Flying Planes Can Be Dangerous: aspects of the technological praxis of Air Traffic Control

R.J. Anderson[1] : J.A. Hughes[2] : D. Z. Schapiro [2]: W. W. Sharrock[3] : with R. Harper[2] and S. Gibbons[4]

The research upon which this report is based was supported by SERC Grant No GR/D/85157. We would like to thank the Civil Aviation Authority and The National Air Traffic Service for their support of our research. We would also like to thank the management and staff of the London Air Traffic Control Centre for their advice and help during the periods of fieldwork. Lucy Suchman, Liam Bannon and a number of anaonymous referees have given us many helpful comments on earlier drafts of this paper.

1 Rank Xerox EuroPARC

2 University of Lancaster

3 University of Manchester

4 Shell (GB) Ltd. PLC

A PROBLEM AND ONE SOLUTION

In their recent thoughtful summary of the contribution which social science might make to the design of computational technologies, Galegher and Kraut (1990) warn that integrating the two disciplinary matrices will be no easy matter. To begin with social science and design derive from different methodological backgrounds, evince different analytic interests and often have entirely different value orientations. In addition, the tendency within the social sciences is to seek to demonstrate just how complex and interrelated the phenomena of our ordinary taken for granted lives are. Time after time, what to the ordinary eye appear to be simple and quite straightforward matters are revealed to be a dense interweaving of norms, values and social processes. Despite these difficulties, though, Galegher and Kraut are hopeful of the possibilities for collaboration between design and social science.

> *We believe that successful design of technology for cooperative work requires both expertise in the underlying computer and communications technologies and expertise in the social and behavioral processes that the technology is designed to support.......(T)he knowledge generated in research labs, or even in field studies, is often too abstract to be applied unamiguously in concrete circumstances. Social scientists working with designers can provide the translation needed between the abstract and the concrete, and, using both formal and casual methodologies, can test whether the translation was successful. (Galegher and Kraut, 1990, pp. 11-12)*

We share Galegher and Kraut's aspiration for the integration of social science and design. We too feel that what the social scientist can contribute is an understanding of and sensitivity towards the complexities of specific sets of social and organisational processes. One way that this can be realised, we believe, is through the provision of what Clifford Geertz, drawing upon Gilbert Ryle, called "thick description" (Geertz 1975) and with it an orientation to "ethnographic realism". To explain what he means by "thick description", Geertz takes the Rylean case of two boys rapidly contracting their right eyelids. "In one, this is an involuntary twitch; in the other, a conspriatorial signal to a friend." (Geetz 1975 p 6). Now add a third boy who "'to give malicious amusement to his cronies', parodies the first wink as amateurish, clumsy, obvious and so on" (ibid).

> *...between what Ryle calls the "thin description" of what the rehearser (parodist, winker, twitcher...) is doing ("rapidly contracting his right eyelid") and the "thick description" of what he is doing ("practising a burlesque of a friend faking a wink to deceive an innocent into thinking a conspiracy is in motion") lies the object of ethnography: a stratified hierarch of meaningful structures in terms of which twitches, winks,*

> *fake-winks, parodies, rehearsals of parodies are produced, perceived and interpreted, and without which they would not (not even the sero-form twitches, which as a cultural category are as much nonwinks and winks are nontwitches) in fact exist, no matter what anyone did or didn't do with his eyelids (Geertz, 1975 p 7)*

In drawing out its object, ethnography engages in a reading of culture. Indeed, ethnography is the *inscription* of such a reading or readings.

> *Doing ethnography is like trying to read (in the sense of "construct a reading of") a manuscript - foreign, faded, full of ellipses, incoherencies, suspicious emendations, and tendentious commentaries, but written not in conventionalised graphs of sound but in transient examples of shaped behavior. (Geertz, p 10)*

In this paper, we wish to outline a frame of reference for reading technologies as cultural artifacts (that is, as culturally significant and the bearers of meaning) in the hope that the thick descriptions so provided may be of direct value in the design of different kinds of technologies. We have chosen construct our frame or reference (or motivate our readings) around the central theme of the experience of technology. In that sense, what this framework hopes to provide is the basis for an adequate phenomenology of technology constructed around the actor's (fn who may only sometimes be a "user") point of view. In the next section we lay out the lineaments of our frame of reference. This will involve treating the metaphysics of engagement with technology under a strict methodological rule whereby the taken for granted artifactual character of technologies are treated as socially organised accomplishments. We will then go on to apply this frame of reference to a specific domain of technological work, Air Traffic Control and the activity of real time controlling. Our motif in this description will be the socially organised character of the distinction between 'a system' and 'its user'. Finally, we will review some of the more obvious design implications thrown up by description of Air traffic Control and hence exp,ore the possibilities for collaboration between social scientists and designers.

A PHENOMENOLOGY OF TECHNOLOGY

Embodiment, engagement and technics

Drawing upon the familiar examples from Husserl, Heidegger and Merleau-Ponty, Don Ihde (Idhe 1990) has explored a number of ways suggestive of what he calls a "phenomenology of technics". He begins by focussing upon "the experientially recognisable features that are centred upon the ways we are bodily engaged with technologies." (Ihde, 1990 p73).

> *At one extreme are those relations that approximate technology to a quasi-me (embodiment relations). Those technologies that I can so take into my experience that through their semi-transparency they allow the world to be made immediate thus enter the existential relation which constitutes my self. At the other extreme of the continuum lie alterity relations in which the technology becomes quasi-other, or technology "as" other to which I relate. Between lies the relation with technologies that both mediate and yet also fulfill my perceptual and bodily relation with technologies, hermeneutic relations. (Ihde, 1990, p 107).*

What Ihde is attempting to capture in this continuum is the signficance artifacts may have for us when, for one reason or another, they move to the centre of our zone of operation in the lifeworld. The hammer in Heidegger's case, is not just an extension of my arm when I am hammering a nail. The hammer and I are one in the praxis of hammering. This is what is meant by the embodiment relation. In the hermeneutic relation, the significance or meaning of, say, the sound a musical instrument is making is understood in and through its playing. We do not hear the notes, see the fingers move and then understand the music. Rather the embodied instrumentalism is found by us in the music being played. With the alterity relation, technology bcomes the other to which we relate and in the actions of which we find design, meaning, intentionality. Here we are not speaking metaphorically. In the flow of our experience of technology we situate their actions as those of an other.[1]

The phenomenology of technics which Ihde sketches is an attempt to reconstruct a metaphysics for technology around the nexus of engagement. The categories which are offered (as well as those associated with horizonal and background relations which he also analyses) define or stipulate how this engagement is to be conceived. It is this which makes his discussion *philosophical* in form. The technology and the user of the technology are first set over against one another in an objective relation as a way of fixing the pluridemensionality and complexity of experience (its plenumn). The dichotomy is gradually decomposed as the continuum of subjective relations is developed until in the alterity relation, we treat the technological object as a subject - an other.

The praxeological rule

Ihde's categories bring out the manifold of our relation to technologies. But how can these sensitivities be reconstrued as the basis of a description of the social organisation of technological

---

[1] In part the alterity relation is what is being struggled with in the much debated domain of human computer inteaction and in particular the tool and conversational metaphors which currently dominate that field. See Luff and Gilbert (eds) 1990

relations?[2] One way of answering this question can be found by adopting a particular strategy which transforms philosophical or metaphysical questions into methodological ones. This strategy involves the application of the application of a “praxeological rule” first worked out in the ethnomethodological studies of Harold Garfinkel.

> *...accounts by sociologists of the conditions under which a phenomenon occurs may be mapped point for point into the terms of strategies that persons follow whereby, knowingly or not, they achieve the pay-off represented in the variable under study. The praxeological rule states that any and all properties of a social system that a sociologist might elect to study and account for are to be treated as technical values which the personnel of the system achieve by their modes of play.(Garfinkel. 1956, p 191)*

There are several complex ideas compacted into this single paragraph. We will tease out only one, the *methodological stipulation*. which Garfinkel claims is definitive of sociology. In Garfinkel's view, all forms of sociology deliberately choose to account for human activity purely as social action. This being so, sociology has no choice but to treat the describable properties of activities in a social setting as the ‘outcomes’, 'accomplishment' or 'achievement' of those participating in it.

We should also notice the importance of the term "treated" in the above quotation. Nothing is being said about the goals, purposes or aspirations of members themselves. All that is being claimed is that as a matter of disciplinary method and hence for or the purposes of sociological analysis, the social features of some setting are to be viewed as the outcomes, the achievements of those in the setting.[3] Thus, if the social phenomenon is 'ordinary conversation', the sociologist treats the participants as "doing conversation". If it is 'the gender division of labour', then this too is treated and analysed as an accomplishment. The purpose is to bring out and analyse how those achievements are organised and hence in what mutual orientations, expectations and understandings they are grounded.

Garfinkel himself applies the praxeological rule by modifying a notion central to Phenomenology, namely that of the "époche" or "phenomenological bracketing" and particularly that of “époche of the natural attitude". The term "epoche" is used to characterise part of the phenomenological

---

[2] The phrasing here is important. We are not saying that technological relations are wholly or exclusively social. Nor are we proposing a hard and fast dichotomy between the social and the technical. Rather, the question is ‘In so far as technology can be treated as social, how might a description of the social properties of technological relations be constituted?’.

[3] It is important to recognise the methodological character of this move. It is, as we saw with Geertz, a way of constructing a reading of activities, not a way of explaining them.

method or turn whereby the appearances or properties of some phenomenon are investigated. The analyst takes some phenomenon, let us say the keyboard in front of me, and brackets off the keyboard as an object in my direct field of action from all other understanding , presumptions which I might have about the keyboard. For example how it works, where it was made, its relationship to other keyboards I have used, and so on. The phenomenological epoche insists that we address this keyboard first as a datum of experience here and now. Thus, the keyboard appears (or is appresented) for analysis as a locale of activity, handiwork, which is ecologically organised by handedness and fingering (Sudnow 1979). What this work is for, how it relates to other work we do, and other questions are not dismissed. They are simply set aside, for now, while our attention is turned to the keyboard as a phenomenon of immediate experience. The aim of this method is to move, step by step, through the levels of our experience, constituting each 'higher' level from those on which it is premised.[4] In this way, what is taken for granted or assumed at one level is, in its turn, subjected to scrutiny.

Garfinkel takes this époche, this natural attitude of the commonsense world, to be his phenomenon.[5] He asks how appearances, the obvious features of a phenomenon, are constituted as the taken-for-granted facticities they are. In line with the praxeological rule, he proposes that the commonsense properties any social phenomenon (just what they are for for those in the setting) be treated for the purposes of sociological analysis as the outcomes, the achievements, of those participating in the setting at hand. What some activity or other phenomenon might be for the participants is "bracketed". In its place an enquiry is initiated into members methods for producing that phenomenon there and then as that phenomenon and recognisably that phenomenon. Hence the commonsense recognisability, meaningfulness, comprehensibility (what Garfinkel calls "accountability") of phenomena becomes the objective of enquiry. At the heart of this enquiry is the proposal to treat this commonsense accountability as an *intersubjective* accomplishment.[6] The methods used to secure the accountability of phenomena, ie to produce"cultural objects" within a setting, are intersubjectively shared. They are the cultural resources which interpretive social actors draw upon to produce the meaningful character of social life.

---

[4] The quotation marks here are important. The phenomenological method begins by withdrawing subscription to an ordering such as this. That we can point to hierarchies of experience reflects our capacity to organise and construct the facticities of the world, not any essential chanracter it must have.

[5] In so doing he draws extensively on Husserl’s later work (Husserl 1965) and its interpretation by Schutz (1967).

[6] The term “interpretive” is used in preference to “interactional” to emphasise the character of these methods. In line with the Geertzian exhortation, our framework focusses on the processes of signification. It also enables us to avoid the possibility of a narrow identification of these methods with face to face interaction or other forms of synchronous co-presence.

Applying the rule to technological work

The studies which Garfinkel and others have carried out within this broad programme have picked out numerous themes. We will relate them to the topic in hand by picking out three and by designating for each some possible 'objects for enquiry' from within the arena of Air Traffic Control for each. It is important at this point to remember that when some phenomenon becomes an object for enquiry, its commonsense status is suspended. Instead the phenomenon becomes a praxeologically achieved "cultural object" (Garfinkel et al. 1981). What it is, what it means, the place it has in the scheme of things etc., is a members accomplishment.

The themes we have selected are:

(1) *The natural metaphysics of a setting*: How are the objects which those active in the setting (eg the system, the controller, the state of the system, a near miss, a conflict, rogue aircraft, a stack, a Standard Instrument Departure (SID)) constituted, recognised, oriented to organised and related to one another? How are they classified, grouped, arranged?

(2) *The situated reasoning of the setting*: How do members recognise similarities and differences, relationships and discontinuities between objects and classes of objects? How are the causal and other sequences of actions produced such that a train of events becomes obviously 'a near miss', clearly 'a routine departure', recognisably ‘stack jumping’ and so on. Here two distinctive clusters of notions have emerged. They are "the local historicity"and "natural accountability" of phenomena. The former refers to the precise course of actions and treatments through which some phenomenon passes. The latter is the consequence of "the documentary method of interpretation" whereby the sense or meaning which some phenomenon has, what it 'really' is, is not decided in the abstract but is constantly reviewed and revised in the context of particular events as the local historicity of the cultural object unfolds.

(3) *The ‘haeccities’ of the work site*[7]: How is the activity in hand produced as just that activity? How do Controllers do " competent controlling" in ways which are recognisable to any practitioner as just that and not mimicry, going through the motions, simulation, or whatever. How can one Controller assemble a picture of the traffic flow which another is organising? The haeccity of the activity-in-the-work site is the specifying details of, for instance, (1) and (2).

---

[7] This neologism was first offered by Lynch and Bogen (198) in a discussion of

We will now sketch one way in which these ideas might be applied to the description and analysis of one aspect of human computer interaction. We then return to them in the specific context of Air Traffic Control.

THE SYSTEM AND ITS USER: AN ANALYTIC SEPARATION

We can begin with what is, or at least appears to be, a natural distinction, that between the user considered as a bundle of social, cognitive and biological properties and the collection of hardware and software properties designated as the system. Design, of the user interface at any rate, is the explicit attempt to manage this natural distinction. Putting it another way, user interface designs which are successful (in whatever ways one wishes to measure such success) seem to be so in virtue of their capacity to solve the problems engendered by this separation. Somehow or other (and because our knowledge of their success remains so vague, it is somehow or other) the user manages to reach across the divide to manipulate the system or to communicate with it in some way. Hence the attractiveness of the two most widely used metaphors for human-computer interaction.

This distinction is grounded in the metaphysics of the natural attitude, a metaphysics which constitutes classes of objects such as pieces of machinery, computer programs, and human beings in entirely different ways and with radically different although not mutually exclusive ranges of properties. For anyone moving around their ordinary world, machines, programs and people *just are* different. Being able to spot the differences and relate to each class in an appropriate way is what constitutes normal social and cultural competency.

If one now approaches the distinction from the position set out above, although that ontology, that certainty about the grounding of the differences just listed, remains fixed and stable, what underpins those grounds is subjected to a particular order of scrutiny. This involves stepping back from the certainties of the situation, those things wich are taken for granted within the situation, and suspending judgement upon them. It follows that the task now becomes the description of procedures (what are termed "members methods") by which the recognisability and naturally accountability of the distinctions such as that between system and user are achieved as part and parcel of doing, for example, air traffic control. In turn, this means focussing upon system and user as *culturally achieved objects*, to use a term we invoked earlier, and the distinction between them as a culturally produced one. In this way, the facticity of the distinction is suspended while at the same time recognising that Controllers depend on this facticity to do their work. Instead, we can ask what, at any particular moment, the system was *for* the user, or where the boundaries of *the-system-in-use* might be located for those working with it. How, for instance, do Controllers recognise the character of their tasks when they are actively enaged in working with the system to manage traffic? Are they manipulating objects within a system, communicating with it, immersed in it, or what? Any of these? All of these? Do they experience the natural distinction between system and user or, is it a feature of workaday ATC that they become part of the working system? How does the distinction between user and system, Controller and managed technology, feature in the endogenous accounts of the way the work is done? Is the distinction evident in working and talking air traffic control? Does the practice of air traffic control rest upon the system-user dichotomy? Questions such as these indicate how we can make the distinction between system and user move from being a premise for practical action in the world to a definition of a topic for analysis.

It is here that the notions of the documentary method of interpretation and the situated organisation of knowledge have their part to play. As numerous studies have demonstrated, the knowledge which ordinary people have and use in their daily lives is not easily summarised as lists of de-contextualised propositions, be they formally specifiable or tacit. It is, rather, highly localised. To use a phrase of Schutz (1967), such commonsense knowledge of the lifeworld is organised by structures of relevance. The massive (and often unappreciated) implication of this view is that actors cannot be treated (for the methodological purposes of sociological analysis, that is) as merely the holders, possessors or repositories of compendia of explicit and tacit knowledge. What they know, what is known, what can be known, what is relevant is a matter of constant enquiry and discovery.

Bringing these ideas to bear upon human-computer interaction throws up a number of possibilities. The first is that any description of the setting/ working system should be a description of the system-as-seen-and produced-from-within. Second, and closely related, the constructs around which that system is organised (such as the separation of and interaction between the user and the system) are treated as resources for the construction of meaning or sense assembly of the working system as its on-going achievement. The working of the 'working system' produces the distinctions. Thus their 'truth', 'facticity', 'veridicality', 'actuality' are bracketed and the methodological practices which produce and re-produce these features of the system (such as the 'fact' of a separation between user and system) can be laid open for investigation. That these practices are essentially reflexive is demonstrated in the ways in which what the system means for the user (say) is at any point open to enquiry by the user and therefore what the system "is" in some sense constantly open to review and revision. From the point of view of this investigative stance, the separation of user and system is an *in-situ* achievement.

The separation between user and system is a conceptual distinction which runs deep in our design of systems. But it is a distinction which we can reflect upon and, should we so choose, abandon.

This is where the implications for of this approach for design become visible. Once we can suspend the natural metaphysics of the setting, we can ask if it is possible to design systems which do not invoke what Merleau-Ponty (1962) described as the *prejuges du monde*.[8] The point is, of course, that having asked where, when, and how the natural distinction is used in particular settings, we can go on to ask if it provides an adequate basis on which to design for those settings. In its place, we might opt for some other tack which does not specify in advance the omni-relevance of the user-system dichotomy. One such might be the approach to design which begins with a deep understanding of work practice.[9]

THE CASE OF AIR TRAFFIC CONTROL

In this section, we will endeavour to demonstrate how the somewhat abstracted considerations we have been discussing can be played out in the description of working with complex technologies. To enable the reader to follow the detail of the work practices we decribe, we begin with a summary of the setting and a sketch of how controlling is carried out. We then turn to two specific aspects of ATC work, managing the strips and stack jumping. Our presentation of these activities will be motivated by the desire to provide an adequate phenomenology of this technological work.

The setting: London Air Traffic Control Centre

The organisation of air space above The British Isles is somewhat complicated and in the medium term likely to become more so as parts are re-configured. The simplest division is that between 'controlled' and 'non-controlled' air space. In the former, aircraft are free to move at will, for the most part. In the latter, all aircraft must be controlled by an appropriate ATC. Controlled air space takes three forms: en-route sectors where planes are at or approaching their crusing heights and speeds, Terminal Manoeuvring Areas (TMAs) where streams of planes seeking to land in Britain are organised, and Aerodrome Approach where planes are taken in to landing. The control of en route sectors and TMAs over England and Wales (except for Manchester) is located at London Air traffic Control Centre (LATCC).

The suite at which a Controller sits is 1960's science fiction dream. (Cf Diagram 1) Panels of buttons and surround the circular luminescent green radar screen. Above them are screens displaying information of various different kinds. Away to each side are clacking line printers.

---

[8] It is precisely this which, we would argue, Buxton and others have been exploring in their development of alternative input and feedback devices for interactive systems. See Baecker and Buxton (1987) and *Human Computer Interaction,* vol , no .

[9] Suchman (1989) and Brown and Duguid (1989) are examples of such an approache. Janik (1988) explores their philosophical basis.

Each Controller is hooked into the suite by a trailing cable from a head-set and microphone. In the centre sits what is possibly the only anomalous feature, a wooden tray holding printed strips with hand written notes scribbled all over them. These are the ‘strips’. To those who work with the technology, the strips are the key to good controlling. As one Controller said: "You have got to have a complete picture of what should be in your sector and what should be in your sector are on those strips." He went on to describe their use:

> *It is a question of how you read those strips......An aircraft has called and wants to descend. Now what the hell has he got in his way? And you've got ping, ping, ping, those three. Where are those three? There they are on the radar. Rather than looking at the radar. One of the aircraft on there has called. Now what has he got in his way? Well, there's aircraft going all over the place, now some of them may not be anything to do with you. It could be above them or below them. Your strips will show you whether the aircraft are above or below, or what aircraft are below you if you want to descend and aircraft, and which will become a confliction. .....You go to those strips. You pick out the ones which are going to be in conflict if you desacend an aircraft, and you look for those on the radar, and you put them on headings of whatever. You find out whether those, what those two are...which conflict with you third one. It might be all sorts of conflicts all over the place on that radar. But only two of them are going to be a problem. And they should show up on my strips.~*

Flight data strips are about 1 inch wide and 8 inches long. They specify the flightpath of an individual aircraft. This includes the aircraft name or "call sign" and type, its departure and destination point, its prefered route, height and speed. In addition, the estimated time of arrival at certain navigation points in the sector is printed at the side. Each sector has three of four key navigation points, strips being printed for each point for each aircraft. The strips are placed in racks or "bays" just above and behind the radar screens. Strips are printed 10 minutes of so before an aircraft is due at a point. The strips are, then, a documentary record of the aircrafts passage through the sector. As each point is crossed, the respective strip is discarded.

This record is what Controllers attend to and use in their work. It is the material instrument and work site of controlling. And yet, the strips do not determine the sequence of actions which Controllers perform in the sense that whatever comes along the production line determines what the line worker has to do next. Rather, the Controller has to organise the strips so they can become a resource which in turn helps to organise the work of controlling. Strips are 'glanced at', 'searched for', 'taken heed of', 'ignored', 'revised', not just when they first arrive but continuously.

This activity is itself work, the work of organising of 'doing' the work. The outcome of this work is that, at any moment, what the strips provide the controller is an 'at hand' and 'in hand' sequence of actions through which to create 'order in the skies'. Management of the strips is, in very large measure, management of the traffic.

The Controller's problem

At its simplest and most general, the Controllers problem is a scheduling one. For any controller controlling any segment of air space, the traffic has to be taken as and when it arrives and threaded together into an orderly pattern before each individual plane is handed off to the next sector or controlling segment.[10] All of this scheduling and traffic management has to be achieved in and through making the traffic flow. Aeroplanes cannot be "parked" for a couple of minutes; nor can traffic jams be allowed to occur. Even when they are put on holding patterns of various sorts, aircraft are still on the move, part of the flow of traffic.

To solve the scheduling problem, the Controller utilises a number of different resources. Two are, in essence, technological givens in the environment since they are related to or constrained by the hardware and associated software located on the suites themselves. These armight be summarised as:

(1) *informational resources*:
  a. the radar screen and its data
  b. -the flight strips
  c. screens of weather conditions.etc
(2) *communication resources*:
  d. radio-telephone to aircraft
  e. telephone links to other controllers etc
  f. face to face communication with the suite team

In addition, there is the Controller's working knowledge of the system itself. This accumulated, know-how, know-what of years of experience is brought to bear on the resources provided by the technology to determine what in any particular set of circumstances appropriate courses of action should be. The point we are making is an obvious and well known one (Reason 1986). The Air Traffic Control system comprises numerous complex sub-systems, instantiated in hardware , software, regulations for controlling, working practices and the like. In the face of the ordinary contingencies of practical working life, conflicts, inconsistencies and incompatabilities are bound

---

[10] We have discussed this in Anderson et al (1990)

to arise. These constitute 'the normal, natural troubles' (Garfinkel 1967) of the Controller's working life. Dealing with these troubles is part and parcel of competent controlling. Indeed, being able to recognise them as the 'normal. natural' phenomena they are, is to some extent anyway, what being a competent controller involves. Since these troubles can occur both with the traffic and with the technology, their solution is achieved by managing the traffic in the context of the technology. The skills required to do this (what earlier we called 'working with the technology') are multi-layered and interwoven. Further, they often seem to lack a sense of deliberation, cogitation, task-definition, specification and solution. Rather, the process is interpretive but in a somewhat different sense than normal. The ATCO just knows what to do. What 'knowing' means here is interpreting the conditions at this suite at this point in time, against a background of what has gone so far, what time of day it is, where everything else is presently, what has yet to arrive, what is going on in neighborouring sectors, etc. The whole is a *gestalt contexture* which provides the meaning of what is to hand. The problem the Controller is faced with is *this* problem here and now where that is *obviously* the appropriate course of action to take. The ATCO experiences problems and their solution, then, as part of a flow of work. The description of that experience is that of the work's phenomenology.

Managing the strips

We can begin to focus down on the achievement of competent controlling by looking at the detail of managing the strips. This move can be made in any number of ways. One is to ask what the strips are as cultural objects in this environment. In a lengthy description of one such exercise, Garfinkel and his colleagues (1981) refer to this move as "extracting the animal from the foliage". By this phrase they mean that, in order to be taken for granted, the organisation which produces a phenomenonhas to be made invisible, unquestionable, routine. The methods used to make the phenomenon invisible are somewhat akin to those which create a "potter's object"; that is an object which has been shaped and given its being through its production process. In describing them in this way, we are able to focus on the ways in which the orderliness of the strips is produced. Part of this will involve explicating how the orderliness of the strips is related to the orderliness of planes in the sky. The attractiveness of adopting this strategy towards flight strips and their management, is first thatn it makes a number of features are made immediately available for analysis. Second, it brings out the importance of the unforlding dynamic of activity. Courses of action emerge and are construed over time.[11]

The features we wish to bring out are:

---

[11] See McGrath (1990) for an illuminating discussion of time in collaborative work.

(1) *Strips are institutionally organised objects*: when a strip appears in a bay, it is not just an object in the environment. It represents a whole set of institutional processes; the scheduling of flights; the filing of flight plans; the application of control regulations and procedures; prior control actions; all of which have shaped the character of the information presented and preserved on the strip. Further, while the Controller takes it for granted that these processes have operated successfully to produce this strip, he or she also takes it for granted this is an appropriate strip for this position on this suite. Obviously, the system does not run faultlessly, but routinely the strips arrive in the right place at the right time and are presumed to do so. The character or location of any strip is not grounds for immediate enquiry.

(2) *Strips are part of a mutually explicating system*: the focus of work is not, s is conventionally thought to be the case, the screen and the R. T. Rather what is seen on the screen and heard over the radio telephone is made sense of via the strips. This process of mutual elaboration and explication is effortless and hence invisible while the various components are routinely co-ordinated. When co-ordination fails discrepancies between what is on the screen, what on the strips and what has been said or heard have to 'normalised'. That is, they have to be treated as objects of 'normal' (that is, routine) enquiries, and thus dealt with as 'normal , natural' troubles of the system. For example, strips state when a plane is due to arrive at a particular point in a sector. Failure of the plane to appear at the appropriate time, or failure of a pilot to contact LATCC when required to is not grounds for the presumption that the system has failed or the aircraft crashed. Rather the presumption is that the data updates are in front of the traffic and that this strip should be set aside for now. There is no need to suspect that the displays are faulty or the RT malfunctioning. Thus, in 'normalising' troubles in this way, the Controller is able to organise events into events which require attention just now and those which can safely, surely, quite responsibly be set aside until later. Exactly the same orientation to the mutual explication of the system, its interdependency as a working system if you like, is brought to bear on those occasions where strips are 'missing' (ie have not been placed in the bays at the appropriate time) or where there is no plane-on-the-screen.[12] Controlling continues *as if* the appropriate conditions were in place. Planes are not sent to holding patterns; enquiries as to the 'state of the system' are not instigated. In the case of the lack of a plane-on-the-screen, a number of routine 'remedies' or 'repairs' may be initiated. Controllers may carry on without the data on the screen, request the pilot to check his transponder, switch to primary radar and so forth.

---

[12] The notion of a plane-on-the-screen designates the radar blip, its ‘footprint’ and the related data block.

Both of the above features reflect what might be called the "Controller`s horizonal structures of relevance". What is taken heed of, what is seen as necessary "to do the job" by the working Controller at the suite. By its very nature, this structure is not predetermined. It is, rather, responsive to the exigencies of the moment and hence contextual in form. What a Contoller feels he needs to know, what the screens, strips or data indicate, are things which are worked out on a moment by moment basis. This "sense assembly~ is, as Garfinkel (1967) points out a version of the 'documentary method of interpretation'. That is, it is a method by which what these events mean and signify is determined 'for now' by what just happened, what can be seen to be about to happen, what is expected, and what is known about typical situations like this. The meaning is found (or better constructed) in context.

(3) *The strips as a proxy orderliness*: The ordering of the strips is usually top-down. Those at the bottom are the strips currently in use or 'live'. Those at the top are those which are pending. Within these global distinctions. strips are organised in terms of the estimated arrival time at the relevant navigation beacons. In another paper (Anderson et al 1990), we discuss the 'proxy orderliness' of the screen and the sky. There we argue, that if considered as the work site of controlling the strip bay and the radar screen *are* the sky. The one goes proxy for the other. While this is so for the strips, some qualifications need to be made. The pattern of movements is what the screen represents. It is a computer generated, two dimensional "picture" of the sky. The strips are the materials through which the patterns on the screen and thus patterns in the sky can be organised. With the strips, what we have is a linear sequence of 'objects for processing'; that is, a sequence of sequence of working out and yet to be worked out courses of action. In this sense, the strips represent a pattern of tasks which as it is gone through and completed produces the orderliness of the 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, etc etc is not just a making up of the record of work, it is doing the work.

The handiwork of working the strips

Much has been made, lately, of Heidegger's (1962) analysis of the "ready to hand" (Winograd and Flores 1986, Dreyfus and Dreyfus 1986, and Ehn 1988). However, while this term has been used to pick out some features of technological objects in the world, their artifactual, constructed characteristics, very little has been said about the ways in which routine work practices produce ready to hand organisation of objects in this way. This organisation we will call "handiwork". Typically, a strip becomes 'live' when a radio message is received from a plane when it enters a sector. The Controller selects the appropriate strip and moves it down the rack to the live strip section. Live and pending strips are separated in the rack by a strip designating the navigational point currently in use. Location in the rack, then, clearly designates status.

However, once 'live', the strip is not just placed anywhere. As with pending ones, live strips are in sequence, with the latest to arrive at the top. There is then a normal flow of traffic down the bay and where any aircraft is on its way through the STAR (Standard Arrival), for instance, can be seen simply by locating the strip in the bay. This sequential organisation both contributes to and reflects the fact that traffic management decision making is an inherently sequential task. The latest additions to the flow will not be finally dealt with until after earlier ones have passed through.

The production of this sequential flow of strips is collaborative work in an obvious way since a number of personnel have the right to arrange strips, mark them, and so on (Harper et al 1989). However, at any point, it is important to be able to determine just who is making what decisions concerning the aircrafts routing and who is responsible for the status of the information being displayed. Thus when a strip is moved it will be colour coded by the Controller who moved it. Sector Chiefs and Air Traffic Control Assistants (ATCAs) may replace strips by more up to date versions displaying revised pending times, destinations etc. In that it is unmarked or marked up in a particular way, enables a Controller to 'see at a glance' what needs to be attended to and what the implications are.

Any command given to an aircraft is marked on the strip. This has a two-fold function. It ensures the strips are an accurate record of decisions taken: it also enables this record to be publically available. The strips are ready to hand for anyone who can read them.

Finally, when the aircraft crosses the navigational point, the strips are not just thrown away. The Controller puts a cross through it. This is particularly important when the plane is handed off to the next sector (Anderson et al 1990). When this happens, the strip is crossed out when the the pilot is told to contact the next sector. In other words, the crossed out strip is both evidence of and the final stage in the completion of the work.

Management of the strips is, then, the creation of an orderly traffic flow. Smoothly moving strips are smoothly moving planes. If we look at this aspect of ATCO routine work from this perspective, it becomes increasingly difficult to support the contention that the strips represent an externalised 'knowledge stack' which Controllers continually build up, maintain and put to use in order to be able to operate the system. Rather from the point of view of the working Controller immersed in handling the traffic, what he knows, what the strips mean, what choices are open, the order of events recorded on the strips, where they are in the bays, where the aircraft are in the sky, the trajectories and routings given to the planes, in short what the strips represent, are all part of the system.[13] The pragmatic justification of the distinction between system and its user disappears. The Controller is in the system: and the orderliness of the system is managing and reproducing itself.[14]

Stack Jumping

Stacks are located in the London Terminal Manoeuvring Area (LTMA), a sector of the airways roughly co-incident with South East England. For controlling purposes, it is divided into a North and a South sector, each of which is further divided as occasion demands. (cf. Diagram 2) The primary task of LTMA controllers is to separate outbound traffic leaving Heathrow and Gatwick and climbing to the levels stipulated by their Standard Instrument Departure patterns (SIDs) from inbound aircraft to the same aerodromes. To ease congestion at busy periods, in-bounds are sent to one of several locations, or holding points, where they circle until space is available for a landing approach. These locations are the "stacks". Heathrow has four (over Lambourne, Biggin Hill, Ocham, and Bovington). Gatwick has two (at Mayfield and Willo). The number of aircraft in each stack and the number of stacks in use varies with how busy the sector is. As airspace fills, they are held at higher and higher levels. Each plane is separated from those above and below by 1000 feet. As planes leave the bottom of the stack, those above are directed to spiral down one level.

The purpose of stacking is, of course, to turn a varied flow of aircraft coming from all directions into a predictable stream of planes which airports can handle. The Controllers only have to direct aircraft to the top of the stack while Approach Control (situated at the airport) takes them out from the bottom. On the other hand, departures consist in a regular stream which have to distributed across the various routeways. The LTMA Controllers receive traffic from the main airports and must direct it around the in-bounds before allowing it to turn away onto its designated routes. What this involves, in practice, is feeding planes around or over the stacks.

There are, as one would expect, sets of procedures for both these tasks. Within LTMA sectors, the most important of these relate to the standard profiles of aircraft into and departing from the airports and related stacks. These procedures are laid down in the manuals and take the form of

1. Standard Instrument Departures (SIDs) which detail the exact trajectory of outbound traffic and are designed to satisfy noise abatement requirements and ensure no flight conflicts with in-bounds.

---

[13] This point is akin to Lynch's (1989) definition of experimental method as an "externalised retina".

[14] We would speculate that the distinction between system and user is visible only when routine procedures are not adhered to. In ATC such occasions are rare. In another study (Anderson et al 1989), we have described ways in which "one offs" are routinely normaliused and proceduralised. We suspect the same processes are at work in ATC.

2. Standard Arrival Procedures (STARs) designed to co-ordinate in-bound traffic. As with SIDs, STARs reflect the destination and route of the aircraft.
3. Agreements between LTMA and neighbouring sectors about which levels aircraft should be handed over on.

However, from both the Controller's and the aircraft's point of view, the standard procedural flight rules may not always be the most expeditious way of controlling. Nor do they necessarily ensure safety. STARs and SIDs are complex and, in many ways, restrictive because they have been designed to weave traffic through but away from all other traffic. In addition, they do not take account of the differences in performance between aircraft. They are such that any aircraft can follow them. They can result in delaying the ascent of an aircraft to its optimal cruising level and speed, thus prolonging flights, increasing costs and creating extra pilot work into the bargain. In addition, they create more work for Controllers, since planes on STARs and SIDs can be in a sector for much longer than they need and so use up airspace, RT time and require extended Controller attention.

Not surprisingly, then, Controllers have developed procedures for dealing with the “troubles” which the conflicting demands of STARs, SIDs and stacks create. These procedures are now seen to be essential to professional competence and controlling skill. They and the techniques associated with them, avoid delay, reduce work-load and contribute to increased safety by reducing the time an aircraft remains in a busy sector. They are ways in which expert controllers display their expertise by working within the system to manage the system. When faced with the possibility of, if not conflict then certainly inconsistency between sub-goals of the system, eg segregation of traffic and expeditiousness, Controllers use the resources provided by the system to achieve working and workable solutions. One of these techniques is known as "stack jumping": that is, climbing an out-bound over an in-bound, the latter being either fixed at a level or descending in the holding stack,. This procedure is used rather than climbing the out-bounds more slowly by passing them under the stack and routeing them to the outskirts of the TMA before allowing them to climb as the SIDs require.

"Jumping" a plane through the stack can only be done because the Controller is at one with the system. Although the aircraft remains at a low speed to satisfy noise requirements, from the configuration of in-bounds, both in the stack and on their way to it, as well as those under the control of the Approach Controller, the Controller senses there is enough "space", for the plane to jump through where "enough" here means "enough to satisfy the requirements of safety and competent handling". The former is defined by the Air Traffic Control Manual and the latter by the cultural practices of the setting. For example, there may be two planes circling in the Biggin Hill stack, one at 7000' and one at 8000'. An out-bound on its way to the Daventry sector would only have to climb to 9000' before or by Biggin Hill to be safely clear of the stack and be able to continue its climb out of the LTMA, even before it has reached the northern geographical boundary.

In this way, the out-bound will have "jumped" all the in-bounds under TMA control and will, almost certainly, be allowed to continue its climb in the relatively empty sector above the TMA much sooner than allowed in the SID.

On the face of it, the practice of "stack jumping" looks to be a relatively straightforward tactic. Just what you might expect experts to do. The issue is, though, not that controllers produce a "simple and easy" solution to a problem (which they do), but the work and skill which allows them to *see and feel* just how and where the system affords such a solution to problems it has itself created, and then to employ it in the ways in which they do. Effortless though it appears, this work, this expertise, is by no means simple to describe nor easy to acquire.

To begin with, stack jumping requires a complex series of judgements about the changing structure of the traffic flow, the performance characteristics of particular aircraft and an awareness of how things are going. This Controllers refer to as their "picture". When jumping a stack, previous out-bounds have to be considered in case they are slow and thus likely to be in the way of subsequent, faster planes. Or there may be too many planes converging on the stack at its top level indicating that it will have to be raised before a possible "jumper" could get there. On the other hand, there may be the possibility of creating space at the top level by slowing down planes which are approaching the stack. Added to this is the fact that the speed of modern planes is such that often there are only moments to notice a slot and decide which out-bound to jump. For this reason, stack jumping is rarely practised by new Controllers or those who have been off duty for some time.

The advantages of stack jumping for Controllers and the system are obvious. It ensures quick exit of aircraft from the sector. It frees RT time, gives the plane to the en route Controller earlier which can ease the handling problems of in-bounds. It enables planes to reach efficient operating height and speed quicker and, since it is a simpler trajectory, often eases passenger comfort. So keen are some pilots to jump that they "offer good climb rates" to Controllers on their first contact with LTMA. Many aircraft do not have the capability to climb as fast as jumping requires. Informing a Controller early that such a climb rate is possible greatly increases the likelihood of being offered one.

There are, however, limitations. Apart from the need for quick assessments of such situations, the most troublesome is the failure of the plane to reach its directed climb rate. There may be various reasons for this, but it has serious consequences. A Controller may be depending on an outbound to climb in front of an in-bound, and if the out-bound does not climb fast enough a possible "confliction"[15] may occur. Other problems relate to the distribution of controlling responsibilities between the Approach and LTMA Controller. Occasionally, one or other might direct an aircraft through the other's airspace without prior co-ordination. For example, the LTMA Controller may choose to take a plane not over the stack but around the middle of it and hence through the flight path of those emerging at the bottom of the stack (remember they are climbing all the time). On others, a plane may be taken from mid-way up the stack In cases such as these, one or both planes may have to be re-directed.

The invisible skills of "stack jumping

Stack jumping requires intimate knowledge of the routines of the sector and of the aircraft currently being controlled, traffic requirements, an awareness of the amount of attention the Controller must give to any one manoeuvre, and much, much more. It is then sector specific. This knowledge has to be applied and honed time and time again to allow the procedure to be effective, smooth and trouble free. It requires exact assessments of the progress of aircraft along their given vectors and where "in the sky" they are in relation to one another. These assessments are based on information 'seen at a glance' and then the appropriate course of action is "executed' immediately, without deliberation, almost nonchalantly (Sacks no date). It also requires trust in the system. For all these reasons novice Controllers generally shy away from it. However, the skill and the work by which it is brought off are , in a sense, made invisible by the very effortlessness of the achievement. That experienced Controllers do not haver, ponder the possibilities; that they act smoothly and efficiently to produce the space for a jumper to jump through with no hiccoughs, finger crossing, drastic changes of mind or direction makes it difficult to see the artful and professional handling of the system which makes its all possible. This is all the more so since such artfulness and skill are not sedimented in a presevable product but, rather, documented only ephemerally in the orderly progression of planes-on-the-screen, strips in the bay, inscriptions on the strips and exchanges with aircraft, controllers, and so on. This skill of competent controlling through stack jumping involves working the system to satisfy the procedural rules of air traffic control "here and now" (Garfinkel and Sacks 1970), where what counts as satisfying the rules is the production of smoothly flowing traffic and demonstrably competent controlling of whatever aircraft are in the sector at any moment.

---

[15] A confliction is the merging of the trajectories of two aircraft to violate the minimal requirement of two miles horizonal and 5000 feel vertical separation.

CONCLUSION

At first sight, to judge from the case of air traffic control, the framework we have outlined seems to make the designer’s task harder not easier. Things seem to be complicated beyond necessity if not belief. However, we would suggest that drawing attention to the specific character of our experience with and of technology can help to frame difficult issues of design in novel and perhaps more tractable ways. By way of conclusion we will summarise four ways in which this might occur. Since the materials we have presented are drawn from air traffic control, the issues which these primarily relate to are ATC ones. However, they are also more general than this.

1. The approach pushes the conceptualisation of the working system beyond the hardware and software components. The system-in-use is treated as a fluid constellation of objects, processes and actors within which the user is immersed. The courses of action engaged in are not defined simply as the implementation of procedurally defined rules but as the contingent outcome of interpretation as to how the rules fit the case in hand.

2. The focus of the approach is the user as a manager of technology, that is as someone involved in getting something very specific done, here and now, with whatever resources are to hand. Activities of others, behviours of the system, are all treated from this praxeological point of view. Thus they are seen as sequences and entrained (McGrath 1990) and their meaning deciphered from their location in a stream of action.

3. The unit of analysis is the embodied course of action which often is best represented by the working team not the individual at the work station. It is the working team which achieves a division of labour, which circulates knowledge and which reproduces the production processes. The individual (controller in the case above) is an individual-controller-in-a-team. Failing to appreciate this aspect can lead to unfortunate design decisions, such as, for instance the fragmentation of teams by isolating controllers and thereby depriving them of many of the invisible but vital resources they need to carry out recognisably competent controlling. Since these resources are relied upon in ways controllers find it hard to articulate, they might well be left out of a requirements analysis. They are not compensated for by providing enhanced data sets on-line. It is precisely what is not on-line that is known in common, read off from the configuration of objects, the ‘picture’ the controller has.

4. The possibilities of re-constituting the flexibility equation are enhanced. Designers are now seeking to make technologies more flexible and more adaptive. Until recently, it has been the user who has had to adapt to the technology not vice versa. What this approach offers is a way of grappling with what the features of this flexibility would have to be - where it is located and the forms it takes. What does a flexible technology have to fit in with? This is where the distinction between user and system might well be re-appraised.

We have no illusions that this paper does any more than hint at these possibilities. Much more needs to be said about almost everything. Further, we are also well aware that many will not see in what we have said anything like an conclusive demonstration of either the distinctiveness or the fertility of the approach we recommend. This would disturb us greatly if it was demonstration of that order we seeking to provide. But it was not. All we wish to do is indicate one way in which the problem of interrelating the social and design disciplines might be addressed and to show what we feel are some interesting possibilities which might then be made available. One upshot might well be a cessation of the indiscriminate use the various formulae or metaphors such as 'dialogue' or 'tool manipulation' offered to depict human-computer interaction and the development of alternative ways of talking about working with and within the system as an environment or ecology of action. One outcome of this might be a reconfiguration of the analytic relationships we stipulate as holding between user and system.

REFERENCES

Anderson, R., Hughes, J., and Sharrock, W. 1990 The Division of Labour. Paper presented at 'Action Analysis:Conversation Analysis', Maison De L'Homme, Paris. September.

Blomberg, J. 1985 Computer technology and social organisation of work. Paper presented at Annual Meeting of American Anthropological Association Washington DC December 1985.

Baecker, R. and Buxton, W. 1987 Readings in Human Comp[uter Interaction. Los Altos, California. Morgan Kaufman

Galegher, J. and Kraut, R. 1990. Technology for intellectual teamwork: perspectives on research and design. In J. Galegher, R. Kraut and C. Egido (eds) Intellectual Teamwork. Hillsdale, N.J. Lawrence Erlbaum. Pp1 - 20.

Garfinkel, H. and Sacks, H. 1970 On formal structures of practical actions. E. Tiryakian and J. McKinney (eds) Theoretical Sociology. New York. Appleton Century Crofts.

Garfinkel, H., Lynch, M., and Livingston, E. 1981 The work of discovering sciences construed with materials from the optically discovered pulsar. Philosophy of the Social Sciences, vol 11, pp 131-158.

Geertz, C. 1975. “Thick description”. Notes towards a theory of ethnography . The Interpretation of Culture. Hutchison. London. Pp 1- 30.

Green, T. Schiele, F. and Payne , G. 1988 Formalisable models of user knowledge. T. Green et al (eds) Working with Computers: theory versus Outcome. London. Academic Press.

Heidegger, M. 1962. Being and Time. New York. Harper and Row

Husserl, E. 1965 Phenomenology and the Crisis of Philosophy. New York. Harper and Row.

1960 The Cartesian Meditations. The Hague. Martinus Nijhoff

Lynch, M. 1988 The externalised retina: selection and mathematization in visual documentation of objects in the life sciences. Human Studies, vol 11, no 2-3, pp 201-34.

Merleau Ponty, M. 1962 The Phenomenology of Perception. London. Routledge and Kegan Paul.

McGrath, J. 1990 Time matters in groups. In J. Galegher, R. Kraut and C. Egido (eds), Intellectual Teamwork. Hillsdale, N.J. Lawrence Erlbaum. pp 23 - 62.

Reason, J. 1986 Recurrent errors in process environments. E. Holhangel et al (eds) Intelligent Support in Process Environments. Berlin. Springer Verlag

Sacks, H., Schegloff, E. and Jefferson, G. 1974 A Simplest Systematics for the organisation of turn taking for conversation. Language, vol 50, no 5, pp 696-735.

Sacks, H. (No date) The Lawyers Work

Schutz, A. 1967 The Phenomenologyof the Social World. Evanston. North Western University Pressˇˇ

Suchman, L. 1987 Plans and Situated Actions. Cambridge. Cambridge University Press.

Sudnow, D. 1979 Ways of the Hand. New York. Knopf.