S. Bødker, M. Kyng, and K. Schmidt (eds.). *Proceedings of the Sixth European Conference on Computer-Supported Cooperative Work, 12-16 September 1999, Copenhagen, Denmark*

Informing Collaborative Information Visualisation Through an Ethnography of Ambulance Control

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

David Martin
Psychology and Computer Science Departments, University of Manchester, UK

Abstract. An ethnographic analysis of an ambulance control centre is presented, specifically investigating the design of information displays and their practical use in this setting. The spatial distribution of the displays around the control room is described and its consequences for cooperative work drawn out. From these analyses, we make several suggestions for information visualisations in virtual environments, including a design concept of multiple displays coexisting within a 3D environment as an alternative to the notion of 'immersive' information visualisation more commonly encountered. The paper closes with a reflection on the relationship between ethnographic analysis and system development that our work here exemplifies.

Introduction

This paper presents an analysis of field work conducted at an ambulance control centre in a large metropolitan region in the North of England. The purpose of this analysis is to inform the development of Collaborative Virtual Environments (CVEs), and in particular, CVEs which are designed to support collaborative information visualisation. In recent years in the research field of Computer

Supported Cooperative Work (CSCW), there has been much interest in CVEs, where distributed virtual reality technology is deployed to support cooperative working activities. While a number of general CVE platforms and development environments exist, (e.g. DIVE, Carlsson and Hagsand, 1993, and MASSIVE, Greenhalgh and Benford, 1995) and have been presented within CSCW, three areas have been the principal focus of application development and usage. *Virtual conferences*, where a CVE is deployed to provide some form of meeting environment as an arena for social interaction, have been discussed by Greenhalgh and Benford (1995) amongst others. *Inhabited television* explores strategies for combining CVEs with broadcast media (McGrath et al., 1998, Benford et al., 1999). *Collaborative information visualisation* has been the subject of applications such as VR-VIBE (Benford et al., 1995) and Q-PIT (Mariani et al., 1995). Here a (typically 3D) visualisation is embedded within a CVE so that, in a sense, users are embodied within the database itself. This might enable them to be aware not only of information but also of other users and the interactions with information and each other that they are engaged in.

An overarching concern in our work is how understandings of virtual reality, and in particular CVEs, might be informed by empirical social scientific analysis. Specifically, we are concerned to deploy ethnographic research methods (which emphasise protracted contact with a research setting and the deep descriptive documentation of the social organisation of activity) alongside video analysis work (concerned with the detailed study of how social interaction is coordinated on a moment-by-moment basis). Throughout, we prefer the study of real-world settings which have an existence independent of our interest in them. There are two strands to our research agenda with respect to CVEs. On the one hand, we conduct studies in real-world settings to inform CVE development—settings which, though often automated, do not typically involve the usage of virtual reality technology. For example, Pycock and Bowers (1996) report a study of work in the fashion industry so as to investigate the plausibility of introducing CVEs to support 3D design and information visualisation in this setting. On the other hand, we conduct studies of the usage of CVEs to explore whether the arguments for their utility are borne out in practice. For example, Bowers, O'Brien and Pycock (1996) and Bowers, Pycock and O'Brien (1996) address questions such as: do (and in what sense) CVEs enable users to deploy their everyday interactional competencies in the 'natural' way that is often hoped for? Do (and in what sense) users achieve a sense of 'place', 'presence' or 'immersion' in a CVE? Do users obtain an awareness of the conduct of others in a CVE which enables a sense of mutuality to come into existence? If so, how?

The present paper is of the first sort—a study of an existing setting: ambulance control. However, it is important that we are clear at the outset of the purposes of our study and exactly how we intend it to inform CVE research. We are *not* proposing the use of CVEs to support real world ambulance control, nor are we here concerned with making suggestions for other forms of cooperative technology in this setting. These are matters more prominently covered in Martin, Bowers and Wastell (1997). Our preeminent concerns in the current paper concern more basic issues in research on CVEs for collaborative information visualisation.

In particular, we wish to inform the design of collaborative information visualisations by offering a social scientific perspective on how databases are used in a complex, time and safety-critical setting. In existing work, the design of 3D information visualisations in the research literature has largely been considered from a computing, indeed algorithmic, standpoint. The distribution of virtual objects, signifying database entries, around a 3D space has been governed by algorithms based on, for example, cluster analysis, physical spring models of attractions and repulsions, or a weighted summing of overlaps between database entries and queries (see Benford et al., 1995; Chalmers, 1994). From time to time, considerations of human perception enter into the design of 3D visualisations. For example, Chalmers (1994) discusses the importance of 'landmarks' and 'regions' for making an information visualisation more easily perceptible and navigable. However, rarely have the actual principles by means of which representations of data are constructed been discussed from a social scientific viewpoint. Largely, an algorithm is selected on the basis of computational (e.g. efficiency) or aesthetic (e.g. the visual elegance of the virtual forms produced) criteria. This means that, ironically, existing CVEs for collaborative information visualisation tend to employ algorithms for constructing information displays *without explicit regard* for how differently designed displays might or might not afford the social interaction and mutual awareness their designers hope for. It is our belief that exactly how objects are distributed in virtual space will matter critically to whether a CVE for information visualisation will effectively support cooperative work. We suggest that designers of CVEs can learn from real-world settings in this regard and that (maybe) even the selection of the algorithms used for the construction of information visualisations might be influenced by empirical social scientific work. The following study is an exploration of this possibility.

Field Study

The ambulance control room is located above one of the 35 ambulance stations that serve the region (population around 2.5 million), centrally, in its main city. The control room operates by answering calls from the public (on the 999 emergency line), medical services and other emergency services. These calls are received by Call Operators (1-6 working depending on time of day/week) who enter basic details onto a computer-based form, their central task being to gain a geographical fix on the incident and some details of the nature of the problem.

As soon as a geographical fix is made and 'enter' is pressed these details are passed directly over to the Dispatchers (4 working at all times), whose job it is to assign each incident to an ambulance in the region (30-70 operational at any one moment). Dispatchers are constantly overseen by two Supervisors who sit behind them. This group work to ensure that, with a finite number of ambulances, incidents can be assigned to them such that adequate cover is maintained throughout the region (at least one ambulance covering every station) and ambulances are dispatched and reach incidents within UK government guidelines.

In order to manage incidents for the whole region it is divided up into four areas known as 'boards' with each dispatcher assigned one of these as their central dispatch area. Supervisors on the other hand have responsibility for the whole region. Dispatchers are presented with the geographically relevant incidents to their board as soon as they are entered. Their job is then to select a proximal, free ambulance for that incident while ensuring that adequate cover will be left consequent upon that selection. Once a selection is made, Dispatchers send abridged details electronically to an alphanumeric display and keypad on the relevant ambulance's dashboard. Supervisors, while not being assigned a particular set of dispatch jobs are constantly involved in monitoring and overseeing work, suggesting and even making dispatches (relieving Dispatchers' workloads) and preparing and planning with a special eye to maintaining adequate cover. The final staff member of importance to our study is the Control Manager, who for the majority of the time has a less hands-on presence dealing instead with general morale, quality of work, contacts with other organisations, and so forth. However, even the Control Manager might actually assign incidents to ambulances at times of extreme workload in the control room.

The central technological resource utilised by control room workers is a computer system containing a series of DOS-based applications (accessible via sequences of keystrokes) for the purpose of collating ambulance information in various ways so operations can be performed on it and distributed to relevant parties. The system is networked to over 20 terminals within the room and has a radio link to the display panels in the ambulances. A Global Positioning System (GPS) links to the database providing near real-time information on, for example, positions of ambulances relative to incidents. We now discuss the three most used applications and their relevance for dispatch.

The Incident Stack

As fixes are made on incidents they are visible through an application that lists them. This is known as the Incident Stack. It is split into two halves. The top displays 'WAITING' incidents (those still to be assigned to ambulances), the bottom 'ACTIVE' incidents (assigned). In each line, a variety of information is displayed including the time the call was received and, for incidents in the 'ACTIVE' category, the call signs of ambulances attending the incidents and their as-the-crow-flies distance (computed by the GPS). At the top of the 'WAITING' list is the oldest unassigned emergency incident while at the top of the lower half is the most recently assigned incident. Each Dispatcher has an Incident Stack for their 'board' accessible via their terminal (indeed this is their default screen setting), while the stack for the whole region is displayed on a large screen on the left hand side of the wall in front of the Dispatchers (see Figures 1 and 2).

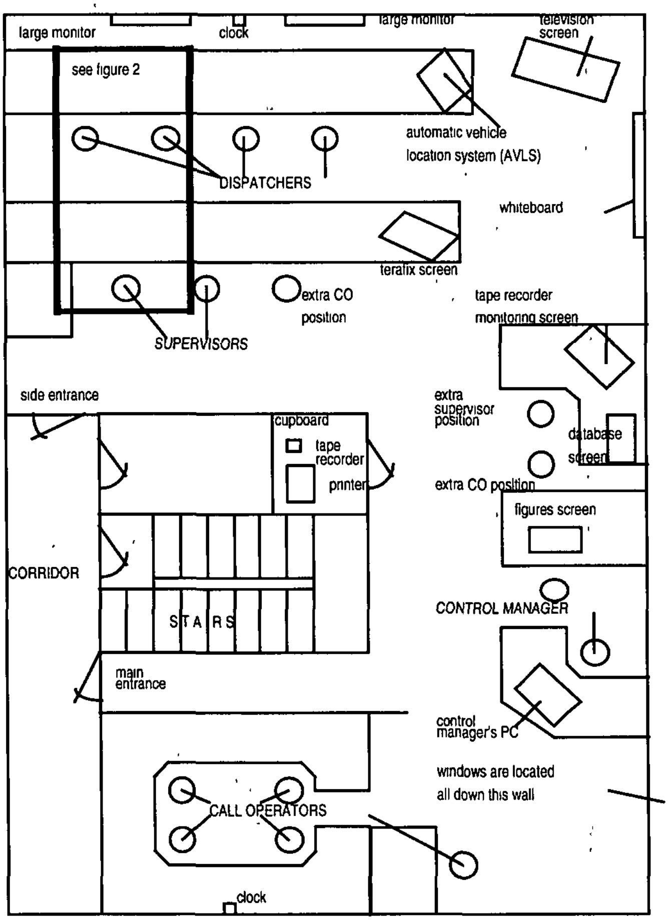

Figure 1· Plan of the ambulance control room

The Dispatch Selection Screen

When a new incident appears highlighted at the top of the Dispatcher's Incident Stack it is selected simply by pressing the enter key. The screen then refreshes to show the basic details of the incident (e.g. its type and location) in the top half of the screen and a list of the nearest as-the-crow-flies ambulances in the bottom half. This list does not take into account the status of these ambulances (whether active or free), however the nearest free ambulance is flagged in blue as a guide to allocation. The Dispatcher then selects an ambulance (often the flagged one), which again leads to the screen refreshing, this time showing the full incident details. These are then checked and abridged by the Dispatcher and transmitted to the ambulance. On receipt the crew press the first button on their keypad which leads to an automatic updating of the Incident Stack (the incident moves to the top of the 'ACTIVE' list) and the Vehicle Availability Map.

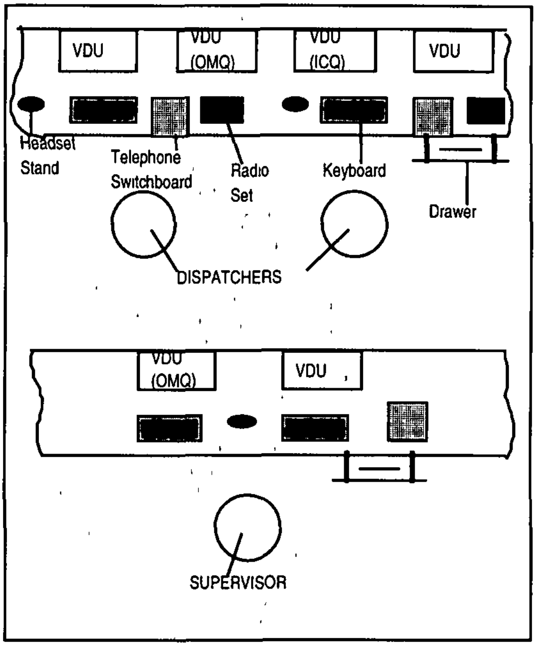

Figure 2 Dispatchers' and Supervisors' local environment

The Vehicle Availability Map (VAM)

As mentioned before, a key role of the Dispatchers and Supervisors is ensuring that adequate cover for the region is maintained. The VAM, which is displayed on the right hand large monitor on the front wall (Figure 1), is the central resource for this purpose. The VAM (see Figure 3) is a set of 15 lists each headed by a three letter abbreviation of the area covered, and comprising the call signs of ambulances from 2-3 proximal stations. The lists are arranged to correspond roughly to the actual geographical distribution of stations in the region (e.g. stations close to each other tend to be part of the same list, adjoining areas are listed proximally on the VAM). The important features of ambulance statuses are displayed by colour coding. Ambulances actively attending emergency incidents are shown in red. Those attending urgent incidents (transporting patients under doctor's orders) are flagged in green. Those placed at standby locations (a designated place to provide cover between two stations) have flashing call signs and those that are free are unflagged. Through examining the VAM, Dispatchers and Supervisors can glean a variety of important information for dealing with cover issues, as we shall shortly describe in more depth.

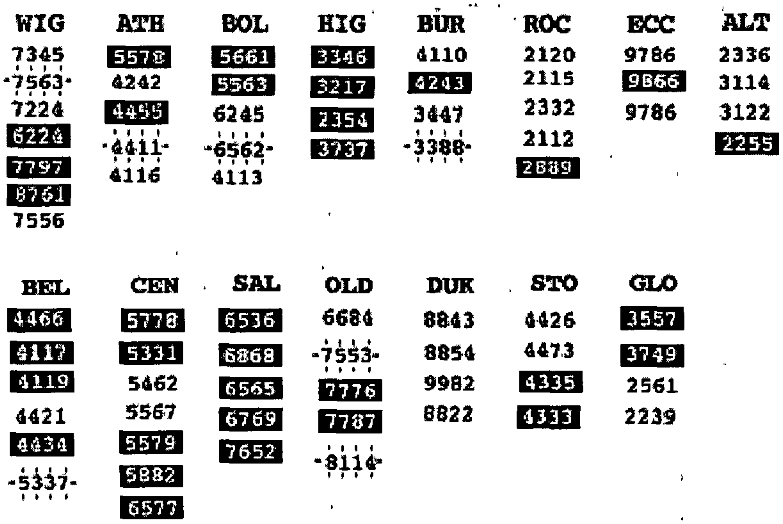

Figure 3 The Vehicle Availability Map (VAM) which lists ambulances by regions and is approximated here in greyscale Ambulances active on emergencies are here shown with their IDs against a black background Ambulances on urgent calls are shown against a grey background Ambulances on standby are shown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

Display Design

We now wish to analytically draw out some features of the work of ambulance control and its relationship to the design of the displays we have just discussed. Focusing on displays is justified both by our interest in information visualisation and because there is a strong sense in which *ambulance control work is display-inspection-and-manipulation-work.* That is, much of the work of ambulance control precisely consists in the inspection and manipulation of displays and it is through such activities that ambulance dispatch and cover maintenance are achieved. Dispatch could, of course, be accomplished by other means and, before the introduction of the system we have discussed, the extensive use of radio for voice communication with crews would have been essential. However, comparatively rarely now is there the need for direct communication with crews in order to accomplish dispatch. Such interactions are now occasioned by 'non-routine' occurrences such as if the crew cannot find the incident or if they are being asked to attend an incident when they are due a meal break or as a back up check to make sure the electronically delivered dispatches have been noticed (as crews can be in a station's staff room when a message is sent to their cab). However, routine dispatch work is accomplished by and monitored through the inspection and manipulation of displays. This centrality of the displays to control room work prompts us to analyse out various features of significance to them.

Distribution of Functionality

A first observation: different screens are used for different purposes. It is *not* the case that any screen in the control room can display any view on incident data from any application. There are dedicated screens for displaying incoming and outgoing message packets (the ICQ and OCQ) to monitor the effectiveness of radio transmission and to troubleshoot possible technical failures. For example, a message which has been transmitted but which does not appear in an ambulance's own display suggests that the latter may be at fault. There is a dedicated screen for the VAM. There is a dedicated, though not routinely used, screen for the display of ambulance locations on a map-like representation of the region (the AVLS display which presents the GPS information in cartographic form, see below). Furthermore, a password system controls access to applications so that, for example, the screens in front of the Dispatchers tend to display only dispatch-related information for the Dispatcher's region. That is, a number, but not all, of the applications making use of incident data will display to these screens. Equally, there are applications available to a Supervisor which allow for incident data from a number of areas to be compared on the one screen. The Control Manager can, in addition, inspect the results of applications which compute running statistics on the performance of the service in getting ambulances to incidents. The overall inspection of such statistics and the monitoring of control room performance is an important feature of the Control Manager's work. Hence, displays of such information are ready-to-hand for him on his screen and not on the screens of Supervisors and Dispatchers. Supervisors and Dispatchers will become aware of such matters through interaction with the Control Manager and not direct from the system to their screens.

This distribution of functionality across displays facilitates workers in that the information they need for their job is at hand, and readily so—it does not have to be found amongst a range of job-irrelevant displays. It also ensures that if workers need become aware, for whatever reason, of information *not* accessible from their own screen that they engage in embodied activity which makes the fact that they are so doing available to others in the control room. For example, if a Dispatcher needs to check on the activity in an adjacent board in order to ensure adequate cover for the service, they may lean over to glance at their neighbour's screen. The fact that they are so checking is thereby available to the Supervisor behind them. This may, in turn, alert the Supervisor to issues concerned with the relationship between ambulance deployments across the two boards—a matter of relevance to the Supervisor's job. In short, *constraining* accessibility and systematically *distributing* the displays and the functionality of the system across the control room *enables* opportunities for work-relevant communication and awareness.

Text

The vast majority of the information displays on view in the ambulance control room is textual in nature. We tend not see graphical objects being used to depict incidents or ambulances. This is *not* because control room personnel are tolerating antiquated non-graphical DOS systems. Rather, it is because the appropriate way of delivering information to control room personnel is in textual form. For example, the primary means for individuating one ambulance from another is by means of its call signal ID—a four number string. If ambulance-specific information needs to be displayed, attaching it to a four letter string seems most effective—especially if any anomalies about its status or whatever may well be resolved through spoken interaction with the crew which will be initiated by keying in the call sign on the radio (this opens a channel to the ambulance). Denoting ambulances with their ID in this fashion compactly individuates them, while giving exactly what is required to initiate interaction if problems occur. A graphical representation (e.g. a little ambulance icon coloured distinctively) might require personnel to *remember* which ambulance is which, how colour is used and so forth. Simple textual displays often give the information required *here and now* (memory not required).

Numerical and compact letter IDs are used ubiquitously in the control room (for incidents, areas within the region and ambulance stations, as well as a key for different types of incident: e.g. RTA for road traffic accident) in ways which are well and commonly understood. This gives text a priority in making displays of information work-relevant. It also makes means of organising text on screen around 'lines' and 'fields' essential. This is not to say that other means of organising text at the interface are not worth exploring. It is not even to say that iconic or symbolic means for representation might not also on occasion have a role. Rather, it is to emphasise that DOS-based textual displays have a lot going for them with a considerable work-relevant rationale behind them.

Colour

Colour is used sparsely in display design and for little other than highlighting purposes. Colour is not used for representational purposes, at least not alone. Colour is always attached to text. Colours do not appear attached to an iconic graphical object or to an abstract symbolic form. The colour 'palette' is restricted to red, blue, green and yellow. Again, this is not merely because control room personnel are being tolerant of or held back by DOS restrictions. In the VAM, for example, red versus green, flashing versus non-flashing, coloured versus plain, are used to indicate primary distinctions of relevance to the work of ambulance dispatch and cover maintenance. It is questionable whether more needs to be shown without introducing confusions which might be caused by reducing perceptual clarity (this restricted colour set allows the same discriminations to be made under different viewing conditions) or difficulties remembering a more complex coding scheme. The restricted colour palette also enables all the greens (say) to be seen as instances of the same class. In other words, at-a-glance perceptual categorisation is possible, something which might be militated against if more varied shadings were used. Again, we do not deny that it might be possible to creatively use colour in displays. Rather, we note that simple distinctions, made with a limited colour palette, seem appropriate for supporting the cooperative work of ambulance controlling practiced in the fashion we have observed.

Dimensionality

We have remarked that, for the purposes of the cooperative work of ambulance controlling, textual representations and displays seem most effective. If text has a kind of priority in display design, then so will lines and lists and fields (for input). Familiar reading practices can then be capitalised upon so that there will arise an expectation that a sense will be found reading a line from left to right and a screen from top to bottom. In the Incident Stack, for example, to the top are the incidents which are yet to be assigned and at the very top are the oldest, the most overdue. To the left of each line, one will find information to individuate the event (e.g. its time, its geographical fix in an abbreviated form), to the right information about what has been done (the ID of an ambulance mobilised, its distance from the event). In other words, while the screen can be considered a flat 2D display, this surface becomes enriched not just with text, colour, lines, fields, lists but also with incidents, locations, times, and a sense of urgency. Skills of normal reading practice give the user of such a display much. A working practical knowledge of ambulance controlling enables yet more to be read in. It is questionable whether other uses of the screen (for example as a graphically ornate representation of information) would give any more. Indeed, again the question arises as to whether such displays might distract through over-elaboration.

Let us give an example to illustrate this last point. Perhaps the most aesthetically interesting screen within the control room is also the least used. The Automatic Vehicle Location System (AVLS) is a map-like display of the region containing representations of ambulances positioned by means of GPS data. This, while situated on the desk near the Dispatchers and Supervisors, is seldom consulted for a number of reasons. First, it only displays a part of the region at one time meaning that it needs to be scrolled through. In contrast, the VAM shows *all* the ambulances *there* and *then*. Second, estimating distances becomes a task of spatial judgement on the AVLS rather than simply reading figures or noticing blue highlighting on the Dispatch Selection Screen. Finally, ambulances which are close to one another geographically tend to appear on top of each other on the AVLS (e.g. if several are together at a station). The VAM, by contrast, separates out different ambulances and gives a clearer picture of the relative workload in different areas.

The VAM shows ambulances with a geographical sensitivity that is *appropriately approximate* for dispatch decisions. The areas in the region are given a quasi-geographical distribution around the VAM screen. This enables comparisons to be made between, say, the commitment levels of adjacent areas, without, for the time being, being distracted by exactly how far apart those areas and the stations within them are. It is this that is important for making judgements about whether to move an ambulance to a standby location. One might well miss opportunities for flexibly deploying ambulances to standby locations if, say, one could not see that one area could help another out because their relative level of commitment had to be discerned from a display rather than seen at-a-glance. Equally, a more literal geographical representation might hide from ready-view *which* comparisons of commitment should be made. Two stations which seem near each other but which are separated by a motorway with entry and exit restrictions may not have a convenient standby location between them. The grouping of stations into areas and the spacing of area-lists around the VAM has been done precisely so that such confusions tend not occur. The VAM captures the 'logic of the region' with respect to ambulance cover maintenance—a very different matter from the logics of more geographically oriented maps or spaces defined by separation distance alone.

This is not to say that geographically oriented maps are irrelevant to ambulance controlling. Copies of the local street atlas (the so-called A-Z) are ubiquitous in the control room and on board ambulances and the page numbering and grid reference system of one publisher have even been adopted to give a representation of the location of incidents internally within the database system. However, the use of such maps themselves is occasioned. They might be consulted if an incident location cannot be found by a crew or if the Call Operator cannot get a fix. Similarly, the (rare) occasions we have seen the AVLS used within ambulance controlling are confined to helping out an ambulance crew who report being lost.

The summary point to make out of these observations is that there is no sense in which any form of display (textual, 2D, 3D) is *inherently* superior to any other. What is important is designing displays so that the information required for the job is ready-to-hand *and no more.* To accomplish this may point to the (principled) selection of textual, list-oriented displays on a 2D screen and may point away from graphical, map-like or 3D projections if those would just confuse the work.

Layers of Seeing

Consider again 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. While such summary impressions are available at-a-glance, both the VAM and the Incident Stacks can be engaged with further 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 engagement they have with it. Indeed, we have observed personnel dramatising the form of engagement they have with a screen so as to ensure that any other personnel who might be attending can pick up on what they are picking up from the screen. For example, a Supervisor might lean forward over their desk conspicuously peering into the VAM. In so doing they are displaying to all who might see that they are engaged in careful scrutiny and, in all likelihood, are focusing on a problematic deployment rather than the status of the service as a whole. Naturally, such understandings of the screen and the activities of persons in relation to it is essentially contextual. The point is that the fact that such variable understandings are possible is in part afforded by the coexistence on the VAM of (quite gross) visual features which enable different kinds of seeing.

We might say then that the VAM and similar displays afford 'layers of seeing' varying from an at-a-glance perception of the load on the service to a more detailed picture of the activities of individual ambulances with more careful scrutiny. This layering is accomplished through the careful and sparing use of the simple interface techniques we have discussed. Indeed, this further testifies to the subtlety of list-based representational formats in this setting. More complex displays may well not permit such readily distinguishable forms of seeing. A complex display may always require extensive engagement, thereby making it hard for others to fathom just what someone is doing or is concerned with.

No End-User Configurability

Above, we noted that the distribution of function across displays in the control room enables personnel to become aware of the status of each other's work in just such a way as to help them with their own. This can, for example, be sustained by juxtaposing different views of the state of the service. For example, a Supervisor might compare the Incident Stack of one Dispatcher with that of another to gain an impression of the relative stress being placed on different regions within the service. This might enable the Supervisor to judge that an ambulance should be moved to a standby location between the regions. Part of what makes views on the service juxtaposable is that the displays are *not* end-user configurable. Use of mouseless, keyboard-driven, window-free interfaces helps ensure that one Incident Stack has the same basic appearance as any other. Thus, visible differences between them will be work-related affairs and not arising from preferences of different users or from any work-unrelated layering of windows. Supervisory activities could well be disrupted by seemingly more sophisticated techniques and principles of user-choice (cf. Bentley et al., 1992).

Display Combination

So far we have been pointing to various features of the information displays employed in the ambulance control setting we have studied. There are good, work-oriented reasons why the displays are predominately textual, use colour sparingly, have standard non-tailorable formats, are often structured around lists presented on a flat screen (rather than graphical objects in any higher-dimensional space), are many in kind (each devoted to some particular purpose or related purposes), and are distributed around the screens of the control room, rather than each being available on every screen. We have already hinted that the *juxtaposition* of multiple displays is important to ambulance control. This and other related features of the work will be further emphasised in this section.

Different Loci and Spatialising What is Variable and What is Not

The different displays of information have specific loci on screens distributed through the ambulance control room. This spatial distribution is not haphazard. Indeed, it was carefully designed and considered by control room staff at the same time as requirements were put together for the current information system. In other words, the information system and its physical, ecological realisation were designed together, the new system requiring new furniture and some rebuilding of the control room itself.

Figures 1 and 2 depict this spatial distribution of information displays. The far wall is the site for the merged Incident Stack, the VAM, and a clock. The ICQ and OCQ sit between the Dispatchers just in front of this wall. The displays are all large (20" plus) monitors and are spaced out evenly. This gives them a clear boundedness and facilitates picking up on which display a control room worker might be glancing at. These displays are not variable under normal operation. The monitor showing the ICQ always shows the ICQ, the monitor showing the VAM always the VAM, and so forth.

The monitors on individual control room worker's desks, by contrast, are variable in what can be displayed upon them. The Dispatchers can access a number of different information displays, the Supervisors yet more, the Control Manager more still. That these displays are located before designated workers gives them the information they need ready-to-hand as we have emphasised. It also enables workers, who may need to know something not normally provided for on their own displays, to know where to look. A Supervisor concerned with performance statistics knows that this information can be accessed from the Control Manager's terminal. The spatialisation of different displays to specific loci enables routine information to be ready-to-hand without interfering confusion from what is needed only occasionally, while the locus of information required only on occasion is nevertheless commonly known.

The seating of control room personnel 'in rows' with Dispatchers nearer the far wall than Supervisors facilitates supervisory activities while minimising the distractions Dispatchers might otherwise be exposed to. There is a similar 'layering' of displays as one follows a trajectory from the far wall of the control room past the Dispatchers, past the Supervisors, and then to the Control Manager. We move from displays which are invariable, larger, generally available and show 'merged' features of the service through displays which are variable and manipulable, smaller, designated for individuals and concerned with area-specific reckoning through again to displays which are still individual, variable and manipulable but add in summaries of the whole service. In short, the disposition of displays and persons in the room comprise a structured working ecology.

Displays and Inspectable Embodied Conduct

We have given examples already of how the distribution of several, separate displays, each with their own specific locus, enables workers' embodied conduct to be intelligible to others. If a Dispatcher leans over to scrutinise another Dispatcher's display, this can enable a Supervisor to pick up on a matter of potential cross-board significance. If a Supervisor stares long and hard at the VAM, this may alert the Control Manager to a problem in maintaining cover in a part of the region. The variety of different displays, each with their own distinct spatial locus, supports discriminability in the understandings that can be obtained through inspecting the embodied conduct of others. The existence of potential issues to do with message transmission, providing cover, a difficult dispatch decision, a particularly feisty crew and so forth can all be distinguished *in part* because relevant displays are in turn distinguished. Naturally, there is no one-to-one mapping going on here which could be captured in crude 'rules' like 'if a Supervisor glances to display D, this means X'. This is why we emphasised *in part*. The variety of coexisting displays in the one space affords such opportunities, it doesn't mandate the attributions made in any mechanical fashion.

The Personal and the Public

In Martin et al. (1997), we emphasise in depth how control room personnel maintain a 'working division of labour' so as to help each other out when the service is stressed and to collectively manage the contingencies which arise in control room work, while nevertheless maintaining a sense of 'my job'. In the terms of Hughes et al (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?). The distribution of information displays in the control room is apt for the management of a working division of labour. There are to be sure 'personal' machines and displays. Personnel have a sense of 'my computer'. However, viewing information displayed on 'my computer' doesn't rule it out being inspected by some other who may be looking over my shoulder or who inserts themselves into a side-by-side view to momentarily share the screen. Equally, there are the public displays towards the far wall. But their being public doesn't necessarily rule out their accessibility from a personal terminal. This distribution of information displays allows routine work to be carried on without disruption from others while not disabling the possibility of making a problem public if needs be. For example, a particularly desperate Dispatcher can draw the attention of the entire control room to the paucity of available ambulances in their board by gesturing towards the VAM and pointing out on the merged Incident Stack an incident which is hard to deal with. Between the extremes of routine work and pleas for help, personnel have a range of displays and other persons they can call upon so as to appropriately make their own difficulties public or attune to the conduct of others.

Conclusions

We now rejoin our initial discussions of information visualisation in collaborative virtual environments (CVEs). This will be structured around our development of *multiple information displays coexisting within a shared virtual space* as a design concept for consideration. We offer this image as an alternative to notions of 'immersion in information' more commonly explored in CVEs for information visualisation.

Coexistence of Multiple Embedded Displays

Scrutiny of our ambulance control study suggests the utility of having multiple displays coexisting and embedded within a shared environment. Multiple displays enable the information required at each 'station of work' to be made readily available there and so be ready-to-hand when required by persons working at that station. This enables some division and distribution of function and work activity. However, the coexistence of multiple displays in a shared environment enables the in-principle accessibility of other information if needs be.

The principles governing the distribution of displays in the control room are subtle. Manipulable displays have different loci from non-manipulable displays. Monitors for variable displays have different loci from those with fixed displays. Displays which are public to all control room workers have different loci from those which are personal. This variation and distribution supports a working division of labour between personnel, a collective sense of the state of the service coupled with personal duties.

The possibility of multiple coexisting displays all embedded in a shared space has not been directly explored in work on CVEs for information visualisation. Most existing systems focus on producing a single information visualisation within a 3D space. Much attention has been paid to different visualisation and manipulation techniques but typically only a single visualisation at a time is computed. We would suggest that this ignores a critical feature of information sharing in a setting like ambulance control and, perhaps, more generally: *information is typically visualised (and acted on) in the context of other visualisations systematically distributed in the same space.*

Recent, yet to be published, work by Taylor (in prep.) is a partial exception to this as he experiments with various techniques by means of which multiple simultaneous displays can be computed, manipulated and distributed within a common virtual space. However, Taylor's systems are for single users and are not yet integrated within a CVE. This means that several features his work (e.g. choice of algorithms for computing displays and their distribution) have not yet been explicitly responsive to the requirement to support cooperative work. What we are attempting here are some first steps at using ethnographic work to inform the design of shared information visualisations of this sort.

Interactional Affordances and Freeing Up Space

We have been at pains to argue that distributing functionally and spatially distinct displays enables the gestures and talk of personnel to be richly yet readily understood by others. Some aspects of this can be captured in the notion of 'interactional affordance' (cf. Martin et al., 1997). Information displays afford interaction in the dual sense that (i) they enable a user to interact with the information displayed, and (ii) such human-computer interaction can be publicly available to 'third parties' and afford them with opportunities for interaction with either of the first two 'parties' (user or display). Information displays can be designed not merely to make them readily intelligible by their users but also to make their users' activities with them be intelligible to others. For example, the VAM not only affords different forms of seeing, it also enables people viewing it to make their forms of seeing public to others. In our ambulance control setting *human computer interaction is a public phenomenon.* Let us give some further elaboration to see how understanding this notion of interactional affordance might shape the development of design ideas for collaborative information visualisation.

For interactional affordances to exist, for the actions of persons with artifacts to be coherently available visually to third parties, *space must be left free.* For example, the disposition of the displays towards the far wall of the control room leaves a great deal of the wall space free. As we have argued above, this separates out the displays one from another to make gesture and gaze to one clearly discriminable by third parties from gesture and gaze to another. There would be no clear advantage in having a very large (and very expensive) monitor across the whole of this wall displaying many information sources when keeping those sources distinct is important to the intelligibility of them and of gesture and gaze with respect to them... and this is most readily accomplished by having a number of cheaper, more conventional monitors!

Just as a 'plane' (like the control room's wall) need not be completely filled with displays, a three dimensional space (like a control room or, we would suggest, a VE for information visualisation) needn't be densely filled either. More densely packing our control room with displays would not merely make for a claustrophobic working environment, it would also make it harder to interpret the actions of others. Perhaps we can say: in a shared space containing multiple displays, the local intensities/densities of information need to be sparsely distributed for interactional affordances to be possible. Put less grandly, if the whole space is cluttered with information sources not only will you have difficulty making sense of it all, you'll also have trouble picking up on what others are doing and making sense of that.

These considerations also have, we believe, novel implications for research on shared information visualisations using virtual reality technology. First, they give us good reasons for exploring visualisation techniques which yield geometrical forms with a dimensionality less than three, even if such forms are embedded within a 3D CVE. The discipline of confining oneself to 2D forms (for example) can encourage us to thin out data to show just what is relevant for the purposes of the display. 2D forms allow 'side-by-side' or 'face-to-face' views between two or more embodied users who have lines of sight (more or less) perpendicular to the 2D plane. This can facilitate the mutual intelligibility of activity, as our gestures and lines of sight can become 'aligned' more readily. There is also a lesson in the ClearBoard system (Ishii and Kobayishi, 1992) which combines two video streams and a shared 2D drawing surface with participants having 'face-to-face' views on either side of the surface—an arrangement which also can support mutually intelligible gesturing. Here, not only are the shared forms confined to a 2D plane, most of the drawing plane is unfilled, indeed is 'clear' (hence the name of the system). This further testifies to the importance of *leaving space free for interactional purposes*. Finally, Chalmers (1994) argues that 2.nD (n small) displays, having a dimensionality suggestive of a 'landscape' enable us more readily as earthbound, perambulatory creatures to utilise everyday perceptual and naming practices (e.g. 'up' and 'down' can have a ready significance). In summary, then, there are good reasons, based on questions of what forms of perception and interaction are afforded by different geometrical forms for (1) exploring forms with a dimensionality less than three even (especially) if they are to be embedded in 3-space and (2) preferring external viewpoints which allow 'side-by-side' or 'face-to-face' mutual orientations towards displays rather than 'immersed' viewpoints which not only militate against obtaining an overview but also might cause problems for the mutual intelligibility of gaze and gesture.

Considered in this way, there are strong practical reasons for flat, more or less two dimensional information displays in VEs. Yet more arguments for flatness accumulate for settings where text is to be visualised, as argued above. These arguments intensify still further, we would suggest, when multiple displays coexist embedded within a shared VE. Flat displays sparsely distributed in separate loci minimise the occasions when one display (or visual forms from it) might occlude another (or forms from it). When they do occur, such occlusion problems can be rectified by (in the real world) leaning around, standing up or 'craning the neck', or (in the virtual world) by a small levitation of the user-embodiment. But this is possible principally only if the displays are themselves laid out so that they do not fill three dimensions evenly. Space has to be left free so that such adjustments of viewpoint counteract occlusions systematically. If 3D space is, in principle, always everywhere equally filled, then an occlusion problem might be remedied this time by a shift up, this time by a shift down, on yet another occasion by a movement to the side. A flatter 'landscape' or other 2.nD (n small) structure will encourage occlusion problems to be resolved and overviews to be sought by means of displacements in the direction of the 'fractional dimension' ('up' in the case where displays gravitate towards the groundplane).

This systematicity lays the ground for an interactional affordance. In the case of a 'landscape' of displays close to the groundplane, displacements up can become intelligible to others as attempts to get a view of a distant display rather than a 'pure' navigational movement. If, as in the control room, displays can be different from each other functionally and not just spatially, then rich inferences can potentially be drawn about the exact significance of a simple move up. An embodiment moving up can be supposed to be one trying to find out about just those things that the distal display depicts. Systematically distributed displays help enable *actions* to be perceived, not just behaviour noted.

When the visualisation of complex dynamically changing data from multiple perspectives is at issue, then, we are developing an argument for exploring CVEs comprised of displays which:

- are multiple, localised, discrete, functionally distinct and embedded within the VE
- have a distribution in the VE which fills out its three dimensions unequally and sparsely
- each have a dimensionality less than the VE itself (i.e. tend to be relatively flat)
- are designed for 'external viewing' primarily (i.e. while user-embodiments and information displays are to be 'immersed' in the VE, the further immersion of users within the displays themselves may not always be called for or desirable).

It is algorithms which generate spatial distributions of this sort that we would commend for exploration for collaborative information visualisation.

Naturally, VEs for particular applications will have to be developed in ways which are sensitive to application-specific requirements. We are not suggesting that CVEs along these lines above will always satisfy all needs. Any one of the features we note could be relaxed in a particular case. What we are suggesting is that the above is a particularly interesting constellation of features for CVEs which is motivated by our ethnographic consideration of ambulance control and is under-investigated in the literature on CVEs for information visualisation to date.

Ethnographic Practice in CSCW

We are aware that this has been an ambitious paper in that we have tried to present ethnographic analyses in some detail while also engaging with technical design issues in CVE application development. Some readers may feel uneasy about this, believing that we are overstepping what can be legitimately concluded from empirical social scientific research. To allay these fears, let us try and be clear again about what we are offering and the limits of our argument.

We are putting forward for consideration a *design image*, a set of concepts which could be instantiated in actual systems. Our design image (of multiple coexisting information displays embedded within a shared virtual space) is proposed as something to *add* to the stock-in-trade of CSCW researchers concerned with shared environments for information visualisation. It is an alternative to an idea of 'immersion within information' and, under some circumstances, may turn out to be more useful, under others, less. As is standardly the case with ethnographic practice in CSCW, we want to add to developers'/designers' resources, not take anything away (cf. Randall, Hughes and Shapiro, 1992). What we are offering is not any abstract theory of information display and manipulation but a design image of potential practical use inspired by our ethnographic description. And like many 'design recommendations' arising from social scientific work, it is *inspired by* field study, not *deduced from* it. The proper evaluation of this would be by reference to particular applications of our design image in plausible contexts of use. It would be too much, we believe, to ask of this now in the current paper but, naturally, this is the subject of our ongoing work.

To be willing to offer up such 'design images' when they present themselves is, for us, *a* response (we have others!) to Plowman, Rogers and Ramage's (1995) question, posed at ECSCW95, *what are workplace studies for?*

Final Word

Gibson (William) once defined cyberspace as "the mind turned inside out". We prefer an altogether more mundane conceptualisation. Virtual environments comprise another arena for social interaction and cooperative work alongside others. Under these auspices, we have argued that virtual environments are to be apprehended in terms of how they display information so that it can be practically acted upon and so that one person's engagement with information can be picked up by others. This is why we emphasise Gibson's (J.J.'s, 1979) conception of affordance. It is also why we grant that there can be good reasons for displays which are textual, list-oriented, non-user-configurable, not very subtly coloured and flat. We do not believe that this is conservatism or lack of imagination on our parts when confronted with the 'new frontier' of virtual reality. Rather it is sensitivity when confronted with the details of lived work practice and a desire to take those seriously as a source of inspiration for technical design.

Acknowledgements

The authors acknowledge support from the eSCAPE (Electronic Landscapes) and eRENA (Electronic Arenas for Culture, Performance, Art and Entertainment) projects of the European Communities' ESPRIT Long Term Research Programme.

References

Benford, S., Greenhalgh, C., Craven, M., Walker, G., Regan, T., Morphett, J., Wyver, J. and Bowers, J. (1999) "Broadcasting on-line social interaction as inhabited television", *Proc. ECSCW99.* (in this volume)

Benford, S., Snowdon, D, Greenhalgh, C, Ingram, R, Knox, I. & Brown, C (1995): "VR-VIBE: A virtual environment for co-operative information retrieval", *Proc. Eurographics'95.*

Bentley, R, Hughes, J., Randall, D, Rodden, T, Sawyer, P, Shapiro, D & Sommerville, I (1992). "Ethnographically-informed systems design for air traffic control", *Proc. CSCW'92*

Bowers, J, O'Brien, J. & Pycock, J. (1996): "Practically accomplishing immersion", *Proc. CSCW'96.*

Bowers, J, Pycock, J. & O'Brien, J. (1996): "Talk and embodiment in collaborative virtual environments", *Proc. CHI'96*

Carlsson, C & Hagsand, O (1993): *"DIVE", Computer & Graphics*, vol 17, pp.663-669

Chalmers, M. (1994) "Information environments", in L MacDonald & J Vince (eds), *Interacting with virtual environments*, Wiley, Chichester.

Gibson, J. J. (1979): *The Ecological Approach to Visual Perception*, Bantam, New York

Greenhalgh, C and Benford, S (1995): "MASSIVE. A virtual reality system for tele-conferencing", *ACM: TOCHI*, vol 2, no. 3, pp 239-261

Hughes, J, Randall, D & Shapiro, D. (1992): "Faltering from ethnography to design", *Proc CSCW'92.*

Ishii, H & Kobayishi, M. (1992): "Integration of inter-personal space and shared workspace", *Proc CSCW'92*

McGrath, A, Oldroyd, A. and Walker, G. (1998): "The Mirror reflections on inhabited TV", *Proc CSCW'98 Video programme.*

Mariani, J., Rodden, T., Colebourne, A., Palfreyman, K. & Smith, G (1995): "Q-PIT a populated information terrain", *IS&T/SPIE Symposium on Electronic Imaging: Science and Technology.*

Martin, D, Bowers, J & Wastell, D. (1997): "The interactional affordances of technology an ethnography of human-computer interaction in an ambulance control centre", *Proc HCI'97*

Pycock, J. and Bowers, J (1996): "Getting others to get it right", *Proc CSCW'96.*

Plowman, L., Rogers, Y and Ramage, M (1995) "What are workplace studies for? ", *Proc ECSCW'95.*

Randall, D., Hughes, J and Shapiro, D (1992): "Using ethnography to inform system design", *Journal of Intelligent Systems*, vol. 4, pp 1-2