1 Technology and social action

> The full introduction of the computer system effectively did away with the radio and telephone calls to stations, with the computer dispatching crews to answer calls. But within hours, during the morning rush, it became obvious to crews and control room staff that calls were going missing in the system; ambulances were arriving late or doubling up on calls. Distraught emergency callers were also held in a queuing system which failed to put them through for up to 30 minutes.
>
> Ian MacKinnon and Stephen Goodwin (*Independent* 29 October 92)

> The medium-term future [1983–90], therefore, will see the first shift towards a decentralisation of commercial and business life. The importance of the city office, with its *mêlée* of agitated human beings passing each other redundant messages of paper and the printed word, will be rapidly eroded. More and more often, office and home will be combined, the public transport system will give way to giant data communication networks, the business motorcar will be traded in for the latest videoconference system. For the first time since Man began to behave as a social animal and gather his kind together into ever larger working and communicating units, a significant trend will emerge. The cities will empty and expensive office blocks will gather dust. For centuries Man has been accustomed to the notion that he must travel to find his work; from the 1980s into the 1990s the work – such as is to do – will travel to meet Man.
>
> Evans (1979: 142)

In the early autumn of 1992, the London Ambulance Service, at the time the largest ambulance service in the world, introduced a system for Computer-Aided Dispatch into the control room. The principal aim of the system was to replace the outmoded and inefficient practice of documenting the details of emergency calls on paper slips. It was recognised that the system would necessitate some change in working practices, not only amongst control room staff but also for the ambulance drivers. Details from the emergency calls entered into the computer system could be matched with the location of ambulances sent from the vehicles, in order to schedule and allocate crews automatically. This would mean that, amongst other things, call takers and dispatchers would not have to

rely on the current paper-based system, and that ambulance stations would no longer have any responsibility for which crew dealt with which call. The day it went 'live', 26 October, problems began to emerge. On not a particularly busy day, response times to arrive at an incident were longer than usual: less than 20% were arriving within the target time of 15 minutes. More importantly, the time taken to answer the calls began to rise alarmingly; the average time a call to the control room was left ringing peaked at 10 minutes. Together, these delays also meant that the number of calls also started to increase, with patients ringing in to find out whether an ambulance was on its way. For the ambulance crews, the automated system was causing some frustration. The system needed near-perfect information about the location of the ambulances which it was not possible to produce. Without this, the system began to allocate crews incorrectly. More than one ambulance would arrive at the scene of an emergency, or the closest ambulance would not be allocated. These problems led to crews transmitting more requests to the control centre and the control centre sending more messages. Unfortunately, as the volume of the messages increased, so did difficulties with the system: messages were lost and the 'awaiting attention' and 'exception message' lists on the computer became so long that items scrolled off the top of the screen and were forgotten. Distraught citizens called saying that they had been waiting for more than half an hour for an ambulance. In the control room personnel began to lose track of which cases had been successfully allocated and dealt with, and in all the confusion crews began to receive incorrect information concerning the location and 'status' of particular cases. Following a second day of problems, the service reverted to a semi-manual system and on the 4th of November the system crashed completely. Personnel returned to the original paper-based system and reinstated some semblance of order.

The difficulties which arose following the introduction of the Computer-Aided Dispatch into the London Ambulance Service are not the most costly disaster to occur when new technology is introduced into a working environment. For example, it has been estimated that the TAURUS project, designed to replace paper certification in the London Stock Exchange with a computerised system, cost the Exchange alone about £75 million and with other firms building interrelated systems the total spent might have been as much as £400 million (Collins and Bicknell, 1997). Despite more than three years work on the project, the system was never introduced, and it has taken more than five years for an alternative and less ambitious technology to be deployed. TAURUS is one of many computer disasters reported in popular publications and the press, where systems either fail to work when they are introduced, or even

never reach a stage where they do get introduced (Collins and Bicknell, 1997; Neumann, 1995; Wiener, 1993) . Putting to one side such dramatic failures, there are numerous examples where seemingly innovative and reliable systems have failed when introduced into organisational environments. In many of these cases, the system in question does not cause severe problems; it simply sits there underused. Indeed, it is hard to find a modern organisation which has not had its problems in attempting to exploit the apparent benefits of new technology.

The official inquiry into the London Ambulance Service fiasco (Page et al., 1993) identified a number of key issues which led to the problems which arose. Perhaps the most important of these is that the project team assumed that the computer system would naturally bring about changes in the working practices of personnel; practices which in part were seen as outmoded and inefficient. The report of the inquiry suggests:

> Management were misguided or naive in believing that computer systems in themselves could bring about [such] changes in human practices. Experience in many different environments proves that computer systems cannot influence change in this way. They can only assist in the process and any attempt to force change through the introduction of a system with the characteristics of an operational 'strait-jacket' would be potentially doomed to failure. (London Ambulance Service Inquiry Report (Page et al., 1993: 40))

This disregard for the ways in which people organise their work, coupled with a disdain for the ordinary resources on which they rely, is a common feature of many projects involving new technology. Management can seem dazzled by the splendours of computers, and terms such as multimedia, the internet and digital age are used to suggest that technology will soon transform our mundane workaday world. The idea of the paperless office might bring a wry smile to the face, but many contemporary organisations still believe that technology alone is the solution. Even the world of fine art is haunted by the banal prediction that 'painting is dead', unable to survive the splendours of the digital age.

Technological failures have led to a growing interest amongst those in both industry and academia in developing new and more reliable ways of identifying the requirements for complex systems. There is a growing recognition that what are unfortunately classified as 'non-functional requirements' need to be taken seriously, and that 'human factors' consist of more than a concern with the interface between an individual and a workstation, but may involve the social and the organisational. Methods from the outer reaches of the social and cognitive sciences are being unearthed, and viewed with regard to whether they alone, or in some curious combination, might provide the key to designing technology which seamlessly

supports and transforms what people do. There is even a growing recognition that more traditional ways of working, and seemingly mundane tools such as pens and paper, may be worth taking seriously in design as resources for thinking about innovative solutions to organisational problems.

More strangely, notwithstanding the growing body of research concerned with the relationship between the 'social and the technical', is how little we know about the ways in which individuals, both alone and in concert with each other, use tools and technologies in the practical accomplishment of their daily work. There is little to which engineers and designers can turn to find out about how technology is used in the workplace. For example, if we consider the problems which arose with the introduction of Computer-Aided Dispatch into the London Ambulance Service, it is curious to realise that we have little idea as to how paper documents are used by personnel to co-ordinate organisational activities, or, for example, how individuals, in the course of talking to others either face-to-face or over the telephone, document relevant details of an event and the encounter. We know even less about the advantages of paper as opposed to computers, and why, despite the onslaught of new technology, our offices remain littered with documents. Our relative ignorance of the use of paper in organisations is complemented by our understanding of the use of complex systems. Despite a substantial body of research, we still have little understanding of the ways in which new technologies feature in practical organisational conduct. Individuals glance at screens, they refer to documents, they discuss plans, they send messages to each other, they turn talk into data, they discover facts and findings; that is, they use these technologies within the practicalities and constraints of their everyday activities. The ways in which these tools and technologies, even basic information systems, are embedded in and depend upon practical activities within the workplace and the practices, procedures and reasoning of personnel, remain largely unknown. Tragedies such as the introduction of Computer-Aided Dispatch into the London Ambulance Service, throw into relief how little we know of the ways in which tools and technologies, ranging from pen and paper through to complex multimedia workstations, feature in day-to-day organisational activities. It is hardly surprising, therefore, that designers turn to the methods of the social sciences to enrich their understanding of the workplace, since we have so little to say about the operation and organisation of even the most mundane objects and artefacts which inhabit our workplace.

There is, however, a growing body of research concerned with the ways in which tools and technologies feature in work and interaction in organisational environments. This body of research, commonly known as

'workplace studies', has emerged within both the social and cognitive sciences over the past decade, and largely consists of naturalistic studies, 'ethnographies', of a broad range of organisational domains. There are, for example, studies of work and technology in air traffic control, emergency dispatch centres, newsrooms, architectural practices, consulting rooms, banks, trading rooms and construction sites. These studies remain relatively unknown in the social sciences, and yet have an increasing influence on research in such areas as Computer Supported Cooperative Work (CSCW), Cognitive Science, and Human–Computer Interaction (HCI). In this chapter we wish to discuss the curious provenance of workplace studies, including our own studies, and show how they have emerged in the light of convergent debates and developments in research on HCI, CSCW, requirements engineering and organisational behaviour. Before doing so, however, it is perhaps worthwhile saying one or two words concerning how technology has formed a topic in social science and in particular sociology.

Over the past decade or so, a substantial body of literature has emerged concerned with the social aspects of technology. Perhaps the most important and wide-ranging contributions are concerned with the ways in which new computer and communication technologies are changing the character of contemporary society and in particular the organisation of work. It is argued that new technology has penetrated almost every sphere of contemporary life, computer networks, digital communications and the like permeating the private, public and market sectors and transforming the activities of individuals and more generally society. These massive changes are characterised in various ways. Bell's (1976) 'post-industrial society' has been replaced with a host of different terms, ranging from Böhme and Stehr's (1986) 'knowledge society', Kreibich's (1986) 'science society' and Münch's (1991) 'communication society'. As Knoblauch (1996, 1997) points out, the most widely accepted scientific and lay characterisation of these changes is the 'information society'. It is argued by Castells (1996) and Webster (1995) however, that the term 'information society' and its counterparts such as the 'information economy' lead to more confusion than clarification, and there continues to be wide-ranging debate as to what changes are taking place and how they should be conceptualised. Attempts to operationalise the concept of the information society and related characterisations have met with some difficulty. It is unclear how a reliable and sensitive measure of 'information' can be identified and transformed into a satisfactory empirical object. As Aldridge suggests, little attention has been paid to the semantic content or the quality of information, rather 'theorists have leapt from quantitative measurements of the volume of information and the velocity

of its circulation to sweeping conclusions about the qualitative changes in culture and society' (Aldridge, 1997: 389). More recent attempts to define the information society in terms of work and occupational structure have not added much light to the concept. As Hensel (1980), Webster (1995) and Knoblauch (1997) suggest, the idea of 'information work' has generated an array of seemingly *ad hoc* distinctions concerning the character of particular occupations, and said little about the ways in which 'information' features in the performance of organisational activities.

From a rather different standpoint, there is a growing body of research concerned with the impact of new technology on the workplace, and in particular the ways in which information and communication systems lend support to, and engender, new forms of organisation. So, for example, it is argued that the movement from more bureaucratic and, more recently, matrix forms of organisation towards 'disaggregation' and 'dynamic networks of firms' is increasingly facilitated by the widespread availability and access to digital technologies (e.g. Barnatt 1995, 1997). The argument, and in particular the idea that new forms of organisation are supported, if not engendered, by technological change, has a long-standing tradition in the social sciences. From the early writings of Marx onwards, it has been argued that technologies have a profound impact on organisational arrangements, and successive 'schools' from socio-technical systems onwards have delineated ways in which tools and technologies shape work and organisations. Whatever reservations are voiced with respect to particular approaches, at least they place the technical and social at the heart of the analytic agenda, even if, as Grint and Woolgar (1997) argue, a technological determinism underpins much of the work. Unfortunately however, despite the important contribution of studies of technology and organisations, research is principally concerned with the ways in which communication and information systems influence, and are influenced by, such aspects as the division of labour, work-force skilling and de-skilling, occupational structure and associated features such as power, job opportunity and unionisation. This is hardly surprising. Such topics and issues are undoubtedly critical to a sociological understanding of new technology and the characteristics of organisational and occupational structure, but unfortunately draw attention away from how technology features in the production and co-ordination of workplace activities.

In a very different vein, we have seen the emergence of a growing body of research concerned with the social shaping and construction of technology. For example, in a series of essays, Mackenzie (1996) and Mackenzie and Wajcman (1985) have powerfully demonstrated how the

meanings of technology shift not only between different socio-historical contexts, but also in the light of the positioned influence of particular social groups. They speak of 'natural trajectories' in technological change and how they are constituted in and through the activities of particular individuals and organisations, and contrast their own approach with the technological determinism which haunts much of the social science literature. At times the argument can appear to reproduce the sorts of arguments we find in certain organisational studies: class relations, gender and the like, shaping people's responses to and sense of particular technologies. However, the underlying theoretical argument takes a more radical standpoint on the relationship between the social and the technical, in which, for example, the technical is constituted by human activity (see in particular, Mackenzie, 1996). As Mackenzie notes, a parallel body of research with related analytic concerns has emerged over the past decade; a body of research which has begun to reconsider many of the key concepts and issues which inform our understanding of technology and social organisation. Whilst this research embodies an array of analytic standpoints, ranging from actor–network theory through to more socio-historical models, in various ways all reflect a concern with reconfiguring the social and the technical, destroying the spurious boundaries between each, and pursuing an 'empirical programme of relativism' (Bijker et al., 1990). Such research has led to a rich and rewarding body of 'thick description', 'looking into what has been seen as the black box of technology', and has powerfully demonstrated how particular groups ascribe, dispute, exclude and cohere the sense and meaning(s) of technologies. Despite the important contribution of such studies, and their powerful demonstration of the shortcomings of technological determinism, their substantive concerns and commitment have directed attention away from how technology features in mundane activities in ordinary working environments (see Button, 1993).

In purely substantive terms therefore, it is interesting to note how the use of technology in practical organisational conduct has escaped the sociological eye. It is acknowledged that new technology is having a profound impact on work and human interaction, and yet the ways in which computers and other tools and artefacts feature in the accomplishment of organisational activities have disappeared from view. Contemporary sociological research concerned with technology seems in various ways to separate systems, both technical and human, from social action, so that we are unable to recover just how tools and artefacts and the 'rest of the furniture' of the modern organisation is constituted in and through the activities of the participants themselves. Even if we ignore debates concerning agency, context and meaning, it would seem unfortunate to rest

with a sociology which treats as epiphenomenal the socially organised competencies and reasoning on which personnel rely in using technologies, whatever they might be, as part of their daily work. Indeed, for the naive, it might appear somewhat peculiar that we know so much about the social organisation of technology in one sense, yet so little about the part it plays in everyday organisational activities and interaction.

1.1 Computers and situated conduct

Despite prevailing trends within sociological studies of technology, we have begun to witness a growing interest in the ways in which complex tools and artefacts feature within practical organisational conduct. Workplace studies are concerned with the work, interaction and technology in complex organisational environments. They are ethnographies, naturalistic studies of domains such as air traffic control rooms, architectural practices, newsrooms, construction sites, banks, dealing rooms and emergency centres. In various ways, they are concerned with how technologies, ranging from complex systems through to mundane tools, feature in the practical accomplishment of organisational activities. These ethnographies are informed by various analytic standpoints ranging from symbolic interactionism through to distributed cognition, but, for reasons which will become increasingly apparent, it is perhaps ethnomethodology and conversation analysis which have had the most profound influence on the emergence the workplace studies.

Workplace studies have emerged in the light of at least three developments which have driven analytic attention towards the *in situ* organisation of technologically informed practical activities. In the first place, we have witnessed a wide-ranging critique of the more conventional models which inform our understanding of human–computer interaction, models which have permeated HCI, Artificial Intelligence (AI), and cognitive science. Secondly, there has been a growing interest in developing technologies to support collaborative activities amongst personnel who may be co-present or located in distinct physical domains. Thirdly, a series of well-publicised technological failures has led to a growing interest amongst computer scientists and engineers in finding new and more reliable methods for the identification of requirements for complex systems. These developments have been accompanied by a growing recognition that technological deployment is more complex than hitherto assumed; work practices do not necessarily change to make systems work. The provenance of workplace studies, therefore, involves a curious mix of academic debate with more practical problems.

As far back as the 1970s, Dreyfus (1972) and others, including Coulter

(1979) and Searle (1980), developed a wide-ranging critique of artificial intelligence and cognitive science; in particular the idea that computers reflected, or provided a model for, the cognitive and social competencies of human beings. The criticisms were not concerned with debunking the technical ambitions of developers of new technologies and systems, despite some initiatives, such as the understanding of natural language, proving highly intractable. Rather, these critiques were concerned with the idea that human intelligence was akin to, and could be modelled by, the operation of a computer. In particular, Dreyfus and others set out to counter the idea that human conduct could be adequately explained in terms of an individual's ability to process information through the manipulation of symbols and by developing appropriate representations, and that intelligible action and interaction is accomplished following predefined goals, plans or scripts. One consequence of these assumptions for HCI has been the idea that, by looking at how individuals use or 'interact' with technology, one might be able to discover the 'grammar of the head' (Payne and Green, 1986) or the 'structure and process of a person's mind' (Carroll, 1984). It has even been argued that, by studying the use of technology in terms of the mental models of the user, themselves based on an analogy with computers, it would be possible to design a system which mirrors the cognitive processes of its users (Norman, 1983).

Assumptions concerning the nature of human conduct drawn from cognitive science have not only permeated the theoretical work undertaken within HCI, but have had a profound influence on the methods which have been adopted. Both specific evaluations of particular technologies and general analyses of the nature of computer use tend to adopt an experimental paradigm focusing on the individual user at the workstation. Analyses of the users' activities on the computer may draw upon such measures as the time to react to happenings on the screen or how long it takes to achieve a predefined task, but these are frequently also considered against some model of the activity, concerning, for example, the 'information processing' required by the user: how the user's conceptions, plans and interpretations are processed to solve problems and execute actions through the interface.

It is not the first time, nor will it be the last, that the operation of technology has provided a model for the scientific characterisation of the workings of the human mind. Earlier this century, for example, the newly invented telephone exchange served as a metaphor for mental processes, with its inputs, outputs, plugs and wires. In the early 1980s, with the emergence of HCI, we witnessed a renewed attack on computational metaphor and its concomitant assumptions. For example, Winograd and Flores (1986) utilised speech act theory in an attempt to underscore the

social when thinking about how individuals 'interact' in and through computers, and Frohlich and Luff and others began exploring conversation analytic models which might be used to design interfaces to complex systems (Cawsey, 1990; Frohlich and Luff, 1989; Norman and Thomas, 1990). It is Suchman (1987) who, perhaps more than any other, has had the most profound impact on the ways in which we might consider human–computer interaction and the emergence of workplace studies in both the United States and Europe. In one sense, Suchman simply takes issue with the idea that action is determined by prespecified plans and goals, and provides an alternative way of exploring how individuals 'interact with' or use technologies. Her powerful critique of AI and HCI, however, coupled with her position within one of the world's leading system laboratories, namely Xerox PARC, transformed the debate and its academic and practical relevance.

Suchman begins her treatise by citing a well-documented distinction between European and Micronesian navigation. Since the Renaissance, European sailors have relied upon highly complex charts and plans to organise their voyages. Navigation consists of following the chart or plan, developed with regard to universal principles. In contrast, non-European sailors, such as the Trukese of Micronesia, use a rather different approach. They have no maps or plans, but rather utilise information provided by 'the wind, the waves, the tide, and current, the fauna, the clouds, and the sound of water'. They navigate in an *ad hoc* fashion, adjusting their speed and direction with regard to the circumstances at hand and the contingencies that arise. Suchman argues that the case provides an analogy to the ways in which cognitive science and related disciplines such as AI and HCI model human conduct. She suggests that the European navigator 'exemplifies the prevailing model of purposeful action found in cognitive science'. Human conduct is goal-oriented, and action is governed according to rules, scripts and plans. In attempting to achieve a goal, an individual may divide it into sub-goals and break down the task into a series of component actions. The individual identifies and invokes the appropriate representations, the relevant rules and plans, to meet the goal. It is this model which is perhaps best exemplified in the influential study of human–computer interaction by Card, Moran and Newell (1980, 1983) where they develop GOMS, a model developed on several layers and utilising a framework based on explicit goals, operators, methods and rules for selecting between options.

Suchman suggests that the goal-oriented, plan-based models of human conduct which inform HCI and cognitive science have a number of shortcomings. In the first place, they diminish the importance of the immediate context of action, and, in particular, the ways in which plans and