



THE INTERACTIONAL AFFORDANCES OF TECHNOLOGY: AN ETHNOGRAPHY OF HUMAN-COMPUTER INTERACTION IN AN AMBULANCE CONTROL CENTRE

David Martin^{1,2}, John Bowers^{1,3} and David Wastell²
Departments of ¹Psychology and ²Computer Science
University of Manchester, UK

³Centre for User-Oriented IT-Design (CiD)
Department of Numerical Analysis and Computer Science (NADA)
Royal Institute of Technology (KTH)
Stockholm, Sweden

ABSTRACT

This paper reports an ethnography of ambulance dispatch work in a large UK metropolitan region. The interplay between control centre ecology, usage of a computerised dispatch system, and cooperative work of control personnel is analysed. The methods by which a 'working division of labour' is sustained to effectively manage dispatch in the face of high workload and manifold contingency are explicated, and contrasted with methods employed by workers in other control room settings known from the literature. The implications of the study for system improvement and for several emphases in HCI research (including discussions of 'affordances') are explored.

INTRODUCTION

Ethnographic research methods are becoming increasingly influential in research fields such as Computer Supported Cooperative Work (CSCW) and Human Computer Interaction (HCI). For many researchers, an ethnographic study involving prolonged participant-observation at a real-world work setting is an invaluable resource in technology design and development. In the CSCW literature, for example, there exist studies in settings as varied as Air Traffic Control (ATC), banks, management training centres, the print and fashion industries (Hughes, Randall and Shapiro, 1992; Randall, Rouncefield, and Hughes, 1995; Rouncefield, Hughes, Rodden and Viller, 1994; Bowers, Button and Sharrock, 1995; Pycock and Bowers, 1996)—studies which have been frequently coupled to system design projects.

The ethnographic work reported in this paper was conducted in the control room of the ambulance service of a large metropolitan region in the UK. It forms part of a long-term collaboration between university-based researchers and the ambulance service but it is only recently that concerted ethnographic research has been added to the mix of methods in this collaboration. Control rooms studies, however, are quite familiar territory for ethnographers. Previous work, for example, on London Underground control (Heath and Luff, 1992) and ATC (Hughes, Randall and Shapiro, 1992) point to a number of issues that seem to be important to cooperative work in control room settings: mutual monitoring and awareness, the affordances of artefacts (some of which

are at first sight quite mundane) which support 'at-a-glance-perception' within an appropriately organised workplace ecology, the maintenance of flexible working divisions of labour which enable 'running repairs' and so forth.

However, those are studies of air traffic and London Underground controlling. We do not know from *them* the inflections in control work in our new setting—ambulance control—which to our knowledge has been little studied in the same ethnographic depth as those others¹. This seems important to redress for a number of reasons. The ambulance control room we have studied differs from other settings in the CSCW/HCI literature in that, *first*, control work is driven by calls from the public, *secondly*, the control centre we have studied is quite advanced in its computerisation, and relatedly and *thirdly*, there have been some highly visible failures to automate ambulance dispatch (e.g. the London Ambulance Service controversy of 1991) though the setting we have studied mostly works well. Thus, we intend to learn from it to see what makes for workable computer systems in such settings as well as to reciprocally contribute to their improvement.

THE FIELDSITE: A UK AMBULANCE CONTROL CENTRE

The ambulance control room in this study serves the whole of a large region of around 2.5 million people and is situated centrally in the region's main city. The site also serves as one of the 35 ambulance stations for the region. The control room communicates with all ambulances (30 to 70 operational at any one time) and the stations from which they operate. Essentially, the centre takes calls from the public (999s), police, fire service, doctors and hospitals and provides ambulances for all incidents arising from these calls according to certain Government-stipulated rules, regulations and targets. The targets specify that for emergency calls 95% of ambulances should be mobile 3 minutes after receiving a call, 50% at the scene by 8 minutes and 95% by 14 minutes. Currently, the region consistently outperforms these figures. The major task for control room workers is to meet the need for ambulances, according to target, while operating with limited resources. The control room is staffed by between one and six Call Operators depending on the time of day, four Dispatchers, two Supervisors and a Control Manager. The present paper will focus on the Dispatchers and Supervisors who organise the distribution of ambulances over the region and responses to incidents. Further work is planned to examine call operation in greater detail.

Central to the job of ambulance control is a computer system which supports control room staff and links to the ambulances themselves via radio. Telephones, a global positioning system (GPS) and other artefacts are also employed intermittently. The computer system principally comprises a relational database containing details of calls, incidents, vehicle availability and so forth, which a number of client applications make requests to. The most important of these applications for ambulance dispatch are discussed shortly. The database server is networked to over 20 terminals in the control room. The system is mainly DOS-based and has been developed by an independent software company to the regional service's own specifications.

Our most intensive contact with the centre took place in the period from April to October 1996, during which about 40 person-days were spent on-site. We ensured that the visits covered both 'quiet' and 'busy' times and sampled the daily and weekly rhythms of the work. The centre's IT Manager was an initial point of contact and informant who provided us with introductions to all centre staff, whereupon we were able to intensively observe control room work as it happened. Much of this study involved 'sitting in' with different staff members, observing their work and occasionally getting them to describe what they were doing. Informal interviews, where we were concerned to clarify our emerging understanding of control room activity, were undertaken opportunistically on the job or during breaktimes. Field notes, ambulance service documents, forms, training materials, system manuals and other related materials were collected. Our methods of ethnographic analysis broadly follow the programme of ethnomethodological ethnography as exemplified in CSCW and HCI by Hughes et al. (e.g. 1992) and discussed by Bowers (1996). Amongst other matters, this involves attending to the real-time details of work and how, in and through them, the orderliness of control room activity is produced as a recognisable social accomplishment.

¹Whalen's (e.g. 1995) work is a partial exception to this though he does not focus, as we do, on ambulance controlling as a species of cooperative work.

Incoming Incidents

Ambulance control work commences on receipt of a call reporting an incident. Incoming incidents are logged by the Call Operators onto computer based forms containing fields for, for example 'name', 'address', 'incident type' and so forth. The minimum information the Call Operator requires from the caller to initiate response is an address and incident type. This is because, on making a 'fix' on the address, the computer system automatically assigns an A-Z (street atlas) grid reference and a zone number to that location, which are displayed on the Call Operator's screen. Once these details have been entered the form automatically passes over to a incident pending list ('Incident Stack') displayed for the Dispatchers and Supervisors, while the Call Operator attempts to gain the information to complete the rest of the form. This feature, whereby Dispatchers can start formulating their dispatch decisions while the call is still being taken, is considered by the service to be a time saving advantage of an electronically based system.




Figure 1: Plan view of the ambulance control room (most of the technologies noted in the figure are discussed in the main text).

The control room also deals with another type of call, 'urgent incidents', which are received from doctors, and cover cases where a patient requires transportation to hospital from either home or surgery. These cases are logged on similar forms, however they can be completed between one and four hours after receipt rather than immediately, and even this deadline can be extended occasionally. Urgent incidents can create problems in the

management of dispatch and cover. For instance, a number of urgent cases may have been allocated due to the service being quiet, when suddenly there is a rush of emergency incidents.

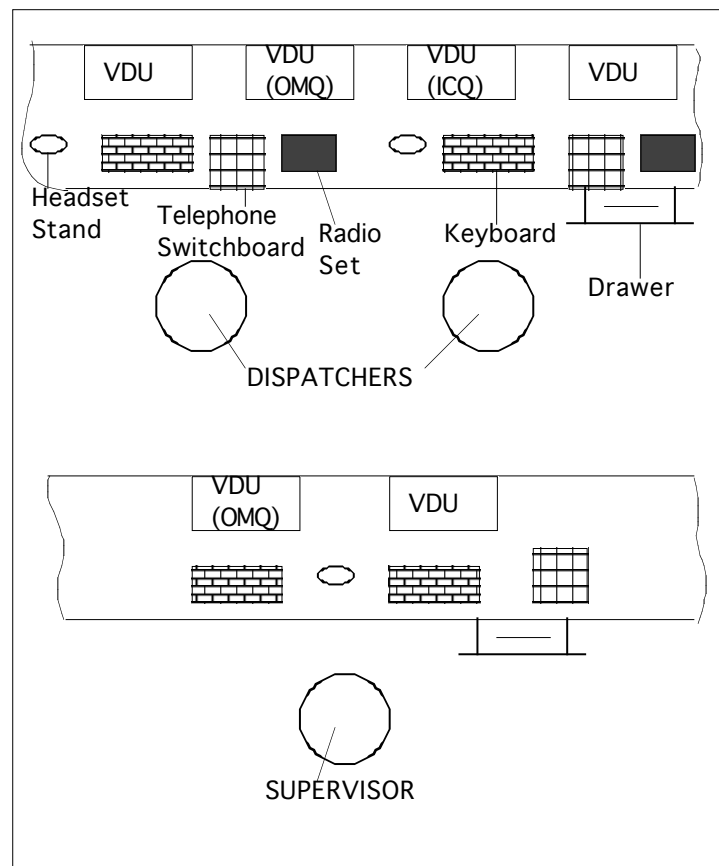


Figure 2: Plan of Dispatchers' and Supervisors' local environment within the control room. (In addition to the technologies discussed in the main text, the ICQ (Incoming Message Queue) and the OCQ (Outgoing Message Queue) screens display the flow of packets of information between the control room and ambulances.)

Dispatchers and Supervisors

Although there may be occasional discussions with Call Operators (e.g. if the details of a form need clarification), for the most part the Dispatchers and Supervisors operate as a self contained group located at the other end of the control room (see Figure 1). The Dispatchers are primarily concerned with assigning incidents to ambulances and sending the details, as an abridged message, through to the ambulances where they are presented on a small panel containing an alphanumeric display and keypad.

Initially, it might seem to be a reasonable strategy for a Dispatcher to send the nearest available ambulance to the incident as soon as a geographical fix has been obtained. Often, this is what a Dispatcher will do. But this cannot be done without consideration of a decision's implications for the provision of *cover*. If the dispatch of a certain ambulance would lead to an area of the region being inaccessible within target times, then an ambulance other than the nearest might be dispatched. Indeed, with limited resources, it is *essential* for Dispatchers to be mindful of the implications for cover of their responses to incidents. In the face of these demands, and with a workload which can rise to several hundred calls a day, how is the complexity of dispatch and cover to be managed?

First, the region is separated into four main areas (or 'boards'—a term which reminds one of earlier non-computer-based systems). Each Dispatcher is assigned the duties of dispatch and cover management for a single board. However, while their board forms their focus, interaction and collaboration across boards is a required and indeed common feature of the work. The Supervisors' job is to ensure the management of dispatch and the adequacy of cover for the whole region. This entails overseeing activities with the Supervisor providing advice and making checks on the Dispatchers' work. Furthermore, it is common for a Supervisor to actually dispatch ambulances in order to lessen the workload for Dispatchers, or deal with emerging problems. In this way, the Supervisors and the Dispatchers maintain a 'working division of labour'. That is, each person has their own job to do but they routinely coordinate with each other and, when necessary, help each other out. We shall return to the question of exactly *how* they maintain their working division of labour later.

Computing Technology

Although Dispatchers and Supervisors deploy a number of non-computer-based artefacts (e.g. street atlases, manuals and personal notepads), the centre's computer system comprises the central technological resource for dispatch. In this section, we describe most important applications within this system. In the typical character of DOS-based systems, moving from one application to another is accomplished by a key sequence whereupon the entire screen changes to show the new application.

The Incident Stack

When an incident is passed over to a Dispatcher it appears at the top of the 'Incident Stack' screen, pushing earlier incidents down one row. Each Dispatcher has an Incident Stack for their own region displayed on the terminal in front of them. An incident's description will be routed to the terminal associated with the Dispatcher who has working responsibility for the board containing the incident's geographical fix. The screen is split into two with the upper half listing all the incidents that are 'WAITING' to be assigned to ambulances and the lower half those that are 'ACTIVE', that is, being attended to. Top of the list in the 'WAITING' half is the oldest unassigned emergency. A variety of information is displayed in each line including when the call was received, the number of ambulances attending and their IDs, and various items provided by the GPS including the 'as-the-crow-flies' distance the ambulance is from the incident. 'WAITING' incidents show similar details but, being unassigned, have no ambulance information. While each individual Dispatcher has an Incident Stack for their board, the stack for the whole of the region is permanently displayed on a large screen on the left hand side of the wall in front of the Dispatchers (see Figure 1).

The Dispatch Selection Screen

When an incident appears on a Dispatcher's stack and they select it by pressing the enter key, the top half of the incident form appears on screen along with a list of the nearest ambulances to the incident. This is defined in terms of the nearest stations to the zone the incident has occurred in. The list presents ambulances irrespective of whether they are active or free. However, the nearest free ambulance is flagged in blue as a guide to allocation. When the selection of ambulance is made the rest of the data which has been entered by the Call Operator appears on screen. The Dispatcher then checks and edits this before transmitting it to the crew, typically without any accompanying voice-radio contact. On receiving the details of incident, the crew must respond by pressing a designated button on the keypad in the ambulance cab. This updates the listing for the incident on the Incident Stack and the listing of the ambulance on the VAM screen.

The Vehicle Availability Map (VAM)

To ensure that cover is adequate or to manage the risk of leaving a station 'empty', Dispatchers must monitor the status of ambulances on the VAM. This is displayed on the wall in front of them on the large monitor shown at the top-centre of Figure 1 and in Figure 3, and can also be accessed from the Dispatchers' individual terminals. The VAM consists of a set of lists of ambulance IDs, each list depicting ambulances from 2 or 3 proximate stations. The lists are arranged on screen to give an approximate representation of the geographical relations between stations (e.g. stations in the same 'board' will be listed close to each other). Colour is used to signify the statuses which are most pertinent to dispatch decisions. If flagged in red the ambulance is active on an emergency, if in green active on an urgent call, and if unflagged the ambulance is available. If the number flashes this indicates that the ambulance is on standby, placed at a designated location between two or more stations to provide emergency cover for not only its home station but another nearby that is low or without cover. Thus, the Dispatchers and Supervisors can see at a glance what the general configuration for the region is

and identify potential problems of cover. (The detail of the VAM is shown in Figure 3 with the colour highlighting given a greyscale approximation.)

| WIG | ATH | BOL | HIG | BUR | ROC | ECC | ALT |
|--------|--------|--------|--------|--------|------|------|------|
| 7345 | 5578 | 5661 | 3346 | 4110 | 2120 | 9786 | 2336 |
| -7563- | 4242 | 5563 | 3217 | 4243 | 2115 | 9866 | 3114 |
| 7224 | 4455 | 6245 | 2354 | 3447 | 2332 | 9786 | 3122 |
| 6224 | -4411- | -6562- | 3737 | -3388- | 2112 | | 2255 |
| 7797 | 4116 | 4113 | | | 2889 | | |
| 8761 | | | | | | | |
| 7556 | | | | | | | |
| BEL | CEN | SAL | OLD | DUK | STO | GLO | |
| 4466 | 5778 | 6536 | 6684 | 8843 | 4426 | 3557 | |
| 4117 | 5331 | 6868 | -7553- | 8854 | 4473 | 3749 | |
| 4119 | 5462 | 6565 | 7776 | 9982 | 4335 | 2561 | |
| 4421 | 5567 | 6769 | 7787 | 8822 | 4333 | 2239 | |
| 4434 | 5579 | 7652 | -8114- | | | | |
| -5337- | 5882 | | | | | | |
| | 6577 | | | | | | |

Figure 3: The Vehicle Availability Map (VAM) which lists ambulances by regions. Ambulances active on emergencies are here shown with their IDs against a black background. Ambulances on urgent calls are shown against a grey background. Ambulances on standby are shown 'flashing'. Available ambulances are just depicted by their ID. See main text for further explanation.

Contingencies

Although often the ambulance suggested on the Dispatch Selection screen (the nearest available one) is chosen, there are multiple varying contingencies that must be considered in the on-going flow of work. Each dispatch decision is in the context of various previous decisions and in turn will influence others. In addition to proximity and availability and a consideration of implications for cover, the kinds of contingencies which must be reckoned with from time to time include:

- Are the crew due a meal-break?
- When does the crew's shift end? When will a new crew's shift begin?
- Have the crew just dealt with one or more harrowing incidents?
- Does the ambulance have the right equipment for the incident?
- Is the nearest (as the crow flies) ambulance on the fastest route?
- Are there road works, traffic problems etc. on a particular route?
- Which side of the motorway is a particular accident on?
- How serious is the incident?
- Are there urgent cases which must be allocated soon?
- And so forth.

Furthermore, ensuring cover is no simple algorithmic matter, as the demands on the service fluctuate dramatically at different times of year and on different days of the week. Public occasions may require special deployment and an emergency call-out at the region's major international airport (which can be actioned by, amongst others, an incoming pilot in difficulties) immediately commits 5 ambulances through a standing agreement with the airport authorities.

It should be clear that it would be negligent to simply select the system's suggestion for dispatch without attending to these many and varied matters. Accordingly, Dispatchers and Supervisors often need to examine whether, for instance, there is the possibility of allocating another ambulance in the vicinity to the job, or whether standby cover can be arranged, or whether another ambulance from the station will become available soon.

Workload and Contingency Management

Ironically, it is precisely when workload is at its highest, and there is the least time for Dispatchers to reason about contingencies, that they most urgently impact upon the work of ambulance dispatch. Consider, for example, a crew which has just attended a particularly harrowing accident. If there are few other incidents outstanding, there will usually be a suitable alternative to this crew for attending any new incident. But difficult problems occur when, for instance, any respite due to a crew has to be weighed up against the disquiet of several others who are overdue on their meal-breaks and there are still new incidents coming in! Plans may be made (e.g. greater cover on a Friday night) but these must be open to constant revision as the situation unfolds. Although proximity, availability and cover are always relevant, there is not a simple hierarchy in operation. Contingencies compete against one another and become more or less important according to the situation. For example, sending the closest available ambulance may leave no cover and mean that the crew must work longer than their allotted shift. However, if the incident is a serious road traffic accident (RTA), and there appears no other feasible option, this will be the case.

The management of contingencies, especially at times of high workload, is facilitated by occasional cooperation between Dispatchers and between Dispatchers and Supervisors. Not only might Supervisors assist with dispatch itself, they will provide advice over vehicle allocation based on their overview of the region, an overview which is hard to obtain for individual Dispatchers especially if their workload is high. Additionally, a Dispatcher whose 'board' requires less attention may assist a Dispatcher who is having to make many dispatch decisions in short order. While knowledge and experience is routinely 'pooled' and exchanged both on the job and during coffee breaks, this becomes most perspicuous at times of high workload when the control room is a flurry of mutual assistance, cross-checking and shouted suggestions.

Dispatch in Action

In order for the group of two Supervisors and four Dispatchers to effectively cooperate and achieve the successful management of dispatch and cover, they have to maintain a *mutual awareness* of one another's work. To do this, the Dispatchers and Supervisors monitor each other peripherally in a continual manner while a certain person who is in difficulty or requires advice may be focused on more directly if the need arises. Let us examine in more depth how this mutual monitoring is achieved.

The Ecology of Dispatch and Supervision

Crucial to achieving mutual awareness between Dispatchers and Supervisors is the interplay between the ecology and technology of the control room and cooperative activity within it. Refer again to Figures 1 and 2. How does the physical layout of the room, the positioning and design of artefacts within it and the placement of workers effect the way workers act and interact?

The Dispatchers and Supervisors share the same computer based artefact allowing them, in principle, to view through their own terminals one another's work as it occurs. However, this facility tends to be only used by the Supervisors as most of the Dispatchers' time is taken up deploying ambulances. The large regional Incident Stack and VAM monitors, on the other hand, are routinely inspected by everyone in the control room, not just Supervisors and Dispatchers, and hence constitute the major shared artefacts in the work in that all personnel are always presented with the same data. Dispatchers visually monitor each other's work mainly through these artefacts and also through directly viewing the screens and actions of those beside them. Supervisors are afforded more opportunity for visual monitoring, as their seating position allows them a more direct view of all of the Dispatchers' screens and indeed, of any activity that the Dispatchers are involved in.

These physical and technical arrangements afford opportunities for becoming aware of critical factors which can inform a dispatch decision. However, they alone do not *guarantee* that information will be picked up by the appropriate recipient of it or attended to for the appropriate details. Consider, for example, the VAM containing its lists of ambulance IDs colour-flagged to indicate their current status. To the experienced worker, this shows at-a-glance how busy the service is. Swathes of red across the VAM screen (as is typical on a Friday night) reveals a service whose resources are stretched. If this is combined with a similar abundance of red on the WAITING portion of the Incident Stack, then we have a service which, for the moment, is finding it hard to cope. While such summary impressions are available at-a-glance, both the VAM and the Incident Stacks can be engaged with to a further degree to find out exactly which individual ambulances are deployed and exactly which incidents are being attended to. While the displays afford the pick-up of both sorts of information, what

exactly a Dispatcher or a Supervisor will get from the screen depends upon the kind of attentiveness they show to it. Accordingly, we often see co-workers explicitly signalling to each other to point out details or to draw attention to anomalies. That is, in the control room the sharing and use of information must be worked at collaboratively in order for the successful management of dispatch and cover to take place.

In short, the ecology of the control room may afford opportunities for information pick-up but it is only through social interaction between personnel and the appropriate degree of engagement with artefacts that these opportunities are realised in actual situated dispatch decisions.

Mutual Monitoring and Occasioned Cooperation

As we have noted, it is common for control room personnel to monitor each other peripherally for opportunities to correct mistakes, offer advice or work collaboratively. For instance, Supervisors often suggest dispatch choices to Dispatchers for incidents they are not yet dealing with, that is, for incidents which are outside a Dispatcher's current 'focus of work'. The degree of shared understanding within the group is readily discernible from the terseness of talk between personnel, talk which is inextricably immersed in and indexically tied to its situation of utterance. For example on one occasion a Supervisor uttered to a Dispatcher (whose back was turned): "I've got that Moss-way job". The Dispatcher then acknowledged this with a small glance towards the Supervisor. The Supervisor had decided to allocate an incident on behalf of the Dispatcher and a mutual awareness of the action was gained simply through a brief statement which could be readily interpreted as a WAITING incident changes its display to ACTIVE.

Here, the monitoring activity of the Supervisor did not disrupt the Dispatcher's activity, and indeed the Dispatcher may have been unaware that it was even occurring. However, in the control room this is not always the case, and, for instance we have witnessed numerous occasions where personnel rendered their monitoring obvious by asking questions of those they are observing either for checking or clarification, and it is common for Supervisors to get out of their seats and stand directly behind the Dispatchers.

Intervention by those who are monitoring the work of a Dispatcher is not the only way in which cooperation within a working division of labour is achieved. We observed a case where a first Dispatcher gained the attention of a second beside her by giving her a small nudge. The first had been conducting a radio conversation with a crew who were having trouble finding a suitable route to an RTA. A discussion then ensued between the second Dispatcher and a third Dispatcher sitting further away (who voluntarily joined in) as to the location of the ambulance and how it should reach its required destination. An A-Z atlas was used as the focal point of their discussion, with the Supervisor positioned behind the first two Dispatchers also volunteering suggestions. The first Dispatcher then asked the second to take over the call, with the second then providing updated information for the ambulance crew.

This compactly demonstrates several of the points made so far. On this occasion, the first Dispatcher established interaction with the minimal effort of a nudge, however during the course of the incident workers other than the initial addressee negotiated their way into the discussion indicating that they were at least peripherally aware of what had been occurring. Furthermore this instance shows how the division of labour in the control room can proceed with great flexibility with different workers entering into direct collaboration when they have the opportunity, before relatively seamlessly returning to their own work.

Negotiating Interaction Opportunities

There are many ways in which Dispatchers or Supervisors will attempt to solicit assistance. We have numerous examples of Dispatchers or Supervisors looking, or gesturing to attract attention to themselves, or sighing, or muttering while they face the screen, or asking questions either generally or directedly in order to solicit assistance. Often these behaviours are used in combination or one after the other until a satisfactory contact has been made or solution offered. For example, if a Dispatcher asks a question about whether to put a certain crew on their meal-break out loud while looking at their screen and there is no response, they may turn and face the Supervisor behind them and ask the question more directly. The Supervisor may indicate with a wave of their hand that they are too busy, meanwhile the Dispatcher beside her may make a suggestion because she has been listening in. The original Dispatcher may act on this suggestion or try to solicit other opinions. While in a similar situation the Dispatcher may simply hear a "yes" from the Supervisor or the Supervisor may reply that they have independently put the crew on a meal-break.

When one control room worker tries to establish an interaction with another they will often first try and do this with minimal intrusion. However, due at least in part to the number of different personnel there are in the control room, their involvement in other duties and also occasional ambiguity as to whom (if anyone specifically) a question or statement is addressed, satisfactory interaction is often not immediately established. Both the addresser and the addressee(s) regularly need to negotiate their positions through qualifying remarks, questions, looks and gestures, in order to successfully engage. In short, a perspicuous feature of this control room is how *persons' availability for interaction has to be itself negotiated*, sometimes by progressively upgrading from (for example) glance to gesture to direct request for help and on occasion to an insistent plea-ful raising of the voice.

DISCUSSION

Let us bring out some points of general interest grounded in our studies of this control room. In so doing, we will make comparisons with related studies, while also giving an impression of the prototype applications we are developing for consideration by ambulance personnel. Necessarily our treatment here will be brief but we hope to indicate how ethnographic work can not only motivate system design but also offer new perspectives on HCI research agendas. We order our remarks around 4 key observations.

(1) Contingencies are Interactionally Managed

We have noted that ambulance dispatch is a contingent affair with multiple considerations impinging upon dispatch selections. Commonly, Dispatchers can resolve these contingencies through interaction with the different resources available to them, juxtaposing, for example, the VAM's representation of the state of the service with their own Incident Stack and those of neighbours. However, oftentimes contingencies will be managed not merely through the interaction of individual Dispatchers with system resources but through *social* interaction amongst Dispatchers and Supervisors. Indeed, on occasion, the significance of a contingency may be resolved through interaction between control room personnel and others outside—a member of the public still on the line or an ambulance crew. For example, we noted earlier that the timing of meal breaks for crews has to be considered in making dispatch decisions. The *severity* of an overdue meal break—and hence its significance for the work of dispatch—is often argued out between Dispatchers and the crew themselves. Indeed, the use of voice-radio is predominantly occasioned by such discussions!

Observations like this make us sceptical about the feasibility of automatic decision making systems in the ambulance service—systems which might, say, dispatch ambulances on the basis of algorithms operating to satisfy multiple constraints. And this for two reasons. One, the list of contingencies which we have begun to document shows no obvious sign of terminating and, two, the practical weight to be attached to considerations is often socially interactively determined and negotiated. Accordingly, our development agenda is to investigate how systems can be designed for a socio-technical setting where *co-operation* is a normal feature of how work is done².

(2) Divisions of Labour are Working Ones

While there exist distinctions between control personnel in terms of what they do, these are defeasible in the light of the exigencies of work. In the terms of Hughes et al. (1992) and Sharrock and Anderson (1992), personnel simultaneously maintain an *egological* orientation to the division of labour (what is there for *me* to do? what's on *my* board?), and an *alteriological* orientation (what can I do to make the work of *others* easier? how can I help with *their* overdue dispatches?). In Hughes et al.'s studies of ATC, control room assistants are

²This is not to say that we see no future for, say, the use of constraint-satisfaction algorithms and other pieces of formal machinery in supporting ambulance dispatch. While not highest on our own research agenda, we do intend to consider the applicability of such computational techniques. We imagine, though, that the recommendations deriving from such algorithms should be just that: recommendations, which are defeasible by personnel much like the results of the 'find the nearest available ambulance algorithm' currently are.

oriented (alteriologically) to create a "protective cocoon" around the controllers so that the controllers have all that they need to do their work ready-to-hand and do not have to search out for it themselves. In the ambulance service, however, there is no such cocoon and the intervention of personnel in each other's work is a much more perspicuous feature. In dispatch³, the ambulance room's working division of labour is differently achieved, through extensive mutual monitoring, oftentimes peripherally (while personnel engage in their own work) and, when required, more focally (as others attract attention to their difficulties).

The kind of mutual awareness and extensive intervention (especially when the tempo of work is high) which is observable in our setting also marks out a difference with Heath and Luff's (1992) findings in London Underground control. While Heath and Luff did observe controllers engaging in exaggerated gestures and the like to gain attention, the abiding impression one has from their research is of two workers seamlessly coordinating their activities without interruption. As we have argued, in the ambulance control room, interaction opportunities sometimes require extensive and more perspicuous negotiation. This is not because personnel are refusing an altereological orientation and resisting giving help. Rather, in the hurly burly of dealing with very many emergencies at one time, others may happen to be engaged in some activity they cannot interrupt, or which they momentarily regard as more pressing, or the level of noise means that an imploring request has been missed. It is also important to note that the existence of *several* others in the ambulance control room means that an exaggerated gesture (or whatever) may not be automatically picked up as being for *a particular* other's benefit. Projecting one's activity so that it can be monitored by others or attempting to attract attention has an element of 'specific-other-selection' in the design of gesture and talk not necessary in settings (like Heath and Luff's) where essentially just two persons are working.

We would suggest that the maintenance of working divisions of labour is a more subtle and variable affair than some methods of 'task analysis' or 'process modelling' in HCI and computer science indicate. The ready defeasibility of the tasks associated with individual jobs (Dispatcher, Supervisor and, even, if required, Control Manager) are not experienced or negotiated as 'exceptions' to normal procedure or to a standard task or process model. Indeed, it would be exceptional (and reprimandable) *not* to show an orientation to others. 'Normal operation' *is* this essential flexibility in conduct. Accordingly, our development strategy is to consider the design of computing technologies with respect to how they can support a working division of labour, rather than, say, encoding a process model or assuming a differentiation between persons with respect to the tasks they perform.

In this regard, there is a lot to learn from our control room. There are artefacts devoted to 'individual views' (the Incident Stacks for each board) and to 'collective views' (the combined Incident Stack and the VAM). These different views are shared in their accessibility and can be juxtaposed the one to the other. Interestingly, part of what makes those views juxtaposable is that the system is *not* end-user configurable at the screen-interface. Indeed, the continued use of a mouseless, DOS-based, keyboard-driven, window-free system helps ensure that one Incident Stack has the same basic design as any other. Hence, visible differences between them will be work-related differences and not due to any idiosyncratic preference of the end-user. Supervisory functions could well be disrupted by seemingly more sophisticated interface techniques and ideologies of user-choice⁴.

However, while current arrangements enable the region to exceed target performance, we feel there is scope for further supporting the awareness that co-workers have of each other through technology. Workers currently complain of the high levels of noise as people shout at each other at peak times. If some of what is achieved through that noise (an awareness of what others are doing) is removed from the medium of speech and supported with enhancements to, say, the individual Incident Stacks, the 'air would be clear' for managing the most difficult contingencies. Our prototype explorations in this area will be made the subject of another paper.

³For there is, indeed, a sense in which a protective cocoon exists around the Call Operators so that they can attend to the business of interacting with an often extremely distressed public.

⁴Bentley, Hughes, Randall, Rodden, Sawyer, Shapiro and Sommerville (1992) make a similar point in connection with supervisory duties in air traffic control, even though, as we have argued, these are differently realised in that setting.

(3) Representations are Work-Oriented

It is possible to regard the VAM, the Incident Stacks and the computer forms completed by Call Operators as representations of (various aspects of) incidents and ambulance service. However, it is necessary to emphasise that these representations are oriented in their design to the effective support of *cooperative work in the control room*. While, naturally, they are intended to be easily perceivable and manipulable for the individuals before whose eyes they fall, their status as *work-oriented artefacts* within the cooperative assembly of the control room must not be missed. Again, *for the purposes of the cooperative work of ambulance controlling*, textual representations and displays, organised around lists, with sparse and carefully selected embellishments and colour highlighting, are most effective.

Let us illustrate this point with an example. While several features of the centre's GPS are frequently used, the 'flagship feature' often promoted by the manufacturers—a map-like display of ambulance locations on a large computer screen—is rarely consulted. The Automatic Vehicle Location System (or AVLS as this display is known) is typically turned away from the sight of the Dispatchers even though it is on the work-surface alongside them, and only easily inspectable by Dispatchers or Supervisors if they come close (see Figure 1). To understand this, it is important to realise that obtaining geographical fixes on the ambulance fleet is only *one part* of judging which to dispatch and how to maintain adequate cover (remember the many other contingencies), and a swiftly visible presentation of the *relevant* aspects of ambulance-location can be given by the GPS's computation of as-the-crow-flies distances and highlighting proximal, available candidates. A literal visual representation of ambulance locations might require a worker to engage in further deliberation to *extract* these details from a cosmetically impressive display⁵. Furthermore, the AVLS screen only shows one part of the region at any one time. Horizontal and vertical scrolling to find visualisations of ambulance locations on a map mostly full of 'empty' streets and countryside is needlessly time-consuming when, in contrast, the VAM compactly shows ambulances with a geographical sensitivity that is *appropriately approximate* for dispatch decisions. Importantly, on the VAM, *all* the ambulances are represented *there* (along with relevant status information), they do not have to be *found* in a 2D map-space⁶.

In short, for the purposes of dispatch and cover, seemingly crude lists may be the most appropriate form of representation, giving an easy impression of how many ambulances are available at every station. However, again, it is possible that more (and more detailed) cover-relevant information could be introduced to the Dispatchers' own resources for ambulance selection. In this regard, in our prototyping work, we are considering re-designs of the Dispatch Selection screen to contain 'at-a-glance-what-if information': e.g. alongside the as-the-crow-flies distance, some representation of the number of ambulances remaining at local stations if the one represented by the row of the list were selected.

(4) Human Computer Interaction is a Public Phenomenon

As we have emphasised, the sharing of artefacts within a shared ecology facilitates mutual awareness and coordination between control room personnel. In this respect, human-computer interaction is itself a *public phenomenon* in that any act of engagement by a person with a computer system, screen or display is potentially available to *third parties* and a resource *for them* in noticing or repairing troubles or gaining an impression of

⁵Indeed, it seems that (currently) one of the prime usages of the map-like visualisation is to impress visitors!

⁶Or a 3D space. While the so-called 'infinite interface' offered by immersive Virtual Reality systems may overcome the limitations of conventional visualisations which are bounded by the borders of the screen, it is hard to see the immediate utility of these forms of 3D computing technology to ambulance control settings (though applications of VR are commonly mooted for control purposes). In a sense, the 'dimensionally challenged' VAM (and other screens in the control room) 'filters out' details from view which would detract from effective ambulance controlling. Pycock and Bowers (1996) make a similar point in the context of fashion design work when they argue that it is 2D representations of garments which are found to be most appropriate for highlighting of significant detail for standard manufacture. This argument, of course, does not rule out the possibility that, for purposes *other* than manufacture (e.g. a customer viewing a garment draped over a representation of their own body through an advanced interactive shopping system), 3D visualisations may have uses.

generally how stretched the service is. This suggests the opportunity to, as we might put it, *design for third parties*, that is, to design interaction techniques, key sequences, screen changes and so forth so that they can be detected by others as appropriate occasions for, say, initiating interaction (cf. Greatbatch, Heath, Luff and Campion, 1993). Accordingly, in the ambulance control room, we are considering re-designs of the screens and interaction techniques used in dispatch to make it more noticeable by third parties when (and which) selections are made and when (and which) changes in view (Incident Stack or Dispatch Selection etc.) take place.

Gaver (e.g. 1991) has done much to popularise within HCI the emphases of J. J. Gibson (e.g. 1979) on the affordances of artefacts and show how everyday affordances can be exploited in system design. We speak of *interactional affordances* to particularly highlight that, in real work settings, *first* many artefacts (and not merely communications technologies) afford *occasions for social interaction* between co-workers, and *secondly* it is often only through *interaction and engagement* that affordances are revealed (cf. Coulter's, 1990, respecification of Gibson's account of ecological perception). We noted earlier that there is a sense in which the VAM affords several different kinds of 'seeing' depending on the engagement personnel have with it, e.g., an at-a-glance perception of the load on the service or a more detailed picture of the activities of individual ambulances with more careful scrutiny⁷. Our point now is that the ecology of the control room makes it available *to others* whether someone takes a swift glance or a more protracted stare. This often enables the act of visual inspection to be taken as an act of information pick-up: whether someone is picking up the details or the overall picture is available in the manner of their looking. That is, the status of human-computer interaction as a public phenomenon often links our two senses (above) of 'interactional affordance'. The fundamental challenge, we believe, arising for an ecologically-oriented HCI is to systematically support design, evaluation and implementation with these senses in mind. We have indicated some of our next tentative steps.

ACKNOWLEDGEMENTS

We would like to thank all ambulance service personnel who generously accommodated us throughout this research. The support of the European Communities' ESPRIT project *eSCAPE: Electronic Landscapes* is also acknowledged.

REFERENCES

- Bentley, R., Hughes, J., Randall, D., Rodden, T., Sawyer, P., Shapiro, D. and Sommerville, I. (1992). Ethnographically-informed systems design for air traffic control. In *Proceedings of CSCW '92*, the Fourth Conference on Computer Supported Cooperative Work. New York: ACM Press.
- Bowers, J., Button, G. and Sharrock, W. (1995). Workflow from within and without. In *Proceedings of ECSCW'95*, the Fourth European Conference on Computer Supported Cooperative Work. Dordrecht: Kluwer.
- Bowers, J. (1996). Hanging around and making something of it: Ethnography. In Haworth, J. (ed.) *Psychological Research: Innovative Methods and Strategies*. London: Routledge.
- Coulter, J. (1990). The praxeology of visual perception. *Inquiry*, 33, 251-72.
- Gaver, W. (1991). Technology affordances. In *Proceedings of CHI 1991* (New Orleans, Louisiana, April 28 - May 2, 1991). New York: ACM Press.
- Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*. New York: Houghton Mifflin.
- Greatbatch, D., Luff, P., Heath, C. and Campion, P. (1993). Interpersonal communication and human computer interaction. *Interacting with Computers*, 5, 193-216.

⁷Indeed, this further testifies to the subtlety of such list-based representational formats in this setting.

- Heath, C., and Luff, P. (1992). Collaboration and control : Crisis management and multi-media technology in London Underground Line Control Rooms. *Computer Supported Cooperative Work*, 1, 69-94.
- Hughes, A., Randall, D., and Shapiro, D. (1992). Faltering from ethnography to design. In *Proceedings of CSCW '92*, the Fourth Conference on Computer Supported Cooperative Work. New York: ACM Press.
- Pycock, J. and Bowers, J. (1996). Getting others to get it right: An ethnography of design work in the fashion industry. In *Proceedings of CSCW '96*, the Sixth Conference on Computer Supported Cooperative Work. New York: ACM Press.
- Randall, D., Rouncefield, M. and Hughes, J. (1995). Chalk and cheese. In *Proceedings of ECSCW'95*, the Fourth European Conference on Computer Supported Cooperative Work. Dordrecht: Kluwer.
- Rouncefield, M., Hughes, J., Rodden, T. and Viller, S. (1994). Working with "constant interruption". In *Proceedings of CSCW '94*, the Fifth Conference on Computer Supported Cooperative Work. New York: ACM Press.
- Sharrock, W. and Anderson, R. (1992). Can organisations afford knowledge? *Computer Supported Cooperative Work*, 1, 143-162.
- Whalen, J. (1995). Expert systems versus systems for experts: Computer aided dispatch as a support system in real world environments. In Ed. Thomas, P. *The Social and Interactional Dimensions of Human-Computer Interfaces*. Cambridge: CU