

Preprint. Final version available from:

BLANDFORD, A. E., WONG, B. L. W., CONNELL, I. & GREEN, T. R. G. (2002)
Multiple Viewpoints On Computer Supported Team Work: A Case Study On
Ambulance Dispatch. In X. Faulkner, J. Finlay & F. Détienne (Eds.) *People
and Computers XVI: Proc. HCI'02*. 139-156. Springer.

Multiple Viewpoints On Computer Supported Team Work: A Case Study On Ambulance Dispatch

**Ann Blandford¹, B. L. William Wong², Iain Connell¹
& Thomas Green³**

¹ UCLIC, UCL, Gower Street, London, WC1E 6BT, U.K.
{A.Blandford, I.Connell}@ucl.ac.uk

² University of Otago, Box 56, Dunedin, New Zealand,
william.wong@stonebow.otago.ac.nz

³ University of Leeds, Leeds, LS9 2JT, UK.
Thomas.Green@ndirect.co.uk

A novel usability evaluation technique, Ontological Sketch Modelling (OSM), was applied to the analysis of systems used within a complex work setting, namely emergency medical dispatch. OSM focuses on the structure of the domain in question and the devices which are applied to that domain, in order to reason about the quality of fit between the two. This analysis shows how OSM can be used to identify misfits between domain (here incidents, ambulance calls and real-time call processing by ambulance service staff) and device (the computer aided dispatch system) in real work settings. We show how OSM can aid additional reasoning about the way in which a new or proposed computer system can both support and enhance existing work structures. The analysis presented here also yields important insights into both the still-developing OSM and the structure of emergency medical dispatch systems.

Keywords: usability evaluation; emergency medical dispatch; conceptual model; domain model; OSM.

1. Introduction

There is still a shortage of effective techniques that are suitable for analysing the systems used within a complex work setting. On the one hand, techniques based on approaches such as activity theory and distributed cognition yield valuable insights, but at a level of description that is not readily translated into implementation requirements. On the other, formal approaches that model tasks or systems provide a clearer bridge to requirements and implementation, but often do not scale well to complex systems. Approaches that bridge this gulf are gradually

emerging – for example, Cognitive Work Analysis (Vicente, 1999) and the Resources Model (Wright, Fields & Harrison, 2000). Ontological Sketch Modelling (OSM: Blandford & Green, 2001a), the approach presented here, also aims to bridge this gulf, focusing particularly on the users' conceptualisation of the domain in which they are working, the device representation(s) that support that work, and the quality of fit between the two. OSM is still under development; it is described briefly below. Here, we report on its application within the Central Ambulance Control room (CAC) at the London Ambulance Service (LAS). A secondary outcome of this work is a focused account of the work of ambulance control and the systems currently used to support that work.

OSM focuses on the structure, or ontology, of the domain and devices under consideration. By simultaneously modelling both the users' view of the domain and the devices that support work in that domain, the quality of fit between devices and users is revealed. As implied in its name, OSM is a *sketchy* approach to modelling; that is: we are investigating the value and limitations of models that are not constrained to be complete or necessarily consistent. This contrasts with most formal modelling techniques, which demand of the analyst that the resulting model should be both consistent and complete – a demand that appears to have limited the uptake of modelling in practice.

Models can be seen either as predictive tools, usually highly formalised, or as less formal, insight-bringing abstractions for assessing a design from a particular viewpoint. OSM belongs in the latter class. In developing a new approach to modelling, there are many key questions to be addressed:

- what is the core ontology of the model to be?
- how is the modelling technique to be applied?
- what kinds of insights does it afford (whether through the process or the product of modelling)?

Established usability evaluation techniques often miss the point that what makes a tool difficult to learn and use is the number of novel concepts to be acquired before productive work is possible. OSM successfully revealed such problems in two contexts studied earlier, namely drawing tools (Blandford & Green, 1997) and diary use (Blandford & Green, 2001b), domains in which tasks are frequently ill-structured. Ambulance dispatch represented an interesting domain because it contrasts with personal tools such as drawing packages and diaries on many important dimensions:

- team working;
- subdomains that need to be described separately; and
- much information that is essential for participants, but not changed by them.

OSM describes *concepts* (entities and attributes that are significant in both the domain and the device), the actions that change the state of entities and attributes, and important relationships between entities, attributes and actions. Problems of misfit become apparent as the model is constructed. Both existing work concepts and existing support systems are taken into account. This contrasts with an approach such as contextual design (Beyer & Holtzblatt, 1998), which considers existing systems including work concepts and support systems, but does not explicitly consider the quality of fit between the two.

Constructing an OSM model involves developing a (semi-formal) description of the users' conception of their work system. This is done by role: if several

individuals perform the same role, the aim is to construct a general model that captures the key features of their shared conceptualisation, rather than focusing on individual differences. The users' model is described in terms of entities and attributes (see Table 1 below for an example), and the analyst may add annotations that expand on the key entities and attributes. This model is derived from an analysis of qualitative data from users – see section 2 for an example of the kinds of data used. In parallel, a model of the device is developed, by reference to existing implementations, documentation, specifications, or other appropriate device descriptions. Again, entities and attributes are described and, in this case, information about actions is also included. These are actions that create or delete entities, and set or change attributes. The analyst is also invited to add annotations. In some cases – though this is not considered further here – there may also be concepts that are 'device-private', which the user needs to know about but cannot easily access or change.

Sources of misfits include key user concepts that are not represented by the device (forcing users to find workarounds) and device concepts that are not a natural part of the users' model (forcing users to make explicit or be aware of information that is not key to their conceptualisation). Misfits of the first type can be translated into design requirements that such concepts should be represented at the interface; misfits of the second type may indicate that the device representation needs to be restructured (to better fit the users' conceptualisation) or that device concepts need to be more clearly represented to the user (so that the users' conceptualisation can better adapt to them). A more detailed classification of misfits and account of their implications for design is under development.

In applying OSM to the LAS our aims were, firstly, to assist ongoing development work by gaining a better understanding of the domain, and, secondly, to prompt reflections on the OSM technique itself. We expected the modelling to be difficult, since ambulance dispatch is oriented towards procedures rather than conceptual structures, but were interested to see what it would yield. In practice, the analysis yielded important insights, about both OSM and the structure of emergency medical dispatch.

1.1 Central Ambulance Control at LAS

London Ambulance Service is one of the largest in the world, receiving an average of 3200 calls per day. The city is split into 7 sectors, each of which is managed by a sector controller with a team of support staff. Each sector team manages the work of about 35 ambulances, operating out of, on average, 10 ambulance stations. This is a much larger scale than any other ambulance service in the UK; for example, the metropolitan service described by Martin *et al* (1997) receives 'hundreds' of calls per day rather than thousands.

The control room is organised into two physically separate sections, for call taking and dispatch. In the call taking area, there is one call supervisor's desk and 24 call takers' positions. Six to eight of these positions are dedicated for doctors' urgent calls. There is little communication between call takers: each is essentially 'on their own', although backup services are available when needed. The dispatch area is primarily organised around the seven sector desks, which are physically distributed around the room to mirror the layout of the city. Additional desks

around the periphery of the room perform particular dedicated tasks such as managing the HEMS (Helicopter Emergency Medical Service) resources, dealing with vehicles that are off the road and providing the overall supervision for the room.

At each sector desk, there is a team of three or four people: the sector controller, who usually performs the role of ‘allocator’, deciding what resources (ambulances and other vehicles or personnel) to send to each incident; the radio operator, who maintains radio contact with crews on the road; and one or two telephone dispatchers who communicate with ambulance stations, hospitals and others (e.g. the Metropolitan Police or members of the public) by telephone. Although radio operators and telephone dispatchers are widely regarded as “doing the same function” – that is, dispatching vehicles to incidents – the different technologies they use have a strong influence over their actual roles. In addition, each has specified additional duties; in particular, the radio operator acts as a ‘second in command’ to the sector controller. The existence of so many clearly defined roles contrasts with the Services described by both Martin *et al* (1997) and McCarthy *et al* (1997), where there are between one and three main roles.

To set the context for this work and describe the basic functionality of the system being analysed, we present a brief history of developments within LAS.

1.2 LAS: A Brief History of System Development

Although London Ambulance Service has moved towards automating many aspects of operations, it has not progressed as far in this direction as many other Services. LAS achieved a certain notoriety when one attempt at computerisation failed in 1992, resulting in a blaze of negative publicity. While that failure is the most well known, it was not the first.

There have been many analyses of the failure (Page *et al.*, 1993; Benyon-Davies, 1995; Finkelstein & Dowell, 1996), and those findings will not be repeated here. Following the major failure, LAS reverted to a manual (paper-based) system, used until 1996. It then adopted a new approach to computerisation, based on evolutionary design and implementation, involving users within the design lifecycle and employing a very small team of developers in-house. LAS went on to win an award for their subsequent slower introduction of a computerised system for call-taking (British Computer Society, 1997; Fitzgerald, 2000), in which call takers typed up details as the calls were being taken. As soon as the relevant sector could be identified (matching the address that had been entered against a computer-based gazetteer), the information was transmitted to the sector desk so that the allocator could start planning what resources to send. When the call was complete, the call details were automatically printed out on a ‘ticket’ at the sector desk. The allocator acknowledged receipt of the printed ticket, then whoever had dispatched a vehicle to the incident acknowledged that dispatch. At that point, the computerised ‘ticket’ would disappear from the display, although it could be recalled by entering the Computer Aided Dispatch (CAD) Number corresponding to the call. The sector staff could view various screens of information, of which the two most important were an overview screen listing all outstanding jobs (except non-emergency calls for which the scheduled time of arrival was more than an hour ahead) and a screen giving details of one call.

Further management of the call was done manually, using the paper tickets. The management of tickets centres around the 'activation file', or 'allocator's box'. This is a slotted metal box with each slot corresponding to a vehicle in the sector. The ticket assigned to a vehicle, representing the job to which it is currently assigned, is kept in the relevant slot. The ticket faces forward while the vehicle is on the call, and is reversed when the vehicle is returning to station but 'free'. The box sits between the allocator and radio operator, where either may easily access it.

Various new developments have recently been introduced, and more are ongoing. Some of the most significant changes have been as follows.

Printing on station: Until early 2001, all information about a call was communicated by telephone to the crew on station. Allocators can now send tickets to be printed at the station, and the telephone dispatcher can simply call to alert the crew to the arrival of a ticket and confirm which crew will be attending the call.

Automating prioritising of calls: Until Spring 2001, call takers followed a flip card-based manual procedure for taking call information. This is known as AMPDS: Advanced Medical Priority Dispatch System. Then ProQA, a computerised version of AMPDS, was installed. This included automatic coding of determinants, codes that categorise the incident type in detail. The installation was done in-house, integrating the system with data from the existing CAD system. Allocators can now access additional information about each incident, and the system bars allocators from dispatch acknowledging a job until they confirm that the ProQA information has been given to the crew.

Mapping: Automatic Vehicle Location System (AVLS) equipment has been fitted in all vehicles. In November 2001, the system was extended to present graphical information about the locations of vehicles and incidents on sector desks.

2. Method

Data collected in LAS is being analysed in various ways (e.g. Wong & Blandford, 2001), including the OSM analysis presented here. The data was collected from staff in LAS over an extended period of time, from August 2000 to September 2001. Consequently, the data covers a period when substantial system changes were introduced; however, it does not cover the introduction of the mapping system, which is the most substantive development in terms of information presented to sector teams. Data collection focused on the roles of staff working in the control room, but also included others:

- Early on, several members of senior management were informally interviewed, to establish their perceptions of key issues in system development. This helped to establish the context for the work.
- Internal LAS documents, including an internal audit report that analysed the roles of the various groups of staff within CAC, were studied. This helped identify the 'front', or official, version (Goffman, 1959) of what staff were expected to do and how their performance was assessed.
- Thirteen 1-hour interviews were conducted using the Critical Decision Method, in which staff recall details of a particularly memorable past incident (Wong & Blandford, 2001). Participants were 6 allocators, 5 radio operators

and 2 telephone dispatchers. This provided data on major incidents (the kind that are, fortunately, rare) and staff's accounts of how they deal with them.

- 16 observation sessions with individual members of staff were conducted in a style similar to Contextual Inquiry interviews (Beyer & Holtzblatt, 1998), such that staff were observed working, and asked about features of that work when they had quiet moments between activities. Each session lasted between one and two hours. 5 call takers, 4 allocators, 4 radio operators and 3 dispatchers were observed. This provided data on routine operations; no major incidents were observed.
- Additional members of staff were interviewed and observed to help set the context of the work; this included a two-hour interview with a key member of the system development team and a day out with a crew.
- At three points during the study, our findings to date were formally presented to staff representatives and management, both to inform LAS of our results and also to verify those findings (by soliciting feedback on any inaccuracies that were detected).

The OSM analysis has been conducted by working through data (notably LAS documentation and transcripts of interviews and observations), identifying key concepts – entities and attributes – that appear in multiple sources or that are clearly important (according to the context in which they appear), plus relationships between those concepts and actions that change the state of the system. For simplicity, we model the system as if it were static, incorporating all the developments outlined above *except* the introduction of mapping.

3. Analysis

We structure the OSM analysis around the different roles played within CAC. However, first we outline some of the core components of a high-level system model. That is: there is a model of the system that a member of the general public would not be aware of, but that is a core part of the domain model for all workers in CAC. This model is very similar to that of systems implemented by other ambulance services (e.g. Martin et al, 1997). Space does not permit a full presentation of the model, but the core entities we can distinguish are as follows:

- Call types: there are three types of call, namely emergency ('999'), urgent (typically requested by doctors, for the patient to be transferred within the next two hours) and 'white', or 'Patient Transport System' calls, which are generally dealt with separately.
- All emergency calls have an AMPDS level allocated, representing the priority of the call. These levels are generally referred to as 'red', 'amber' and 'green'.
- Call tickets for emergency and urgent calls are all indexed by Computer Aided Dispatch Numbers (CAD Nos).
- The city is divided into physical and logical sectors, as described above.
- 'Virtual ambulance' identifiers relate a physical vehicle (with a fleet number) to a call sign which identifies both the station out of which it works (e.g. C5) and vehicle number (e.g. 01). The call sign generally identifies a particular crew for any given shift.

More details of the domain model are discussed below, as we present details of each role. The computer-based implementation (the CAD system) includes:

- Electronic tickets, which may be shown, listed, accessible by any individual, and may be complete or incomplete.
- The details of individual calls, as discussed below.
- Pending call list (shown on the allocator's screen).
- Information on nearest ambulance stations and hospitals.

There are constraints on this system: in particular, only one electronic ticket can be viewed at one position at one time; only certain roles can access or update particular information; information can only be updated on the electronic tickets while it is incomplete, and only by call takers.

For illustrative purposes, we present detailed extracts from the OSM model for the call taking role. Space does not permit full presentation of the model for other roles, so we discuss the most important points that emerged from modelling.

3.1 The Model of the Call Taker

The role of the call taker is primarily to take details of emergency calls and enter them in the computer-based form (or ticket). Questions are asked according to the AMPDS protocol, which is based on the chief complaint. Of the four roles considered, that of call taker is the least complicated. The call and its attributes are summarised in Table 1, and the corresponding device concept (also called a 'call' by staff, but referred to here as a 'call record') is outlined in Table 2.

Entity	Notes
Call	The call is a domain object, which is not 'known' to the computer system.
Attribute	Notes
Incident or query?	Incidents have call records; queries have to be dealt with on a one-by-one basis. The most common queries are about estimated time of arrival (ETA) of an ambulance.
Chief complaint	The chief complaint determines what questions are asked. For some calls, the call taker stays on the phone and gives pre-defined medical advice on care of the patient.
Caller features	Calls from the police or a doctor's surgery are sometimes treated slightly differently from calls from the public.
Instructions to caller	May include pre-determined care information for the patient until the ambulance arrives, or instructions to prepare something (e.g. note of GP) for the crew.

Table 1: the core attributes of a call (as a domain concept)

The location of an incident is important; this may include information on the major road from which a minor road leads, or detailed landmarks for a call from a large place (e.g. Heathrow airport). There are two reasons for this focus on location. The first is that the computer-based gazetteer needs to be able to match the address to a map reference which determines the sector to which the call information is sent and which is also a key component of the directions given to crews. The second is that call takers seem to form a picture of the situation in order to get enough information for the crew to be sure of locating the incident quickly.

AMPDS is central to the way call takers work. Once they have identified the location to their satisfaction, they work through the chief complaint and other information in the prescribed order, asking only questions they expect the caller to be able to answer (depending on whether or not the caller was on scene at the time of the incident and how well (s)he knows the patient(s)). Call takers are aware that

the determinant derived from AMPDS (e.g. 27D1S = serious stabbing) is matched by the system to ratings (e.g. 'red 3'), but many of them are unaware of how this information is subsequently used: "There's only red, amber and green. I don't think they take much notice of the other numbers, do they?"

Call takers are generally aware of other key domain concepts, such as tickets and sectors, but these have little impact on their work.

Entity	Add by	Delete by	Notes
Call record	Opening new record	Not possible	Call records are created for calls reporting (or giving up-dated information about) incidents.
Attribute	Set by	Change by	Notes
CAD Number	Set automatically by system	Not possible	Each call for the day has a unique CAD number. Numbers start from 1 at midnight.
Caller's number	May be transmitted automatically by exchange or entered manually	Overtime	
Caller's sex	Enter M / F	Overtime	
Caller's relationship to patient	Type in	Overtime	May be unrelated; may even be elsewhere – e.g. a care service.
Address information (of incident)	Typed in	Overtime	If the gazetteer does not recognise the information, there is extended dialogue between caller and call taker to re-formulate the address.
Map reference	Automatically generated by gazetteer when address is recognised	Not possible	See comment above. Many staff are dependent on the system to provide this information.
Phone number for incident	Typed in	Overtime	Only if different from caller's number.
Patient's age	Typed in	Overtime	Caller may not know accurately.
Patient's sex	Enter M / F	Overtime	
Chief complaint	Selected from list	Re-select	Initiates AMPDS protocol: determines which further questions are asked (there are 32 possibilities).
Conscious? Breathing? Alert?	Select y / n / unknown	Re-select	Call takers sometimes have low confidence in the response as terms are hard to understand.
Determinant	Select from series of lists	Re-select	Questions are pre-determined, depending on chief complaint.
Access information	Type in	Overtime	Call taker tries to ensure crew have good information for reaching the patient.
Safety information	Type in	Overtime	There may be reason to have concern for the safety of a crew.
Call category	Set by system based on determinant	Not possible	Of little relevance to call takers.

Table 2: the core concepts of a call record (as entered into the CAD system)

There are no essential constraints or relationships that impact on the work of call takers. Because the CAD system is tailored to the call taking function, and much of the call taker's job is determined by the system, there is a good fit between system and role.

3.2 The Model of the Telephone Dispatcher

The primary role of the dispatcher is to dispatch crews from stations to calls at the request of the allocator. A secondary task specified in the Dispatch Audit Report is to maintain accurate vehicle availability records.

There are clear procedures to follow in giving out call information. When the dispatcher is asked to send a vehicle to a job, the job is referred to simply by CAD number, and the information is taken from the corresponding ticket. The ticket is usually on the screen, rather than on paper. The dispatcher is responsible for noting which crew from a station has responded to a call, and whether or not there is a paramedic in the crew. There is a good fit between the system and the task of giving out call information, but a less good one for the task of recording which crew has gone, because the dispatcher cannot enter information on the electronic ticket, and must note it on a separate form until the paper ticket appears.

One important attribute of call information for the dispatcher is how much of it has been communicated to a crew before they have left the station (if the call taker has not completed the call, which is common). The dispatcher has to mark on the paper ticket any information not yet given out, so that the radio operator can update the crew before they arrive on scene.

Dispatchers occasionally check the computerised list of nearest stations to check that they have heard the instructions correctly (if the station is nowhere near the incident, they may have misheard either the station name or the CAD number). Apart from this, the system is used primarily to retrieve electronic tickets before they have been printed in order to confirm the details to the crew.

As well as the tasks for which there are clear procedures, dispatchers are involved in various other tasks involving the telephone. These include:

- Taking calls from crews about various matters such as broken down vehicles.
- Calling hospitals to try and locate any crew that has 'gone missing', or to alert them to a 'blue call' (a very serious injury) that will require a trauma team to meet the ambulance when it arrives.
- Calling back to the address from which a 999 call had originated to get more precise directions on how to find the building.

In all of these cases, the dispatcher's role is to deliver a message, with content determined by someone else, to some destination. In order to do this effectively, dispatchers need to understand enough about the jobs of other team members to be able to make sense of sometimes cryptic messages. This role also involves some running around, as messages are typically recorded on bits of paper that need to be delivered to the destination: this role is not yet well supported by the system.

Some dispatchers try to maintain a similar kind of overview of workload to allocators by typing in a succession of CAD numbers and seeing which upcoming jobs are likely to be for their sector: "Some people like to scan through them and... I don't really worry about that." This task, while not core (it is done only at quiet times), is also not well supported by the system at present.

3.3 The Model of the Radio Operator

The primary role of the radio operator is to keep crews fully informed of all available information relating to calls at the request of the allocator. A secondary task specified in the Dispatch Audit Report is to maintain accurate activation files. The radio operators also receive calls from crews on the road.

The radio operator has to be aware of device concepts related to the communication medium being used – e.g., radio channels, vehicle fleet numbers and call signs, all directly represented on the radio operator’s workstation: there is a good fit between this aspect of the role and the device used to support it.

One key concept the radio operator has to work with is the queue of pending calls, which is represented in two places: calls from crews appear as buttons on the screen, while calls requested by the allocator are represented by sheets of paper on the desk. At times of low workload, the existence of two streams to be integrated is unproblematic, but when workload is high, prioritisation across two sources is difficult. One extreme example is the beginning of a major incident: because radio calls are heard by all vehicles in the sector, during a major incident “they’ll hear and go”, and radio in, and CAC have to keep track of which crews have gone.

An additional challenge when workload is high is that crews have a concept of the radio operator as being busy or free: the only way for crews to assess this status is by the messages on the radio, so they may consider the radio operator to be free when in fact (s)he is busy in some other way – e.g. completing paper work. Radio operators may signal their status to crews by announcing “stand by unless priority”. Workload problems are likely to be alleviated as more vehicles are equipped with data links, reducing the time needed for voice communications.

3.4 The Model of the Allocator

The obvious and most recognised role of the allocator is to allocate resources to incidents within their sector, aiming to keep within government-defined time targets. In practice, this entails many additional tasks, including:

- Maintaining adequate coverage within all areas in the sector;
- Catering for the needs of crews – allocating meal breaks, organising appropriate support after crews have dealt with a difficult case (e.g. fatal road traffic accident), etc.;
- Keeping track of the locations of all crews to ensure their safety and track their availability;
- Liaising with teams in other sectors and other services (fire brigade, police, etc.) to deliver the required service.

Resources need to be appropriate to the nature of the incident. Because the allocator’s role involves rapid real-time decision-making (Klein, 1998), it is the one that is least well supported by the current system. Here, we outline some of the more substantial misfits between the allocator’s conception of the task and the device model. The current OSM model for the allocator is 11 pages long, and therefore not included here. As with all other roles, the process of constructing the model involved working through transcripts and other allocator-related documentation, identifying core domain and device concepts and relationships between them, plus any misfits between domain and device.

3.4.1 Relationship between incidents and tickets

Of necessity, information about incidents is received – at least initially – through calls. For reference purposes, each call is allocated a CAD number, as described above. With the increasing use of mobile phones, the number of calls being received is rising rapidly: according to one member of staff, “the average six years ago was something like 2000 calls a day, maybe 2200.” Now it is over 3000, and expected to rise further. It is believed that the numbers of both incidents and calls per incident are rising.

In the terminology of CAC staff, both calls and incidents are widely referred to as “calls”. For example, one radio operator explained about “two calls which we linked to the same call ... I’ve had 7 or 8 bits of paper for one call”. The action of deciding whether or not multiple calls refer to the same incident often involves allocators recognising superficial similarities (incident location, description, etc.) and then placing tickets side by side, maybe even showing them to a colleague, to judge the probability of them referring to the same incident. The current system provides no support for identifying duplicate calls, although it is planned as part of future developments, as described by one of the developers (here, ‘RTA’ is a Road Traffic Accident):

“Since we’ve started using ProQA, or having determinants, there’s always been the issue of what might be a possible duplicate call. Time is obviously one factor. Distance is another. The condition, or complaint, should be a key as well. But before we had determinants, all that was entered was free text, so there was no mechanism for doing a comparison. Now that we’ve got determinants what I’m going to do is effectively to get the users to provide me with a matrix of what are possible duplicates. Because a 29D2 might be a 29D5: an RTA vehicle versus pedestrian as against vehicle versus vehicle. [...] A snake bite probably is not likely to be a duplicate of an RTA. [...] We’ll never make the decision automatically.”

Allocators currently gather together information on an incident by grouping the paper tickets of all calls corresponding to that incident. Where information conflicts, they work with the best information available.

The current computer system does not support ticket grouping. Also, when there are multiple calls for one incident, allocators cannot clear them off their screen (until each one has been print- and dispatch-acknowledged) in order to distinguish them from calls for different incidents. One allocator, referring to the start of a major incident, described it as follows:

“Whenever something like that happens, the first call...obviously the first call is gonna come up on the screen, and you see train crash, and you first think to yourself, is it real, is it not? Then you flick back to the first screen, and the screen’s full with calls, so maybe you think, this is a real one. [...] Well when they are on my screen, I press F1 to view the calls [...] and I’ll view it and I go “That’s the same”,. Flick on through it, “That’s the same”, ’cos they have all got a train crash or the same location, but in all of that I might have had someone in a house that has fallen down the stairs, that’s why it is very important to check through and make sure that the tickets are all the same.”

This is an example of a misfit – in this case, between the call and the incident. The computer system only supports processing in terms of calls, and does not support either the identification of calls that relate to the same incident or the kind of grouping of information that is done using the paper tickets.

3.4.2 The “story” of one incident

Just as multiple calls may (unwittingly) give information about the same incident, so successive calls may make reference to the same event. For example, a patient’s representative may ring in again with further information, or querying the estimated time of arrival of an ambulance. Matching items to one stream of activity is poorly supported at the moment, relying largely on the memory of the allocator.

3.4.3 Stack of outstanding job sheets

Tickets for outstanding urgent jobs are arranged on the desk in front of the allocator. Paper tickets appear to be used in preference to the computer-based tickets because they include all outstanding jobs, they can be ordered by time, and the most important information can be made permanently visible by writing it at the top. As Mackay (2000) notes in a different context, the physical actions of writing information at the top of the ticket and of manipulating tickets both appear to contribute to their effective use in supporting planning behaviour. The (partially ordered) stack of outstanding jobs is not well supported by the computer system implementation. Put another way, the allocators process emergency and urgent calls in substantially different ways, and this difference is not fully implemented within the current computer system.

3.4.4 A ticket for a vehicle

If multiple vehicles are needed for one incident, multiple copies of one ticket are printed. One ticket is then allocated to each vehicle ‘slot’ as it is dispatched. The location of tickets in slots serves to relate the vehicle to the incident. The act of printing out duplicate tickets helps the allocator to plan how many vehicles to send, and to track how many have been dispatched. Paper tickets are a ‘mediating representation’ that support allocators in their primary task of dispatching vehicles to incidents. In particular, if vehicles from another sector are being sent, tickets are physically handed from one allocator to the other, signifying a transfer of responsibility. In OSM terms, domain relationships are represented through the system as follows:

- Incidents occur in sectors, represented within the CAD system and revealed when a computerised ticket appears at a particular sector desk;
- Vehicles ‘belong’ to sectors, represented in the slotted box at the sector desk and also at the radio operator’s workstation;
- Vehicles are sent to incidents, represented by the location of the paper ticket.

One consequence of the current implementation, with ‘hard’ sector boundaries, is that these cross-sector transfers cannot be easily organised until the paper ticket has been printed.

3.4.5 Incidents and stations

As noted above, information on the nearest stations (and also nearest hospitals) is provided by the computer system. We have not observed this being used ‘in earnest’ by allocators. They appear to rely more on their knowledge of the area and of traffic patterns in the area – as well as availability of vehicles at each station – when identifying the most appropriate station from which to send an ambulance.

Because the computer system does not have information about the locations of vehicles – for example whether there are any currently at a particular station – the information about nearest stations appears to be of little value to allocators although, as noted above, it is sometimes used by dispatchers.

One of the allocators, describing a major incident, discussed details of traffic movement between a station and the incident:

“Traffic’s gonna be mayhem around the stations, even though you gotta send all them ambulances there, you know, you gotta start thinking of the time of day as well, traffic, you know, what’s the best ways in and things like that.”

The knowledge allocators are using when managing the allocation of resources to incidents is sophisticated, exploiting past experience, and cannot be neatly packaged in terms of easily identified parameters such as ‘nearness’, particularly in a complex and busy environment such as London.

3.4.6 Vehicles and crews

Allocators appear to think largely in terms of unit of resources. Information about a particular crew will come in when there is an exceptional situation to handle, but generally the vehicle and crew are considered together. So, for example, an individual becoming unavailable at some time is thought of in terms of the impact that has on resources. Vehicles are also thought of in terms of their crew type – for example “hotel crew” (i.e. paramedic crew), first responder unit, etc. – particularly in terms of their appropriateness to a particular incident. However, for ‘standard’ requirements, the allocator will simply request a vehicle from a particular station, or give permission for any crew from a station to have a meal break, delegating the detailed allocation to the station itself and simply noting what choice is made.

Allocators keep a paper record of crew shift times, meal times, etc. They use work-arounds like marking vehicles as unavailable if a crew needs time to recover after a difficult job. The current computer-based system provides no support for crew management.

3.4.7 Vehicle locations

Experienced allocators have developed the ability to track the probable locations of the vehicles under their command – certainly in situations where they plan to intercept a vehicle on its route back to station and send it elsewhere. For example, an allocator sending an ambulance on an urgent (i.e. not emergency) call when it was part-way back from a hospital to its station explained:

“... if I give him a SW2 call at Mayday, I could have another vehicle come up nearer. [...] So I let him get back a while into his own area, and then give him the call. Because between Mayday ... you’ve got Mayday and you’ve got Thornton Heath. He has to go through Thornton Heath, then he goes through Norbury, then he goes that way, then they go Lower Streatham, to Streatham Vale End, then they come up to Streatham Hill. So when they get up near that area ... it’s just by judgement, of knowing your areas.”

As noted above, mapping has recently been introduced, so that this information is readily available to allocators, but was not at the time of this study: at that time, this could have been viewed as a misfit between the computer system and the information needed. Whether allocators shift to using the external representation of

vehicle locations, or whether they will continue to work with their ability to track this information mentally, is a question for further research.

3.4.8 The allocator's mental picture and level of cover

Experienced allocators have a very detailed knowledge of their area. They know without looking how many vehicles are normally based at each station, where normal stand-by points are, where stations are, which hospitals normally have particular facilities. When they come on duty, they simply have to adapt their general knowledge to local variations – for example “that’ll be a vehicle down”. While on duty, they are updating their mental picture, or situation awareness (Wong and Blandford, 2001), with real-time information such as the locations and availability of particular vehicles. In maintaining their mental pictures of their sector, allocators ensure that all areas within the sector have adequate cover. Selection of ambulance(s) for a particular incident is made with this in mind. During a major incident that draws all vehicles from an area to one spot, other ambulances are moved in to provide cover. Again, the recently introduced mapping system should help allocators rapidly check that cover is being maintained.

Another aspect of the ‘big picture’ is keeping an overview of new incidents in the sector. Allocators flick rapidly between overview and detail screens to see the details of a call while also maintaining an overall picture of the level of demand within their sector. The current computer implementation does not indicate the arrival of a new call while the allocator is viewing the details of a call – this can only be seen when in overview mode. A revised implementation would ideally address these twin needs of needing overview and details in a seamless way.

4. Discussion and Conclusions

Modelling the LAS with OSM has successfully generated many insights into the quality of fit between each group of workers and the supporting computer systems. While in many cases a good fit was found, in some cases the fit was less good; for example we have discussed eight important features of the allocator’s task that are poorly supported by the current implementation. As noted above, however, the computer support system is still evolving, so it will be possible to address these issues as part of the ongoing development work.

Using OSM also allows comparison between different work groups, e.g. allocators and call takers. Not surprisingly, nearly all important concepts are used by the allocators; call takers share the smallest subset of those concepts, but the set that is best matched by the currently implemented computer system. The remaining two groups, telephone dispatchers and radio operators, gave us a surprise. Although they are described as ‘doing the same job’, these two groups share a relatively small set of key concepts with each other, though each shares a substantial set with the allocator. The existing computer implementation supports dispatchers and radio operators reasonably well, but supports allocators less well. Overall, the current system works well because the paper system is manipulated expertly by highly skilled individuals.

We believe it unlikely that these insights could have been obtained from analysis techniques based more directly on tasks, procedures and actions. In this

case, for roles other than call-taking, the lack of computer-system-related actions rendered the explicit description of actions largely fruitless. Although most of the work of ambulance dispatch consists of standard physical actions which can indeed be modelled using traditional task-centred approaches – making telephone calls, handing over pieces of paper, etc. – modelling these actions reveals little about how well the computer and other systems support the primary tasks of the staff. As in most complex work domains, the prescribed and implemented systems work well for standard tasks, but various workarounds have evolved to deal with unusual situations, so that *in practice* work flows efficiently, effectively and reliably.

The exercise has also increased our knowledge of OSM and its applicability. In the first place, we have learnt that the technique successfully scales up from our earlier domains of study (drawing tools and diary systems) to large, multi-person work systems. Moreover, the investment of effort was not unreasonably high. Whereas it would have taken weeks to produce a fully formal model of this domain, if it were possible at all, the OSM model was created in a few days. The approach of sketchy modelling does, subjectively, appear to yield useful results for reasonable effort.

We were particularly pleased to find that sub-systems of the LAS, reflecting the different roles, could successfully be separated out for analysis. Inevitably problems of boundary demarcation and inclusiveness arose: for example should we model roles within the organisation as part of the system? Should we model the HEMS desk or other specialist desks? We invariably opted for the simpler approach. However, as a sketchy approach to modelling, OSM would allow either decision to be taken.

Unlike the domains previously studied, ambulance dispatch contains plentiful many-to-many relationships between domain terms and concepts. As noted above, the term ‘call’ can refer to at least three different concepts: the actual telephone call, the electronic record of that call, or the incident to which a call refers. Similarly, ‘red’ is used to refer to a category of call (highest priority), to the colour of a ticket (which means ‘emergency’ rather than ‘urgent’ or ‘patient transport’), or to the underlining on a ticket to show information that has not yet been communicated to a crew (“it’s got red on it”). Conversely, the highest priority of call is often referred to as ‘category A’ (rather than ‘red’). Consequently, extracting concepts from verbal protocols was far from straightforward, and required validation with LAS staff. Once noted, this was readily accommodated in the OSM model.

Overall, our first large-scale exercise with OSM has proved very successful. Modelling the domain and device in terms of entities, attributes, actions and relationships has proved feasible, and has yielded a model that helped provide useful insights on the domain.

Ongoing work is focusing on refining the model, clarifying the approach to modelling, testing in other domains (with other groups of users and developers) and developing tool support for model building.

Acknowledgements

We are very grateful to all staff in CAC and elsewhere in LAS who gave time to contribute to this study, and particularly to Avril Hardy, Development Manager for

CAC, for comments on an earlier version of this paper. The work has been supported by grants from Middlesex University and EPSRC (GR/R39108).

References

- Benyon-Davies, P. (1995). Information systems 'failure': the case of the London Ambulance Service's Computer Aided Despatch project. *European Journal of Information Systems*, 4(3), 171-184.
- Blandford, A. E. & Green, T.R.G. (1997) OSM: an ontology-based approach to usability evaluation. *Proceedings of Workshop on Representations*, Queen Mary & Westfield College (July 1997).
- Blandford, A. & Green, T. (2001a) From tasks to conceptual structures: misfit analysis. In J. Vanderdonckt, A. Blandford & A. Derycke (Eds.) *IHM-HCI2001 Vol. II* (short paper). Pp. 113-116.
- Blandford, A. E. & Green, T. R. G. (2001b) Group and individual time management tools: what you get is not what you need. *Personal and Ubiquitous Computing*. Vol 5 No 4 Pages 213-230.
- British Computer Society. (1997). London Ambulance Service wins British Computer Society Excellence Award, www1.bcs.org/docs/01000/1097/pastwin.htm, access date 3 Sep 2001.
- Finkelstein, A. & Dowell, J. (1996). A Comedy of Errors: the London Ambulance Service case study, *Proceedings of the 8th International Workshop on Software Specification & Design IWSSD-8 (2-4)*: IEEE CS Press.
- Fitzgerald, G. (2000) *IT at the heart of Business*. The IS Management Series Vol. I. British Computer Society.
- Goffman, E. (1959) *The Presentation of Self in Everyday Life*, New York, Doubleday.
- Klein, G. A. (1999) *Sources of Power: How people make decisions*. Cambridge, MA: The MIT Press.
- Mackay, W. (1999) Is paper safer? The role of paper flight strips in air traffic control. *ACM Trans. CHI*. 6. 311-340.
- Martin, D., Bowers, J. & Wastell, D. (1997) The Interactional Affordances Of Technology: An Ethnography Of Human-Computer Interaction In An Ambulance Control Centre. In *Proceedings HCI'97*. Springer. 263 – 281.
- McCarthy, J. C., Wright, P. C., Healey, P., Dearden, A. & Harrison, M. D. (1997) Locating The Scene: The Particular And The General In Contexts Of Ambulance Control. In *Proc. GROUP'97*. 101 – 110.
- Page, D., Williams, P., & Boyd, D. (1993). *Report of the inquiry into the London Ambulance Service*. London: South West Thames Regional Health Authority.
- Vicente, K. (1999) *Cognitive Work Analysis*. Mahwah, NJ : Lawrence Erlbaum.
- Wong, B. L. W. & Blandford, A. E. (2001) Situation Awareness and its Implications for Human-Systems Interaction. In W. Smith, R. Thomas, & M. Apperley (Eds.), *Proceedings of the Australian Conference on Computer-Human Interaction OzCHI 2001*, 181-186. Perth: CHISIG, Ergonomics Society of Australia.
- Wright, P. C., Fields, R. E. & Harrison, M. D. (2000) Analysing Human-Computer Interaction as Distributed Cognition: The Resources Model. *HCI Journal*. 15.1. 1-41.