“What a f-ing system! Send ’em all to the same place and then expect us to stop ’em hitting”: Making Technology Work in Air Traffic Control

R.H.R. Harper and John A. Hughes

Technical Report EPC-1991-125

Published in:
Technology in Working Order: Studies of Work, Interaction, and Technology,
Graham Button (ed),
Routledge, London, 1993,
pp 127-144.

Rank Xerox Research Centre
Cambridge Laboratory
61 Regent Street
Cambridge CB2 1AB

Tel:+44 1223 341500
Fax:+44 1223 341510

"What a f-ing system! Send 'em all to the same place and then expect us to stop 'em hitting": Making Technology Work in Air Traffic Control

by

R.H.R.Harper

Rank Xerox Cambridge EuroPARC

and

John A. Hughes

Lancaster University

1. Introduction

Despite the long prominence of technology in social life, especially since the industrial revolution and most saliently celebrated, for social theory, in Marx's disquisitions on *homo faber,* the social study of technology has not been similarly prominent, until recently that is, with the development of social constructionist views on technology drawing heavily upon similar perspectives on the sociology of science.[1] While there have been no shortage of macro characterisations of the shifts and changes technology has wrought in work and play, health and wealth, life and death, in social and political division, there has been relatively little detailed study of how technological artifacts are understood and used as sociocultural objects within the world of living actors. Much of social science's interest in technology has rarely been an interest in technology itself but much more concerned with how it has affected work, leisure, socioeconomic change, the media, demographics, and so on, rather than with what we would want to call the sociality of technology; that is, a concern to explicate how technology as a set of cultural artifacts come to be incorporated as features of socially organised activities.

Of course, such an interest is exceedingly wide but what has helped to focus it in the last half decade or so is a recognition by system designers that a greater understanding of the sociality of work and its technology is a necessary condition of their art. (Suchman, 1987, 1989; (Grief (ed), 1988; Gallagher, et al, 1990; Wilson, 1990; Bannon and Schmidt, 1991; Bowers and Benford, (eds), 1991). In the development and design of complex technologies, the systems metaphor is not only a useful heuristic: it is also a beguiling one.[2] The interrelationship of parts each smoothly intermeshing with others to produce an overall flow of system generated actions captures the sense of a well-oiled mechanism of controlled power; an efficacious arrangements of parts designed and directed toward a goal that each part separately cannot attain. However, such an overly roseate conception underplays the fact that complex systems are also some mix of the human and the machine in which the human elements are very often the grit in the designer's vision. Whereas it might be said with not too much exaggeration that the last two hundred years or more of the history of technology it is humans that have had to adapt to technologies, for many reasons this is no longer something that designers can assume. Further, and as more automated technologies become essential to the control and performance of complex systems, nuclear power stations being but the most notorious of examples, how humans and machines resonate together becomes of paramount importance both for efficiency and safety. Moreover, it is also recognised that designers can no longer rely on 'imagined' representations of human activities as an adequate input to informing their design conceptions. Work is, among other things, fundamentally a socially organised activity and understanding in some detail how and in what ways this is manifested in the varieties of work in human social life is coming to be seen as a prerequisite of more effective system design: system design which produces systems which 'work with' and support human operators as *an instrument* of their work rather than as a recalcitrant and mysterious dominatrix.

Air traffic control (ATC) belongs to that class of communication and control systems which typically depend on a multi-user data base of relevant information which is capable of reporting on specific states of the system, either as a result of interrogation by human operators or by continuous automatic updating. The operators rely on this data base for information pertinent to their particular tasks. In the case of ATC this is much of this information is provided by a computer which both drives the radar and the flight data information which the controller, along with other members of the controlling team around the suite, use to coordinate air traffic. In the early days the movement of what few planes there were could be left to little more than the pilot's eyesight to ensure that one did not crash into another. The inexorable growth of air passenger travel, in the last decades much accelerated with the advent of jet aircraft, required traffic regulation by means of flightway procedures and a corpus of personnel to directly coordinate the flow of aircraft as they travelled from point to point, and flew from much the *same* points through much the *same* airspace toward much the *same* points. It is this which remains the responsibility of the ATC system, namely, to prevent aircraft flying through the same airspace from colliding: a task that has its absurd side as captured by the quotation that is the title of the chapter. It is a task which has over the years accumulated, and necessarily so, a considerable amount of technological support for the controller. (See Hughes et al, 1988 for further details and Harper et al, forthcoming)

Contemporary ATC is a complex of subsystems, technological, procedural, human, bureaucratic, regulatory, legal, and more. As far as the technology is concerned, it is a highly sophisticated arrangement of computer-driven radars, drawing upon an elaborate nationwide network of radar and navigation beacons, computer processed flight progress data, control suites for communicating with other facets of the system, and so on; all arranged to facilitate the work of controllers whose task it is to direct aircraft into an orderly flow of traffic. The orderly flow of traffic is the outcome of the working practices that deploy the technology 'to hand'; a weave of practices and skills that constitute the work of air traffic controllers. The system is also a regulatory one in that aircraft are expected to follow certain 'rules of the road', rules contained in the *Manual of Air Traffic Services* which specify in great detail the routes, procedures and regulations to be followed by air traffic and by controllers in directing aircraft. Thus, there is an important sense in which the work of controlling consists in applying this immense corpus of rules to aircraft entering UK airspace using technology to achieve this.

Rules as Instructions for Seeing

Sociological interest in work and its relationship to rules was provoked by Weber's (1947) reflections on bureaucratic organisation, a concern developed and extended by ethnomethodology, in particular in its resistance to the notion that rule-use, as a feature of the social organisation of activities, is straightforwardly one of stable compliance with what the rule stipulates. (Wilson, 1974) The orthodox view, very briefly put, is to regard rules, be they the cultural norms and institutions of Parsons (1958), or the proprieties of interpersonal conduct (Goffman, 1959), or the formal rules of rationally conceived organisation (Weber, 1947) as 'external' factors governing, even determining, the behaviour of those subject to the rules. Such a view fails to acknowledge that rules have to be applied within a setting such that what a rule or a procedure means, what actions fall under it, is a matter that has to be decided, judged, determined on occasions of its application. Social actors, that is, have to make judgements as to whether *this* rule applies *here* and *now* in respect of *these* circumstances. Action in accord with a rule is situatedly accomplished by actors and done 'in light of' the particularities of the setting. (Coulter, 19..). Rules do not stand as disembodied regulations 'mindlessly' applied, but are constitutive of the situation itself and can, across the endless variety of work settings for example, exhibit a variety of relationships to the 'things' and the practices of the work.

Baccus, for example, describes the regulatory rules that provide virtually no guidance and little relationship to the work of truck wheel maintenance mechanics because workplace personnel 'know better' than "any standard provisions what safe work conditions consist of as in fact they know their job as a reasonable account" (Baccus, 1986: 20-56). Wieder (1974), by contrast, describes how an informal 'convict code' of the kind regularly identified in studies of organisations of all types, regulated in fine detail staff-inmate interaction in halfway house. The 'code' was never written down as a set of regulations, though it could have been, but was attended to, cited, mentioned, discussed and generally manifested through the particularities of talk and action. Nonetheless, it was made intelligible and itself furnished intelligibility to the actions and activities within the community of halfway house. Harper (1988:297-306) describes how the predictable features of routine accounts are used as rules to determine accounting errors. The point is that rules, procedures, regulations, endemic to work of all kinds, do not, can not, stand independently of the activities of work but furnish those who do the work with ways of seeing and recognising things and activities as relevantly features of the work. And it is this interweaving of rules and activities that is also characteristic of ATC. (Shapiro, et al, 1991) The *Manual of Air Traffic Services* does not stand in some disembodied relation to the work of controlling but is integral to its activities and work as socially organised. The rules of controlling have to be imbibed *as activities* that constitute controlling as a situated and artful skill. They constitute a plan, to use Suchman's (1987) phrase, an enablement of the work of organising air traffic and, as such, furnish controllers with 'instructions' for seeing the information to hand, much of it technologically furnished, as proxy for the current state of air traffic. But, and to stress, the plan in its actualities is a continually evolving one in that though, in general, a controller knows what s/he will be doing, it is not known in advance in detail. It is, we might say, in dealing with the unfolding details of air traffic that the work of controlling consists; in a word, it is a discretionary system. (Shapiro, et al, 1991).

The interest in the relationship of rules and action also raises the matter of skills and competencies incorporated in working activities, the skills and the knowledge involved in 'doing the work'. (Lynch, 1985; 1982: 499-533; see also Garfinkel et al, 1981: 131-158; Livingstone, 1986 for example) Controlling air traffic is managing events in real time rather than following the kind of plan which determines in advance what it is the controller should do next. There are always inevitable and contingent factors, including technical troubles of various kinds, which have to be dealt with. The rules, along with technically provided information 'to hand', are resources reflexively deployed such that 'competent use' is founded on the controller's practised grasp of what particular actions are necessary on a given occasion to provide adequate aircraft separation *according to the rules and procedures.* Controller's culture is saturated with the recogition of the artfulness of the craft; an artful and skilful use of rules and procedures, the anticipation of 'likely problems', the smooth control of an aircraft's flight, the organisation of the pattern of traffic to ease the work of controller's on adjacent sectors, and more, that is exercised within a 'working division of labour' around a controlling suite.

Flight strips and the picture

What we particularly want to focus on in this chapter is one crucial element of controllers' working practices, namely, how they make Flight Progress Strips, or 'strips', *useable* within the flow of work.[3] 'Strips' represent the movement of an aircraft through an airspace which is configured into sectors; that is, blocks of sky which are distinct areas of control responsibility. As an aircraft flies through controlled airspace it is passed from sector to sector under the direction of Air Traffic Control Officers (ATCOs) whose main task is to ensure the effective separation of aircraft one from another as well as the expeditious flow of the traffic. Not surprisingly, the means of achieving both of these goals requires accurate information about planned routes, flight levels and directional headings. Coordination between sectors is one of the core tasks of controlling and one that can, potentially and in actuality, take up a great deal of a controller's time and attention. Configured airspace is the 'stage' upon which the patterns of air traffic are 'drawn' using the information and the rules and procedures of ATC. Air traffic flows into and out of controlled airspace in a variable stream which is the task of ATCOs to produce. The passage of an aircraft into, through and out of airspace is, initially, indicated through a Flight Data Processing (FDP) application filed prior to departure. Updating these is done through limited radar links with the database through Radar Data Processing (RDP) and by inputting flight data. The information from the data base is distributed throughout the operations room according to sectors.[4]

At its simplest and most general, the controller's problem is a scheduling one. For any controller the traffic has to be taken as and when it arrives in the segment of airspace for which s/he is responsible and threaded together into an orderly pattern before aircraft are handed over to the next sector. The scheduling has to be achieved in and through making the traffic flow. Aircraft cannot be parked for a couple of minutes nor can 'jams' be allowed. Even in holding patterns aircraft are still on the move, part of the flow of traffic and must, therefore, be taken into account. It is the 'strips', a formatted strip of paper containing boxes of information relating to individual flights, next reporting point on route, time due to pass that point, call sign of the aircraft, its 'squawk code', aircraft type, planned flight path, requested cruising height, departure and destination airports, and more, which are crucial elements in the melding of the information available to the controller, which also includes radar and R/T, which constitute the mutually elaborating processes by which 'order in the skies' is assured.

The local culture describes a controller's orientation to the configuration of traffic in a sector and the likely problems it poses, as 'getting the picture'. This is not a detached description but one used and understood within the controllers culture to refer to various practices in controlling work, not least the regular habit of an incoming controller spending anything up to 10 minutes watching over the shoulder of his or her predecessor to 'build up the picture' before taking over the position.[5] An important part of this is "looking about 5 to 10 minutes ahead all the time.... sometimes a little further than that", in order to get an idea of the actions needed to be taken 'now'. As one controller described the process at some length,

> ".. the kinds of check you do is check your information which is your strips which will show you whether an aircraft is in there or not, and that means then that you don't have to look all over the radar... it would be an impossible job to sit down and look at the radar and look at all the different blips and try to avoid them by putting the aircraft into blank spaces on the radar, so you have got to have this information to tell you what traffic is coming into and out of the sector. From your strips you can find out whether there or not there is a possible confliction.... and what you can do about it, then you go to your radar and look for that particular aircraft and see where it is in reference to the, outbound from Heathrow, for example, and you want to climb it , to see where it is in relation to that. Then you decide what you are going to do with it, whether you are going to go underneath it, whether you are going to wait until its gone past it, whether, its one on your frequency, whether you can put them on parallel headings and then you can climb it up to the other aircraft levels. Same as inbounds, its the same sort of thing".

The controller's 'picture' has been of interest to ATC researchers for some time (Whitfield,1979: 19-28). Its potential significance was first appreciated during trials on Interactive Conflict Resolution, a computer-assistant which extrapolated aircraft movements to detect potential conflictions. (Whitfield, Ball and Ord, 1980: 569-580). This led to the development of various techniques to illuminate aspects of the 'picture' as a 'mental model', or 'internal representation' and its relation to measures of workload (Jackson and Onslow, undated). The analytic idea of the picture, best summarily characterised in Whitfield and Jackson's words as the "overall appreciation of the traffic situation for which they are responsible", recognises that controlling work has a subjective dimension to it in the sense that the controller has to think and make decisions about the situation before him or her and, it is argued, that this process is characterisable as the matching of the 'picture' or 'mental model' with the information provided by screen, R/T and strips. What is also important is that the idea recognises, albeit tentatively, that controlling is not simply a matter of unreflectively applying rules and procedures but a matter of applying rules with respect to an ongoing configuration of traffic; in short a matter of using the rules and procedures along with the provided information for interpreting, *or making sense of,* what is going on 'now'.[6]

The importance of the 'picture' was emphasised by the growing possibilities for providing the controller with more automated assistance and the possibilities inherent in this of reducing the ATCO to supervisory control: an issue which raises the difficult question of deciding the appropriate mix of decision capacity between human operators and automated assistance given that the ultimate responsibility resides with the human operator. A major problem for designers, given that human beings are notably bad monitors, is to devise ways in which any automated facility could be elegantly overriden by the human operator; a matter also of keeping the operator 'warm', or engaged, throughout the period of duty so that if required to override the automated system time is not wasted finding out what needs to be done.[7] Thus, the examination of the components of the controller's 'picture' seemed to offer a useful route to a fuller understanding of controlling skils which could inform the design of automated systems more adequately adapted to human capacities (Whitfield and Jackson, 1982; Hopkin, 1979).

Interviews with controllers suggested that although they often found it difficult to describe the picture verbally, certain themes did emerge as important ingredients, such as, 'getting the picture' before taking over, the value of experience in handling information, the division of traffic into foreground and background, and planning and prediction and, of course, the frightening possibility of 'losing the picture'. (Whitfield and Jackson, op cit). Further studies using the Verbal Protocol technique (Bainbridge, 1974; Rasmussen and Jensen, 1974: 293-308) where an operator is asked to give a running commentary describing his or her sequence of thought and action prior to taking over a position, which is recorded and analysed, provided the following counts of protocol elements:

- strips were predominant elements establishing the picture. Expected aircraft were not yet on radar and strips alone contain the detailed information about each flight;

- radar was mentioned less frequently, and then often in conjunction with information provided by strips;

- mental activities inferred, such as time check, memory reference, predictions, calculation, decisions, were small in frequency, search activities prior to the reference to the strip or to radar were relatively infrequent.

Further analysis of the strip category suggested that controllers appeared to use route information frequently when thinking about aircraft and obtaining a general sense of the traffic situation. The next most important one was flight level, then beacon time. However, it was noted that there are variations between protocols, even though the most frequent procedure when taking over a sector is to examine the flight strips in time order. Departures from this included exploring potential conflicts first, inserting a new flight strip and considering that first, overhearing RT, among others. Also, some controllers apparently organise their picture in terms of inbound and outbound flows for various airports, relying on their knowledge of typical routes and procedures, while yet others organise it in terms of flight levels. Some also focussed on the picture on unusual and non-routine flights.

Yet, it seems to us that this approach does not treat as particularly problematical what it is controllers are doing when they perform their working tasks. That is, the reference to the 'picture' as the *analytic object* of inquiry, independent of its contexts-of-use, detracts from specification of the skills that are in use when the 'picture' is referred to, talked about, or even, and here we think use of the term is being pushed to its limits, when it is being depended upon[8]. From our point of view, these inquires oversimplify situations of choice as they are encountered by those who make choices, in this case the controllers themselves. It also underates the knowledge available to controllers as decision making agents and ignores the socially organised, and thereby situated, features of the environment within which the choice is being made. While such an approach to the examination of controlling may lead to a generalised picture - if the pun is excused - of the activity in question, what is not clear is how that picture relates to particular cases, particular instances and competences as used and understood by those doing the specific job in hand: in short, its quiddity (Garfinkel, et al, 1981: 131-158). What is lost is the realistic character of the situation in which the practical decision making takes place. It is to avoid this that we propose one needs to look at controlling in detail, as it takes place.

From this perspective the first thing to note is that controlling is not an individualised cognitive activity, but a social one in that it involves teamwork.[9] Controllers are part of a formal division of labour around the controlling suite normally consisting of the controllers for each of the two sectors, the chief and 'wingmen'.[10]

Although the distribution of positions is work specific and directed toward serving the activities of the radar controller, the actualities of the 'working division of labour' (Anderson, Hughes and Sharrock, 1989) are such that all members of the team attend to the information and the tasks and activities it supports in distinctive but closely related ways. All explicate and interrelate their respective activities as aspects of a working division of labour which they trust and rely upon. To effect this activities and the information that serves them are publically available as 'accountable and witnessable' features of the work for those who 'need to know' and have the 'know how'.

From the ethnomethodological point of view, the visibility of actions, and therefore their 'witnessability', is a key feature of social action. Within the operations room at LATCC for example, among those familiar with it, what a controller is doing is visible by looking at the strips, the screen, listening to the talk and so on, because of its familiarity to those doing the work involved. The incoming controller, in getting the picture, does not routinely have to interrogate the outgoing controller to determine what is happening but simply sits and observes the strips, the radar and the talk, and from these resources, is able to determine what the active controller is doing and why.[11] From this point of view, ATCO work becomes not so much the accumulation and construction of an 'internal representation' as the production of an observable state of affairs at the console - that array of radar screens, VDU's, radio and telephone equipment controllers work with - this being done not by a sequence of inferred cognitive operations but by a succession of witnessable social actions done so that their recognisability is evident to those 'in the know'.

As said earlier, 'strips' are pieces of paper about one inch wide and eight inches long that are formatted into 'boxes' containing information about individual flights. The information is derived from the original filed flight plans for each aircraft contained in the computer database and is, in this sense, historical rather than a real time record of flight progress. Although some updating may have taken place this is not a continuous process. Each sector will have three or four key navigation points, strips being printed for each point for each aircraft from the FDP computer up to 40 minutes prior to its arrival at the navigation point, though the time can be much shorter than this. These strips represent the aircraft at each stage of its journey through the sector. As each point is crossed so the respective strip is discarded by the controller.[12] The strips are placed in racks or bays, just above and behind the flat radar screens and this is another way of making the strip informative. That is, information *on* the strip is not the only way in which it is made relevant for the controlling activities of the team. On the 'wing' positions are 'pending racks' containing the strips of traffic due in the sector which provide an early indication of incoming traffic, its routes and likely problems. They also help some anticipatory planning of flight levels at which aircraft can enter the sector.[13] When an aircraft enters the sector the strips become 'live' and placed in a separate rack in front of the radar controller.

The strips provide the template for what is and will happen in the sector. They are the material instruments that the controllers attend to and use in their work. They are also institutionally organised objects representing whole sets of institutional processes such as the filing of flight plans, the application of control regulations and procedures, prior control actions, and more. Yet the strips do not determine the sequence of actions of controllers in the sense that what comes along a production line determines what a production line worker has to do next. Rather, the controller has to organise the strips so that they can become an instrument that helps organise and so make possible, controlling work. Strips are manipulated, glanced at, taken heed of, ignored, revised, written on, and so on. And not just when they are first placed on the racks but continuously all the time they are in use. The end result of these activities is that at any moment in time, what the strips indicate and create is the sequence of controller actions that results in order in the skies. Thus, management of the strips constitutes a large part of the work that underscores controlling competence.

Working the strips

As pointed out earlier, the controller's problem is scheduling aircraft within a sector taking each aircraft as it arrives and threading it into an orderly pattern of traffic before handing it on to the next sector.[14] The scheduling is achieved in and through making the traffic flow using the information resources of radar, strip, telephone and R/T, together with specific information about weather patterns, for example. In addition, there is the accumulated 'know how', much of it tacitly held and understood, relating to the system itself, its affordances and its problems, which is brought to bear on the information resources provided by the system's technology, to determine what, in respect of any configuration of aircraft, should be done.

Strips are not some detached record but a vital instrument of the work. (See Harper et al, 1989a, 1989b for further details). 'Working the strips' is a means of achieving a solution to the scheduling problem. As noted, strips are activitated up to 40 minutes before an aircraft is due at a navigation point. They are placed in the racks in front of the controllers. As already suggested the controller does not treat these as determinants of his or her behaviour, but rather as a resource, among others, that can be used to organise the controller's activity. The next step is to order the strips in ways that reflect the work that needs to be done. Planes fly through a sector in a sequence reflected in the succession of time an aircraft are estimated to pass various navigation points. The controllers uses these estimates and the navigation point to order the strips on the racks so that the next plane due at any point is at, say, the top of the rack, the last to arrive at the bottom [15]. Strips are not related to the real time events in the sky as represented on the radar screens, since at this stage, the aircraft the strips refer to are not yet the business of this sector controller. The actions described are very much preparatory. Though subsequently the strips may come to be seen, in as yet unforeseen and unpredictable ways, as in need of revision, supplementation or may, indeed turn out to be correct, what the business of the controller at this stage is, is ordering this resource in such a way that what he or she may require from them is available in useful ways. To use a phrase made much of by Heidegger, it is an activity of making the strips '*come to be at hand'.* In fact the bulk of time estimates are fairly accurate and so ordering the strips in time sequence usually turns out correct and hence, a useful way of ordering them. The controller will know immediately where a needed strip is, whether he or she should take account of any special features in relation to it, (i.e., in respect of those circumstances that the controller thought worthy of being marked out) and so on.

This ordering enables the controller to get a clear idea of what decisions he or she is likely to have to make in the near future. Ordering the strips shapes the controller's attention in terms of what is likely to happen in the sector, for example with respect to standard traffic patterns but also toward any special problems that need to be anticipated. Special problems that need to be taken into account include such things as two aircraft estimated to reach the same point simultaneously and at the same height (the kinds of problems that the air traffic routing system can be thought of as causing itself). Although this may well be of no immediate concern, controllers mark out such problem strips by slightly lifting them, or 'cocking them', out of the rack. In this manner when the strips become live, the controller will have already prepared them so that they indicate to the controller and others around the suite how they need to be read, specifically, and 'at a glance', in relation to that potential problem. The way in which the controller organises the strips, for instance according to arrival time over a reporting point, or flight level, or possible confliction points, organises information about the state of the sector. Moving the strips is to organise the 'picture' by organising the information in terms of work activities and, through this, organising the traffic.

'Working the strips', or making them 'at hand', continues once they become live. Typically a strip becomes live on the receipt of a radio message from the respective plane when it enters the sector or nears the navigation point. The controller selects the appropriate strip and moves it down the rack to the live strip section: live and pending strips being separated by a strip designating the navigation point being used. The live strip is not placed in just any position among the already live strips. As with pending ones, it is placed in sequential order: for example, the latest at the top, old and finished strips exiting from the bottom. This order reflects, and helps organise, the fact that the controlling decision making is a sequential matter. So the latest addition to the sequential order will not be finished with ordinarily until ones beneath, hence before it, have been finished with.

The strip is also used to display information not configured as part of the computer output. Once outputted by the computer the information on the strip is 'frozen' unless altered by the suite team. Updating is provided through other members of the team around the suite, such as chiefs, assistants as well as controllers themselves, using the strip as a notepad. As a controller remarked, the strips are "like your memory, everything is there" and, as a Chief Sector Controller put it, "Get used to writing it down as you say it. Every time you tell the airplane what to do, write it on the strip". For example, as controllers instruct pilots to climb or descend, follow particular headings, and so on, these instructions are written on the relevant strip, as are the pilot's acknowledgement and their attainment. Attention-getting information may also be written on the strip, such as arrows to indicate unusual routes, symbols designating 'crossovers, joiners and leavers', circles around unusual destinations, and so on. Such 'notes' can include information indicating coordination, changes to eta or to route, changes in call sign, and so on. The strip, that is, conveys to members of the team what actions have been done with respect to a particular aircraft. They embody its control history.

The strip provides a facility for rapidly updating information by members of the suite team as they record their actions in ways relevant to them, and in ways mutually visible to them. This is formally recognised in the colour protocol used to make notes on the strips to show whether or not it is the chief, radar controller or assistant who has written on the strip and using, but not exclusively, conventional signs to denote actions taken or about to be taken. This preserves not only *what* changes or decisions have been made but by *whom* so adding to 'history' responsibility and accountability. Chiefs will normally write coordination agreements on the strips and the 'wings', or assistants, updates of eta's. Any of the team may 'cock out', or slightly lift out from the racks, strips in order to draw attention to them. In other words, management of the strips is very much a collaborative activity exhibited in and through using them as integral features of the work.

Finally when a plane crosses the navigation point represented by a strip the controller does not just throw it away but puts a cross through it. This is especially important on those strips which represent the last point through which an aircraft passes in a sector. The controller puts the mark through the strip when he or she directs that plane to contact the next sector controller. In other words, it is a physical mark to demonstrate that the controller's work has been properly completed and indicates that the strip has not 'just' been thrown away.

Thus strips play a key role in enabling controllers to use the radar quickly and effectively and achieve 'good technique'. As one controller put it, "You've got to have a complete picture of what should be in your sector and what should be in your sector should be on those strips". He went on to describe their use:

> "It's 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, these three, where are those three, there they are on the radar. Rather than looking at the radar, one of the aircraft in 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 or below them, your strips will show you whether the aircraft are above or below, or what the aircraft are below you if you want to descend an aircraft and which will become a confliction... you go to those strips and pick out the ones that are going to be in conflict if you descend an aircraft, and you look for those on the radar and those, what those two are.... which conflict with your 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."

The strips and their organisation are a proxy orderliness for the configuration of the traffic flow. While the radar is a computer generated 2-dimensional picture of the sky, the strips are the means whereby the patterns on the screen, and thus the sky, can be seen as the patterns that they are. Srips are not just placed anywhere but are organised so as to give a sequence to them which provides a sequentiality to the traffic flow through the airspace.

Strips, Radar and R/T

Though we have concentrated on strips, and only referred to other information sources briefly, in particular R/T and radar, we suggested at the outset that all these resources come to be useable in being *mutually determinative*. That is, these sources of information are not distinctly separate items of information which are, as it were additive, but mutually explicate the sense of each other. What each source of information means is reflexively determined by the sense made of the others: a mutuality that is premised upon 'learnt through experience' knowledge of controlling and the trust in the teamwork of others. These interpretive practices, the commonsense reasoning they embody to make sense of these resources, are not only not described in the manuals, but could not be; they are, in short, tacit features of controlling activities.

Whereas much of controlling is routine, problems can also arise which are regularly treated as 'routine problems'. The mutual determinacy of strips, radar and radio is reflected in the ways in which discrepancies in the information are dealt with. Strips, as mentioned earlier, state when a plane is due to arrive in a sector. Failure of an aircraft to appear on the radar screen at the time stated or failure of the pilot to contact the controller at the appointed time, is not normally treated as a reason for the controller to think that the strip is spurious, that the aircraft never existed or that there is a technical malfunction in the equipment. Routinely a controller in a situation such as this will assume that the radio and radar are functioning normally and that the strip should be set aside until further information clarifies its status. Controllers know, in other words, that there are many 'good reasons' why an aircraft does not arrive at the time indicated on the strip. Mutually checking and rechecking available information is a function of the 'working division of labour' and the 'working knowledge' of how information is relevant to work activities.

The mutual determinacy of strips, radio and radar means that each is 'made sense of ' by what the others indicate; a process of sense making that, for most of the time, is 'at a glance', and reflects what Garfinkel identified as the salient feature of common sense practical reasoning, namely, its documentary character (Garfinkel, 1967). It is this which makes 'what is happening' transparent. That is, that on the basis of the information the resources provide and through the experience of aircraft doing-these-sorts-of-things-in-these-typical-ways, a controller can recognise or know where a plane will be in the future - in controlling terms a matter of a few minutes away - and hence, where it will be vis a vis all other aircraft in the sector. It is this ability to make transparent what is happening on a sector which is the quiddity of the controlling work. Team members make practical assumptions about what each datum represents or indicates and on the basis of, among other things, the predictable ways aircraft fly and in so doing, can judge whether the absence of one or other of these resources represents a serious problem or not.

The functionality of the strips resides in their status as one of the embodiments of the working division of labour around the controlling suite and as its work site. As we have said, the printed strips quickly amass a great deal of written information, information that constitutes an evolving history and unfolding plan of the controller's intentions and decisions, recorded as they are made. The negotatiation of the information on the strips by members of the team 'trustably' incorporates them into work activities. Controllers are only too aware of the fact that mistakes are made, errors that can be alarming to the inexperienced outsider. But the trustability of the strips does not lie in any technical failsafe but in their accessibility to members of the team and the mutual checking that goes on. A routine part of the work is checking information and this begins as soon as the strip is printed. It involves talk about routes and destinations, questions to each other about procedure and the coordination between sectors, confessions about the need to do it this way rather than that, and so on. Thus, although coordinating rules and procedures are laid down in the *Manuals*, what is striking are the richly varied ways in which they are applied in the activities of controlling.

The picture as a sequence of working tasks

Newcomers to an ATC console cannot see there what the experienced controller can, though they can, very quickly, be given sufficient guidances to how to see on the screen what the experienced ATCO is seeing. They can see this in at least the sense in which they are shown where the outlines of the coast are[16], the flight lanes, the major airports, the blips which represent the aircraft, how the data blocks identify the flight number and its height. Similarly with the strips: their left to right (navigation marker separation) and up to down (temporal) arrangement can be shown, what the letters and the figures in the cells refer to, what the coloured marks mean, and so on. However, none of these things can be apprehended in the fluent way in which the controller does. The ATCO of course, does not just see a series of identifiable units and patterns of movement, but also the history and the likely future of these units and patterns and furthermore, can see that these are not only a coherent and coordinated series of movements, but how they came to be coherent and coordinated.

It is this quality which we suggest is sought by the notion of the 'picture'. In vernacular terms, the picture refers among other things, to the controllers capacity to 'keep it all together'; to see and give coherence and organisation to the patterns of aircraft movements under varying conditions. As we discussed earlier, to date the marked tendency has been to treat the picture as an 'internal representation', the controller taking time to sit observing his or her predecessors work, observe what is going on and, through this, build up a sense of 'where things are' and both generally and specifically, what is happening in the sector. At a certain point the incoming controller takes over, and presumably, begins to work in terms of the synthetic appreciation of the information available putatively matching his or her internal picture with the developing information from screen, strips and radio.

However our case here is directed at viewing the 'picture' from a different angle. Our interest in the ATCO is not so much as an information processing device, so to speak, but as a worker and hence, an interest not so much in the ATCO as 'cognitive information processing machine' but in the ATCO as someone working out the organisation of a set of tasks. To repeat an earlier comparison: ATCO work is not like an assembly line in which a recurrent sequence of steps has to be followed through, but one in which the work consists of putting the tasks to be done into a sequence of steps that can be followed through. Any current transaction between an ATCO and a pilot, or any other member of the team on and around the console, is not an isolated one but part of a developing sequence whereby current steps are built upon previous ones and shaping subsequent ones. The activities, that is, are treated as 'entrained' sequences. (McGrath, 1990) and their meaning derived from their location in stream of action. From this view, rather than thinking of the air traffic control situation as one which involves a pictorial representation of the state of affairs up in the sky, it is better to think of it as *a display of a set of task requirements*. This is how the ATCO is looking at the information presented in the strips, the radar and the R/T; to see what needs doing 'now', 'in a moment', 'sometime later on', and so on. Accordingly the working picture is a term for the 'stuff' which is already dealt with, or which will take care of itself, 'stuff' which now or soon requires ATCO intervention, things that are coming along and will or not need dealing with and so on.

Conclusion: The Sociality of Work and Technology

In the context of ATC, the technology embedded in the suite becomes not so much a material object the performance of which is governed precisely by the technical specifications which its design incorporates, but an intergral part of the organisation of things which a controller needs to do in his or her work. This includes such things as 'serving up the information when needed', 'giving information that is not needed just at the moment' 'providing facilities that are useless', 'telling me what needs to be done', all of which we have illustrated with strip use; meeting, in short, to a greater or lesser extent the practical needs of the ATCOs as experienced in doing the work of controlling. In this respect, conceptions of problems, ways of dealing with them, of likely exigencies, of efficiencies, and so on, are part of the technology as experienced by those doing the work. Controllers respond to technology and its operation in sensible and rational ways predicated on the motivations induced by the need to know 'what the technology can do as part of my work'. In this sense, the workings of the system as known yields determinate causes for occurrences.[17] A key element of this is working with the information that is available, reading and evaluating it. This is the kind of activity that anyone familiar with the work will be able to know without saying or describing. Working with the strips, with R/T, the radar, is the routine stuff of controlling and the 'bottom line' feature of the ATCO's task is the continuing consultation of these resources.

Of course there is more to controlling that we have reviewed here. Our focus on how the strips are interwoven with the work activities around the suite has largely precluded other than passing reference to other important aspects of controlling work done by 'wingmen' and chiefs. Nevertheless what we hope to have shown is the situated character of the work in which, routinely and mundanely, the meaning of the technology and the information it provides, the rules of controlling, the activities of others around the suite are all interwoven and organised as a system-in-use.

What we have been trying to illustrate here is a way of approaching the study of the sociality of technology which abandons the classical, one could say 'natural', distinction between the technological system and the user. Describing the system-as-seen-from-within means declining to accept the analytic omnirelevance of the strict separation of the technical system and the user. From an ethnomethodological point of view the constructs around which the system is organised are treated as resources for the construction of the working system as an on-going achievement. Far from the technology and the user being set against each other they are mutli-layered and interwoven into a motile configuration. (Anderson et al, 1990; Shapiro et al, 1991) The system-in-use is treated as a changing constellation of objects, activities, procedures and actors within which the user is immersed in and through the specifics of the work. Controlling actions are not then to be looked at simply as the following of procedurally defined rules but as the contingent outcome of processes of interpretation as to how the rules fit the case to hand. Thus, solving the sequencing problem is a mattter of moment-to-moment using the information, the technology, the rules to achieve, in and through the the specifics of the flow of air traffic, an orderliness to that flow. The user, in this case the controller, is a manager of the technology and involved in getting specific things done through that technology using the procedures 'here and now'.

But, and this is an important point, the focus is not the individual but the individual-in-a-team. As we have pointed out and elaborated, the controller at the screen is part of a division of labour around the suite and thus part of an embodied collection of courses of action achieved by the working team. It is the team which circulates knowledge, which reproduces the production processes, checks 'how things are going', doing their respective jobs, and so on. It is in the work of the team that the invisible but vital skills, and other resources furnished by the system-in-use, are 'made to hand' for controlling. It is in the practical, accountable, visibly rational working of the team in the local circumstances wherein lies the sociality of work and its technology.

By way of final comment: as we said at the beginning, there is an increasing movement in computer system design that argues for a greater recognition of the sociality of work and technology. Ethnomethodological studies, too, are being examined with this end in mind: a move which occasions no little controvery within ethnomethodology itself.[18] But whatever the outcome of this interest, there is no doubt that ethnomethodological analyses of work do, on the face of it, make the task of the designer much harder - the 'grit in the designer's vision' - not least because the user ceases to be some 'static entity' specifiable in terms of a list of cognitive capacities. Examining work, and technology, in terms of the situatedly specific activities as seen 'from within' raises the awesome spectre of a variety that is inimical to a design process which is increasingly becoming industrialised. Moreover, ethnomethodological analyses emphasise the skilful knowledge and resources which saturate the skill of the work but which are largely 'invisible', because routinely familiar, even to the parties to the work. It is these relied upon skills which, in this case controllers but the point is generalisable, workers find hard to articulate. Nevertheless, in being vital to the work they cannot be substituted for by abstracted, decontextualised versions of knowledge and skill. But this is to simply restate the challenge for designers in trying to meet what has always been their ambition if not always their achievement, namely, a better 'goodness of fit' between the technology and the working environment in which it is located. Designers are increasingly seeking to make their technologies more flexible and adaptive. What the approach here at least offers is one way of throwing light on what kinds of flexibility are needed, what kind of technological support is appropriate, in what kinds of work; in a word reappraising the distinction between the system and the user. Needless to say, whether such efforts will lead to better design has yet to be shown.

References

Anderson, R.J., Hughes, J.A., and Sharrock, W.W. (1989), *Working for Profit: The Social Organisation of Calculability in Entrepreneurial Decision Making,* Averbury, Aldershot.

Anderson, R.J., Hughes, J.A., Shapiro, D.Z., Sharrock, W.W., Harper, R., Gibbens, S. (1990), 'Flying Planes Can be Dangerous: Work Skills and Traffic Management in Air Traffic Control', Unpublished manuscript, Rank Xerox EuroPARC.

Baccus, M. D. (1986), 'Multipiece truck wheel accidents and their regulation' in Garfinkel, (ed) *op cit*, pp. 20-56.

Bainbridge, L. (1974), 'Analysis of Verbal Protocols from a Process Control Task', in Edwards, E. and Lee, F. P. (eds) *The Human Operator in Process Control,* Taylor and Francis, London.

bainbridge, 1974

Bannon, L. and Schmidt, K. (1991), 'CSCW: Four Characters in search of a context' in Bowers and Benford, (eds) *op cit*, pp 3-17.

Bowers, J.M. and Benford, S. D. (1991), *Studies in Computer Supported Cooperative Work: Theory, Practice and Design.* Amsterdam: Elsevier, North- Holland.

Douglas, J. (1971) (ed), *Understanding Everyday life.* Routledge and Kegan Paul, London.

Gallagher, J., Kraut, R.E., and Egido, C. (1990) (eds) *Intellectual Teamwork: Social and Technological Foundations of Cooperative Work.* New Jersey, Lawrence Erlbaum Associates.

Garfinkel, H. (1967). *Studies in Ethnomethodology*. Prentice Hall, New York.

Garfinkel, H. (1986) (ed), *Ethnomethodological Studies of Work.* Routledge, London.

Garfinkel, H. and Sacks, H. (1970), 'On formal structures of practical actions'. in McKinney, J.C. and Tiryakian, E.A. (eds) *Theoretical Sociology,* Appleton Century Crofts, New York, pp. 338-366.

Garfinkel, H., Lynch. M. and Livingstone, E. (1981) 'The work of a discovering science construed with materials from the optically disovered pulsar', *Philosophy of Social Science,* 11, pp. 131-158.

Goffman, E. (1959), *The Presentation of Self in Everyday Life,* Penguin Books, Harmondsworth.

Greif, I., (1988) (ed) *Computer-Supported Cooperative Work: A Book of Readings*, Morgan Kaufman, San Mateo, California, 1988.

Harper, R.R. (1988), 'Not any old Numbers: an Examination of Practical Reasoning on an Accoutancy Environment', *The Journal of Interdisciplinary Economics*, Vol 2, pp. 297-306.

Harper, R.R., Hughes, J.A. and Shapiro, D.Z. (1989a) "Working in Harmony: An Examination of Computer Technology in Air Traffic Control" in Bowers, J.M. and Benford, S. D. (1991), *Studies in Computer Supported Cooperative Work: Theory, Practice and Design.* Amsterdam: Elsevier, North- Holland.

Harper, R.R., Hughes, J.A. and Shapiro, D.Z. l989b) *The Functionalities of Flight Data Strips.* CAA, Kingsway, London, February.

Harper, R.R., Hughes, J.A., Randall, D, Shapiro, D., and Sharrock, W.W. (forthcoming), *Order in the Skies*, Routledge and Kegan Paul, London.

Hopkin, V.D. (1975), 'The Controller versus Automation', in Benoit, A. (ed), *A Survey of Modern Air Traffic Control,* AGARDograph, No. 209, AGARD.

Hughes, J.A.., Shapiro,D.Z., Sharrock, W.W., Anderson, R.R., Harper, R.R., and Gibbons, S. (1988) *The Automation of Air Traffic Control,* SERC/ESRC Grant No. GR/D/86157

Jackson, A. and Onslow, G.J. (undated) 'The Replay Technique: The Concept, Initial Experience and Proposed Developments', *RSRE Memorandum*, no 3827.

Livingstone, M. (1986) The Ethnomethodological Foundations of Mathematics, Routledge and Kegan Paul, London.

Luff, P., Gilbert, N., Frohlich, D. (1990), *Computers and Conversation,* Academic Press, London.

Lynch, M. (1985) *Art and Artifact in Laboratory Science.* London, Routledge & Kegan Paul, London.

Lynch, M.E. (1982) 'Technical work and critical inquiry: investigations in a scientific laboratory', *Social Studies of Science,* 12, PP499-534.

McGrath, J. (1990), 'Time matters in groups', in Galagher, Kraut, and Egido, *op.cit.* pp. 23-62.

Parsons, T. (1958), *The Social System,* Tavistock Publications, London.

Rasmussen, J. and Jensen, A. (1974) 'Mental Procedures in Real-life Tasks: A Case Study of Electronic Trouble Shooting', *Ergonomics,* 17, pp293-308.

Shapiro, D.Z., Hughes, J.A., Randall, D. and Harper, R. (1991) 'Visual Re-representation of Database Information: The Flight Data Strip in Air Traffic Control', *Proceedings, 10th Interdisciplinary Workshop on 'Informatics and Psychology': Cognitive Aspects of Visual Language and Visual Interfaces,* Scharding, Austria, May.

Spiegel-Rosing, I. and Price, D. de Solla, (eds), *Science, Technology and Society,* London, Sage.

Suchman, L. (1987) *Plans and Situated Actions: The problem of human-machine communication,* Cambridge University Press.

Suchman, L. (1989) 'Rediscovering Cooperative Work'. Keynote Address, *Proceedings of EC-CSCW Conference,* Gatwick, England, September.

Weber, M. (1947), Theory of Social and Economic Organisation, The Free Press, New York.

Whitfield, D. (1979) 'A preliminary study of the Air Traffic Controllers Picture', *Journal of the Canadian Air Traffic Controller's Association,* ii (i), pp19-28.

Whitfield, D. (1980) *Discussion paper: The Air Traffic Controller and Sector Capacity,* Ergonomics Unit, University of Aston, Birmingham.

Whitfield, D. and Jackson, A. (1982) 'The Air Traffic Controller's 'Picture' as an example of a Mental', *Proceedings of the IFAC-IFIP-IFORS-IEA Conference on Analysis, Design and Evalutation of man-machine Systems,* Baden-Baden.

Whitfield, D., Ball, R.G. and Ord, G. (1980) 'Some Human Factors Aspects of Computer-Aiding Concepts for Air Traffic Controllers', *Human Factors,* 22, pp. 569-580.

Wieder, D.L. (1974) *Language and Social Reality: The Case of Telling the Convict Code,* Mouton,The Hague.

Wilson, P. (1990) *Introducing Computer Supported Cooperative Work*, CCTA Publications.

Wilson, T.P.(1974), 'Normative and Interprative Paradigms in Sociology', in Douglas (ed), *op.cit.*

Wittgenstein, L. (1952), *Philosophical Investigations,* Blackwells, London.

Zimmerman,D. (1971) 'The Practicalities of Rule Use' in Douglas, (ed) *op cit,* pp. 221-238.

1 It is perhaps to history that we owe the renewed interest social science is beginning to show in technology, its innovation and development as a social process. (Hughes, 1979; Laudan, 1984) In a number of case studies on, for brief examples, radio, railways, steam engines, chemical plants, among many others, it has been shown that the forms, innovations, the paths of technology are governed through and through by social, economic and political factors. (See, for example, Bijker et al, (eds), 1987; Mackenzie and Wajcman (eds), 1985) Important in this was the broad interdisciplinary grouping known as the Study of Science, Technology and Society. (See contributions in Speigel-Rosing and de Solla Price (eds), 1977).

2 See, for example, Hughes (1979)

3 The first stage of this research, completed in 1989, funded jointly by the ESRC/SERC Joint Initiative, involved ethnographic observation and interviews at London Air Traffic Control Centre, or LATCC. The second phase, currently underway and funded by the HCI Initiative, involves additional ethnographic research directed toward building a prototype design tool for computerising flight progress strip interfaces. Fellow researchers on these projects included D. Shapiro, D. Randall and S. Gibbons, all of Lancaster University; W.W. Sharrock of Manchester University, and R. Anderson of Rank Xerox, Cambridge Europarc. The computer scientists at Lancaster are I. Sommerville, T. Rodden and R. Bently.

4 There are typically eight suites, each dealing with 2 sectors, in the operations room at LATCC. However, under certain traffic conditions sectors can either be merged or split further.

5 This time is not mandatory but depends on how busy the sector is. Occasionally on very busy sectors the time can be much longer. Normally, controllers work two hours on the 'tube' followed by two hours rest.

6 What is going on 'now' is also a matter of 'scrolling forward' to see where things might be some minutes ahead.

7 A US National Transportation Board report on an accident in February, 1984, when a DC10 ran off the end of the Kennedy Airport runway, cited the crew's "habitual reliance on the proper functioning of the airplane's automatic system" as a probably cause. *The Independent on Sunday,* 23rd June, 1991, 'Who is Really Flying the Plane?' Trusting automated assistance very often means that operators, in the case of failure, have to make inquiries to determine what needs doing; a 'warming up' which can be fatal.

8 One of the problems of studying social action is the relationship between the ways actions are described by those doing them and the actions themselves. As Wittgenstein noted, it is all too easy to be confused by the use of the same word in different contexts. Here we suggest may be a case in point: the word picture is used by controllers to mean a certain kind of thing specific to their work; the same word is used in other contexts to mean things like a graphical representation. One needs to be careful, therefore, not to get muddled up or conceptually confused when one uses the word in analysis, or treats the word as descriptive of an analytic object. See Wittgenstein, (1952).

9 We are not, of course, denying that cognitive operations are involved.

10 There are variations to this depending on the density of the traffic. It the sector is subdivided because traffic loads are heavy this complement will be increased.

11 Of course, not everything the ATCO does is intelligible in this way. There are occasions when the activities of a controller are puzzling even to the most acculturated of observers. At such times interrogation might be necessary. But these are exceptions rather than routine.

12 The strips are in fact collected and used to calculate the cost of the ATC services. The agencies responsible for each aircrtaft are charged accordingly.

13 They can also be used *in lieu* of live strips when, for whatever reason, the latter are not output in time. They can also give advance warning of any input errors, such as

incorrect flight plans and, in the event of computer failure, serve as a basis for manual calculation.

14 Controllers try to minimise the amount of communication between sectors, since it can take up valuable time, by using 'silent handovers', that is, using procedurally agreed flight levels for the transition between sectors.

15 Some controllers put the latest at the bottom. It depends entirely what is preferred by the individual controller. Nonetheless, assistants, who provide the new strips, need to be aware of this preference so that they can orient their activities to ensure that they sustain the preferred order when placing new strips in the racks.

16 Though as it happens most controllers elect for the system not to display such data as it complicates the image presented on the radar screen.

17 Of course, and especially in complex discretionary systems such as ATC, what are diagnosed as 'routine' faults may, in the end, turn out to be so. This, if one likes, is the price paid for systems of such flexibility.

18 One of the more arguable efforts in this regard is the use of Conversation Analysis as informative for HCI design. See Luff, Gilbert, Frohlich (ed) (1990).