

From Ethnographic Record to System Design

Some experiences from the field

JOHN A HUGHES, DAVE RANDALL & DAN SHAPIRO*

Department of Sociology & Centre for Research in CSCW, University of Lancaster, Lancaster LA1 4YL, UK

(Received 11 March 1992; in final form 15 June 1992)

Abstract. This paper explores the issues involved in moving from ethnographic explorations of work in context to a practical contribution to system design. It does so using the example of an interdisciplinary research project involving sociologists and computer scientists in the domain of air traffic control systems. It forms a pair with another paper (Sommerville et al., 1992) exploring these questions from the perspective of our computer science partners. We characterise ethnography as a research method, and consider the differences between undertaking it for strictly sociological or anthropological purposes by contrast with interdisciplinary and design purposes. We summarise some of our results in ethnographic explications of the work of air traffic controllers, and the sociality which it manifests. We describe the dialogues involved in rendering these observations 'informative' for systems design, and the mutual translations implied in attempting to reconcile sociological with software engineering questions about supporting the work. We conclude by specifying some features of cooperative work which an engineering approach is in danger of overlooking; the ways, and limits, in which ethnographers can form a 'bridge' between users and designers; and some of the conflicts of interest entrained in generating technical change.

Keywords. Air traffic control, Command and control, CSCW, Ethnography, HCI, Interdisciplinary design, Social organisation of work, Sociality of technology, System design, Teamwork

One consequence of the growth of interest in CSCW has been to bring sociologists, and their methods, into the field of HCI and, presumably and eventually, into system design. In significant respects this is a consequence of what Grudin (1990) refers to as the "outward step" of the interface, into the social or work setting and a recognition that computers will support work more effectively if they "implicitly or explicitly incorporate social and organisational knowledge". However, one of the main problems is that the corpus of sociological theory is, by and large, ill-suited to the design process since its disciplinary interests have

* John A. Hughes is Professor of Sociological Analysis, Dave Randall is Research Fellow on the Air Traffic Control project referred to, and Dan Shapiro is Senior Lecturer in the Department of Sociology, all at Lancaster University. All are members and Dan Shapiro is Director of the Centre for Research in CSCW at Lancaster. The research is funded by the UK MRC/SERC/ESRC Cognitive Science/HCI Initiative with the title 'Social Analysis of Control Systems for HCI Design', grant number SPG 8931598. The computer science partners are Professor Ian Sommerville, Dr Tom Rodden, Dr Pete Sawyer and Mr Richard Bentley. Thanks are due to Val King, Kjeldt Schmidt, Ian Sommerville and Tom Rodden for very helpful comments on earlier drafts of this paper.

Please address enquiries to: *John Hughes*, Department of Sociology, Lancaster University, LANCASTER, LA1 4YL, UK, Tel: +44 524 65201 ext 4174, Fax: +44 524 844788, Email: J. A. Hughes@lancaster.ac.uk.

not been directed to any kind of detailed applied role to the extent that this is required within CSCW. Certainly if sociology is to emulate the contribution made by the cognitive sciences to interface design, for example, then it is not too outrageous to suggest that sociological studies of work will have to begin anew in order to develop the necessary body of knowledge and resources upon which designers can draw.¹

A key virtue of ethnographic studies is their focus upon the rich and varied 'real world' sociality recovered through a fieldworker's participation in the social life of some setting. Directed toward system use and system design, this implies placing an emphasis on studying the functionalities of a technological system as they evolve from their incorporation into the socially organised work activities of those who use them; rather than, as in many cases, functionalities as the system's designers might imagine them to be. One of the clarion calls of CSCW is the claim that not only *ought* systems to be resonant with the human world of work, and to be of benefit to that world, but that they are more *effective* if they are designed to be so in the first place (Hirschheim and Newman 1988, Schmidt 1991). Although such a claim, clearly, engages with a number of problems that will require researching, it does provide a space for ethnography within CSCW not only as a *post facto* means of showing why systems have failed but as a potential resource for informing system design.

The emphasis here is on 'potential' since there are important questions about how such studies might be incorporated into the design process. This is not a complaint but a recognition of the fact that the alliance between ethnography and systems design is, as yet, but the nomination of an agenda rather than signifying any body of achievement. In this paper we want to review some aspects of this problem using our experience of an ongoing research project which is using an ethnographic study of air traffic controllers to inform the design of an electronic flight information display configuration tool.² We do not intend to offer widely applicable panaceas but some practical solutions to practical problems that can arise in a specific interdisciplinary collaboration between systems design and sociologically inspired ethnographic studies.

Typically, the work of the system designer is directed toward, as a crucial step in the building of a system, a requirements definition, which "is an abstract definition of the services which the system is expected to provide and the constraints under which the system must operate" (Sommerville 992:56). Inevitably, the requirements document is the outcome of a process of negotiation between various interested parties (though rarely *all* interested parties) such as, in the case of commercial system design, management, engineering, sales and marketing, systems designers, and, increasingly, domain experts, in which the services to be provided by the system are reconciled and agreed upon. Requirements capture is about evolving a system specification that can deliver the required services and software development and thus take its place in a progression from a natural, or quasi-natural, language specification of a system's requirements through to the

design of an executable representation or programme.

A key objective of requirements capture is maintaining, by negotiation if necessary, consistency throughout the various stages to avoid incompatibilities arising in services the system is expected to deliver. Thus, and to put it simply, for the purpose of system design the richly textured, often inchoate, features of the 'real world' activities of the relevant domain have to be distilled in order to meet the consistency criterion: a distillation that inevitably involves making compromises among competing alternatives that may be proffered. Many of these will be related to the various, and often competing, interests of groups within the organisation. This is particularly relevant in deciding the most effective, begging the question of quite what this means, mix of human being and automated support: a problem whose solution is likely to vary as between organisational settings and organisational purposes as well as between interested groups within organisations. Seeking computer support in order to replace costly labour is one obvious objective firms within a competitive market might wish to pursue, but this is not necessarily 'effective' in cases, such as air traffic control, where safety is the overriding concern. Of course, in 'real world' situations such objectives are unlikely to be as polarised as stated here (not to speak of the various forms of obfuscations that can be used to conceal some of the purposes of innovating in computer supported work) but they do nonetheless raise an important and practical issue that CSCW design faces as does, our special concern, the role of ethnographic studies of work within it.

To summarise, if we are to take seriously the claim that more effective systems will result when their intervention 'resonates' with existing work practices, a method is required which both elaborates and explicates those practices. There is a *prima facie* case for considering ethnography to be particularly appropriate for this purpose.

1. The character of ethnography and problems of interdisciplinary working

Ethnography has a long and honourable tradition as a research instrument within sociology and anthropology.³ As a mode of social research it is minimally concerned to produce detailed descriptions of the lived social experiences and social activities of social actors within specific contexts. In anthropology, for example, such contexts have generally been aspects of the social life of pre-industrial peoples, but have also included more modern communities as well as studies of occupational cultures. In sociology, such contexts have ranged from juvenile gangs, innumerable occupations, aspects of urban and rural life, religious groups, processes of justice administration, and more. The method relies upon material drawn from the first-hand experiences of an investigator in some natural setting; a research engagement which can sometimes involve concealment of the research role. The emphasis of ethnographic work is to present a portrayal of life as seen

and understood by those who live and work within the domain concerned. It is this objective which is the rationale behind the insistence of the method on the direct involvement of the researcher in the setting under investigation. Indeed, one of the main arguments for ethnographic methods consists of an attack against variable analytic, quantitative social investigation methods, mainly survey and experimental research. Such methods, it is argued, by the structured character of their data collection processes impose the researcher's assumptions on the data. They also reify social phenomena by treating them as more defined and distinct than they really are, and presupposes that those phenomena enter into law-like regularities of association. Further, in making claims about what happens in natural settings on the basis of data produced in artificial situations, or by using prestructured questionnaires, they engage in a highly questionable generalisation procedure (Hammersley 1990, Benson and Hughes 1991).

For the purposes of this paper the following points are important to note. First, the research material produced by ethnography usually consists of the fieldworker's notes and records taken more or less at the time of the fieldwork. These will note and record what is seen, heard, talked about, questioned, participated in, and so on. In the early stages, such fieldnotes may look little more than random jottings of odd remarks, sketches of activities, quotes, and so on. As indicated earlier, it is very much part of ethnography's rationale as a research tool that it begins with a minimum of assumptions as to the detailed character of the activities, understandings and practices that are to be found in a particular setting. Its naturalistic stance requires that its analysis begin with what it is persons do within the natural settings of their ordinary lives. In other words, *how* the setting is understood by and, through these understandings, socially organised by the participants, is not presumed in advance of inquiry, but is the task of the ethnographer to discover. In an early statement of the principles of ethnography, Becker (1958) described the process of analysis as a 'sequential' one in which much of the data collection phase *and* its analysis is carried out while the fieldwork is in progress and, as such, forms an integral part of that phase of the research. That is, it uses sociological interpretations, albeit provisional ones, to inform the direction of subsequent fieldwork.

A second and important point is that although we have so far been emphasising ethnography as a method, it is not devoid of theoretical purpose. The term 'ethnography', like the terms denoting all disciplines, is a gloss for a diversity of interests, concerns and theoretical presumptions that the method can be used to serve.⁴ While these are not always clearly distinguishable, it is far from being the case that ethnography, even when simply regarded as a method, is devoid of analytical purpose. It is not as if the ethnographer simply collects 'facts' for, inevitably, he/she has to be selective in order to collect facts relevant to a research problem, even though ethnography, unlike for example the experiment, entertains a minimum of selective principles. One principle it has to embrace as a *sociological* method is that the work site of the study is socially organised, and

ethnography's task being to uncover *how and in what ways it is socially organised*. In other words, for sociology the sociality of work, the sociality of the design process even, is not a discovery, an additional variable to add to reduce the "unpredictable source of variation" (Rasson, Maas and Kellog 1988) in the HCI equation, so much as the exploitation of a point of view which, as an analytic practice, can be applied to each and all topics of sociological interest.

Although there is a large and disparate literature on the methodology of ethnographic inquiry, much of it is, not unnaturally, directed at the methodological and theoretical concerns of sociology and anthropology. Relatively little attention has been paid to the problems of using the method within an interdisciplinary context, such as serving a discipline, that of computer science, which has its own distinctive ways of working. On the face of it, the very virtues of ethnography for some kinds of social inquiry, such as its attention to the diversity of 'real world' social life, its activities and its settings, the remit to uncover that social life as constituted in and through the understandings and activities of its participants, and its reluctance to presume much about the character of that life in advance of inquiry, would make the task of informing system design a very difficult one. Although ethnography, like any method, has to be selective in its focus, apart from the general presumption that the setting under investigation is socially organised, further focussing is shaped up during and in response to the investigation itself. The ethnographer's task is to gain access to and knowledge of the social practices, knowledge, beliefs, attitudes and activities, etc., as exhibited by participants in some 'natural setting', and to present these in terms of a sociological account of a 'way of life' as organised by its participants. However, the rich, highly detailed, highly textured, but nevertheless partial and selective descriptions associated with ethnography would seem to contribute little to resolving the designer's problem where the objective is to determine what should be designed and how. In the requirements capture phase of the design process this is, normally, resolved through negotiation with and among different classes of users and system procurers whose often differing needs and perspectives have to be reconciled in some way rather than celebrated, as the ethnographer might do, in their diversity.

Nevertheless, such a problem, and it is really an agglomeration of problems, are not objections in principle to using ethnographic methods to inform system design, but more a statement about the importance of trying to identify, as adequately as possible, the lineaments of the problems and the role that ethnographic methods can serve.⁵

2. The design problem of the electronic flight strip

We now try to illustrate these issues by drawing lessons from a particular interdisciplinary research project in the domain of air traffic control, involving sociologists and computer scientists working in collaboration. We had the advantage of

having conducted a prior study of air traffic controllers (Hughes et al. 1988, Harper et al. 1989, Harper et al. 1991). On this occasion, however, the study focussed on the work of the controllers in using the paper 'flight progress strip' – a strip of card measuring about 20 cm by 3 cm containing formatted information about a particular flight – and with the objective of contributing to the design of an electronic equivalent. The airspace over England and Wales is typically divided into 16 sectors⁶ and these sectors are controlled from eight radar suites at the London Air Traffic Control Centre (LATCC). The suite provides the technical facilities for communications with aircraft and with other suites, and provides the radar and flight progress strips for displaying the position and the relevant characteristics of the aircraft within the control responsibility of the sector. The radar is a 2-dimensional, 'real time' representation of the sector's airspace and the present position of aircraft within it. Each aircraft is indicated by a 'blip' and an adjacent data block showing its call sign and height. The radar can display other information about the sector including sector boundaries, coast lines, major and minor airports, etc., which are under the discretion of the controller. However, the radar does not constitute a complete picture of the traffic situation since 'the situation at any point in time' is not simply what is happening 'now' but also what 'might be happening' in a few minutes time and this is immensely relevant to what actions the controller needs to do 'now' by way of anticipating likely problems. It is in this respect that the flight strip is important, along with the radar, as part of a mutually explicating process of making sense of the traffic pattern and, through this, directing it as appropriate.

Briefly, the flight strip is used to represent the passage of an aircraft through controlled airspace.⁷ It is produced from a database consisting of flight plans submitted by pilots through a Flight Data Processing application and filed prior to departure. Updating this information *in the database* is done through limited radar links with the database through Radar Data Processing (RDP), by inputting flight data via the Strip Input and Display Device (SIDD) attached to the 'wings' of every suite or, on occasion, by the radar controller through input devices on the radar suite itself. The information for each flight is printed on a paper strip via the computer up to 40 minutes prior to the aircraft reaching navigation points on the relevant sector. The strip contains formatted information relevant to the flight including call sign, flight level (height), destination, next reporting point on route, eta at next navigation point, radio 'squawk' code, aircraft type, planned flight path, and so on. This information is 'historical' and not a real time record of flight progress in that updating the database which originates the strip is not a continuous process. The real current condition of the sector is therefore represented not in the database but in the manual annotations and symbols written by the controller and the team on the strips themselves.⁸

There are some fairly obvious bonuses that could be achieved by the electronic displays of flight strip information, the main ones having to do with updating the

data base with information about the controller's actions and, through this, offering more effective ways of linking radar and strip information (Shapiro et al. 1991). In particular, this could be useful in building a more effective Conflict Alert Facility.⁹ However, given the incremental style of innovation which necessarily dominates in Air Traffic Control, any technical change would need to involve a minimum of changes in working practices not only to avoid the huge cost of retraining but also because the strip system, to call it that, is only a part of the complex of systems that is air traffic control. In effect this means that any electronic representation of the strip would need to capture its current functionalities in order not to impoverish the accomplishment of controlling by misrepresenting and misspecifying the tasks and activities involved (Hughes et al. 1988).

Thus, the first task of the research was to provide a description of the activities performed by controllers during the work itself. This required the ethnographer to spend prolonged periods of time observing a team working a sector at LATCC. In effect, the fieldworker was required, from the start, to 'learn the ropes' of controlling by watching, asking questions, talking with controllers.¹⁰ Meanwhile, broad design principles of the electronic display system began.

The interests of the computer scientists were concerned to illuminate some of the problems of database visualisation in interactive command and control systems using air traffic control as the exemplar. One of the main problems is designing a user-interface that does not overload the user with information but presents only relevant items in a readily comprehensible fashion. In this case the database is a 'reactive' one in that it can be updated in real-time by both users and external sensors. As previous experience with electronic display systems in ATC had shown, there are immense subtleties to and variations in the way controllers work, and so an important objective was to design a prototype tool which would allow controllers to tailor the system to support different working practices and personal preferences, yet be consistent with the database information stored in real-time. Also, the environment should allow for the construction of a multi-user, multi-screen system where the same database entities could be simultaneously displayed on separate screens and manipulated by different users (see Bentley et al. 1992 for further details). However, apart from these main specifications, it was up to the ethnographic research to describe the working practices to inform the rapid prototyping of the tool.

From the sociological side, and the main concern of this paper, at the outset we made the methodological choice not to take on board the concerns of the system designers. Not because these were unimportant, on the contrary, but to try to conform to one of the strategic injunctions of the method in entertaining as few preconceptions about the setting and its activities as possible in order to better display the 'real world' character of its actual working practices. We felt, rightly or wrongly, that it was important not to prejudge what might be found by taking on board the designer's constraints and problems from too early a

stage. While both sides agreed that an effective strip display system would have to resonate effectively with existing work practices, it was the task of the ethnographic study to specify what these were without any presumptions about how they could or could not be supported by any subsequent computer support system.¹¹

3. Explicating the work of controlling: the first run through

The theoretical-methodological posture guiding the fieldwork was the decision to explicate the sociality of work as it is understood, experienced and organised by those who do the work. It was Garfinkel (1967) who stressed that sociology's disciplinary choice was to account for human activity purely as social action; that is, to treat the describable properties of activities in a social setting as the 'outcomes', 'accomplishments', or 'achievements' of those participating in it using their practical commonsense, mundane knowledge of how the work and its activities are organised.¹²

It was our contention that only by describing the controller's activities as part of a social system in the above fashion could we possibly talk about them in a way which begins to capture their identity and to convey the sense that these activities have both for those doing them and those otherwise involved in them, as activities *within* the setting itself. The controller at the console is carrying out 'controlling activities' and, therefore, it is only through a description of the details of these *as activities of controlling* that we can give other than behaviouralised and reductive descriptions of the work. It is only through detailing the features of the 'air traffic control system' as recognisable and familiar features of the working environment of the controllers, that we can spell out the sense and depict the 'working organisation' of those activities. In other words, the adequate identification of the state of affairs in the sector has to be determined by a relatively protracted study of the 'system from within' as a 'real world' socially organised working environment.

After the first period of fieldwork, of approximately one month, lengthy debriefing sessions took place based on the fieldnotes and the experiences of the researcher. This required going into the work activities in some detail and although it is not the purpose of this paper to go into the kind of detailed accounts produced, a number of important features of the work of controllers emerged.

Controlling is a teamwork activity. A typical radar suite which normally deals with two sectors is manned by a chief, two radar controllers, and the assistants, or 'wings'. Sometimes London Joint Operations (LJO) officers of the RAF, who liaise with civilian controllers concerning military aircraft, are in close attendance. Each of these are responsible for particular tasks, though there is a very important informal sharing of some of them. 'Wingmen', for example, though

performing what might be regarded as the more routine servicing tasks, such as attending to the printed strips on their delivery, tearing them off from the printer, placing them into their correct coloured holders, ensuring that each strip is placed in the correct rack ('Live', 'Pending', or 'Transfer'), also check the information on the strip. Errors stemming from filed flight plans are not uncommon and the 'wingman's' initial checking of that information while they are performing other routine tasks is an important first level check. They also prepare substitute strips where, for whatever reason, the controller has to deal with an aircraft before a 'live' strip is available. As time allows, 'wingmen' will also input route and eta changes via the SIDD.¹³ The chief's responsibilities normally include coordinating traffic between sectors by telephone calls or by visits to the relevant sector suite. The main task, however, is monitoring the flow of traffic at a strategic level, making judgements about the allocation of flight levels, and so on, to facilitate the flow and act as a backup to the radar controller by dealing with unscheduled traffic, drawing attention to potential problems by 'cocking out' strips, or by writing on them.¹⁴ The chief also coordinates military aircraft with the LJO.

The work of the radar controller is the focus of the teamwork on the suite. The controller issues instructions to aircraft and verifies on the radar that the instructions have been carried out. It is their task to coordinate the air traffic safely and expeditiously by maintaining separation levels, directing course changes, ascents and descents. An important feature of this work is that controllers have to take the traffic as it comes and blend incoming aircraft into a steady flow. To this end they spend time organising strips, especially 'live' strips, into a sequence of aircraft, using the radar to check and assess the current 'state of play' and, again in conjunction with the strips, anticipate likely problems. They also spend time coordinating activities with other members of the team, arranging flight levels, discussing 'best' procedures to expedite the flow of traffic, sometime passing flight level changes to 'wings' to coordinate with adjacent sectors, and so on.

However, what is important about the teamwork is not so much that the respective team members have individual tasks to perform, which they do, but that they are woven into a 'working division of labour'. The individuals are individuals-in-a-team, and much of their work consists in the ability to organise the distribution of individual tasks into an ongoing assemblage of activities *within* the 'working division of labour' (See Anderson et al. 1989 for further explication of this notion).

The teamwork character was manifested in myriad ways in the ongoing work itself. It was not merely an adjunct to the work but integral to it. There was no one thing, or set of things, that manifested teamwork since it saturated the work itself – in the talk around the suite, in the activities performed, in the stories and anecdotes told, and more. Indeed, this was much the way in which the ethnography recorded the work, by stories, descriptions of activities, anecdotes, vignettes of events, and so on. But teamwork is very much an aspect of its members' tacit

knowledge which, though it can be and often is articulated in the ways indicated, is not normally explicitly expressed except on the occasions when 'normal and routine' working breaks down.

The above is, of course, only a sketch of the details of the working practices of controlling. The debriefing of the ethnographer provided a much richer picture than the fieldnotes themselves could possibly have furnished. Further, and this is important, the sessions were not intended to provide the rest of the research team simply with an overall picture of the controlling process but with a very detailed one; what Geertz (1973) in reference to ethnography, referred to as 'thick description'. In particular, and not surprising given the objective of the research, much of the debriefing questioning focussed on how the flight strips entered into the teamwork and how they were used within the 'working division of labour'.

4. The strips in the working division of labour

Given the conception of controlling work as teamwork, we could not see the strip simply as a fixed technical artifact but as an instrument whose functionalities derive from the work itself; as part and parcel of a working division of labour.

For those who work the technology of the strips, the strips are the key to 'good controlling'. As one controller put it, 'You have got to have a complete picture of what should be in your sector and what should be in your sector is on those strips'.¹⁵ Strips become a documentary record of an aircraft's passage through the sector. The initial information on the strip contains relevant data concerning an aircraft, information which is amended as instructions are given to and confirmed by the pilot. The strip, thus, accumulates information about actions taken and recorded on the strip according to laid down protocols. In this respect the strip is a note pad conveying to members of the team what actions have been taken with respect to particular aircraft and, importantly, with respect to *this* aircraft in its relation to others in the traffic configuration. Such amendments may be done by the controller, by the chief or, less often, by one of the 'wings'. 'Attention-getting' information may also be written on the strips, such as arrows indicating unusual routes, symbols designating 'crossers, joiners and leavers' (that is, aircraft crossing, leaving or joining the major traffic streams), circles around unusual destinations, and so on. Information indicating that coordination between sectors has taken place, changes in eta or to route, changes in call sign, etc., can be indicated using standard notations, each member of the team using a different coloured pen to annotate the strips, thus adding responsibility and accountability to the record. The strip, in other words, is a public document for the members of the team; a working representation of an aircraft's control history and a work site of controlling.

But information *on* the strip is not the only way in which the strip is informative for the controlling team. The active organisation of the strips themselves on

the Flight Progress Board, or 'racks', by the controller is an essential feature of the work. At its simplest and most general, the controller's problem is a scheduling one. For any segment of airspace the traffic has to be taken as and when it arrives and threaded into an orderly pattern before each individual plane is handed on to the next sector. This scheduling has to be achieved in and through making the traffic flow. Aircraft cannot be parked and even when they are 'on hold' they are still on the move and part of the traffic flow. Organising the strips into bays on the racks is a way of achieving a solution to the scheduling problem by organising the strips as a scheduling of work tasks. 'Live strips' are set out on the rack in front of the controller and represent aircraft which are, or are about to be, under control. These are the current work tasks. On the 'wings' are 'pending' strips which are early indications of traffic due. They provide a picture of the traffic due and the requested routes. They can also be useful in resolving scheduling problems in planning flight levels at which aircraft can enter the sector.¹⁶ A second bay on the 'wings' holds the 'transfer' strips which depict aircraft at their last reporting point and which will be transferred to other controlling centres.¹⁷

Live strips on the Flight Progress Board are actively organised such that their array displays features of the work. Strips are aligned and organised under a number of yellow plastic Designator strips representing the reporting points over which a flight will pass. To the accomplished controller this display is an 'at a glance' means of showing the flow of traffic through the sector and its characteristics. To the experienced and knowledgeable eye, different indicated eta times and different heights may represent problems, for example, 'climbing through' or 'catching up' given the relative performances of aircraft. The way in which the controller organises the strips, for instance according to arrival time over a reporting point, or according to flight level, or to possible confliction points, provides information about the state of the sector. Moving the strips is to organise the information in terms of work activities and, through this, accomplishing the work of organising the traffic.

The importance of the strip to the controlling process is difficult to overestimate. With the radar, and depending on individual controlling styles, the strips are the means by which controllers see and note what is happening, what they have already done, what needs to be done. They are an essential feature of 'getting the picture', 'organising the traffic', which is the means of achieving the orderliness of the traffic. The strips and their organisation are a proxy orderliness of the configuration of the traffic flow. While the radar is a computer generated 2-dimensional picture of relevant aspects of the sky and its traffic, the strips are a means by which the patterns on the screen, and thus in the sky, can be seen *as the patterns that they are*. Strips are organised so as to give a sequence to them which both contributes to and reflects the fact that the management of the traffic is an inherently sequential task, no more so than in the fact that the traffic *flows* through the airspace.

5. Towards system requirements: the specifying dialogue

The above account of 'working the strips' is, to repeat again, only a truncated description of the organisation of work around the suite, and one which particularly focusses on the role of the strips. It is also a 'worked up' account heavily informed by our own sociological interests. However, this is not the way in which it emerged in the discussions with our computer science colleagues. While more than readily appreciative of the idea of the social organisation of work, they were less interested in the *sociological* issues surrounding the notion. For them, and reasonably so, the issues centred around the *detailed and specific* explication of the ways in which the strips were worked with, the significance of strip markings, how the connection between the strips and the radar was effected, and more. Such concerns were not, of course, antithetical to our sociological interests, rather they were interests motivated by different relevances, namely and broadly, their implications for the requirements the electronic display tool would have to serve. Further, even at this early stage there was a sense in which the ethnography, whilst relatively undirected, was already impacting on discussions about the scoping of the system.

The important point we want to stress here is the process by which the experiences and materials of the fieldworker were directed toward the concerns of the system builders. It was very much a dialogic process in which, initially, the fieldworker was interrogated by the rest of the research team. As indicated, this took the form of questions such as 'what happens if...?', 'how is this done...?', 'what would be the consequences of doing X?', and so on: dealing very much with the practicalities of the work rather than disciplinary issues about the work. Scenarios were evolved, such as following an aircraft and its strip through a route, noting how departures from the 'routine' were dealt with; trying to distinguish 'normal' and 'exceptional' circumstances; imagining what might be the result of changing the strips, the working practices in specific ways; estimating the wider constraints on the system, such as airspace configuration, and more. It would be wrong to portray the discussions as systematically organised topic by topic; it was much more inchoate than that. But, and significantly, it was very much a way, as mentioned earlier in reference to Becker (1958), of sequentially analysing the material and the experiences of the fieldworker and directing these toward the aims of the research. The dialogue achieved a number of things. First, it acquainted members of the research team with the work of controlling and did so in some detail. Second, it gave a direction to further fieldwork by identifying matters that needed checking and matters that needed further inquiry. Third, it informed a number of design decisions for the further prototyping of the system. In this respect, the ethnography generated problems but ones germane to reviewing design possibilities and options. Later, the computer scientists, after some deliberation on the debriefing sessions, posed four 'guiding questions':

- what characteristics of the existing manual system are unimportant and need not be supported in a computerised system?
- what are important manual activities which need not be supported in a computerised system because the activities are a consequence of the fact that no computer support is available?
- what characteristics of the manual system must be replicated without change in a computerised system?
- what activities from the manual system may be supported in a way which is different from that used in the manual system?

From our point of view as sociologists these questions gave more directed attention as to how the ethnographic analysis of the work should be examined and analysed and, also, raised in a practical way some of the problems of interdisciplinary collaboration acknowledged in principle in the CSCW literature (see for example, Schmidt 1991). For us, such questions were not easily answerable by reference to work which is as subtle and complex as our ethnographic analysis had shown controlling to be. Our methodological point of view of treating the system as a 'motile configuration' of working practices, activities and technology, made it difficult for us to draw the distinctions necessary to provide answers to such questions. For example, and as discussions with controllers later confirmed, the actual performance of some 'manual' activities, such as writing on the strips, manipulating them in the racks, coordinating by telephone with adjacent suites, and more, serve to keep the controller, and other members of the team, 'geared into' the work. Though some of these could be automated, the important question of whether this would reduce the controller to a more passive role of system monitor rather than an active participant in a system-in-use, and whether or not this is a good thing, is hard to determine from the ethnographical analysis alone.

Further, and a point relevant to interdisciplinary working, the questions posed *seemed* to reduce issues about the social organisation of controlling work to a series of apparently individualistic activities giving scant justice to their *social* character. The emphasis on the sociality of work requires treating that work as embodied courses of action done by a working team which achieves a division of labour; courses of action not adequately decomposable into a set of discrete tasks. And, furthermore, there is a world of difference between designing a system which simply removes some tasks from the manual realm by providing automatic alternatives which facilitate the work done, and a system which reconfigures the work as a whole while preserving the crucial involvement of the user *within* the system-in-use. As we have indicated, decision-making in controlling is a cooperative process evolving in and from the understandings achieved by members of the team. To the extent to which the work displays its apparent seamlessness, this is due to the smooth operation of the 'working division of labour' that is tacitly understood and trusted in and through the active involvement of team members in the 'doing' and the organisation of tasks. But, from the

point of view of the system designers, decomposing the work into a set of discrete tasks is what is required in order to build the system, and its software, so that it performs consistently and efficiently. Of course, such differences between the ethnography and the tasks of the system designer are not antithetical even though they did mark, in our case, an important stage on the learning curve of interdisciplinary working.

6. Conclusion

In trying to review the lessons of this experience for system requirements, it is important to remember that the 'product' is small in scale. Designing and building the system is essentially the work of two or three persons rather than, as in the case of large scale systems, teams of designers and programmers. Thus, the system design has been able to go hand-in-hand with the ethnography through a process of rapid prototyping which will eventually culminate in a clearer requirements understanding. A larger scale system would present other kinds of problems. Within these limitations, however, there are a number of features of the research worth noting.

First, although design inevitably involves making choices – and the further into the design process, the more the consequences of choices initially made become constraining – the ethnography, with its emphasis on identifying controlling activities as activities entrained and sequenced within a 'working division of labour', has meant that some design options have been informed by their likely consequences for working practices. Thus, and for example, controllers encounter a variety of problems that obstruct the expeditious 'flow' of the traffic. Among these is the need for a substantial amount of coordination, especially between sectors, which is compounded, on occasion, by the inadequacies of the information on the printed strips because of the limited updating facilities. Such problems can relatively easily be tackled by building in the appropriate technology. However, although coordination does 'slow up' the work, it is a channel through which much of the artful work of scheduling and maintaining safe separation is done. Moreover, and importantly, it is a means of forming the very essential attentiveness of the controller to the specific tasks 'in hand'. 'Doing things', such as marking strips, organising them in the racks, glancing up and down them, using a pencil on the screen to gauge a vector, effecting coordination with adjacent sectors, and more, are all ways of 'embodying' attentiveness to the tasks of the work and essential to avoiding the controller becoming merely a monitor. To put the point more generally, the ethnographic portrait of the activities as part of a socially organised setting avoids some of the pitfalls in treating tasks as discrete, isolated chunks of behaviour as if they were representations or descriptions of how the work and its tasks is actually done. Identifying the skills, how they are deployed, how work activities are sequenced and how they are

made to connect accountably and recognisably as 'controlling activities' is important to any aspiration to blending systems with working practices. The sensitivity to the place activities have within the totality of activities that constitute controlling work highlights their interdependencies in ways that are not always obvious.

This brings us to a second and related feature of ethnographic work of this kind, and often regarded as one of the prime virtues of the method, namely, the potential it offers for uncovering the tacit knowledge and understandings that are embedded in activities; knowledge and understandings which by virtue of their very familiarity to those performing the tasks, they find it very difficult to articulate. The ethnographer is an 'incompetent' in the culture and can ask the kind of naive questions that would not be permissible to experts. While the capacity to do this is dependent on the setting and the roles it furnishes for the fieldworker, it is also crucially dependent on the kind of relationship the fieldworker can establish with the relevant parties. Of course, the fieldworker cannot be expected to acquire knowledge of the domain in ways that, in this case, controllers know it. Indeed, it can be argued that it is essential that he/she does not since this would be tantamount, to use anthropological argot, to 'going native' and thereby losing the analytic and theoretical attitude that is the point of the method. Thus one of the roles an ethnographer can perform is to act as a bridge between the domain and the system designer; a task that can, in part, be achieved through supporting a dialogue between users and designers, either as part of a design team which could include different kinds of relevant expertise or, as in our case, a very small team of developers. Though it is the ethnographer's task to acquire knowledge of the skills and expertise of the domain, determining the relevance of these for the system designer has to be a matter of mutual interrogation. Fairly obviously, immersion in any social situation for a period of months will mean acquiring knowledge of lots of things about that situation as well as, through growing familiarity, becoming less and less sensitive to the how and the why of things that go on within the setting. The ethnographer will begin to drift toward becoming a 'native'. Accordingly, periodic distancing from the field is important not only to gathering the information, understandings and knowledge the fieldworker has acquired but also to establish the relevance of that knowledge for the research. In a way, the ethnographer is a surrogate for the domain and – paradoxically – can only effectively serve as such if sufficient distance is sustained. It is for the designer to probe and explore possibilities and have them confirmed, doubted, rejected through the dialogue.

A third feature concerns the relationship with that somewhat mythical category, and certainly not a homogeneous group, 'users'. Some years ago, Becker (1966) raised a question for sociology, 'whose side are we on?', in a response to issues about whether sociology should express a commitment or should emphasise neutrality and value freedom in its inquiries. Although it is not an essential part of ethnography's remit that it has to be partisan in quite the way that Becker

sets up the issue, in the context of system design it is eminently suited to exploring and explicating the working tasks and the socially organised world of actual users, and can do so in a number of ways some of which we have already reviewed. However, and pulling on the second feature noted, no matter how competent the ethnographer, he/she can never know the domain as users know it. While it is important to involve users in the design process, and there are other aspects to this we will return to in a moment, this is by no means a straightforward process. As indicated earlier, users often find it difficult to articulate what it is they know since the knowledge that enters into the skilful execution of working practices is not easily summarised as lists of decontextualised propositions, be they formally specifiable or tacit, but is highly localised and a matter of constant enquiry and discovery (Schutz 1967). It is not that users cannot talk about what it is they know, how things are done, but it needs bringing out and directing toward the concerns of the design itself. In this respect, the ethnography can serve as another bridge between the users and the designers. In our case, controllers have advised on the design of the display tool with the ethnographer, as someone knowledgeable about but distanced from the work, and, on the one hand, able to appreciate the significance of the controllers' remarks for their design implications and, on the other hand, familiar enough with the design problems to relate them to the controllers' experiences and comments.

An aspect of the relationship with the user which is also of some significance, and noted at the beginning of the paper, is the fact that the user is normally situated within wider organisational structures which generate their own interests and concerns, and in which any system will become implicated. In the case of ATC, the successful design of an electronic flight data display system, while having every chance of increasing efficiency and safety, would also threaten jobs, particularly those of the 'wingmen'. Further, although controllers are sceptical of technological innovations until well and truly proven, they also realise that in order to cope with current and predicted traffic densities, innovations of this kind will be necessary. For management, the problems are to bring on such innovations as fast as possible. In other words, 'real world' organisations are full of divergent interests, conflicts, inefficiencies, all of which can cross-cut each other in ways that are often difficult to predict. These are also factors which the ethnographer cannot help but be aware of and given the close relationships he/she often establishes in the field, there is a temptation to take sides. While this does not necessarily compromise the ethnographer's objectivity, it does highlight another matter to do with the partiality or the selectivity of the ethnographer's research. This issue has been an enduring one in sociology; one in which there have been many misconceptions, misunderstandings and confusions. It is not our task to go into these here, but merely to point out that selectivity is an inevitable feature of all research; the important question is about the relevance of the research for the question being researched. Not that this is a straightforward matter. To use our own example, we can be sure that we have not covered everything of relevance

in explicating controllers' working practices, every nuance, every subtlety, every possible contingency such as air crashes, 'hijacks', and more which a system must be able to cope with, however exceptional. But then, nor can any other piece of research make such a claim and nor can any system design.¹⁸ It is a commonplace to note that design inevitably involves making choices and there can never be any guarantee that the choices made are the right ones in an ultimate sense. What they can be are choices made in light of the best information available at the time. What we do claim is that within the constraints of the research, within its modest goals, we have evolved a way of interdisciplinary working which provides a practical way of incorporating ethnography into system design and building. Of course, and as already indicated, it is unlikely that the method is scalable upwards to larger systems which require industrial methods to build and which, because of this, involve a very different level of granularity. But on this matter we have no hard evidence either way.

References

- Ackroyd, S. and Hughes, J. A. 1992. *Data Collection in Context*, 2nd ed. London: Longmans.
- Anderson, R. J., Hughes, J. A. and Sharrock, W. W. 1989. *Working for Profit: The Social Organisation of Calculability in an Entrepreneurial Firm*. Aldershot: Avebury.
- Bannon, L. J. 1985. *Extending the Design Boundaries of Human-Computer Interaction*. ICS Report 8505, University of California, San Diego, USA.
- Becker, H. (1958), Problems of Inference and Proof in Participant Observation. *American Sociological Review* 23: 682–90.
- Becker, H. S. (1970), Whose Side Are We On? In *Sociological Work*, 123–39 Chicago: Aldine.
- Benson, D. and Hughes, J. A. (1991), Evidence and Inference. In *Ethnomethodology and the Human Sciences*, ed. Button, G. Cambridge: Cambridge University Press.
- Bentley, R., Rodden, T., Sawyer, P., Sommerville, I. (1992), A Prototyping Environment for Dynamic Data Visualisation, *Proceedings of the 5th IFIP TC2/WC2.7 Working Conference on Engineering for Human-Computer Interaction*, ed. Unger, C. and Larson, J.A., Elsevier Science Publishers.
- Blumer, H. (1955), Sociological Analysis and the Variable. *American Sociological Review* 21: 683–690.
- Garfinkel, H. (1967), *Studies in Ethnomethodology*, Prentice Hall: Englewood Cliffs.
- Geertz, C. (1973), *The Interpretation of Cultures*. New York: Basic Books.
- Grudin, J. (1990), The Computer Reaches Out: The Historical Continuity of Interface Design. *ACM*.
- Grudin, J. (1991), CSCW: The Convergence of Two Development Contexts. *CHI*.
- Hammersley, M. (1990) What's Wrong with Ethnography? The Myth of Theoretical Description. *Sociology* 24: 597–615.
- Hammersley, M. and Atkinson, P. (1983), *Ethnography: Principles in Practice*. London: Routledge.
- Harper, R. R., Hughes, J. A. and Shapiro, D. Z. (1989), *The Functionalities of Flight Data Strips*. Report for CAA, Department of Sociology, Lancaster University.
- Harper, R. R., Hughes, J. A. and Shapiro, D. Z. (1991), Working in Harmony: An Examination of Computer Technology in Air Traffic Control. In *Studies in Computer Supported Cooperative Work: Theory, Practice and Design*, eds. Bowers, J. and Benford, S.D., Amsterdam, Elsevier, North-Holland.
- Hirschheim, R. and Newman, M. (1988), Information Systems and User Resistance: Theory and Practice. *The Computer Journal* 5, Vol. 31.
- Hughes, J. A., Shapiro, D. Z., Sharrock, W. W., Anderson, R. J., Harper and R. R., Gibbens, S. (1988), *The Automation of Air Traffic Control*. Final Report, SERC/ESRC Grant No. GR/D/86157, Department of Sociology, Lancaster University, USA.
- Rasson, M. B., Maas, B. Kellogg, W. A. (1988), The Designer as User: Building Requirements for Design Tools

- from Design Practice. *ACM Communications* 11, Vol. 31.
- Schmidt, K. (1991), Riding a Tiger, or Computer Supported Cooperative Work. In *Proceedings from the 2nd European Conference on CSCW*, Dordrecht: Kluwer Academic Publishers.
- Schutz, A. (1967), *The Phenomenology of the Social World*. Evanston; North Western University Press.
- Shapiro, D. Z., Hughes, J. A., Randall, D. and Harper, R. R. (1991), Visual Re-representation of Database Information: The Flight Strip in Air Traffic Control. In *Proceedings, 10th Interdisciplinary Workshop on Informatics and Psychology. Cognitive Aspects of Visual Language and Visual Interfaces*, Scharding.
- Sommerville, I. (1992), *Software Engineering*. New York: 4th Ed. Addison-Wesley.
- Sommerville, I., Rodden, T., Sawyer, P. and Bentley, R. (1992), *Sociologists Can Be Surprisingly Useful in Interactive System Design* Cambridge University Press. To appear in *Proceedings of HCI Conference*, York University, Sept.
- Stanley, L. (1990), Doing Ethnography, Writing Ethnography: A Comment on Hammersley. *Sociology* 24: 617–627.

Notes

1. This is in part a consequence of perfectly legitimate choices of interest and focus. To study working life from, for example, the perspective of whether or not it is conducive to the formation of a radical class consciousness, is quite properly different from a detailed mechanics of the social organisation of work. In part, however, sociologists have been remiss in taking this social organisation for granted as something easily achieved – other than politically – on the part either of workers or of management. In addition, it is worth noting that one of the problems of interdisciplinary working is that disciplines are likely to bring their own specific problems along with them. In sociology, for example, sociological disputes between, say, Marxist and Ethnomethodological approaches are not likely to prove endearing or enlightening to the members of other disciplines. There is no easy resolution to such problems, of course, except by working them through within their disciplinary home which might have to take on board, as a sociological issue, studies of work done within the interdisciplinary context of CSCW!
2. See Bentley et al. 1991.
3. See, for example, Hammersley and Atkinson (1983), Ackroyd and Hughes (1992) for reviews of this tradition.
4. In connection with ethnography see Stanley (1990) as a reply to Hammersley (1990).
5. In the face of such problems, some have argued that interdisciplinarity must reside in practitioners in the human sciences becoming designers and vice versa as the only effective way of determining respective disciplinary relevances. We ourselves are doubtful about the practical prospects of achieving such a symbiosis. While there is clearly a need for some awareness of each other's way of working, as well as a tolerance for them, effective mastery of bodies of knowledge involving years of apprenticeship would seem to require skills of a neo-Renaissance or post-Faustian kind!
6. Depending on traffic densities sectors can be split or merged.
7. UK airspace is divided into blocks of sky which are distinct areas of control responsibility. They have a floor of 5,500ft. except around airports where the floor descends to the ground. In upper airspace, that is, above 29,000 ft. aircraft are still under ATC control but certain procedures, such as horizontal separation, are different.
8. See Harper et al. 1989 for further details.
9. The 'problem' with the current Conflict Alert Facility is that it cannot take controllers' actions into account. That is, on occasion a controller may deliberately leave or place aircraft on conflicting headings in order to make more effective use of scarce airspace. Of course, the intention is to change the headings well in advance of any violation of separation requirements. However, it is worth pointing out that although initially controllers were reluctant to use this facility, they have now come to rely on it rather more despite the problems just noted.
10. There is a great deal of discussion in the methodological literature on ethnography to do with the appropriate role to adopt when in the field. For our part, we made no attempt to conceal the research role but did insist that the field worker worked shifts with the controllers to establish 'street credibility'.
11. In part, this constraint of the research was a consequence of a realisation that innovation in ATC has to be

incremental rather than revolutionary since any innovation has to be adopted while the whole ATC system is in operation. Although our research is not directed at designing a system which will actually be used, it is a realistic assumption to take on board and one which, we suspect, applies to many cases of system design.

12. This is a different approach to the attempt to specify task description in isolation from the 'real world' settings in which the work is done.
13. This is only done when specific and important pieces of information become available.
14. The strips are arranged in upright racks and can be slightly pulled out of position horizontally, 'cocked out', to bring them to attention.
15. There are signs of generational differences here. Although many controller's expressed various degrees of frustration with the strips, one or two the desire to 'get rid of them altogether', most controllers do see their point within the current organisation of work. A number of younger ones, however, are more receptive to innovations, such as electronic strips, provided that they are effectively designed.
16. They can also be used in lieu of live strips when, for whatever reason, the latter are not output in time. They can also give advance warning of input errors, such as incorrect flight plans, and, in the event of computer failure, serve as a basis for manual calculation.
17. Coordination of these flights is usually done by 'wingmen' since it is a relatively routine procedure though time consuming.
18. This is not, of course, to dismiss the problem of attempting to anticipate exceptional circumstances which must be tackled as a crucial part of the design process.