent and situated nature of organisational activities. However, it is apparent that these demands placed upon CSCW, HCI and requirements engineering cannot be met by a pre-existing set of tools and techniques. It requires that practitioners rethink their current conceptions of everyday work activities in order reconsider the frameworks underpinning current methods both for the analysis of conduct associated with new technologies and for the design of novel systems. The unpicking of how collaborative activities are actually accomplished in workplace settings can thus been seen as a resource for such a reconsideration. It can suggest both ways of reconceptualising key concepts in the analysis of technology-oriented activities and the design of computer systems.

The possibilities afforded by workplace studies for analysis and design have not only been informed by, but also motivated a range of naturalistic case studies of everyday workplace activities, many of which have had a particular concern with the mundane uses of artefacts and technologies in real-world settings. These empirical case studies have stood in stark contrast not only to previous studies of technologically mediated work but also to prior methods for design. They have concentrated on revealing the complexities of everyday social interaction, emphasising the relevance of particular analytic orientations for the examination of empirical materials and drawing out the implications these analyses have for the critical conceptions underpinning the study of work activities.

However, in beginning to reveal the complexity of everyday, collaborative work activities they also reveal the paucity of our current understanding of everyday technologically mediated work. So although workplace studies have been utilised to propose novel designs for technology, extensions to existing design methods and even some possibilities for new approaches to system development, perhaps their more immediate contribution is in outlining the conceptual and methodological innovations required in the social sciences to understand the ways in which artefacts and technologies are utilised in everyday workplaces.

Workplace studies

Workplace studies have arisen in the light of a number of convergent issues and concerns. First, they have been driven by a growing concern, among those in both academia and industry, with the design and deployment of advanced technologies, particularly with a recognition that problems and failures of technologies often derive from our lack of understanding of how ordinary people, in conventional organisational environments, do the things they do. Second, they have arisen in the light of the changing nature of technology, not only the shift towards complex communication systems, but also the growing ability to provide sophisticated support for collaborative activity. Again, in part, the turn towards the social has reflected the possibilities of supporting, in complex ways, people working together. In large part, however, workplace studies have been driven by a concern to develop an understanding of technology which is free from the incumbencies of certain forms of cognitive science, which takes the social and situated seriously, and which drives analytic attention towards the ways people use technologies to accomplish and coordinate their day-to-day practical activities. The practical concerns and implications of workplace studies derive from an analytic agenda, an agenda which is attempting to respecify technology with regard to human practice and social organisation.

Workplace studies have emerged from within various disciplines, including sociology, social anthropology, cognitive science and to some extent computer science. To a large extent many of these workplace studies remain relatively unknown in their original discipline(s), but have had an important impact on interdisciplinary fields such as HCI and more particularly CSCW. Indeed, major interdisciplinary colloquia and conferences such as Computer–Human Interaction (CHI), Computer Supported Cooperative Work (CSCW), and European Conference on Computer Supported Cooperative Work (E-CSCW) are increasingly dominated by papers which interweave workplace studies with more technical concerns. It is interesting to note that despite the diversity of approaches which inform workplace studies, and their wide-ranging concerns in terms of substantive domains and analytic disposition, they retain a number of common characteristics. First, they are principally concerned with the situated organisation of collaborative activities, and the ways in which tools and technologies, objects and artefacts, feature in practical action and interaction in the workplace. Second, they are overwhelmingly naturalistic, ethnographic studies, involving, to use Geertz's (1973) well-used term, 'thick description' of human conduct and cooperation in complex technological environments. Many of these environments, control rooms, newsrooms, financial institutions and the like have been characterised by Suchman (1993b) as 'centres of coordination'. Third, many workplace studies are concerned with reconsidering and re-specifying the concepts and theories which currently infuse our understanding of technology. For example, the idea of the 'user' has been reconsidered in recent years to demonstrate its embeddedness not so much in individual cognitive competencies, but rather in socially organised practice and reasoning which is inseparable from the socially organised activities in which tools and technologies are used. Parallel reconsiderations have been applied to such concepts as 'information', 'communication', 'awareness', 'cooperation' and so forth. Their concern with the contingent and situated character of practical action serves to generate a body of empirical observations and findings concerning the practices and reasoning in and through which participants accomplish and coordinate their actions and activities in the workplace. In this way, they break from the long-standing 'technicism' which Grint and Woolgar (1997), for example, argue pervades both sociological and cognitive studies of technology, and systematically attempt to examine tools and artefacts with regard to the indigenous courses of action and interaction which gives technologies their occasioned, yet determinate sense.

Within this broad set of concerns, however, there are a number of distinct approaches to practical action in the workplace. One of the most

surprising developments is found not within sociology or anthropology but within cognitive science. The growing recognition of the importance of the situated and collaborative character of practical action has led to the emergence of 'distributed cognition'. While there is some debate as to the provenance of the approach, and what it actually involves, there would seem to be strong commitment to exploring the ways in which tools and technologies enable participants to develop common understandings and representations of actions and objects, and thereby facilitate cooperation and collaboration. For example Salomon (1993) suggests:

People appear to *think in conjunction or partnership* with others and with the help of culturally provided tools and implements . . . The thinking of these individuals might be considered to entail not just 'solo' cognitive activities, but *distributed* ones. In other words, it is not just the 'person-solo' who learns, but the 'person-plus,' the whole system of interrelated factors. (Salomon, 1993: xii–xiii, original italics)

The approach has interesting parallels with symbolic interactionism, with its emphasis on the ways in which shared definitions and understandings, themselves the products of social interaction, provide the basis to practical action and collaboration within the workplace. The work of Agre (1997), Hutchins (1995), Rogers (1992) and others who, in various ways, have contributed to our understanding of distributed cognition, especially in organisational environments, also has certain similarities to an approach which developed independently in France, indeed (some suggest) foreshadowed distributed cognition. Commonly known as course-of-action analysis, a number of researchers in ergonomics in France developed an approach not dissimilar to distributed cognition but which emphasised more the ongoing coordination of workplace activities. The approach, emerging in the light of the work of Pinsky and Theureau (Pinsky, 1979; Pinsky and Theureau, 1982, 1992; Theureau, 1991, 1992), is naturalistic, and is principally concerned with explicating the use of tools and technologies from within the courses-of-action in which they are embedded. The approach preserves a commitment to the cognitive, while explicating the ways in which individuals interweave distinct courses of action, often through the assistance of tools and technologies. Like certain forms of distributed cognition, course-of-action analysis preserves the primacy of the individual, and individual cognition, but powerfully demonstrates how representations and action are assembled and disassembled through cooperation and coordination.

A number of other analytic orientations have also informed these workplace studies and begun to generate findings concerning technologically mediated collaborative activity. As suggested, symbolic interactionism, and in particular perhaps the work Strauss (Strauss *et al.*, 1964, 1985),

has informed a range of empirical studies and provided a number of conceptual distinctions which have permeated discussions of the social and technical in CSCW (see e.g. Bowker and Star, 1994; Star, 1989). As in other fields, such as education and literary criticism, there has also been a growing interest in drawing on activity theory, as a methodological and conceptual framework for the analysis of workplace activities (see, e.g. Engeström and Escalante, 1996; Kuutti, 1996). Unlike other approaches used in workplace studies, it is sometimes seen as offering a solution to the vexed problem of the 'micro and macro', which even haunts CSCW, a conceptual vehicle for interweaving the fine details of interaction with the broader organisational constraints and circumstances.

However, it is perhaps ethnomethodology and conversation analysis, more than any other analytic orientation, which had the most prevailing influence on these workplace studies and more generally, social science research in CSCW. This is hardly surprising. Suchman's (1987) original critique of cognitive science and HCI drew on ethnomethodology and conversation analysis, and brought such work to the attention of an audience largely unaware of these analytic developments. They also offered a collection of analytic commitments, and a substantial body of empirical findings, which could provide a vehicle for a distinctive approach to the 'interaction' between human beings and computers within the circumstances of the workplace – an approach which placed the occasioned sense of practical action at the forefront of the analytic agenda. It is interesting to note, however, that ethnomethodology and conversation analysis do not themselves offer a unified analytic orientation to practical action and interaction, and as elsewhere, they have provided the basis to a wide diversity of workplace studies which, while all concerned in general with the collaborative production of technological informed organisational activities, reflect a complex array of interests and commitments.

If workplace studies embody a wide range of analytic orientations, they also address a diverse variety of substantive issues and domains. The substantive domains addressed in workplace studies cover a broad range of organisational settings. There is for example a growing corpus of research concerned with what Suchman (1993b) has characterised as 'centres of coordination'. These include studies of air traffic and ground control (C. Goodwin and M. Goodwin, 1996; M. Goodwin, 1990; Harper and Hughes, 1993; Hughes *et al.*, 1988; Suchman, 1993b), of emergency dispatch centres (Whalen, 1995b; Zimmerman, 1992) and the control rooms of rapid urban transport systems (Filippi and Theureau, 1993; Heath and Luff, 1992a, 1996a). In different ways such studies have examined how the members of such settings, in interaction both with each other and those outside the domain, use various tools and technologies to

preserve a mutually compatible sense of constantly shifting circumstances and events, and maintain a coordinated response to disparate, and to some extent unpredictable, problems and difficulties with which they deal. There is also a growing body of studies concerned with the use and deployment of technologies in financial institutions. So for example, Harper (1998) has undertake a wide-ranging study of the use of documents in the International Monetary Fund and in a rather different vein, Jirotko and colleagues have examined collaborative work in trading rooms in the City of London (Jirotko *et al.*, 1993; Heath *et al.*, 1994–5).

This emphasis of the practicalities of cooperation, the use of documents and the management of ‘normal, natural troubles’ are also reflected in studies of rather different domains. So for example Button and Sharrock have undertaken projects on software engineering and examined how personnel order a complex array of concurrent and serially related activities (Button and Sharrock, 1994, 1996). Bowers and Button have examined the deployment and use of ‘formal systems’ such as workflow models to coordinate and make sense of activities on the shop floor in the printing industry (Bowers and Button, 1995). This interest in the deployment of both social and technical systems is also addressed in studies of the use of email in the civil service (Bowers, 1994), the introduction of information technology for customer services into high street banks (Randall *et al.*, 1995) and the use of medical systems in general medical practice (Greatbatch *et al.*, 1993; Heath and Luff, 1996b).

These and other workplace studies are all naturalistic, and involve extensive fieldwork in the respective settings. Some of the studies, especially those concerned with the interactional organisation of workplace activities, are primarily based on the analysis of video recordings, and direct attention towards the moment-by-moment, collaborative accomplishment of visual, vocal and material conduct. Among workplace studies there is also a growing body of video-based research derived from quasi-naturalistic experiments. These can range from investigations of the deployment of prototype technologies into organisations, such as research laboratories, through to short-term exercises in which subjects are requested to undertake particular tasks. Such experiments have proved invaluable for exploring the use of more advanced, experimental systems, and have helped provide insights into communication and collaboration in more conventional circumstances. So, for example, Heath and Luff (1992b), Gaver *et al.* (1993) and Dourish *et al.* (1996), among others, have undertaken a series of studies exploring interaction, sociability and work among personnel using media space technologies (computing and audio-visual infrastructures designed to support distributed collaborative working) and Bowers *et al.* (1995) and Hindmarsh *et al.*

(1998) have examined the use of virtual environments. In various ways these studies reveal how advanced technologies designed to support distributed collaborative work transform visual and vocal conduct, and can undermine the abilities of participants to establish a mutually coherent sense of objects and their respective environments. In turn these findings have begun to clarify certain practices and presuppositions found in collaborative work in co-located environments such as control centres (cf. Heath and Hindmarsh, 1997).

Despite the varied domains and issues addressed by such analyses these workplace studies share a number of analytic commitments and interests. Aside from their methodological orientation, we find a principal concern with the contingent and situated character of practical organisational conduct, and in explicating the resources on which participants themselves rely in producing and coordinating their actions and activities with each other. They also share an interest in technology, not only an applied concern in using naturalistic research to inform the design, evaluation and deployment of new tools and artefacts, but more importantly perhaps in using studies of work to reconsider the ways in which we understand the relationship between human action and objects, whether highly complex or truly simple. Indeed, perhaps the most significant contribution of these workplace studies are the ways in which they are placing tool or artefact mediated conduct at the heart of the analytic agenda, and attempting to reconceptualise technology, and in particular reveal how the use and intelligibility of objects is produced and constituted in and through social action and interaction. Finally, in various ways, such studies are providing a vehicle to question a common body of issues and concerns which underpin more traditional research concerned with technology and organisational conduct in both the cognitive and social sciences; that is that they are throwing into relief the empirical and theoretical shortcomings of studies in HCI, organisational behaviour, and CSCW, that indigenous social action can be adequately accounted for with regard to rules, plans and procedures which are insensitive to the ways in which participants themselves accomplish their conduct in particular situations and circumstances. It is their concern with the tacit, seen but unnoticed, indigenous resources which inform the production and coordination of *in situ* technologically mediated conduct which unites these workplace studies.

Implications

The growing dissatisfaction with more traditional research in HCI and cognitive science, coupled with the emergence of CSCW and the growing

interest in developing systems to support collaborative work, have played an important part in the emergence of workplace studies. These academic developments have been fuelled by the problems which have arisen with the design and deployment of major technologies, reflected in press coverage of various organisational disasters (e.g. TAURUS, see T. Collins and Bicknell, 1997). These systems failures have contributed to the widespread recognition that we need to abandon the long-standing assumption that complex systems will stimulate efficient changes in work practice. Rather, there are calls to develop approaches more sensitive to the settings and situations in which new technologies are deployed.

The emergence of naturalistic studies of the workplace has begun to reveal the essentially situated and contingent character of collaboration and technology use. Such a sensitivity has been used to respecify understandings of foundational (and taken-for-granted) concepts in the areas of HCI, CSCW and even the practical world of systems design. For example, Suchman's (1987) study of photocopier use reconsidered the plan-based model of human conduct extremely prevalent within expert systems design, HCI, artificial intelligence and cognitive science. Workplace studies have similarly attempted the respecification of concepts such as 'user', 'task', 'context' and even 'collaboration' itself. As a result the findings from these studies can be seen to have made a powerful contribution to the ways in which these fields conceive of the workplace and the activities therein. They have highlighted the complexity of the everyday, tacit resources and activities that underpin collaborative work. In turn, they have provided directions for designers who attempt to envisage and develop new technologies to support (collaborative) work.

These alternate conceptions and observations of work drawn from workplace studies have begun to be utilised within the general development of workplace technologies. For example, studies in a range of domains have suggested to designers how it may be necessary to provide individuals with peripheral awareness of another's conduct. Hence, developers of a variety of collaborative technologies, including audio-visual infrastructures, collaborative virtual environments and even more asynchronous kinds of support have explored ways of providing different types of awareness and participation in their systems. Thus, findings from workplace studies are having an influence on such diverse developments as CSCW systems, ubiquitous computing, virtual reality technologies and mobile devices.

As well as generating new concepts to inform more general design concerns, workplace studies have also been used to examine particular technologies and domains. There have been numerous attempts to generate

requirements for new technologies or comment on proposals for technological change from a detailed understanding of the interactional production of activities in specific domains. From these naturalistic studies, it is often possible to outline key features of the work which appear relevant to the participants and which may have consequences for any future deployment of technology.

For example, from a study of the work of traders in a financial dealing room, Jirotko *et al.* (1993) were able to comment on proposals by a technology development company to introduce a new 'deal capture' system for traders. The company was considering developing a voice recognition system to overcome current problems associated with the hand-written recording of deals, where deals were recorded on paper tickets and inputted into the computer by other personnel some time later. The proposed system aimed to 'capture' the contents of a 'deal' as it was made on the phone. Jirotko *et al.*'s study not only revealed some shortcomings of the developers' assumptions concerning the nature of dealing, such as the amount and nature of the activities carried out on the phone, but also highlighted the collaborative nature of financial trading – both between traders on the phone and with colleagues in the same trading room. These comments led the company to rethink its proposed development. From the analysis, Jirotko *et al.* were also able to propose alternative technologies (e.g. digital pen-based systems) which might support the traders, suggestions that would be more sensitive to the collaborative production of the work.

Workplace studies facilitate a close look at current practices in a domain, which can usefully inform the development and introduction of new systems. This kind of sensitivity to current practices would seem critical if designers are to avoid impeding staff in their everyday activities, indeed to avoid the kinds of disaster that can develop from technologies which impose an 'operational straitjacket' on workers (cf. Page *et al.*, 1993). This is not to suggest that these studies of work merely attempt to replicate current work practices. Rather, they provide understandings of the essential user requirements from which to base thinking about the potential impact of candidate technologies. Moreover, and as mentioned above, they are also being used to inform the development of completely new, 'virtual' workplaces. Understandings of collaboration drawn from workplace studies are contributing to, and shaping the development of, advanced technologies to support distributed work, namely media spaces and collaborative virtual environments (Bowers *et al.*, 1995, 1996; Hindmarsh *et al.*, 1998).

In various ways then, researchers undertaking workplace studies have sought to show how their ethnographies and fieldwork reveal issues of

direct relevance to any system that would be implemented for that workplace or for the general development of technologies to support collaborative work. It is inevitable, then, that practitioners have begun to consider how findings from workplace studies or the accomplishment of fieldwork can be transformed into design guidelines or 'methods' for design (e.g. Randall *et al.*, 1994); a related concern being how to present ethnographic materials and findings in such a form as to be both clear and useful to designers. The traditional 'thick descriptions' of ethnography, though revealing the important detail of workplace activities, are not amenable to the transformation into design guidelines. Indeed, such transformations are likely to gloss the very nature of the activities that the orientation makes visible. Similarly, though designers may seek ways of abstracting and generalising user requirements, methods that seek to formalise these may end up replicating the stipulative character of those approaches they are seeking to replace.

It could be that workplace studies themselves suggest ways in which system development processes could be transformed. For example, Button and Sharrock's (1994) study of software engineers reveals how the developers of a system contingently made use of design principles, methodologies and tools in their work in order to accomplish the organisational demands of the setting in which they worked. To get the work done, the engineers reversed the prescribed order of activities, producing requirements following from a design, used tools contingently for their particular purposes rather than within the confines of the methodology and were oriented to other demands of professional practices. Conventional development methods, particularly those that are highly prescriptive and force a rigid adherence to a sequential process, often overlook such organisational contingencies. In proposing any tools, techniques and methods that draw from workplace studies it may be worth considering how such tools would be utilised for the practical purposes of software designers, that resonate with their own work practices and that are sensitive to their own particular organisational demands.

The influence of these varied contributions to the broad field of CSCW and systems design can be gauged from the calls by computer scientists themselves, for their colleagues to take seriously the tacit features of human conduct (e.g. Brown and Duguid, 1994; Goguen, 1994; Moran and Anderson, 1990; Potts and Newstetter, 1997; Sommerville *et al.*, 1993b), from the growing number and strength of links between various academic computer science departments and social scientists and from the number of large industrial research laboratories that are utilising naturalistic studies of work as part of the systems development activity. This reveals another aspect of the ways in which workplace studies are