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