

Preprint. Final version appears as:

SMITH, P., BLANDFORD, A. & BACK, J. (2008) Questioning, exploring, narrating and playing in the control room to maintain system safety. *Cognition Technology and Work*. DOI 10.1007/s10111-008-0116-1

Questioning, exploring, narrating and playing in the control room to maintain system safety

Penn Smith, Ann Blandford & Jonathan Back

UCL Interaction Centre, University College London, Remax House,
31-32 Alfred Place, London WC1E 7DP, U.K.

Communicating author: Ann Blandford

A.Blandford@ucl.ac.uk

Tel.: +44 20 7679 5288

Fax: +44 20 7679 5295

Questioning, exploring, narrating and playing in the control room to maintain system safety

Penn Smith, Ann Blandford & Jonathan Back

UCL Interaction Centre, University College London, Remax House,
31-32 Alfred Place, London WC1E 7DP, U.K.

A.Blandford@ucl.ac.uk

+44 20 7679 5288

ABSTRACT

Systems whose design is primarily aimed at ensuring efficient, effective and safe working, such as control rooms, have traditionally been evaluated in terms of criteria that correspond directly to those values: functional correctness, time to complete tasks, etc. This paper reports on a study of control room working that identified other factors that contributed directly to overall system safety. These factors included the ability of staff to manage uncertainty, to learn in an exploratory way, to reflect on their actions, and to engage in problem-solving that has many of the hallmarks of playing puzzles which, in turn, supports exploratory learning. These factors, while currently difficult to measure or explicitly design for, must be recognized and valued in design.

Keywords

Human error, reflection on action, safety, problem-solving, control rooms.

INTRODUCTION

The work reported here started out as an investigation into Distributed Cognition (Hutchins, 1995; Hollan *et al.*, 2000) in control rooms. The intention was to study the system designs, working practices and cultures in different control rooms that all performed the same core function – namely to control parts of the London Underground rail network, to understand how the different system configurations all supported safe and effective working. To this end, observational studies were conducted in five control rooms, with the intention of comparing and contrasting the system designs (layouts, technologies used, job functions etc.), to reason about what made each set-up function effectively. As the study progressed, however, it became apparent that the commonalities across the five study settings were much greater than the contrasts, and that features of the settings that have not previously been reported were essential to the safe and effective operations of the systems. These features include managing uncertainty, exploratory learning, reflecting on actions, and engaging in play-like interactions that are discussed in detail below.

BACKGROUND

Railway control room activities and management have been described from a variety of perspectives. In parallel, a

substantial literature on safety and organizational systems has emerged. Here, we draw on these literatures to set the context for this study, focusing on socio-technical studies which are most relevant to the approach we have taken.

Understanding railway control room activities

Safety has been a concern in many railway control room studies. For example, Stearn *et al.* (2005) considered ergonomic issues in the design of a new control room; Pledger *et al.* (2005) present an account of how Human Factors considerations were formally accounted for in the design of a new control room; and Kauppi *et al.* (2006) present a simulator-based study of a novel train traffic planning tool. Of more direct relevance to the work reported here are the studies of Garbis (2002) and Heath & Luff (1992; Luff, & Heath, 2000), as discussed below.

Garbis (2002) studied a control room in Stockholm which is organized in a broadly similar way to the London control rooms that are the focus of this paper. Garbis focused on the roles of “cognitive artefacts” in supporting the coordination of work within the control room. In particular, he considered the artefacts used – the fixed line diagram, the timetable, a database (e.g. of maintenance work) and pen and paper – and the degrees to which they support situation assessment and mutual awareness within the control room. He highlighted the contrasts between the day shift (which represents intense but largely routine work) and the night shift (which involves non-routine work such as maintenance, but which is of much lower intensity). The day shift work described by Garbis has many features in common with the day shift work of LUL, and the value of openness of artifacts was also found to be important in this study.

Heath and Luff (1992) presented an ethnographic study of one London Underground control room (for the Bakerloo Line). They focused in particular on the relationship between two key members of staff in this setting: the controller and the divisional information assistant. They highlight the importance of these two individuals maintaining a good understanding of the other’s activities while also specializing in their own roles. For the information assistant, the key responsibility is communicating with staff and passengers at stations, while

for the controller it is managing the train timetable and reforming trains – an activity that we describe in more detail below. They highlight the importance of there being fluidity between the private actions of individuals and the public presentation of actions to help others maintain appropriate awareness of the situation to inform their own actions. Luff and Heath (2000) present a similar, but more detailed, analysis of a Docklands Light Railway control room.

The studies of Garbis (2002) and Heath and Luff (1992; Luff & Heath, 2000) provide a foundation for the work reported here: they apply similar, observation-based methodologies; Garbis focuses on the roles of artefacts in supporting communication and shared understanding; Heath and Luff consider the role of room layout and conduct a more detailed analysis of the interactions between controller and information assistant. Our study builds on these findings and extends them.

Maintaining safety

The focus of this paper is on how the control room design and behaviours ensure system safety. There is an extensive literature on human error and system safety (e.g. Reason, 1990; Hollnagel, 1998). As Hollnagel (2005) argues, much of the work starts from the premise that system safety is compromised by individual (“erroneous”) human actions, and this approach fails to account for the richness of human behaviour and the contexts within which people perform. Also, a high proportion of the literature focuses on why systems fail, rather than how they remain resilient in the light of unanticipated events of more or less seriousness. In this context, we use the word “resilient” to mean that mechanisms exist for the overall system to remain under control despite perturbations and unanticipated events (Hollnagel, Woods & Leveson, 2006).

Relative to this study, the work of Rochlin (1999) provides appropriate background. Rochlin draws on evidence from a range of studies of organizations in which safety is paramount to argue that safety is a constructed human concept, and that the evaluation of system safety cannot be reduced to the systematic evaluation of sources of error and risk. He identifies properties of organizations that contribute to safety as including:

- Learning: the organisation supports the learning of all individuals. Each individual learns according to his or her situation.
- Duality: teams within the organisation are able to simultaneously maintain multiple representations of the system state. While to the outsider these representations may appear contradictory, to the system actors they are complementary viewpoints that together support reasoning and safe operations. “Confidence in the equipment and the training of the crew does not diminish the need to remain alert for signs that a circumstance exists or is developing in which that confidence is erroneous or misplaced.” (Rochlin, 1999, p.1553).

- Communication: both formal and informal communications (including “chatter”) are recognised as being important for maintaining awareness of system state, and enabling experienced operators to pick up early signals of potentially problematic situations.
- Locus of responsibility: Rochlin discusses the contrast between organisations where individual “heroes” are valued and those where responsibility is shared. In the former, “hero stories” are used to support organisational learning, whereas in the latter cumulative knowledge is encoded in formal procedures.

Rochlin also notes that safe organizations typically value the reporting of error, and regard breakdowns as being the responsibility of the organization rather than the individual. He argues that a “collective commitment to safety is an institutionalized social construct”, noting that one means for maintaining a safety culture is through “rituals and stories that serve to orally transmit operational behaviour” (p.1556). After we present the findings of our study, we relate them to the resilience criteria presented by Rochlin.

Aims of this study

The aim of this study was to investigate aspects of London Underground control rooms, considering layout, communications and operating practices, that contribute to resilient performance. The initial focus was on understanding how differences in both physical structure and organisation of work might lead to different behaviours that achieve broadly the same outcomes (in terms of operating effective services on different train lines). However, using an inductive approach to data gathering and analysis, it rapidly became apparent that the commonalities in strategies and practices were much more significant than the differences, so in the tradition of Grounded Theory (e.g. Charmaz, 2006), subsequent data gathering and analysis further explored factors that contributed to system resilience.

METHOD

Five London Underground control rooms were selected for study so as to gain a broad picture of the different equipment and work practices used to manage the running of the trains. These centres vary in the age of the technology being used, the size of the teams, and the organizational arrangements. The five rooms were selected to represent a variety of system designs: control room design is an evolutionary process, whereby each new design incorporates and adapts design features from earlier system designs, so control rooms share common features while also differing on some important dimensions.

A day was spent in each of the control rooms (each covering a change of shift), observing work practices and technology use. Extensive notes were taken, including diagrams and photographs of equipment, artefacts, and the layout of the rooms. Immediately after each visit the notes were reviewed and further details and comments added

(clearly separated from the original text). These were then labelled with the themes that were emerging ('open coding' in the terms of Grounded Theory). This systematic process sensitised the observer to the emerging themes, and this developing understanding guided subsequent observations, making the observer alert to both commonalities and contrasts between the different sites. Once all the sites had been visited, summaries of each of the encounters were produced, allowing different aspects to become prominent: a similar process to using affinity diagrams (Beyer & Holtzblatt, 1998). Charmaz (2006) discusses the value of coding data at different levels of abstraction, enabling the analyst to engage with (or make sense of) the data in different ways; this process yielded a richer understanding of the emergent themes, but also highlighted further questions that needed to be probed.

To verify the findings and conclusions that were surfacing, a semi-structured interview protocol was devised to further probe the emergent themes. The interview questions sought further evidence both for and against the themes that were emerging from the data, as well as clarifications of some points. Follow-up interviews were carried out in three of the control rooms. These interviews took place at the team member's desk so as to obtain answers arising from the actual situation. The interviews were intended to take an hour each, but interruptions meant that they took longer. They were recorded, and subsequently analysed in relation to the themes already identified, noting both evidence that reinforced the themes and evidence that might have contradicted them, and also developing causal accounts and probing consequences. Direct quotations from these interviews are included in italics to illustrate findings; names have been changed to ensure anonymity for individuals. Once the analysis had been completed, the final report was reviewed by a senior member of staff in a control room that had not been part of the study; he confirmed that the analysis conforms to his experience and understanding and provided further examples of the kinds of behaviour reported here.

THE CONTROL ROOM CONTEXT

Here, we summarize key features of the study settings and operator practices, and the work of controlling the trains, to set the context for our findings.

Control room organisation

The control rooms studied varied in size, configuration and roles allocated to staff. The two oldest control rooms had a hierarchical division of work, reflected in the layout of the room, with signal operators sitting at the lowest level at their control panels facing a large diagram of the line mounted on the wall, and with the controllers and information assistants sitting on an upper level behind them, "like a submarine control deck". Within the more modern control rooms, the large line diagram is still present (e.g. Figure 1), but adjustment of the signals or train routes are controlled via a bank of monitors and computer

displays, set up around curved desks, using number pad, mouse and keyboard (see for example Figure 2). One such control room is described in detail by Heath and Luff (1992).

The operators control and gain their understanding of the system being controlled (i.e. the train system) via a diverse range of interfaces, including the line diagram, signal controls, radio, telephone, CCTV and timetable.

The line diagram shows the track layout (see Figure 1). The track is divided up into track blocks which light up when a train passes over them. The diagram also shows the state manage the movements of trains; under normal conditions, much of this control is determined by the train timetable; this is overridden whenever events demand it.

The radio broadcasts to all the drivers' cabs, while the phone's primary purpose is for safety critical instructions between controller and driver or signaller and driver. Within the rooms, radio calls can be listened to via a handset or, more usually, switched to the speaker so that others in the team can hear.

CCTV shows activity at selected points on the network – e.g. station platforms. The cameras are generally positioned to give a view of the driver's cab as well as the front portion of the platform. Only a limited number of views are available, and control room staff have no control over the views presented.

The timetable is programmed into the system and also available as a file of paper printouts, placed in plastic envelopes so that alterations can be marked on them in chinagraph pencil (Figure 3). These marks indicate where a train is to be withdrawn or switched to take the place of another train. Using solutions such as these, the controller attempts to keep the service running to the expected programmed schedule with a minimum of intervention, as described in more detail below.

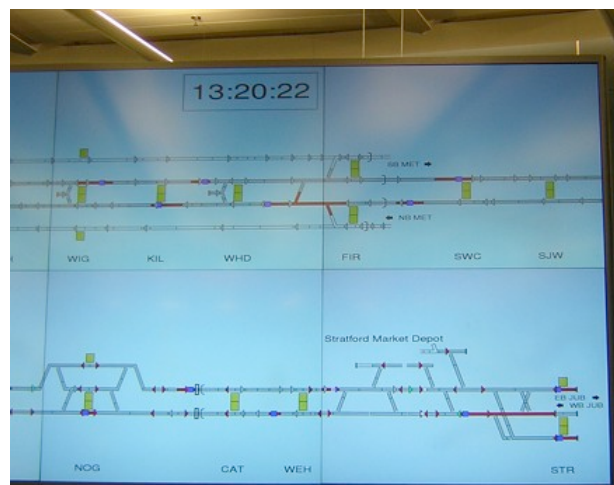


Figure 1. Example of line diagram (some sections are 'lit up' – e.g. by a train passing through that section).



Figure 2: a controller's desk

Events are recorded in a log book, which serves as both a record of incidents and actions, and also a 'to do' list of outstanding actions. The log book is used in the hand over procedure between shifts.

Other equipment provides information about both static and dynamic aspects of the system configuration, but details are not relevant to the present paper.

The London Underground control rooms' equipment may vary between the different rooms but the focus of all the teams' work remains the same – i.e. to keep the trains moving and ensure the train system is operating in a safe manner.

The work of controlling the trains

All the lines have some form of automation that can set up the routes and direct the trains according to a programmed schedule, and that ensures trains do not collide. Therefore, a large part of the control rooms' work is waiting. The teams



Figure 3. Example of the timetable with changes marked.

chat, read, watch TV, make tea, and provide explanations and demonstrations to trainees and visitors, all the while alert for the moment when they will have to intervene: *“All the time we're not doing anything very much, we're watching”*.

They are waiting for a trigger event. Sometimes this is a phone call – e.g. someone ringing with a question about when an engineering train can be allowed through, or a message to be passed on to a driver concerning a change of shift. In more serious situations, the caller will be providing information that needs further investigation and maybe a fast reaction: a report of smoke, a suspect package, a passenger emergency alarm pulled, a cable hanging down beside the line, etc. Automatic alarms also have to be attended to, ringing when a train is stationary for longer than two minutes or when a train identification code is contradicting the timetabled route. On two occasions during the study, the traction power lights lit up, warning of anomalies in the power supply to the tracks.

Often the team themselves are the people who first realize that something is happening, by noting on the line diagram that unexpected gaps are developing between the trains. A common reason for this is that a train has been slow leaving a station due to the number of passengers trying to get on or because of a delay in switching to a new driver. Mechanical problems, such as doors sticking, can also produce delays.

These gaps may also be an indication that a more serious incident is starting. The first judgment the team have to make is always what the effect will be on trains behind the problem situation. If the incident is going to last more than a few minutes then following trains need to be notified – the top priority being to avoid trains entering the area or sitting in tunnels (because of the effect it has on the passengers and because of the difficulty of getting them out if that becomes necessary).

If not dealt with promptly, even short delays can create a serious knock-on effect, especially during peak service: *“During the peak it is 2 1/2 minutes, you get a train with a sticky door for 2 minutes and you've got them backing up. At this time of day its almost 4 minutes and it gives you that second's breathing space”*.

This spacing out of the service is dealt with by slightly delaying the trains in front and behind. If the signaller needs to make a train wait, he or she can access a particular signal and put a hold on it, meaning it remains at red. Twice these procedures resulted in drivers being forgotten and radioing in asking why they were having to wait so long. E.g.:

Driver: Wonder if you're going to hold here for much longer for regulating purposes? I think that customers are getting a bit err

Signaller 1: Yeah, tell him to go. I held him there, [expletive deleted] yeah, tell him to go...

Signaller 2: If you have clear signals, you're okay to proceed ... I didn't know you'd held him.

Signaller 1: Yeah sorry mate, my fault.

This interaction also illustrates an important point to which we return later: that staff are aware that human error is always a possibility and that they should be able to admit to their mistakes.

If delays are longer, or sections of the line have to be closed, then the routing of trains is adjusted. To ensure that a reasonable service can be maintained in unaffected areas, controllers use a variety of methods: holding trains at signals to control their timings; looping the trains at reversal points; thinning the service by withdrawing trains into sidings and depots; bringing out spare trains to fill gaps; and reforming trains. Reforming trains means changing the physical train's logical identifier so that it occupies a different slot in the timetable. One important factor in deciding how to reform trains is the locations and schedules of drivers: the paper timetables show each driver's pick up time and place coloured in by hand for each trip (the crew and train databases are incompatible so cannot be combined automatically). Trips which are the drivers' final duties are also highlighted to remind the controller of which trips are less likely to be a problem if they run late because they will not have a cumulative effect by delaying the next picked up train.

If a situation has become too complex for minor adjustments to be effective, the schedule is switched to emergency service. This is where a number of trains are withdrawn and the remainder are given emergency numbers not listed in the timetable and the service is controlled manually.

If a train is defective, or if it has not been possible to get a driver to the pick-up point, the train may be withdrawn. If this creates an unbalanced service a train intended for another route is switched to take its place, by reforming the train (i.e. changing its ID).

When there are many trains out of schedule the ideal method of dealing with them is as a pair. This is where, if two trains are running late so that the front train is actually traveling in the time slot of the following train, the IDs of the trains are swapped. This produces one train running on time, leaving only one much delayed train to be dealt with by short-tripping or withdrawal. For example:

"205 Mate, make it onto 240, 240 to 205 – I'm seeing that there's an on-time [station P], there's a late running train so I'm making it up on time. And the other one, which is here, I'm making it up afterwards and then turning it short at [station Q]".

The controller has one, usually two, backup plans prepared and will closely monitor what is now happening with the trains.

Throughout any reorganisation, the controller and information assistant disseminate information. External organizations such as the bomb squad or fire service sometimes need to be called, pick-up drivers have to be sent to different platforms, technicians have to be dispatched, and the public has to be continually informed of every change.

Finally, once the incident has been dealt with, the team will bring trains back into service and adjust the train identifiers to match the timetable, returning to timetable as quickly as possible. And then the logbook and the fault and delay sheet can be filled in.

RESULTS

What stood out across the different rooms was the constant assessing of what was actually happening outside the control room, in the world of the drivers and trains, and the deep concentration that the controllers brought to their work of re-organizing the train running order. Three key themes emerged in the study: the way operators deal with uncertainty and establish appropriate trust in their technology; the ways new operators are trained; and the particular mode of working while reforming trains.

Uncertainty and Trust

The control room staff are well aware that the interface with the external world that the control room provides may not accurately represent what is actually happening. Occasionally this is due to faults in the control room technology but mainly it is a realistic acknowledgement of the difficulties of interacting with a train system comprising temperamental machines and unpredictable humans, plus an awareness that unless the interaction is by direct manipulation the action and the reaction will always be at one remove from each other,

"It's just a representation, you're not getting a true picture of what's actually going on outside. Not like in a signal cabin, where it is hardwired so you know exactly what you do happens out on the track, whereas here it goes through a computer then through an IMR, and so on. This is just a reflection of what's happened".

This understanding of the system comes from training and experience; as one operator explained: *"We're paid for what we know, not what we do"*.

One feature that gives rise to uncertainty is that elements of the interface equipment have to be interpreted. When asked to explain the equipment, the operators were regularly translating what they were seeing and making judgments on what this might mean; there were known problems which could occur if elements of the equipment were taken at face value. Line diagrams give a good example of this. The line is in sections with a current running through each of the two rails within that track section. When a train is on a particular track section, the circuit is completed by the connection the train creates between the left and right rails;

this information is transmitted to the control room where it lights up the corresponding track section on the line diagram. However, the circuit can be completed by other events, such as faulty wiring or flooding, so although the line diagram is being monitored for the positions of trains, the lighting up of track sections does not necessarily show the presence of a train.

As all the wheels of a train make this connection, when a train is moving across two sections of track, both track sections light up on the diagram, and if trains are travelling close together this can mean a long string of tracks are lit up:

“That’s three trains – they breach tracks as they go along, so you see two tracks per train”.

Equipment can seem to be controlling an external object when it is not; for example, setting a signal to remain at red (putting a hold on it) does not ensure it will not switch to green. An operator explained this:

“With computers, you can’t trust them, they’re not completely fail safe. Got two computers running the railways, one running the railway and one doing everything in parallel. So if the main one goes down the other steps in and takes over, while doing this it releases the hold, the signal will clear. We can’t guarantee holds on this system”.

Sometimes the confusion is not caused by the interfacing equipment in the control room giving a misleading picture, but because no feedback system can be provided for all that might take place in the external system. For example, when a driver sees a green signal he or she does not necessarily move forward. The rule is that only the person on the ground can decide if it is safe to go, and they may decide not to for many reasons – e.g. objects or people on the track, problems on the train, or overhearing an incident being reported by a driver in front.

Another source of uncertainty is human behaviour. This may be behaviour of London Underground staff, as in the following example (here, DMT is the Duty Manager for Trains):

“He rang wanting drivers at [station X], told them what trains to pick up and make into their trains. They picked up the wrong trains so we’re seeing them as one thing on the track – we put in the number of what the train is into the timetable. They get on the radio and we see their train number come up and then we’ll start knowing something’s wrong ‘cause we’ll be thinking they’re at [station Y] when they’re up at [station W]. Then we’re saying to the DMT at [station Z], ‘course you got the driver he came in on 233, but it isn’t, it’s 245 that come in. It’s easily done”.

Passenger behaviour is also a source of unpredictability; for example:

“Just carried a gentleman into the sidings. I’ve got an over carry, Al. At [station X]. They detrain there and obviously all the passengers are supposed to be taken off. But he

wasn’t taken off. You wouldn’t know till you were told, basically”.

Events elsewhere in the rail network can also have an effect. Within the same line some parts will be under the control of signallers looking after other lines that overlap, and whose priority is their own trains, with the result that it is not possible to predict the running of the trains; the operators just presume the worst scenario:

“It’s seven minutes late when it comes out of the tunnel and it will slow up the next [train from other operator], so they’ll hold it so [train from other operator] goes first. So it’s even more late: sixteen minutes”.

Altogether these uncertainties about the external world means the job of an operator requires them not only to control the trains, but also to be continually making judgments about the true state of the system, not relying just on what the interface is telling them.

London Underground control room staff’s trust in their interfaces is appropriately low. There is an awareness that no equipment provides a continuously accurate interface to the external world: the picture it is showing may be misleading, and messages sent out from the control room equipment may not be acted on as expected. Understanding of the state of the system is achieved from a combination of room culture, the design of equipment and work procedures that have developed over time.

It is important that the control room teams remember the limits of their systems as the consequences of mistakes run from inconvenience (trains sitting in tunnels, gaps in the service) to serious safety issues (train collisions, people injured on the line, fire). The constant awareness of these limits is ensured by the culture within the rooms which is first encountered during training.

The teams are primarily made up of staff with previous experience in the external world of trains (signallers, drivers, station managers) and this broader knowledge improves understanding of the external world they are controlling,

“Being an ex driver you’d know that, you see, come from an engineering background, Ted’s from a signalling background – we all come from different places in the railway, builds up the building blocks to make a good control room”.

This awareness of the limits of an interface has resulted in some good design solutions. The teams’ need to constantly monitor the situation is dealt with using the large line diagrams. These are well placed so as to be seen from any point in the room and have information from which to infer the present train behaviour (although, as discussed above, they show the presence of powered tracks rather than, necessarily, the presence of trains). As discussed by Heath and Luff (1992), this, together with their knowledge of the timetable and their ongoing situation awareness, enables the

operator to maintain a picture of the real world objects and quickly diagnose which of a number of causes might be creating the resultant patterns.

The monitoring task is also supported by the radio, phone and alarms being given distinctive ring tones, and a speaker being provided for the radio. The speaker is provided so team members can listen in to what the controller is arranging with the drivers, and so synchronize their own activities with the controller, but it also makes double checking of his decisions possible. This eavesdropping on each other is assisted by the layout of the modern control rooms (the old control room design does not enable the signallers to see the controllers or the line information assistants).

The team are able to cross check the data they are receiving. The main method is by speaking directly to a driver at the scene via the radio. All the control rooms have some CCTV equipment, and this can give a fast response without the need for explanations to the staff member contacted or interrupting their work, though unfortunately it is not always well positioned. The CCTV is used as a diagnostic tool when needing to find out more about particular situations on the platform or checking what is happening in a driver's cab when a train is stationary at a platform.

In acknowledgement that safe operation of the uncertain world of the trains sometimes requires a duplication of practices, many 'belt and braces' practices have developed. One example is that, as noted above, the operator never tells a driver to go forward past a signal at danger (known as 'apply the rule'): they can only give them permission and the driver must judge the situation for themselves. These safety processes have often been developed through learning from earlier accidents.

This theme of uncertainty and trust has outlined some of the challenges operators face in keeping the trains running safely and effectively. We have identified some of the organisational practices and individual strategies that have been developed to cope with the limited trust operators have in their equipment. Many of these strategies are acquired both through initial training and through working as a member of a control room team.

Training new team members

All trainees receive a twelve week introductory course but they then return to their own control room where, even if the line has a simulator, they are expected to work alongside experienced staff on the active desks. This in-situ training is of different durations in different control rooms. Experienced staff not only provide ongoing explanations of what is happening, but also alter the settings of equipment to illustrate their explanations and hand the system over to the trainee. Many trainees drop out before the end; this was explained as being because they are unable to cope with the necessary concentrated multitasking work (a point also made by Heath and Luff (1992)).

All the control rooms observed had trainees present, shadowing the work, and many visitors, such as drivers and station managers, passed through. The team obviously felt a sense of ownership of the equipment and were happy to demonstrate setting a signal to hold, or switching the setting of the programme machines at the junctions. They sometimes drew other people into the exercise: for example, one called up a driver and asked him to send through an emergency call to demonstrate the differing alert tones.

The staff were adamant that the skills and knowledge required in a control room cannot be taught, but have to be gained by experience in the actual situation. They would not even accept that the simulator could provide this. They shared many stories with trainees; these typically emphasized the importance of vigilance, team interaction, and understanding that the equipment could not guarantee to match the real external train system – elements that are more easily communicated in context than through formalized training.

During the study the researcher was often told that people can only learn by being given a level of responsibility that may result in them making mistakes. This is the way everyone had been taught themselves and they claimed the lessons that stuck came from where mistakes had been made. This meant that even if a trainee suggested an incorrect action they would be allowed to carry that action out (provided that it would only impact on efficiency, and not on system safety). One instance that caused a 5 minute delay in a peak time service (during which there are approximately 2 minute gaps between trains passing through stations) was a decision to take a defective train off a platform backwards. In theory this would have been faster than going forwards as the train could only travel forward slowly, but the trainee had forgotten that to take a train in the opposite direction from the usual running of that track required the presence of a station manager to ensure all the correct procedures were carried out - and station managers can be off dealing with other matters, as was the one at this station.

Mistakes also happen when trainees are given control of rescheduling the trains, and can mean the trainee has to immediately re-schedule again to compensate for the consequences of their first decision. Sometimes the practical details of putting the decision into practice can be what causes it to fail; for example, one trainee made what at first seemed a good choice of trains to reorganize, but he had overlooked the time it would take to get a driver to the new position for the pick-up.

Every decision has consequences for drivers, station managers, technical operators (who travel out to fix faulty trains), depot managers, etc. – plenty of opportunities for a trainee to forget to inform people – and yet this was never complained about. In one of the control rooms where the relationships between the team members is relatively

formal, the mistakes the trainee was making became the opportunity for a lot of banter, relaxing the atmosphere of the room. Making mistakes is accepted as the healthy way to learn.

A focus on reforming

A third theme that emerged was the centrality of reforming trains to the operation of the service. We discuss this activity in particular because it is both a common practice and a key one, involving particular skill and concentration, and hence is a major contributor to system resilience.

At peak times a controller is working against the clock. With trains passing through junctions at as little as two minute intervals, decisions have to be made on which adjustment of the train times will provide a regular service, and which reforming of train destinations will ensure the slots in the timetable are filled - all the while checking the correctly numbered trains are in the right places to be picked up by waiting drivers.

When searching for the solution to a particularly complex disruption in the service, controllers have to concentrate hard, handing other duties over to the rest of the team: *“Right, let’s bash this out - get it sorted before the peak starts”*.

The goal of the task is quite clear and simple, it is to get trains back on schedule, and the success or failure of the moves selected are quickly seen:

“Some people can do it and some can’t. I’ve been here a long time so it’s easy. When you first come in it’s a load of numbers. People do find it quite hard. There isn’t much you can do wrong, that you can make a mistake with. All you’re basically trying to do is marry up trains with drivers. ... So that’s what you’re really doing, it’s not all that complicated as it looks. It’s just making sure you’ve got a driver in there to take the train out”.

The team claim that this craft cannot be taught in text books and it is clearly something that people become more proficient at over time. When a trainee is being introduced to what a number of times was referred to as the *‘black art’*, they will be shown the alternative solutions that could have been used and a quick explanation for the choice made, but the emphasis is that there can be many combinations of switching trains, reforming, looping and short-tripping, all of which might be as successful a plan as another, especially as the train pattern can be continually changing when the trains are seriously off schedule:

“I got a little pair for you. 201 to make 235, and then 35 to 201. And then when you get 204 that can make 234. Yes, 204 is running early, in fact, is running on time, in fact. That’s why. Yes, 204 to make 234. ... Hi Joe, slight change, 34 will stay as 34, and then 204 will make 245, mate. Yes, then 245 can make 246. That’s it”.

The tools for this task are the paper timetable placed in a file of plastic pockets, and a chinagraph pencil with which

to cross out withdrawn train routes and add in arrows to indicate the new combination of trips being suggested (see Figure 3). On the left page can be seen the trains travelling in one direction, and on the right the trains operating in the opposite direction. Often the goal is to ensure that, although the schedule has been lost in one direction, it is corrected by the time the trains have turned - especially if approaching the morning or evening peak hours. This produces a lot of turning of pages back and forth as the controller searches for a solution. No paper is used but the chinagraph pencil marks can be rubbed off with a duster, and many of the controllers talked to themselves as they were working, sounding as if they are solving a puzzle: *“204...12...all the twos, and that 38, lovely”*.

Although the controllers are taking their job seriously there are aspects that are reminiscent of games playing. As when playing a challenging game, the operators become engrossed in the task, resulting in time passing quickly:

“You can come in, and it can be as quiet as the grave, and then two minutes later, it’s chaotic. I like the chaotic bits, because of the day goes quicker for a start”.

As noted above, observing a controller finding the best way for the trains to be brought back on schedule is like watching someone solve a particularly engrossing puzzle. When asked what the experience was like, one of the suggestions offered was that it was comparable to doing a crossword:

“I suppose you could say it’s like, it’s like doing a crossword, because in theory, if it all goes well, everything comes out at the end of the rush hour - if you haven’t got any spare drivers, or spare trains. It goes wrong - you end up having to work. If it’s all gone horrendously wrong - then you’ve no driver, or no trains - and there shouldn’t be trains cancelled. So there’s that little bit of satisfaction at the end, when doing reformation, when it all comes out all right”.

Within the constant requirement to monitor the real-world situation by translating the data being received in the control room, it was a strong contrast to hear them describing the task of reforming the trains as a separated out activity, with the controllers, once they have seen a train situation they want to alter, then standing back from it and turning their full attention to the print-out of the timetable:

“Whatever’s happening out there it doesn’t matter, I’ve got to get on with turning my train - you’re not really turning train, you’re turning numbers and lights”,

“Some people come here and find it hard to equate those red dots with trains packed with people, outside. We lose sight of the fact that that little red dot might have 5 - 600 people on it. I think that’s good, your main aim is just to get trains moving around as quick as possible”,

“Not visualising. Thinking of the numbers”.

One controller even confessed that he felt it could become the equivalent of playing children's games:

"Sometimes it's like PlayStation, you can set up different routes, it's quite fun, a bit childish I admit but it's quite fun. Drivers ringing up to say, Oh, I haven't seen that route before, and I say, Yeah, nothing wrong with that route driver - sad isn't it but keeps us little boys happy".

Once trained, individuals are responsible for developing their own skill at reforming the trains and consequently some of them choose to take action when it is not strictly necessary:

"Different signalmen have different approaches. I was taught similar to this - If not ten minutes late, leave it - but then decided it was better if you act immediately, four minutes, in case it might get worse. Sometimes nicer for the driver if it is all the same way but nice to think about it - do it your own way".

They therefore develop their own styles, using methods they know have worked previously but also adjusting their approach as alternative solutions surface:

"There's always other ways of doing it. Someone might turn it at [station R], whereas someone else might keep stepping up around it and bring it in and turn it at [station S]. Like someone might run it up to [station S] and reform it up at [station S], where other people might say, Well, I'll not bother with that I'll put it into sidings at [station T] until next trip round and then I can bring it out on time. Sometimes you spot a better way of doing it as you are looking at it".

As controllers have to pick up each others' work when changing shift it would be expected that they develop an opinion on other people's skill. Within one of the control rooms where the team seemed particularly comfortable with each other and reflective about the work they were doing it, individuals were pointed out as having an impressive ability: *"He can do ten reforms in one go, can't you?"*

There are certain qualities to control room staff that seem to dominate. They must all be able to multi task extremely well, keep constantly alert to each other (as also noted by Heath & Luff, 1992), plus they can be very outspoken. There is a lot of friendly banter that includes not only teasing but also comments on what a fellow team member is attempting. This atmosphere of honesty and acceptance of human fallibility could be seen by the way a number of times team members interrupted descriptions being given to the researcher with a confession, saying that their colleague was tactfully omitting the way their mistakes had contributed to the event, e.g.:

"A really good Controller would have held it at [station M], I did it instead on the radio at [station N]".

By this method they checked and challenged each other's work and assumptions - reinforcing the need to be alert to the possibilities of mistakes.

DISCUSSION

We have presented our findings under three headings:

- Uncertainty and trust: operators place little trust in individual pieces of equipment, always seeking confirmation of their interpretations of the system state, articulating and reflecting on their assumptions.
- Training: in particular, we noted that on-the-job training is essential, that the training process is highly selective (contributing to the resilience of the resulting system), and that mistakes are accepted and reflected on.
- Focusing on reforms: when doing complex reforms, controllers focus on that one activity, usually protected from distractions by their colleagues, resulting in high engagement with that activity; the activity has important properties including focusing on a world of "numbers and lights" and solving a puzzle.

In this section, we relate our findings to work on resilience (as discussed earlier, when we set the context for this study). We further discuss key features of the practices observed that contribute to ongoing learning, as an important contributor to resilience. In particular, we highlight a finding from this study that has not, as far as we can ascertain, been previously discussed in relation to system resilience – namely the deep engagement, particularly with the reforming task, that has many properties of play.

Resilience in the underground

Our findings are consistent with Rochlin's (1999) at the organisational level:

- Individuals learn according to their own situations – a theme we return to below.
- Duality is important: staff are constantly having to maintain multiple representations of the situation in the outside world; they defer to drivers to make some decisions; they have both formalized procedures and local practices (Wenger, 1999); and they co-ordinate with other rail organizations that use the same resources.
- Communications are facilitated. The majority of control rooms are designed so that everyone can overhear what is going on; there are shared external representations (e.g. the line diagram); and there are both formal and informal interactions (particularly when mistakes happen).
- Responsibility is shared across the team. Individuals are open about their own mistakes and accepting of others', while also recognizing and valuing others' skills. There will often be dynamic reallocation of roles to support individuals engaged in focused work (notably train reformations).

However, our findings go beyond those reported by Rochlin (1999): resilience is a constructive process rather than an attribute since there is often more than one way to accomplish a given task. We have found that individuals learn through reflecting on their own actions and

understanding (e.g. in recounting tales and explaining the system state to visitors), making that understanding open to critique by others and learning by reflecting on their own mistakes and those of others (by which we mean sub-optimal decisions rather than errors that compromise safety). At times, while reforming, there is focused engagement with the task that acquires many of the properties of play. In the remainder of this section, we consider learning within the work setting, the nature of engagement and its relation to play and learning.

Learning at work

As researchers such as Wenger (1999) and Kilgore (2001) argue, knowledge is socially constructed. Wenger (1999) focuses on the role of communities of practice in nurturing situated learning. He argues that learning within any domain is more than a formal acquisition of knowledge or information: it occurs through a process of participation in 'communities of practice'. New members are brought into knowledge communities, and those communities both transform and reproduce themselves. This participation is at first peripheral, but gradually increases in both engagement and complexity. We see this gradual induction into the community of practice within each control room in the process of on-the-job training.

However, the behaviours observed have a quality that goes beyond the rather sober notion of learning as a process of becoming part of a community: there were levels of deep engagement with the ongoing activity, particularly while reforming trains.

Engagement

When someone is working on a task it can be at different levels of involvement, which range from cursory, through interested, to engagement, engrossment and total immersion (Brown & Cairns, 2004). These last three are aspects of an experience (or behaviour) that has various names: being in the zone, peak performance, immersion, flow.

'Flow' (Csikszentmihalyi, 1990) is a common experience described by people such as rock climbers, artists, and chess players. The characteristics of flow are deep concentration, all the person's skills being brought to bear on the task, a feeling of control and belief in one's ability to succeed, a lack of focus on self, and a distortion in the perception of time passing. While it is unlikely that controllers experience all aspects of flow, due to the strong cultural re-enforcement of the need for constant double checking, many of these features were apparent in their behaviour, particularly while reforming trains.

The variety of different solutions that are possible for the same problem, and the acceptance that people need to be left to make mistakes so as to learn, ensures that each individual defines his or her own level of achievement. One service control manager explained that he sees it as part of his job to ensure that others on the team are blocked from

interfering or attempt to influence whoever has the role of controller that day.

Within the control rooms there are many distractions, with the controller needing to instruct the signallers, monitor what is happening elsewhere on the line, hand on messages, update management and passengers. But when a complex incident kicks off or when a severe backlog has developed, everyone on duty plus those officially on meal relief pitch in to, as far as possible, remove any distractions from the controller, leaving just the activity of working with the train number patterns shown on the paper timetable.

The controllers have a task that is challenging. As this is balanced with extensive training and with the freedom to set their own level of achievement to aim for, this results in concentration and satisfaction. This enjoyment of exercising a skill can be an encouragement to practice that skill and the London Underground control rooms are set up to support this. Controllers do not have to justify their decision to adjust the running of the trains as long as it does not impinge on safety and they can show that their intention was to improve the service. The controllers and the trainees can therefore practice their skills, and keep themselves alert, by checking whether any re-scheduling would produce a better regulated flow of trains.

As with games, the reformation task is set against time and requires quick sifting of information and fast choices. It has simple goals that can be soon assessed as to whether they have been achieved or not. This takes place within a culture where experimentation is seen as the acceptable way to learn. All this produces an activity that encourages learning, in a similar way that interactive educational resources are being produced that create motivation and concentration.

Play

Reiber (1996) discusses the value of well-designed microworlds for supporting learning through play. Reiber argues that play is not the opposite of work (and thus contrasts it with leisure). Rather, play is an important means of learning. This understanding of play can be contrasted with that of, for example, Costea *et al.* (2005), who analyse organizations that have an ethos of "work hard, play hard", where the play activity contributes to the sense of "wellness" of individuals, and hence of the organization as a whole, but is not explicitly linked to learning.

In Reiber's (1996) view, play has four key attributes:

- It is voluntary.
- It is intrinsically motivating.
- It involves active, often physical, engagement.
- It has a make-believe quality.

Perhaps surprisingly, all of these qualities can be discerned in the activities of the operators in London Underground, particularly when engaged in reforming trains. As discussed above, reforming is typically voluntary, in the sense that the

operators are choosing to optimize the train service. It is intrinsically motivating to those who choose to engage in this activity, in the way that solving a puzzle is motivating. It involves a high level of engagement. And even though it is clearly concerned with real trains and people, the operators abstract away from those into a world of numbers and lights.

Reiber (1996) discusses several views of play; in particular (of relevance to this study), he discusses play as progress – as a means of learning something useful. The reforming of trains appears to have this quality: of learning about how the system works – in ways that are not properly embodied within the system simulators – by engaging with the system in an exploratory, playful way.

Sengers *et al.* (2005) discuss ludic design (designing for play) as a means of promoting “engagement in the exploration and production of meaning” (p.51). However, this misses the essential role, highlighted by Garris *et al.* (2002), of reflection in learning. Garris *et al.* present a view of games as having six important attributes: fantasy, rules and goals, sensory stimuli, challenge, mystery and control. The activity of reforming trains does not have all these qualities: while it has rules and goals, challenge and control, it does not have fantasy, strong sensory stimuli or mystery. On these criteria, the activity is not a game even though, as discussed above, it does have many of the properties of play. Garris *et al.* also emphasise the importance of motivation in learning; they distinguish between intrinsic motivation (arising from features such as challenge, curiosity and fantasy) and extrinsic motivation (arising from achieving desired outcomes). In these terms, train reforming is both intrinsically and extrinsically motivating to the controllers in this study. Garris *et al.* argue that motivation and engagement are essential to learning. They also argue that both experience and reflection are necessary. These are all features of the activities and culture in the control rooms studied – features that together contribute to an environment in which operators can learn to perform in more anticipatory, and hence more resilient, ways.

CONCLUSION

This study has highlighted some important features of the control rooms that contribute to their resilience. It has confirmed the earlier findings of Heath and Luff (1992) and Garbis (2002) on the importance of communications and the roles of artifacts in facilitating those communications. It has also provided further support for the findings of Rochlin (1999) on system features that promote resilient performance, such as acceptance of mistakes and duality of system views.

All control room operators work with a level of uncertainty, dependent on the quality of information available about the system being controlled. We have identified many of the sources of uncertainty in London Underground, from unreliable equipment to unpredictable people, and

highlighted strategies that operators have developed for managing that uncertainty.

The importance of individuals learning (creating their own personalized understanding of the system) is widely recognized. In most safety-critical industries, simulators are used to provide opportunities for staff to explore the boundaries of system operation in a safe setting, since even the smallest mistake might lead to catastrophe in the real system. Exploration allows individuals to develop their own understanding of the system (e.g. McCarthy and Wright, 2004, Chapter 7). We have discussed ways in which operators are encouraged to explore the system, often making (non-safety-critical) mistakes in the process, which are accepted and used as learning opportunities.

The value of reflection – on both the meanings of system outputs and the effects of actions – has also been highlighted in this study. Control room staff have evolved a broad repertoire of techniques for encouraging reflection on action, including recounting tales to trainees and visitors as well as discussing situations among the team. Reflective activities enable the monitoring of dynamic factors such as the extent to which the system can be trusted, and whether controllers have sufficient expertise and training to solve persistent problems.

Finally, we have highlighted control room behaviours, particularly around reforming trains, that have many properties of play. In our earlier studies of ambulance control (Blandford and Wong, 2004; Furniss and Blandford, 2006), we did not detect any playfulness of the kind described here. Neither have we seen it discussed by other researchers who have studied control rooms or other safety-critical settings. This probably reflects some of the particular properties of train control rooms: that the system is inherently resilient to certain kinds of mistakes; that there can be complicated puzzles to solve (unlike ambulance control, where major reconfigurations of the system are relatively rare); and that time-frames are relatively large (compared, for example, to air traffic control). Play is also less likely to occur in simulators, certainly in situations (e.g. for pilots) where use of the simulator is compulsory and tasks or scenarios are defined by others. Overall, there is no evidence that play is necessary to support learning for resilient performance; our evidence simply suggests that it can be valuable in some situations. We are not aware of previous studies that have recognized the value of play in promoting resilience.

Play is only one element of a rich web of behaviours and attitudes that ensure resilience of the overall system. The study has shown the importance of other factors including establishing an appropriate understanding of the ability of system components to reflect complex real world situations; learning through mistakes; sharing, recounting and reflecting on experiences; and rehearsing understandings with each other and with visitors. These elements cannot be readily understood by decomposing the system into it

individual components, but result in emergent behaviours that make the overall system resilient in a wide range of unpredictable situations.

ACKNOWLEDGMENTS

We are very grateful to all the staff at London Underground who contributed to this study in any way, and to anonymous referees of an earlier version of this paper for constructive criticism. This work was partially funded by EPSRC grants GR/S67494 and GR/S67500.

REFERENCES

- Beyer, H. & Holtzblatt, K. (1998) Contextual Design. San Francisco : Morgan Kaufmann
- Blandford, A. & Wong, W. (2004) Situation Awareness in Emergency Medical Dispatch. *International Journal of Human-Computer Studies*. 61(4). 421-452.
- Brown, E., Cairns, P. (2004) A grounded investigation of immersion in games. *ACM Conf. on Human Factors in Computing Systems, CHI 2004*, ACM Press, 1297-1300.
- Charmaz, K. (2006) *Constructing Grounded Theory*. Sage.
- Costea, B., Crump, N. & Holm, J. (2005) Dionysus at work? The ethos of play and the ethos of management. *Culture and Organization*. 11.2. 139-151.
- Csikszentmihalyi, M. (1990). *Flow: The Psychology of Optimal Experience*. New York: Harper and Row.
- Furniss, D. & Blandford, A. (2006), Understanding Emergency Medical Dispatch in terms of Distributed Cognition: a case study. *Ergonomics Journal*. 49. 12/13. 1174-1203.
- Garbis, C. (2002). Exploring the Openness of Cognitive Artifacts in Cooperative Process Management. *Cognition, Technology and Work*, 4, 9-21
- Garris, R., Ahlers, R. & Driskell, J. (2002) Games, motivation and learning: a research and practice model. *Simulation and Gaming*. 33. 44-467.
- Heath, C. and Luff, P. (1992) Collaboration and Control: Crisis Management and Multimedia Technology in London Underground Line Control Rooms. *Computer Supported Cooperative Work*, 1, 69-94.
- Hollan, J. D., Hutchins, E. L. & Kirsh, D. (2000) Distributed cognition: toward a new foundation for human-computer interaction research. *ACM Transactions on CHI*, 7.2, 174-196.
- Hollnagel, E. (1998) *Cognitive Reliability and Error Analysis Method (CREAM)* Oxford : Elsevier Science.
- Hollnagel, E. (2005) Human Reliability Assessment in Context. *Nuclear Engineering and Technology*. 37.2.159-166.
- Hollnagel, E., Woods, D. D. & Leveson, N. (Eds) (2006) *Resilience engineering: Concepts and precepts*. UK: Ashgate.
- Hutchins, E. (1995) *Cognition In The Wild*. MIT Press, Cambridge, MA.
- Kauppi, A., Wikström, J., Sandblad, B. & Andersson, A. (2006) Future train traffic control: control by re-planning. *Cognition, Technology and Work*. 8. 50-56.
- Kilgore, D. (2001) Critical and postmodern perspectives on adult learning. *New directions for adult and continuing education*. 89. 53-61.
- Luff, P. & Heath, C. (2000) The collaborative production of computer commands in command and control. *International Journal of Human-Computer Studies* 52(4): 669-699.
- McCarthy, J. & Wright, P. (2004) *Technology as Experience*. MIT Press.
- Pledger, S., Horbury, C. & Bourne, A. (2005) Human Factors in LUL – History, Progress and Future. In *Rail Human Factors: supporting the Integrated Railway*. J. Wilson, B. Norris, T. Clarke & A. Mills. (Eds.) Hampshire UK: Ashgate. 497-507.
- Reason, J. (1990) *Human Error*. Cambridge : Cambridge University Press.
- Reiber, L. (1996) Seriously considering play: Designing interactive learning environments based on the blending of microworlds, simulations, and games. *Educational Technology Research and Development*. 44.2. 43-58.
- Rochlin, G. (1999) Safe operation as a social construct. *Ergonomics*. 42.11. 1549-1560.
- Sengers, P., Boehner, K., David, S. & Kaye, J. (2005) Reflective Design. In *Proceedings of the 4th Decennial Conference on Critical Computing: between Sense and Sensibility*. O. W. Bertelsen, N. O. Bouvin, P. G. Krogh, and M. Kyng, Eds. CC '05. ACM Press, New York, 49-58.
- Stearn, M., Clarke, T., & Robinson, J. (2005) Baseline Ergonomics Assessment of Signalling Control Facilities. In *Rail Human Factors: supporting the Integrated Railway*. J. Wilson, B. Norris, T. Clarke & A. Mills. (Eds.) Hampshire UK: Ashgate. 262-271.
- Wenger, E. (1999) *Communities of practice: Learning, meaning and identity*. Cambridge: Cambridge University Press.