Designing with Ethnography: Making Work Visible

One of the main challenges that has come to face system designers in the last decade had been to accomplish an adequate characterisation of the social properties of the work settings into which systems have to fit. Grudin (1990:261) describes this as an "outward movement of the computer's interface to its external environment, from hardware to software to increasingly higher-level cognitive capabilities and finally to social processes". This is particularly apparent in the development of groupware and, potentially, thoroughly CSCW systems. While applications to support cooperating individuals and groups have met all the usual problems of interface design they also have to face up to the additional problems of designing systems which, to quote Grudin (1990) once again, "incorporate organisational and social knowledge".

The rationale behind this shift is not hard to understand. One is an increased emphasis on usability to avoid the huge training and familiarity overheads associated with new systems. A second is that the move toward the development of systems which are intended to support cooperating individuals, relating design to the subtleties of working practices introduces the vital importance of understanding the social context into which a system will have to fit.

Because of the complexity and subtleties and the inherently dynamic nature of working practices, the approach taken in some CSCW systems where a formal model of an activity is produced and used to support the activity is of very limited utility. For example, the COSMOS and AMIGO projects (Wilbur and Young, 1988; Danielson, 1986) developed activity modelling notations but their examples have been limited to very simple situations such as voting. While we accept that this approach may have a limited utility in supporting some simple office procedures, in general we believe it reflects an over-simplistic view of the nature of cooperative work practices.

Some approaches to the challenge identified by Grudin attempt to bypass such problems altogether through a strategy of 'participatory design' (Bjerknes et al, 1987; Bodker et al, 1988; Grudin, 1990; Muller, 1991) in which eventual users of a system enter into a partnership with system designers in the hope of enhancing the effectiveness, acceptability and accessibility of the system.[1] This strategy is perhaps the most well developed way of bringing into the design process the knowledge and expertise of users. However, although it is clearly essential to involve users in the design process, user-centred design is not enough. For while users often understand the nature of their own tasks very well, they are not always aware of how these fit into the overall organisation of working activities. Further, and an aspect of this, often the very familiarity of working activities to those who perform then can make the essential cultural knowledge in which they are embedded 'invisible', a 'seen but unnoticed' feature of the work.

Another strategy, albeit one less developed than the one just mentioned, is represented in Suchman's (1987) seminal work using ethnographic studies of naturally occuring activities to show, in fine detail, how relevant users make sense of a technical artifact in order to make it perform its functions. Generalising this approach is to argue for the detailed ethnographic study of the naturally occuring settings, their activities and the organisational and social knowledge that constitutes the setting for what it is, in order to inform the conceptions and design of appropriate computer support.

Ethnography is a research method which is directed at studying the social character of groups and the activities of their members in their natural settings. Its sociological emphasis means that it examines activities as socially organised *from within* their natural settings by participants to those settings. The ethnographer's task is to reveal and explicate not so much the fact of cooperation and collaboration, for the socially organised character of social life is a presumption of its inquiries, but how, in a natural setting, this is achieved by parties to that setting using *their* knowledge and *their* understandings of the domain. In this respect, the objective of the fieldworker is to uncover the understandings and meanings used by participants to produce the sense the setting, its activities and its objects have for them.

Of particular importance here is uncovering the tacit knowledge and implicit practices that are normally 'invisible' to the casual observer, or are taken for granted but unexamined by task analytic approaches which seek to breakdown work activities into discrete, 'free standing' components without regard to how they are interwoven into a set of socially organised work activities.[2] As a number of sociological studies have shown, in settings ranging from laboratory science (Lynch, 1985; Latour and Woolgar, 1986), to expert systems (Collins, 1987), to police patrol work (Sacks, 19..), and many more, and as studies of organisations have long shown in their use of the distinction between the formal and informal aspects of organisational conduct, skills and activities are embedded in cultural understandings that are rarely explicit but still vital ingredients of the work. Some of these are short-cuts in working practices that are unrecognised in formal job descriptions; others are subversive of the formal patterns of work; others constitute an 'unofficial' version of the organisation; and yet others are integral to the accomplishment of the work itself. Further, such knowledge is not adequately summarised as elements of 'real world' activities as lists of decontextualised propositions, be they formally specified or tacit. Rather such knowledge is highly localised and organised around structures of relevance (Schutz, 1967) which are to do with the specifics of the work site itself. Although much of our cultural knowledge is widely shared, a conditions of our being members of a community at all, much of it is specific to domains and constituting the 'know how' and the 'know what' for those enculturated within them. Disentangling all of these, and more, in order to make judgements about the design of appropriate computer support is not straightforward, not least because they are not discrete categories, but overlap and cross-cut in complex and very particular ways.[3] Nonetheless, gaining some understanding of how they do so is an important problem of designing systems "incorporating organisational and social knowledge".

While various ethnographic studies, including our own work with air traffic controllers, suggest that such studies can inform the systems design process, perhaps surprisingly, few studies have been carried out in conjunction with related software systems development. Our work on air traffic control is part of a research project concerned with investigating user interface designs for systems with a reactive data base; that is, a multiuser data base whose contents may be changed by external agents, either users or automatic sensors, while they are in use.

User interface issues which are being investigated include the need to highlight changes in a way which is not overly distracting, the need to provide facilities which allow different users to tailor the display to their personal style of working and the need to display a large amount of information in a restricted display screen yet still provide the ‘at a glance’ understanding which is important for controllers using the information. The system development studies are being undertaken in three phases: first, the development of a prototyping environment specifically tailored to generate interfaces for reactive databases with end-user participation; second, the development of a number of prototype user interfaces based on flight strips and exploring some of the problems just identified; third, evaluation of the system with the participation of ethnographers and air traffic controllers.[4]

In this paper we are not concerned with the details of the software, which are discussed elsewhere (Bentley *et al*, 1992), but with some of the problems we have experienced in integrating the contributions of the ethnography into the systems design process. Many of these had to do with evolving methods of work between researchers training in disciplines which exhibited very different styles of working and ways of approaching their respective problems. But, and as important from the point of view of this paper, we are beginning to learn more about the role ethnography can play in the system design process, as well as learning about some of its limitations.

The first part of the paper presents a brief characterisation of the work of controlling as constituted from the fieldwork. It focusses especially on the paper flight progress strip which, though output from the database, has no automated facilities for updating the information it contains. Among the advantages of automated updating is the possibility of designing a more effective Conflict Alert Facility.[5] However, it was also clear that any automated facility would have to preserve existing functionalities derived from current working practices.[6]

Air Traffic Control: A Brief Overview

The control of UK airspace is centralised at London Air Traffic Control Centre (LATCC) with sub-stations at Manchester and Prestwick. The airspace is divided into sectors, each of which is the responsibility of a controlling team around a radar suite situated in the operations room at LATCC.[7] The typical manning levels of a sector consists of a chief, two controllers, two 'wingmen' or assistants.[8] ATC belongs to a class of command and control systems in which human operators issue instructions to controlled entities, in this case aircraft, using information provided by a common database. The data base provides two kinds of information for controllers: radar and flight progress strips. The former is a 'real time' representation of the traffic through the sector covered by the radar, each aircraft being identified on the screen by a 'blip' and its call sign.

The latter is a paper strip, approximately 8" long and 1" wide, containing formatted information of each flight, including call sign, next reporting point on route, time due to pass that point, aircraft's 'squawk code', aircraft type, planned flight path, requested cruising height, departure and destination airports, and so on (see Figure 1). Much of this information is filed prior to departure and output as a strip by the computer anything up to 40 minutes prior to the aircraft's arrival at the relevant navigation points along the routes within the sector. The strip is an 'historical' not a 'real time' representation of the state of an aircraft. Any updating of the information on the strip is done by various members of the team as changes in height, direction and speed are directed to the aircraft by the controller. Thus, in the passage of an aircraft through the sector the strip will accumulate a record of its control history.

Figure 1 Flight strip

The work is governed by a set of rules set out in the *Manual of Air Traffic Services* which specifies the details of routes, separation requirements, communication procedures, standard procedures for particular sectors, and more; all of which the controller is required to apply in the course of the work. However, such rules do not say anything about *how* the work is accomplished as an actual course of work activities. As far as the work is concerned, the rules do not stand as disembodied regulations which have to be 'mindlessly' applied, but are integral to controlling activities and inbibed as *activities* that constitute controlling as a situated and artful matter. They 'enable' the work of organising the traffic because they furnish 'instructions' for seeing the information to hand on the radar and on the strips, as indicative of the current state of the traffic in the sector.

Controlling Work

Controlling air traffic is to manage events in real time since a procedural plan cannot be produced which would enable controllers to determine in advance, except in very general terms, what in particular they should do next. There are always and inevitably contingencies, incidents, special flights, technical troubles of various kinds, sudden increases in traffic flows, and so on. At its simplest, the controller's problems is a scheduling one. For any controller the traffic has to be taken as and when it arrives in the sector of airspace for which s/he is responsible and threaded together into an orderly pattern before handing over to the next sector. The scheduling has to be achieved in and through making the traffic flow. Aircraft cannot be parked and even 'holding patterns' they are still on the move, still part of the flow of traffic and, therefore, taken account of in the moment-to-moment management of the traffic.

It is in resolving. moment to moment, the scheduling problem that the 'invisible' skills of controlling are deployed. It is the accumulated 'know how', much of it tacitly held and understood, that is brought to bear on the resources provided by the system's technology to determine what, in any particular circumstances, the appropriate controlling actions are to be. The radar, while providing a 'real time' representation of the present position of aircraft and their configuration, does not provide a complete picture since 'what is likely to happen' in a few minutes is immensely relevant to what actions the controller needs to do now by way of, for example, avoiding conflictions, maintaining separation requirements, expediting traffic effectively, and so on. Information relevant to these matters is provided by the flight strips. Both radar and flight strips are elements in a mutually explicating system of making sense of 'what is going on'. Thus, the solution to the scheduling problem is managing the traffic in the context of what the technology furnishes, and the skills of 'working with the technology' are complex and interwoven.[9] They often seem to lack a sense of deliberation, cogitation, task definition, specification and solution. Rather, the controller 'just knows' what to do, where 'knowing' here means being able to interpret the meaning of *this* set of conditions at *this* suite at *this* point in time against a background of what has gone on so far, what time of day it is, where everything is presently, and so on. The whole is a 'gestalt contexture' which provides the meaning of what is 'to hand'. The problem is defined as *this* problem requiring *this* course of action, and experienced as part of the flow of the work.

Part of the art of controlling is managing the information as part of a sequence of working tasks in and through the 'working division of labour' around the suite: a notion which is intended to direct attention to the *actual course of activities* as they are socially organised and understood by parties to the work, in this case that of controlling, rather than to idealised versions of those activities.[10] Thus, although formally each member of the suite team have their separate tasks and responsibilities, it is essential that they also attend to, as a trustable matter, those of others in the team. It is in this respect that the strip is made into the work site of the activities of members of the team. It becomes a note pad for recording, and thereby updating the information it exhibits, the actions done with respect to the aircraft the strip represents. For example, as instructions are given to pilots to climb, descend, follow particular headings, these are simultaneously written down on the relevant strip as is the pilot's acknowledgement, as is the attainment of the instruction. Coordination between sectors, changes in ETA, routes, and other relevant information are also recorded. In this way, the strip is a facility for the rapid updating of information by members of the team as they record *their* actions in ways relevant to them.[11]

But information *on* the strip is not the only way in which the strip is informative for the controlling team. The controller's problem is a scheduling one because s/he is not dealing with a single aircraft but with several occupying the sector's airspace; a configuration which is continually evolving and which it is the task of the controller to shape. To meet this end, strips are organised into bays. 'Live strips' are set out on a Flight Progress Board in front of the controller and represent aircraft which are currently, or about to be, the controller's responsibility. On the wings of the suite are the 'pending strips' which are indicators of the traffic due in the sector. They provide an overall picture of likely tasks, mainly through the routes aircraft have selected which can be useful in resolving scheduling problems and planning the flight levels at which aircraft will enter the sector. A second bay on the wings holds the 'transfer strips' which depict aircraft at their last reporting point. The 'live strips' are actively organised such that their array displays the schedule of work tasks. They are aligned under Designator Strips representing the reporting points over which a flight will pass. Within these divisions the strips are further and continually organised, according to preference by such as times over reporting points, according to flight levels or possible confliction points. The point being to achieve, through the organisation of the strips, information about the state of the sector that is relevantly available 'at a glance'. It is a way of making work tasks, possible problems, visible to members of the team.

"Working the strips' in the ways indicated is organising them in terms of tasks that need to be done and, through this, accomplishing the work of organising the traffic. the strips become a linear sequence of 'objects for processing'; that is, a sequence of working out and yet to be worked out courses of action. The strips and their organisation are a proxy orderliness for the configuration of the traffic flow. While the radar is a computer generated two-dimensional picture of the relevant sector and its traffic, the strips are the means by which the patterns on the screen, and thus in the sky, are seen as the patterns that they are and what needs to be done in terms of coordinating work tasks. Strips are not just 'placed anywhere' but are organised so as to give a sequence to them reflecting the fact that the management of the traffic is inherently *constructing* a sequence of tasks that will achieve the flow of traffic through the airspace. In this sense the strips represent a pattern of tasks which, as it is gone through and completed, produces the orderliness of traffic in the sky. Keeping the strips straight is keeping the planes straight. Hence, the work on the strips, the marking up of route changes, height changes, coordination actions done, and so on, is not just making up a record of work: it is doing the work.

The Implicit Cooperative Features of Working the Strips

Although we speak of a sequencing of the tasks of controlling, it is important to bear in mind that this is not the kind of sequencing that is, say, characteristic of a production line in which tasks are predetermined and governed by the speed of the line itself. A controller has, in fact, a great deal of discretion in how the traffic is dealt with, and how, within the rules and procedures, the flow is to be managed. We have also been making the point that information on the strips, and of course provided by the radar, is not passively related to the work of controlling but actively organised by various members of the team so as to organise the working tasks. Managing the information is, one could say, managing the traffic.

An important feature of this management of the information is its 'at a glance' availability for those around the suite. This feature is no mere adjunct to the work, but is one oriented to *as a feature* of the cooperative work around the suite itself and, thus, integral to the controlling of the traffic. When traffic levels are high, and this is becoming increasingly the routine, controllers have little time for reflection and deliberation. Decisions have to be made, and the appropriate action taken, quickly, almost nonchalantly. This requires exact assessments of the progress of aircraft along their given vectors and where 'in the sky' they are in relation to one another: assessments based on information seen 'at a glance'. The seeming effortlessness of the achievement makes invisible the skill and the knowledge by which controlling is brought off.

Much of the interaction around the suite is unspoken, routinely limited to gestures such as pointing to strips, or to the radar, 'cocking strips out', reviewing them, and so on. Strips are 'glanced at', 'taken heed of', 'ignored for now', 'revised', not just when they arrive but continuously, so that information is 'at hand' and 'in hand'. Moreover, the activities within what we have referred to as the 'working division of labour' are governed by structures of relevance in that what is taken heed of, what is necessary for the task in hand, what is ignored for now, has very much to do with each member of the team's responsibilities but not in any predetermined fashion but rather in a way responsive to the exigencies of the moment.[12] What a controller feels s/he needs to know, what the screens or the strips indicate, is worked out on a moment to moment basis. For additional example, a 'wingman's' responsibilities are mainly to do with preparing the strips as they are output from the computer, checking the information, and making them ready for when needed.and particularly for aircraft first entering the sector, by attending to the controller's current work and the display of 'live' and 'pending' strips, 'wingmen' can work out appropriate routings, levels, and so on, to facilitate the tasks of the controller. It is knowing what others in the team are doing and how one's own responsibilities can be made to fit into that cooperation. As indicated, much of this intra-suite coordination work is unspoken, relying on the 'know how' and the 'know what' and, importantly, the trust among members of the team. 'Cocking out' strips is another way in which mutual monitoring goes on. What this action does is bring to the attention of the radar controller, though it is also an action the controller sometimes performs, aircraft that need to be 'noticed' for whatever reason, such as a possible confliction or a special flight. By utilising a contingent property of the racks, namely, that they can support strips propped out from the bay, information can be conveyed to the controller routinely without the need for further communication. Such activities, and there are many more instances that could be cited as illustration, are all aspects of the implicit social organisation of the work, learned, understood and continually worked at and relied upon by members of the team. Much of this is taken-for-granted in, but is essential to, the performance of the work as seamless sequence of tasks that have evolved as ways of working with the technology. Much of the cooperation of the team members around the suite consists in the implicit understandings by which separate responsibilities and tasks are smoothly integrated into an overall flow of work that not only 'gets the work done' but also provides for that vital mutual checking of the information 'to hand' and its potential consequences.

Conclusions

What we have tried to do in this paper is illustrate the importance of the implicit understandings of work activities for the design of appropriate computer support. It has to be admitted that ATC is not typical of many contexts of work in that it is done by a relatively homogenous group of highly trained personnel whose work is tightly focussed around the locale of the radar suite. Skill and, above all, experience, means that although the tasks, especially those of the radar controller, require intense prolonged concentration and, to the naive observer seem radically impossible to perform, they are sharply focussed for the periods 'on the tube'.[13] This is unlike, say, police beat work which tends to be fragmented into the routine and the emergency but not in any predictable way (Ackroyd et al, 1992). However, unlike production line workers who have to perform similar concentrated tasks, ATC is much more discretionary. Further, the work is done to achieve the safe passage of air traffic through the skies, where safety is an imperative which dominates the skill of controlling and its performance.

Noting these differences is important not just to specify the distinctiveness of ATC work, but to stress the more general point that designers need to acknowledge the variety of work, its settings and its practices and, by implication, the need to study these in their natural settings in order to inform the design of appropriate computer support. While such study is the rationale of ethnographic inquiry, it is by no means a guaranteed solution to Grudin's challenge, least of all in commercial contexts where requirements capture often has to be done within time scales which most ethnographers would regard as a joke. Nevertheless, such limitations do not make the problems go away.

In our own case, it was realised from an early stage of the research that many of the implicit work practices were not just incidental adjuncts to the cooperative work around the suite, but highly relevant to its successful performance. The active 'doing' of marking the strips, manipulating them in various ways such as 'cocking them out', organising them, pointing to them, all serve to 'gear' the members of the team into the ongoing work. They became, as it were, 'at one' with the system. Thus, any electronic version of the flight strip would need not only to offer additional functionalities but also preserve many of the existing ones deriving from the implicit interactions of the team members. Thus, and for example, any automatic updating of information, one of the putative advantages of an electronic flight data system, would not only need to be brought to the attention of the controller, but also provide a means whereby s/he could acknowledge the 'noticing' of that change. More generally, providing the 'at a glance' visibility of the state of the sector is important not just for the radar controller but also for other members of the team in terms of their relevances so that their work can be smoothly 'geared into' the work of the controller without the need for much deliberation, inquiry or discussion.

From the point of view of interdisciplinary working, and as indicated in the beginning, we have had to overcome a number of problems, both cultural and practical. One of the basic objectives of the ethnographer is to elucidate the details of settings and their activities rather than to rely on asbtracted and general characterisations of them.[14] By contrast, the software engineer is continually seeking out abstractions which can be modelled in the software. While these differences are differences to do with the specific problems of the respective disciplines, and in this sense not antithetical, the different styles of working did lead to misunderstandings.

The software developers on the project hoped that the ethnographic studies would inform both the design of the prototyping environment and the design of experimental system interfaces. They wanted the ethnographic study to identify activities, objects, attributes and relationships which were fundamental to air traffic control and to separate these from the details of the flight strip 'user interface' to the underlying flight plan database so that fundamental entities would be supported directly in the prototype with user interface 'details' left until later phases of the research.

After more than a year of working together, we now realise that the expectations of the system designers were unrealistic. Whilst the system designers continue to maintain that there are 'core' air traffic control activities, there is a greater realisation that ethnography does not present its analyses in ways designed to furnish such a description. The sociologists, for their part, did not see the need for such a distinction or, more accurately, felt that any description of such 'core' activities could only be a gloss for the complex work of controlling. Nevertheless, despite such differences a style of working evolved based in extensive debriefings after a period of ethnography where the fieldworker discussed the fieldnotes, highlighted what seemed to be the key points and answered questions posed by the system developers. These questions, typically, raised other issues which directed further ethnographic study, and in this way we iterated toward a better understanding of the work process of air traffic control.

While we are confident that the ethnographic studies have been valuable in the context of system design, we must admit that this is as much a matter of faith as it is backed up by evidence.[15] The nature of ethnography and (we stress that this is very much from the point of view system design) its 'unstructured' observations is such that it has been difficult to organise and direct the ethnographic record so that we can correlate observations and system requirements. Indeed, one of our activities in this project is investigating methods of organising fieldnotes to make them more accessible and more useful to system designers.

The problem of system's designers dealing with the 'unstructured' nature of ethnographic fieldwork data is, of course, only one of the problems facing researchers interesting in informing the design process through ethnography.[16] Other problems, which we simply note here without any pretense of offering solutions, include:

- the cost-effectiveness of ethnography - typically an ethnographic study will take several months of fieldwork yet the timescales imposed on developers are such that the system specification must be completed within a much shorter time period than is usual for an ethnographic study. Thus, if ethnography is to be used at all in systems development ways need to be found of foreshortening the periods of fieldwork.
- again as pointed out earlier in the paper, the system we are currently researching is unusual inasmuch as end users are relatively homogeneous in their requirements. More typically, user communities are much more diverse, often to the extent of having mutually conflicting requirements. Although this is a general problem of system design, we do need a better understanding of how ethnographic studies can contribute to the identification of these different groups, their different requirements and how, if at all, such divergencies may be resolved.
- finally, but not exhaustively, there is the problem of determining the 'completeness' of the ethnography. This is related to the first of the problems noted but is more fundamental in that irrespective of how much time is spent on system studies, of whatever kind, we can always be sure that they are incomplete. For example, and referring to our own research, we have never observed a terrorist incident, any kind of air accident or even the onset of very severe weather conditions. Nevertheless, any workable automated system must be able to cope with the stresses

> imposed by exceptional situations and while this is a general problem, for which we can never guarantee that it is solved, we need to think about how to integrate ethnographic observations with often hypothetical scenarios of exceptional situations.

At the start of the project it was believed that the principal role of the ethnography would be to help clarify system requirements. This has happened to some extent. We also believe that ethnography has an alternative and equally important role as the 'user's champion'. While this does not necessarily mean defending the user come what may, the understanding gained by ethnographers in the course of their studies makes them not only knowledgeable about user's tasks but also sympathetic if only in the sense that they have no especially interest in the success or otherwise of automated projects.

One of the most significant problems in system evaluation and, indeed, in user-centred design is that typical system users are often too busy doing their job to participate in the design and evaluation process. One result of this is that it is often atypical users, such as those with an interest in and sympathy for automation, who will participate. Accordingly, one valuable role ethnography can play is 'standing in' for the end-user for at least part of the evaluation, using the ethnographic record to validate that essential features of the 'working division of labour' are supported by the automated system so reducing the probability of exposing end-users to an inappropriate system which may be immediately rejected.[17]

On its own ethnography cannot be a 'complete' method of requirements capture, but in this is equivalent to many of the other methods proposed. Further studies are needed to discover how it may be integrated with other approaches to system requirements derivation such as interviewing, viewpoint analysis, soft systems methodology, object-oriented analysis, data flow analysis, and so on. Our experience has shown that the understanding of social, cooperative processes gained through detailed ethnographic studies is effective in informing the process of system design. However, we are a long way from understanding how to integrate such studies routinely into the process of system specification. The differences in culture between the disciplines involved, the timescale required for ethnography, the 'unstructured' character of the ethnographic record, and more, limit it for the time being to very specific types of system where the social interactions and understandings are strongly focussed and recognised and accepted by everyone, not just sociologists, involved in the development process.

The software development team must have some understanding of, and must be sympathetic to, the problems of the ethnographer. Equally, the sociologists studying the setting involved need to realise that economic and other practical constraints placed on development may not allow the 'best' solutions to problems they have identified to be adopted. More effective tools to support the ethnographer and which help organise, structure and browse the ethnographic record are needed. Integrating ethnography into the software development process is itself a significant cooperative activity with challenging requirements for computer support.

Bibiography

Ackroyd, S., Harper, R., Hughes, J.A., Shapiro, D., and Soothill, K. (1992), *New Technology and Practical Police Work,* Milton Keynes, Open University Press.

Anderson, R.J., Hughes, J.A., Sharrock, W.W. (1989), *Working for Profit: The Social Organisation of Calculation in an Entrepreneurial Firm,* Aldershot, Avebury.

Becker, H.S. (19..), 'Sociological Analysis and the Variable',

Benson, D. and Hughes, J.A. (1991), 'Evidence and Inference', in Button, G. (ed), *Ethnomethodology and the Human Sciences,* Cambridge, Cambridge University Press.

Bentley, R., Rodden, T., Sawyer, P., and Sommerville, I. (1992), 'A Prototyping Environment for Dynamic Data Visualisation', Submitted for presentation at the *5th IFIP Working Conference on User Interfaces,* Ellivuori (Finland), 10-14 August.

Bjerknes, G., Ehn, P. and Kyng, M. (eds) (1987), *Computers and Democracy - a Scandinavian Challenge,* Aldershot, Gower.

Bodker, S., Ehn, P., Knudsen, J., Kyng, M., and Modsen, K. (1988), "Computer Support for Cooperative Design', *Proceedings of CSCW '88 Conference.*

Bucciarelli, L.L. (1984), 'Reflective practice in engineering design', *Design Studies,* 5, pp. 185-190.

Collins, H.M. (1987), 'Expert Systems and the Science of Knowledge', in Bijker, W., Hughes, T. and Pinch, T. (eds), *New Developments in the Social Studies of Technology,* Massacheussetts, MIT Press, pp. 329-348.

Crawley, R. (1982), *Predicting Air Traffic Controller reaction to computer assistance: a follow up study,* AP Report 105, Applied Psychology Department, Aston University.

Danielson, T. *et al.,* (1986), 'The AMIGO Project: advanced group communication model for computer-based communication environment', *Proc. CSCS'86*, Austin, Texas.

Grudin, J. (1990), 'The Computer Reaches Out: The Historical Continuities in Interface Design', *ACM.*

Harper, R.R., Hughes, J.A. and Shapiro, D.Z. (1989), *The Functionality of Flight Strips in ATC Work: A Report for the Civil Aviation Authority,* Department of Sociology, Lancaster University.

Hughes, J.A., Shapiro, D.Z., Randall, D., and Harper, R.R. (1991), 'Visual Rerepresentation of Database Information: The Flight Strip in Air Traffic Control', Co-Tech Working Group Conference, Scharding, May.

Hughes, J.A., Shapiro, D.Z., Sharrock, W.W., Anderson, R.J., Harper, R.R. and Gibbens, S. (1988), *The Automation of Air Traffic Control,* Research Report SERC/ESRC, Department of Sociology, Lancaster University.

Latour, B. and Woolgar, S. (1989), *Laboratory Life:The Construction of Scientific Facts,* Princeton, Princeton University Press.

Lynch, M. (1985), *Art and Artifact in Laboratory Science: A Study of Shop Work and Shop Talk in a Research Laboratory,* London, Routledge and Kegan Paul.

Muller, M. (1991), 'Participatory Design in Britain and North America: Responding to the "Scandinavian Challenge"', in *Proceedings of CHI '91 Conference: Reaching Through Technology,* ed. Robertson, S.P., Olsen, G.M. and Olsen, J.S.

Schmidt, K. (1990), *Analysis of Cooperative Work: A Conceptual Framework,* Roskilde, Riso National Laboratory.

Schutz, A. (1967), *The Phenomenology of the Social World,* Evanston, Northwestern University Press.

Shapiro, D., Hughes, J.A., Randall, D., and Harper, R. (1991), 'Visual Re-representation of Database Information: The Flight Data Strip in Air Traffic Control',

Suchman, L. (1987), *Plans and Situated Actions: The Problem of Human-Machine Communication,* Cambridge, Cambridge University Press.

Whitfield, D. (1980), 'Discussion Paper: The Air Traffic Controller and Sector Capacity', Ergonomics Development Unit, University of Aston in Birmingham.

Wilbur, S.B. and Young, R.E. (1988), 'The COSMOS Project: a multi-disciplinary approach to design of computer-supported group working', *Proc. EUTECO88: Research into Networks and Distributed Applications,* Vienna, Austria, April 1988.

---

[1] Of course, the motivation for 'participatory design' is not that of avoiding *sociological* problems, though it is to achieve effective design by incorporating user's knowledge of the relevant domain.

[2] In the context of ATC research see, for example, Crawley (1982) and Whitfield (1980),

[3] Such features are, of course, endemic in the design process itself. See, for example, Bucciarelli (1984) and Schmidt (1990) for an overview.

[4] The first phase of this research is almost complete.

[5] The current facility does not distinguish between spurious and actual conflictions. For example, the former can arise when controllers deliberately place aircraft on conflicting headings 'for now' knowing that in a short space of time they will take appropriate action. The reasons for this are various but mainly to do with freeing airspace to solve other problems. However, having said this, most controllers, despite some initial misgivings, are happy enough with the current system since it does provide important backup even though many of its warnings are unnecessary.

[6] See, for example, reports based on earlier work in Hughes et al (1988), Harper et al (1989), Hughes et al (1991).

[7] Each sector is a block of airspace, generally from flight level 5.5 up to flight level 25. They are highly variable in their horizontal dimensions, however. Flight levels are expressed in thousands of feet, adjusted daily according to air pressure. Around airports sectors descend to the ground. Traffic within controlled airspace is obliged to maintain contact with ATC, obey its instructions and carry transponders which 'inform' the radars of the call sign of the aircraft and its height. Both these are displayed on the radars alongside the 'blips' and, along with other information, on the flight progress strips.

[8] This can vary widely depending upon how busy the sector is. Sectors, for example, can be 'split' during busy periods or 'bandboxed' when traffic is light.

[9] Controllers work within a technologically rich environment. The radar and the strip information is provided by a common database stored on the computer. In addition, there is RT, telephone links, weather information screens, and so on, all built into the radar suite.

[10] See Anderson et al, (1989) for a further explication of this notion.

[11] This is formally recognised by a colour protocol by which different members of the team use different coloured pens to write on the strip using, but not exclusively, conventional signs to denote actions taken or about to be taken, or information to be noted.

[12] In recent discussion with controllers it is becoming clearer that the information on the strip can be grouped into static (route, squawk identifier, aircraft make, etc.) and dynamic (current level, position, speed, etc.) information with the latter, when under direct control, of more relevance that the former. Currently, whenever strips are printed for each navigation point en route, both types of information are retained. Reconfiguring the types of information on separate screens,

yet retaining the flexibility for the controller to expand or limit each type of configuration, is a scheme we are exploring.

[13] Normally controllers work for periods of approximately 2 hours followed by a 2 hour rest.

[14] Within social research, this is one of the major justications for ethnography and one of its central criticisms against other styles of social research. See, for example, an early statement of this by Becker (19..) and Benson and Hughes (1991).

[15] From a sociological point of view as exploring a method of studying work and its activities, it has been immensely valuable. The issue we are addressing here is how far they can be valuable in the system design process.

[16] One of our current activities in this project is to investigate methods of organising the ethnographic records to make them more accessible to and more useful to designers.

[17] As it happens, the occupational culture is full of 'war stories' about ill-thought about innovations which have succeeded in making the work more difficult.