

Modelling of human alarm handling response times: a case study of the Ladbroke Grove rail accident in the UK

Stanton, N; Baber, Christopher

DOI:

[10.1080/00140130701695419](https://doi.org/10.1080/00140130701695419)

Document Version

Early version, also known as pre-print

Citation for published version (Harvard):

Stanton, N & Baber, C 2008, 'Modelling of human alarm handling response times: a case study of the Ladbroke Grove rail accident in the UK', *Ergonomics*, vol. 51, no. 4, pp. 423-440.
<https://doi.org/10.1080/00140130701695419>

[Link to publication on Research at Birmingham portal](#)

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.


When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

This article was downloaded by:[Neville.Stanton@brunel.ac.uk]
On: 21 March 2008
Access Details: [subscription number 791605627]
Publisher: Taylor & Francis
Informa Ltd Registered in England and Wales Registered Number: 1072954
Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



Ergonomics

Publication details, including instructions for authors and subscription information:
<http://www.informaworld.com/smpp/title~content=t713701117>

Modelling of human alarm handling response times: a case study of the Ladbroke Grove rail accident in the UK

Neville A. Stanton ^a; Christopher Baber ^b

^a Ergonomics Research Group, School of Engineering and Design, Brunel University, Uxbridge, Middlesex, UK

^b School of Electrical, Electronic and Computer Engineering, University of Birmingham, Edgbaston, Birmingham, UK

Online Publication Date: 01 April 2008

To cite this Article: Stanton, Neville A. and Baber, Christopher (2008) 'Modelling of human alarm handling response times: a case study of the Ladbroke Grove rail accident in the UK', Ergonomics, 51:4, 423 - 440

To link to this article: DOI: 10.1080/00140130701695419

URL: <http://dx.doi.org/10.1080/00140130701695419>

PLEASE SCROLL DOWN FOR ARTICLE

Full terms and conditions of use: <http://www.informaworld.com/terms-and-conditions-of-access.pdf>

This article maybe used for research, teaching and private study purposes. Any substantial or systematic reproduction, re-distribution, re-selling, loan or sub-licensing, systematic supply or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The accuracy of any instructions, formulae and drug doses should be independently verified with primary sources. The publisher shall not be liable for any loss, actions, claims, proceedings, demand or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

Modelling of human alarm handling response times: a case study of the Ladbroke Grove rail accident in the UK

Neville A. Stanton^{a*} and Christopher Baber^b

^a*Ergonomics Research Group, School of Engineering and Design, Brunel University, Uxbridge, Middlesex, UB8 3PH, UK;* ^b*School of Electrical, Electronic and Computer Engineering, University of Birmingham, Edgbaston, Birmingham, B15 2TT, UK*

The purpose of the paper was to address the timeliness of the signaller's intervention in the Ladbroke Grove rail incident in the UK, as well as to consider the utility of human performance time modelling more generally. Human performance response time modelling is a critical area for Human Factors and Ergonomics research. This research applied two approaches to the same problem to see if they arrived at the same conclusion. The first modelling approach used the alarm initiated activity (AIA) model. This approach is useful for indicating general response times in emergency events, but it cannot comment in detail on any specific case. The second modelling approach employed a multi-modal critical path analysis (CPA) technique. The advantage of the latter approach is that it can be used to model a specific incident on the basis of the known factors from the accident inquiry. The results show that the AIA model produced an estimated response time of 17 s, whereas the CPA model produced an estimated response time of 19 s. This compares with the actual response time of the signaller of 18 s. The response time data from both approaches are concordant and suggest that the signaller's response time in the Ladbroke Grove rail accident was reasonable. This research has application to the modelling of human responses to emergency events in all domains. Rather than the forensic reconstruction approach used in this paper, the models could be used in a predictive manner to anticipate how long human operators of safety-critical systems might take to respond in emergency scenarios.

Keywords: response time; alarms; control room; performance modelling; emergency; rail system

1. Introduction

1.1. The Ladbroke Grove rail incident

On 5 October 1999 at 08.06, a light commuter train left Paddington station from Platform 9 on route to Bedwyn in Wiltshire. Approximately 3 min later the train collided with a high-speed train coming from the opposite direction at a combined speed of 130 mph causing the deaths of 31 people, including the drivers of both trains. At the subsequent inquiry led by Lord Cullen (Cullen 2000), one of the key questions was: 'why did it take the signaller in the Slough control room [*responsible for that section of track*] some 18 s to

*Corresponding author. Email: neville.stanton@brunel.ac.uk

respond to the alarm advising him of an unauthorised track occupation?’ (italics added). In his report, Cullen (2000) criticised the time taken by the signaller to respond, although he did concede that an earlier response ‘could not have prevented the crash’ (Cullen 2000). The point at issue was whether 18 s constitutes a reasonable time or whether one might anticipate a much faster time. Certainly, if the response was simply a matter of hearing an alarm and then pressing a button to cancel the alarm, then one might expect human performance to be around 3–4 s. However, such a response time might miss crucial aspects of signaller activity, e.g. determining an appropriate form of response as well as performing that response.

In a recent paper addressing the underlying causes of the accident, Lawton and Ward (2005) took a systems analysis perspective to identify potential latent conditions contributing to the Ladbroke Grove accident. They reviewed witness statements and the two reports produced by the enquiry that followed the event. The focus of the paper was largely upon the driver of the light commuter train, to explain why he might have driven through a red signal. Many factors were identified, such as those associated with driver psychology (e.g. perceptual, attentional and psycho-motor abilities), the design of the rail system (e.g. track layout, display interfaces, feedback and communications), the organisation (e.g. safety management, training and system design) and system defences (e.g. safety devices, signalling, policies and awareness of hazards). Using Reason’s (1990) ‘system pathogen’ model, Lawton and Ward (2005) suggested how these latent conditions gave rise to the active failure (i.e. the driver passing the red signal). Taking their lead from the Cullen report, the authors also questioned the timeliness of the signaller’s intervention in the incident.

There are at least three approaches that could be taken to determine whether or not Lord Cullen’s criticism of the signaller’s response time is justified. The first approach would be to recreate the conditions in a signal control room environment, to determine how a range of signallers deal with the event (Farrington-Derby *et al.* 2006). This approach could be criticised because, while participants may be encouraged to follow standard operating procedures and thus perform to similar standards, the conditions in an experiment may not reflect those of the real-life incident. The second approach would be to use Klein *et al.*’s (1989) critical decision methodology, to interpret the activities of the signaller in question (Walker *et al.* 2006). This approach could be criticised because the event took place some years ago and therefore the recall of the signaller is likely to be unreliable (particularly as the signaller has already had his account challenged and debated in Court). A third approach would be to model response times of the signaller. It is this latter approach that will be followed in this paper. In particular, the paper focuses on two ways of achieving model response times: the first considers response times for emergency handling across a range of industries; the second develops a model of the signaller’s probable action, given the known contextual factors and the description of events provided by the Cullen Enquiry Reports. The two modelling approaches have the benefit of offering independent data on how long it could reasonably be expected for people to respond in an emergency situation.

1.2. Summary of the Ladbroke Grove incident

A light commuter train (with the headcode of IK20 on the signaller’s screen and called train 1 in this paper) had passed a red signal called SN109 (indicated by the black circle at the centre right of Figure 1) on the track out of Paddington station as if signal SN109 was set to green, indicating that the train could proceed safely. A high-speed train (with the




Figure 1. A schematic of the relevant part of the track layout at Ladbroke Grove.

headcode of 1A06 on the signaller's screen and called train 2 in this paper) was coming into Paddington from the opposite direction. An illustration of the direction of travel for both trains is presented in Figure 1. The Ladbroke Grove rail junction has six tracks, three of which are illustrated in Figure 1. The tracks are divided into segments, called track circuits, which indicate the position of the train. Train 1 travelled through track circuits GD, GE (passing signal 109 – the red signal), GF and GG as indicated by the arrow going left in Figure 1, whilst train 2 travelled through track circuits MX, MY and MZ as indicated by the arrow going right in Figure 1. The collision occurred where the black rectangle is placed over the two intersecting sections of track.

The timing of events is taken from Lord Cullen's report and presented below. The timeline represents the events that were logged by the signaller's workstation and, in this extract, picks up the journey from the moment that train passed signal SN109, which was set to red. The train driver should have waited at this signal until it turned green. The paper by Lawton and Ward (2005) contains more details on the activities of the driver of train 1, whereas the present paper focuses on the signaller.

08:08:29 – A 'track circuit GE occupied IK20' warning message is presented on the alarm screen in the signaller's workstation (see Figure 2 for a general picture of the workstation) and auditory 'tweet' sounds (an auditory track occupation alarm referring to the same track circuit GE occupation by train 1); at the same time a red line appears on the track layout on the track display and the train headcode of IK20 stays at signal 109, the red signal.

08:08:32 – The oncoming train 2 occupied track circuit MZ and a red line appears on the appropriate track display with the headcode 1A06 (the number associated with train 2).

08:08:34 – Auditory alarm 'tweet' sounds as rear of train 1 clears track circuit GD (i.e. the track circuit before GE) and the track circuit is shown as cleared on track display.

08:08:36 – Track circuit GF occupied message displayed and auditory 'tweet' sounds (track occupation alarm referring to the occupation of track circuit GF by train 1 – at the same time a red line appears on the track layout on the track display).

08:08:41 – Rear of train 1 clears track circuit GE (i.e. the track circuit before GF) and track circuit shows as cleared on track display.

08:08:42 – The rear of the oncoming train 2 clears track circuit MY and track circuit MY shows as cleared on track display.

08:08:49 – Track circuit GG occupied by IK20 alarm message displayed and auditory 'tweet' sounds referring to track circuit GG by train 1 – at the same time a red line appears on the track layout on the track display.

08:08:50 – Train 1 and train 2 collide.

The photograph in Figure 2 shows a signaller's workstation, which comprises six screens, a trackball and buttons, a keyboard and four telephones. The track displays can be seen on the four screens from far right. These are similar to the track schematic shown in Figure 1, only with increased complexity to reflect six tracks and the interconnections. Notice also that the 'alarm list' is presented on the screen to the far left, which can easily become filled with messages during an incident.

Unless the signaller happens to be looking directly at the appropriate point on the alarm screen (on the far left screen in the workstation, shown in Figure 2), the first he or she will know of a new alarm is an auditory warning. There are four categories of visual alarm and these are colour-coded yellow, blue, green and red, presented in that order from top to bottom of the screen. Only those alarms that currently apply are displayed. All four categories have the same auditory 'tweet', i.e. there is no differentiation of auditory information, which means that the response to a 'tweet' is to search the alarm screen in order to determine the type of alarm. The track occupation alarms are colour-coded in red (the only warnings discussed in this paper) and appear at the bottom of the screen (the reasons for this layout are historical and beyond the scope of this paper). When the signaller had read the track occupation alarm, he would be aware that a train has overshot the point at which it was supposed to stop. This may occur for a number of reasons. For example, the track occupation alarm can also sound when trains are shunting (e.g. the manoeuvring of carriages and train cabs along the track into and out of sidings). In this case the alarm is false and the signaller will be expecting the alarm and be able to acknowledge and ignore it. There are a few occasions on which the track occupation alarm will be presented due to slight incompatibilities in the track circuit design and the direction



Figure 2. The signaller's workstation.

of train running. This means that occasional false alarms are displayed and the signaller will be able to acknowledge the alarm and ignore it. On some occasions, the driver might simply misjudge the stopping distance required by the train. This is a genuine alarm (i.e. an alarm requiring intervention by the signaller – Stanton 1994); the driver normally calls the signaller after the train has come to a halt. Most tracks have a safe overshoot area to cope with this. For example, in the case being studied here, the signaller probably expected the driver of train 1 to pull up within track segment GE in Figure 1, as was stated in the testimony during the enquiry by the signaller: 'In every other SPAD [signal passed at danger] incident I have been involved in, the train involved has stopped within the overlap. At the first instant in the situation on 5 October 1999, I monitored the workstation, expecting the train IK20 to stop...' (Cullen 2000). In fact, if the driver of train 1 had stopped at any point up to and including track segment GG, the accident would have been avoided. Given this background information, it is possible that the signaller (on noting an overshoot) could have sought confirmation of the event, i.e. wait for an additional alarm in order to confirm that this was a 'real' train runaway event, or a call from the driver to say that he or she had come to a halt within the safe overshoot area.

On only a tiny percentage of occasions the track occupation alarm will refer to a real train runaway (typically only 1% of alarms will require any action by the signaller), when a train has continued to run on a line for which it has not been cleared. This is a real emergency. Fortunately, these occurrences are rare and it is possible for a signaller to have never encountered a real event before. The effects of the low positive predictive value (PPV) of the alarm system have been shown to slow down the human reaction times considerably, particularly below a PPV of 0.25 (Getty *et al.* 1995). If this does happen, the signaller has to first find the track occupation on one of his four track displays and then decide what course of action to take. If the signaller decides that this is a real case of a runaway train, then they may have several alternative courses of action. These courses of action may include sending a stop message to the train or any oncoming trains, switching the points to send the train or oncoming trains onto an unoccupied track or changing the signalling for oncoming trains.

All of these decisions, i.e. sending a stop message, switching the points and changing the signals, require the signaller to read the track ahead of the train in order to decide if there is an oncoming train and if either of the trains can be stopped in time. Good situation awareness by the signallers, through information communicated via the displays and/or interaction with other persons, is essential to help them reach the right decisions (Stanton *et al.* 2006). Sending a stop message would require the assessment of the stopping distance and the trajectory of the train(s). Moving the track points ahead of the train could divert the train onto track 4 instead of track 2 (see Figure 1) to prevent an impending collision, but the points can lock up when they detect a train is in the circuit in order to prevent an inadvertent point change. Changing the signals from green to red requires the signaller to anticipate the likely stopping distance of the oncoming trains. These are complex decisions that have to be made by the signaller, who will be under considerable time-stress. The signallers' rule book does not give unequivocal guidance about which is the best strategy; it just requires the signaller to 'act immediately'. The question addressed by the current paper was: 'could the signaller have responded quicker than 18 s?'

2. A model of human alarm handling from other domains

The model of alarm initiated activity (AIA) shows potential pathways for human activity involving alarms (Stanton 1994, Stanton and Baber 1995, Stanton and Edworthy 1999).

The term ‘activities’ refers to the ensuing behaviours triggered by the occurrence of alarms. It is assumed that these activities would not have been triggered without the alarm being present; thus, they can be thought of as ‘alarm initiated activities’. This may be deduced by asking people why they pursued one course of action over another in a particular situation (assuming they know why and are being truthful); for example, if an individual is observed to interrupt a task or change a strategy within a task. The AIAs are linked to other human supervisory control activities that can be typified as continuous tasks (such as visual scanning of track monitors) and discrete tasks (such as manual routing, dealing with phone calls, updating logs, etc.). In these tasks, alarm information may be used instead of, or in conjunction with, other information (e.g. data regarding weather conditions, reference to track state, comments from other signalmen or reports from engineers who are trackside).

The main stages of AIAs are as follows:

- Observation: the initial detection of the alarm.
- Acceptance: the act of acknowledging the presence of an alarm (usually pressing an ‘accept’ key to remove the highlighting as a physical acknowledgement that the signaller has noted the presence of the alarm).
- Analysis: the initial assessment and prioritisation of the alarm. There are five potential pathways that the activity can follow, as illustrated in Figure 3.
- Investigation: the activity directed at determining the underlying cause for the alarm.
- Correction: the stage at which the system controller implements their response to the alarm condition.
- Monitoring: the assessment of success of the analysis, investigation and correction activities.
- Resetting: extinguishing the alarm, returning it to its inactive state (usually pressing a ‘reset’ key to remove the alarm from the display, as a physical acknowledgement that the alarm condition has passed).

The stages and pathways of the AIA model are shown in Figure 3. The AIA model has been used by other researchers. Shorrock and Scaife (2001) report on the use of the model in the development of design principles for alarm systems in air traffic control. At a high level, the air traffic controller’s tasks have some commonalities with the signaller’s tasks, as both sets of tasks are concerned with the safe and efficient movement of transportation systems. Shorrock and Scaife comment that their work has ‘demonstrated the usefulness of the model of AIA in the design and evaluation of alarm systems’ (Shorrock and Scaife 2001). More recently, in an overview of the whole field of alarm handling, Bliss and Fallon (2004) argue that the AIA model is particularly useful because it acknowledges cognitive activities as well as physical actions. Other researchers have remarked on the relatively small amount of observable actions, commenting that most of the time control room operators are engaged in covert, cognitive activities (Wilson and Rajan 1995, Moray 1997, Riera and Debernard 2003).

The alarm handling model in Figure 3 shows stages and pathways in alarm handling. It is recognised that not all alarms are dealt with in exactly the same way, as the activities are context dependent. After an alarm has been triggered, it may be observed and accepted by the signaller, then some analysis of the likely trigger conditions is undertaken (which will be highly dependent upon the situation awareness – Stanton *et al.* 2006). Depending on the outcome of this analysis, one of five pathways could be followed, as indicated in Figure 3.

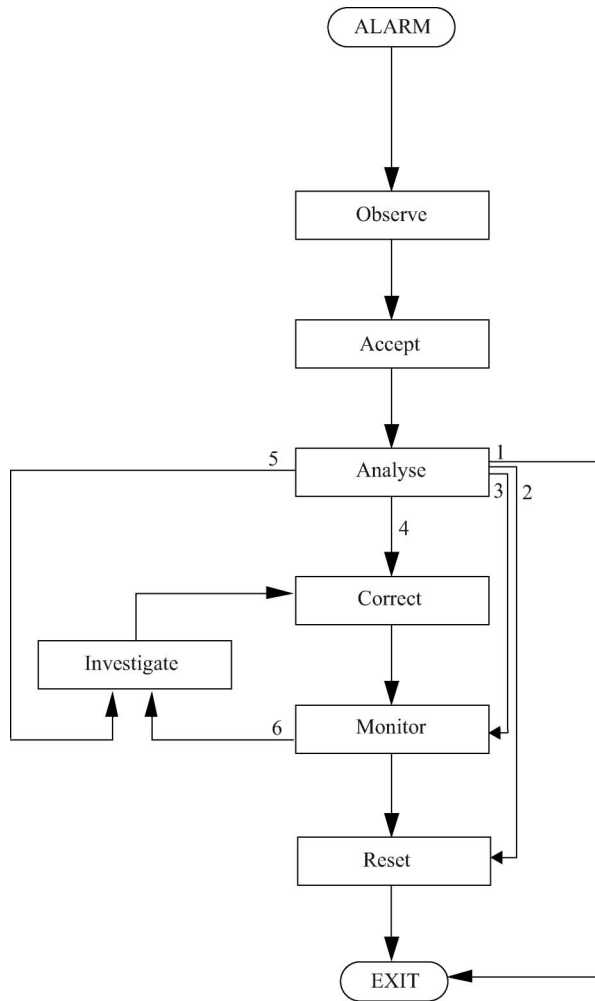


Figure 3. Stages in alarm initiated activities.

These pathways include correcting the situation, monitoring the condition, resetting the alarm, investigating the condition and exiting the analysis.

A literature search on human alarm handling in emergency events identified 13 research papers pertinent to the study in hand. The review revealed seven research papers that dealt with reaction times to alarm events in aviation (Butcher *et al.* 1993, Bliss *et al.* 1995, Getty *et al.* 1995, Bliss 1997, Tyler 1997, Singer and Dekker 2000, Trujillo 2001), two papers that reported reaction times in the nuclear industry (Roth *et al.* 1992, Hollywell and Marshall 1994) and four papers that presented reaction times in process control (Kragt 1983, Moray and Rotenberg 1989, Stanton and Baber 1997, Stanton and Stammers 1998).

The studies reviewed ranged from simple tracking tasks with secondary alarm handling, through simple simulations and alarm categorisation tasks, to full-scope task simulation with critical events. Thus, the literature was not found to contain data directly relevant to rail operations nor specifically to signal handling. It is proposed that, despite obvious variation in context, it is feasible that for the elementary activities within the AIA

framework, human performance can be compared across domains. There is considerable variation in the reports of response times for human intervention in alarm initiative events, from 1 s to 90 s.

The data from the review of the alarm handling response time literature (which presented experimental data rather than real-world studies) were compiled into a table of response times, based on the AIA model. The definitions of observation, acceptance, analysis, investigation correction, monitoring and resetting were used to identify which AIA(s) the research papers were referring to in their results. Although the authors of the works may not have used these terms exactly, it was possible to deduce to which stages(s) the response data they collected referred by using the AIA definitions. This made it possible to identify the amount of time people take for different aspects of the alarm handling task, as summarised in Table 1. The minimum and maximum values of these data are shown in Table 2, together with their respective cumulative totals. In order to produce the timing estimates for each AIA reported in Table 2, the times for the previous stages were subtracted. Thus, to get the time taken for the 'accept' stage, the time for the 'observe' stage was subtracted, and so on. A cumulative total was produced for each of the columns of data. The figure in parentheses refers to the percentage that each activity represents as a proportion of the total response time. Of the 13 studies, only six report data on the completion of corrective actions (i.e. the studies of Kragt 1983, Moray and Rotenburg 1989, Butcher *et al.* 1993, Stanton and Baber 1997, Tyler 1997, Trujillo 2001). Only one of these studies reported a mean response time that was quicker than the signaller, of 17.12 s compared to 18 s. Most of the reported response times were between 20 and 60 s. On this basis, the signaller's response time of 18 s is comparable to the shortest response times in the literature.

Table 2 illustrates the minimum and maximum times reported taken by each stage in alarm initiated activity. The cumulative totals for minimum and maximum response times are 17 s and 90 s respectively. The signaller's response time of 18 s in the Ladbroke Grove incident is at the quicker end of this range. The answer to the question why the signaller took as long as 18 s rests on the idea that alarm handling involves a series of activities that must be completed. If successful alarm handling consisted of only observing the presence of an alarm and pressing an 'alarm acceptance' key then this could take as little as 2 s. These activities only represent about 10–12% of the total response time. If successful alarm handling required the analysis and corrective actions as well, then this could take 11 s. If the alarm handling activities also require investigative activities to be performed, as indeed the railway rule book prescribes, the minimum response time will be 17 s. The reason for the increases in response time are due to the analysis, investigation and correction phases being considerably more demanding in terms of time requirements, which accounts for the remaining 88–90% of the response time. This means that observing the presence of an alarm is not the same as interpreting it and understanding it in the broader context of the system. Deciding upon, and implementing, an appropriate course of action takes the bulk of the time in alarm initiated activity in emergency events.

3. A multimodal critical path analysis model of human alarm handling

The multimodal critical path analysis (CPA) method has its roots in two traditions. CPA is based in project management literature (e.g. Lockyer and Gordon 1991) and is used to estimate the duration of a project in which some activities can be performed in parallel. It is possible to apply these ideas to any time-based activity, such as human performance.

Table 1. Analysis of human response time data from other domains.

Author Stage	Bliss <i>et al.</i> 1995	Bliss <i>et al.</i> 1997	Butcher <i>et al.</i> 1993	Getty <i>et al.</i> 1995	Kragt 1983	Hollywell and Marshall 1994	Moray and Rotenberg 1989	Roth <i>et al.</i> 1992	Singer and Dekker 2000	Stanton and Baber 1997	Stanton and Stammers 1998	Trujillo 2001	Tyler 1997
Observe							1.18 to 2						
Accept		2.65		1 to 8						2.93			
Analyse	6 to 6.5					3.5 to 4.7					5.02 to 7.58		
Investigate								30 to 45	10				
Monitor													
Correct			41		25.35		26.6 to 32.3						
Reset										17.12		20 to 60	60 to 90

Note: Grey blocks indicate that the time taken to undertake the previous activities are also subsumed in the response time data reported. Empty cells indicate that these data are not reported.

Table 2. Estimates of minimum and maximum response times (RT) for each stage of alarm initiated activity with the percentage accounted for of total response time in brackets.

Alarm initiated activity	Minimum RT (s)	Cumulative minimum (s)	Maximum RT (s)	Cumulative maximum (s)
Observation	1 (6%)	1	2 (2%)	2
Acceptance	1 (6%)	2	7 (8%)	9
Analysis	2 (12%)	4	6 (7%)	15
Investigation	6 (36%)	10	30 (33%)	45
Monitoring	Variable*	10	Variable*	45
Correction	7 (40%)	17	45 (50%)	90

*No figures are entered for the monitoring task as it is difficult to estimate these from the data – it is assumed that monitoring is a continuous activity that can be performed in parallel with the other actions.

In order to calculate CPA, one needs to know the order in which tasks are performed, their duration and their dependency. The notion of dependency is, for traditional CPA, based on the question of what tasks need to be completed before another task is allowed to commence. When applied to human performance models, however, dependency offers a richer conceptual framework in that it allows consideration of parallel activity. Traditional methods for modelling human response time are constrained because they do not represent parallelism. For example, the keystroke level model method offers a simple additive method for calculating response times in computing tasks (Card *et al.* 1983). Using concepts from CPA, it is possible to demonstrate how some activities can be performed in parallel, which can provide more accurate estimates of performance time (Schweickert 1978, John 1990, John and Newell 1990, Gray *et al.* 1993, Baber and Mellor 2001). As Gray and Boehm-Davis (2000) point out, the use of critical path models provides an account of human behaviour that has the benefit of a graphic notation (making it relatively easy to see what is happening) and minimalist theory (making it easier to focus on what they term ‘micro-strategies’ in task performance).

The question of how to address parallel activity is usually based on a multiple attentional resource model. Wickens (1984) amalgamated a considerable amount of research on multiple task performance to propose a theory of multiple attentional resources. The theory proposes a general pool (or reservoir) of attentional resources, which contains two sub-pools, one for spatial resources and one for verbal resources. While the notion of multiple resources has not been without its critics (e.g. Spence and Driver 1997, Vidulich and Tsang 2007), the implementation in this paper draws on the assumption that physical actions can be assumed to be performed separately. Thus, the present model uses a ‘light’ version of multiple-resource theory, in which physical actions are treated as separate, but that cognitive actions are performed under a general level (which might function like a ‘central executive’). Such a model would help to determine the possibility of tasks being performed in series or parallel, i.e. two ‘visual’ tasks would need to be performed in series (for the simple reason that one cannot look in two places at the same time), but ‘auditory’ and ‘visual’ tasks could possibly be performed in parallel, e.g. the (visual) monitoring of displays could be performed in parallel with the (auditory) hearing of an alarm. Although the notion of different resources provides some insight into dependency, it is still essential to know the sequence of tasks and their relationship in order to construct the CPA model.

The modelling of the signaller’s response times in the events from the presentation of the first signal passed at danger (SPAD) alarm is undertaken using cognitive CPA based

upon a method initially developed by Gray *et al.* (1993) and further refined by Baber and Mellor (2001). The method may be proceduralised as follows.

- (1) Analyse the tasks to be modelled: The tasks need to be analysed in fine detail if they are to be modelled by multimodal CPA. Hierarchical task analysis can be used to specify the tasks, but it needs to be conducted down to the level of individual task units. This fine-grained level of analysis is essential if reasonable predictions of response times are to be made. The result of this analysis is an initial sequencing of the tasks in terms of their order of occurrence. This ordering defines the temporal dependency of the model.
- (2) Allocate sub-tasks to input/processing/output modality: Each unit task then needs to be assigned to a modality. For the purposes of control room tasks (analysed by the first author), examples of these modalities are as follows:
 - (a) visual tasks: looking at the track displays, looking at the alarm screen and looking at written notes and procedures;
 - (b) auditory tasks: listening for an auditory warning or listening to a verbal request;
 - (c) central processing tasks: making decisions about whether or not to intervene and selecting intervention strategies;
 - (d) manual tasks: typing codes on the keyboard, pressing button and moving the cursor with the tracker ball;
 - (e) verbal tasks: talking on the phone, talking to another signaller in the control room.
- (3) Modify the sequence of tasks in terms of modality: Tasks that employ the same modality, e.g. vision, cannot be performed in parallel; rather, one task in the modality must be completed before the next can commence. This defines the modality dependency of the model. Figure 4 shows a diagram for a simple activity. Notice that the tasks are sequenced over time (from left to right) and defined by particular modalities (the diagram only shows four modalities but this is due only to space constraints and one would normally show all relevant modalities in the diagram).
- (4) Allocate timings: Timings for the tasks are derived from a number of sources. For the purposes of this exercise, the timings used are based on the human-computer interaction (HCI) literature and are presented in Table 3. Just as standard times have been taken for component AIA tasks in the previous section to represent elemental human performance, so it is believed that unit tasks can be assigned generic values. These times are based on underlying aspects of human performance that are independent of context. In this model, the shortest time that can be found in the literature has been taken because the study's interest is in defining a minimum response time for the specific signal activity under consideration.
- (5) Construct the CPA diagram and calculate the critical path: Figure 4 shows the preferred method of representation, i.e. action-on-node. In this format, each node represents a task, as shown by the task name. The task has a defined duration (taken from Table 3). The early start time is defined as the time for all preceding tasks to be completed, i.e. it is the longest of the durations of previous tasks. The early finish time is the sum of the early start and duration. The late finish time is the time for all succeeding tasks to have started and the late start is the difference between late finish and duration. Slack is the difference between the two start times.

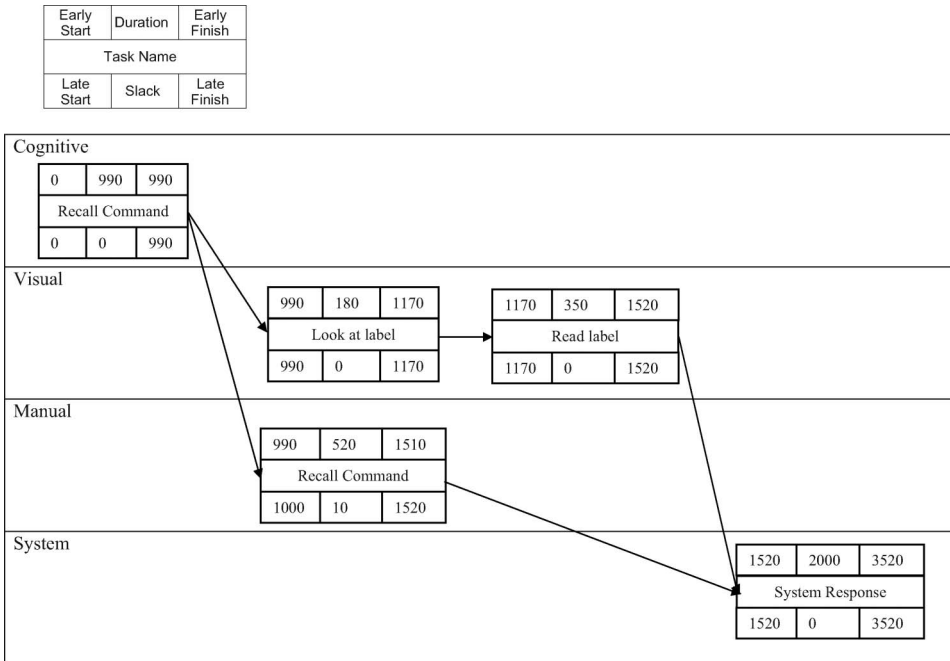


Figure 4. Critical path analysis representation.

If slack = 0, then the task lies on the critical path, otherwise the task is free to start any time between the early and late start times. The total time is, therefore, the sum of the durations of all tasks that lie on the critical path. A representation of a node is also shown in Figure 4.

Based upon the CPA presented in Figure 5, the predicted time to undertake the activities was around 19 s (i.e. 18.625 s). The model constructed for this analysis, and shown in Figure 4, assumes that the signaller would seek to change the signal first. Recall that, in section 1.2, the possible causes of action for the signaller would include changing the points, changing the signal and sending a stop message, and that the option for changing the points was deemed unlikely in this case. An alternative model was constructed (not shown in this paper due to space constraints) in which the actions of ‘change the signal’ and ‘send a stop message’ are exchanged in the time sequence. This second model resulted in a time of around 24 s and the model that resulted in the shortest time was chosen. According to Lawton and Ward (2005): ‘...there is no evidence that the signaller sent an emergency stop message to the Turbo (train 1) driver’ – although there is an indication that the request for such a message was recorded on verbal recordings. This would seem to corroborate the assumption that the signaller opted for the fastest course of action in this case.

This timing, of around 19 s, is derived by following the critical path task, which is summarised in Table 4. These analyses seem to support the evidence in the previous section, that it would take a signaller around 19 s to change signal SN120 from green to red in response to the alarms associated with the runaway train at Ladbroke Grove. The

Table 3. Estimates of activity times from the literature on human-computer interaction.

Activity	RT (ms)	Source	Times used in the model
Read (alarm message, headcode, etc)			
Glance at simple information	180	Olsen and Olsen (1990)	Look at screen 180
Read short textual descriptions	1800	John and Newell (1990)	Interpret information 340
Recognise familiar words or objects	314-340	Olsen and Olsen (1990)	
Hear (auditory warning)	300	Graham (1999)	Hear auditory 'tweet' 300
Search (screen for alarm or train(s))			
Checking or monitoring or searching	2700	Baber and Mellor (2001)	Search VDU for new information 2300
Scanning, storing and retrieving	2300-4600	Olsen and Olsen (1990)	Primed search (faster than search) 1300
Primed search	1300-3600	estimated	
Diagnosis or decision			
Mental preparation for response	1350	Card <i>et al.</i> (1983)	
Choosing between alternative responses	1760	John and Newell (1990)	Diagnosis activity 990
Simple problem solving	990	Olsen and Nielson (1988)	
Basic Cognitive Operation	50	Gray <i>et al.</i> (1993)	Manipulate information 150 (3 steps)
Response			
Speak (e.g.: 'We've got a SPAD')	100 per phoneme 1112	Hone and Baber (2001) Average from speaking the phrase 10 times	1112 for 'We've got a SPAD' 1500 for 'Send Message - Stop HST'
Move hand to tracker ball or keyboard	214-400 320	Card <i>et al.</i> (1983) Baber and Mellor (2001)	Move hand to trackball 320
Move tracker ball to target item	1500	Olsen and Olsen (1990)	
Move cursor via tracker ball 100 mm	1245	Baber and Mellor (2001)	Move trackball 1245
Press key (e.g. ACK or CANCEL key)	200 80-750 230	Baber and Mellor (2001) Baber and Mellor (2001) Card <i>et al.</i> (1983) Olsen and Olsen (1990)	Press button 200 [total 520 with 'move hand' time]
Type headcode			
Average typist (40 wpm)	280	Card <i>et al.</i> (1983)	
Typing random letters	500	Card <i>et al.</i> (1983)	
Typing complex codes	750	Card <i>et al.</i> (1983)	
Auditory processing (e.g. speech)	2300	Olsen and Olsen (1990)	
Switch attention from one part of a visual display to another	320	Olsen and Olsen (1990)	

VDU = visual display unit; SPAD = signal passed at danger; HST = high speed train.

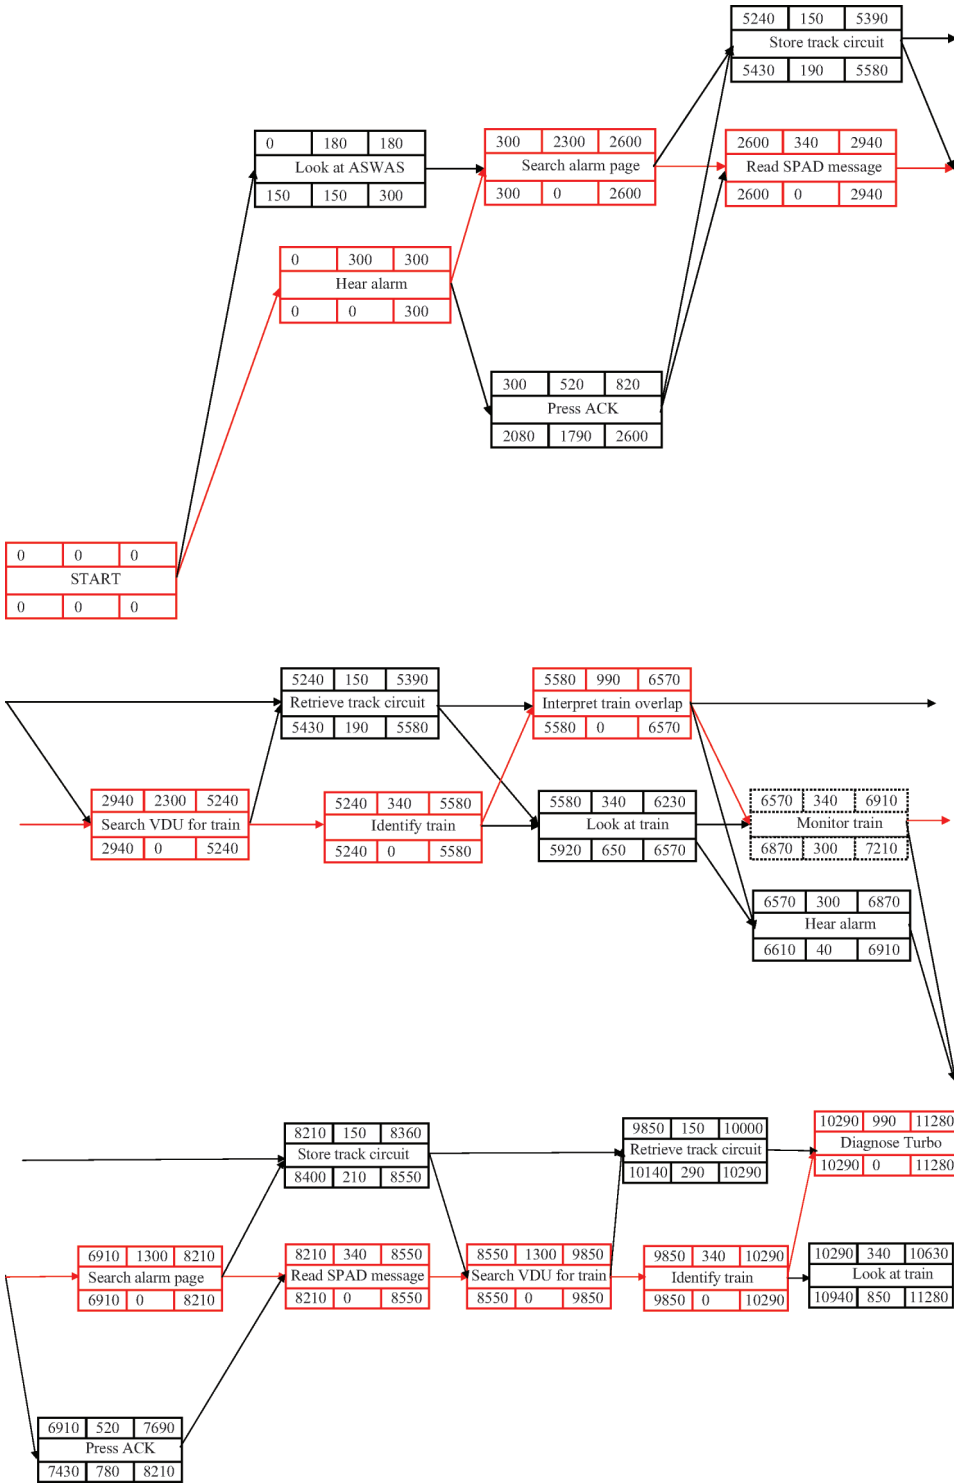


Figure 5. Critical path analysis of the signaller's activities. SPAD = signal passed at danger; VDU = visual display unit.




Figure 5. (Continued).

CPA timings were based on published data of human response times in particular activities, with the exception of the vocalisations. As with the AIA analysis in the previous section, the CPA analysis demonstrates that the response to the SPAD warnings is rather more involved than simply pressing a few keys. Most of the activities are visual and cognitive, which account for 24 out of a possible 30 activities. Manual tasks account for only four of the activities, which reinforces the ideas that alarm handling is predominantly a cognitive activity. Interestingly, the pressing of the ACK key is not on the critical path on either occasion and it has a relatively large float time associated with it, suggesting that this is not a crucial factor in the alarm handling performance of the signaller. Analysis of the 'micro-strategies' (cf. Gray and Boehm-Davis 2000) suggests that the performance of the signaller is less to do with key pressing and more to do with cognitive and visual aspects of the tasks, as there are more of these activities on the critical path. Thus, it can be

Table 4. Critical path tasks and timings.

Tasks	Start time	Duration	End time
Hear alarm	0	300	300
Search alarm page	300	2300	2600
Read SPAD message	2600	340	2940
Search VDU for train	2940	2300	5240
Identify train	5240	340	5580
Interpret train overlap	5580	990	6570
Monitor train	6570	340	6910
Search alarm page	6910	1300	8210
Read SPAD message	8210	340	8550
Search VDU for train	8550	1300	9850
Identify train	9850	340	10290
Diagnose Turbo	10290	990	11280
Search for conflict	11280	2300	13580
Diagnose conflict	13580	990	14570
Plan to change SN120	14570	990	15560
Locate SN120	15560	1300	16860
Move cursor	16860	1245	18105
Press CANCEL	18105	520	18625

SPAD = signal passed at danger; VDU = visual display unit.

concluded that delays in the visual and cognitive activities are likely to have a bigger impact on performance than delays in the alarm acknowledgement activities. The low probability of the SPAD alarm relating to an actual SPAD (cf. Getty *et al.* 1995) might explain why the signaller's initial diagnosis assumed that the driver of the Turbo would come to a halt after passing signal SN109. All previous experiences of the track occupation alarm had shown this to be the case. It is likely, therefore, that the signaller revised this diagnosis on receipt of the second alarm.

4. Conclusions from the modelling of alarm handling

In conclusion, the research presented in this paper suggests that it is possible to consider a wide range of studies under a unifying model of human alarm handling. The data in the model present the range of response times that may be expected. The CPAs show that a conservative estimate of the signaller's response time would be around 19 s. Evidence from alarm response time data in the published literature (i.e. reports on simulated and laboratory studies in other domains) suggest a minimum response time of 17 s. Unlike the CPA, the AIA is a generic model of activity rather than being specific to any particular circumstance. Taken together, the CPA and AIA models provide convergent evidence from divergent modelling approaches. It is interesting to note that they arrived at approximately the same response time, despite their differing methodologies. The AIA model presents a more general model of human activity than the CPA model. Both models use data from the academic literature, rather than data specific to the railway industry. Whether or not the models would perform in a similar manner in another domain remains to be seen. On the basis of the evidence from the cognitive CPA and the published literature in other domains, it is concluded that the signaller's response time of approximately 18 s appears to have been not just reasonable, but commendable, under the circumstances. This research offers some support for the use of time-based modelling when evaluating the reaction time of people in emergencies. Future research should

consider whether these approaches could be used in a predictive manner, in order to anticipate likely reaction times under defined sets of circumstances.

References

- Baber, C. and Mellor, B., 2001. Using critical path analysis to model multimodal human-computer interaction. *International journal of human-computer studies*, 54, 613–636.
- Bliss, J.P., 1997. Alarm reaction patterns of pilots as a function of reaction modality. *International journal of aviation psychology*, 7, 1–14.
- Bliss, J.P. and Fallon, C.K., 2004. Active warnings II: False warnings. In: M. Wolgalter, ed. *The handbook of warnings*. Hillsdale, NJ: Lawrence Erlbaum Associates, 231–242.
- Bliss, J.P., Gilson, R.D., and Deaton, J.E., 1995. Human probability matching behaviour in response to alarms of varying reliability. *Ergonomics*, 38, 2300–2312.
- Butcher, R., Grashel, M., and Lufaila, T., 1993. The effects of boredom on reaction time in the cockpit. In: R.S. Jensen and D. Neumeister, eds. *Proceedings of the seventh international symposium on aviation psychology – volume 2*, 22–29 April. Columbus, OH: Ohio State University, 1045–1048.
- Card, S.K., Moran, T.P., and Newell, A., 1983. *The psychology of human-computer interaction*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Cullen, Rt Hon Lord, 2000. *The ladbrooke grove rail inquiry. Part 1 Report*. Norwich: HSE Books, HMSO.
- Farrington-Derby, T., Wilson, J.R., Norris, B.J., and Clarke, T., 2006. A naturalistic study of railway controllers. *Ergonomics*, 49, 1370–1394.
- Getty, D.J., Swets, J.A., Pickett, R.M., and Gonthier, D., 1995. System operator response to warnings of danger: a laboratory investigation to the effects of the predictive value of a warning on human response time. *Journal of experimental psychology: applied*, 1, 19–33.
- Graham, R., 1999. Use of auditory icons as emergency warnings: evaluation within a vehicle collision avoidance application. *Ergonomics*, 42, 1233–1248.
- Gray, W.D. and Boehm-Davis, D.A., 2000. Milliseconds matter: an introduction to microstrategies and to their use in describing and predicting interactive behaviour. *Journal of experimental psychology: applied*, 6, 322–335.
- Gray, W.D., John, B.E., and Atwood, M.E., 1993. Project Ernestine: validating a GOMS analysis for predicting and explaining real-world performance. *Human-computer interaction*, 8, 237–309.
- Hollywell, P.D. and Marshall, E.C., 1994. An experiment to support the design of VDU-based alarm lists for power plant operators. In: N.A. Stanton, ed. *Human factors in alarm design*. London: Taylor and Francis, 31–44.
- Hone, K. and Baber, C., 2001. Designing habitable dialogues for speech-based interaction with computers. *International journal of human computer studies*, 54, 637–662.
- John, B.A., 1990. Extensions of GOMS analyses to expert performance requiring perception of dynamic visual and auditory information, *CHI '90*. New York: ACM Press, 107–115.
- John, B.A. and Newell, A., 1990. Toward an engineering model of stimulus-response compatibility. In: R.W. Proctor and T.G. Reeve, eds. *Stimulus-response compatibility*. Amsterdam: North-Holland, 427–479.
- Klein, G.A., Calderwood, R., and Macgregor, D., 1989. Critical decision method for eliciting knowledge. *IEEE transactions on systems, man, and cybernetics*, 19, 462–472.
- Kragt, H., ed. 1983. A comparative simulator study of annunciator systems. In: *Operator tasks and annunciator systems*. The Netherlands: Eindhoven University of Technology, 44–59.
- Lawton, R. and Ward, N.J., 2005. A systems analysis of the Ladbrooke Grove rail crash. *Accident analysis and prevention*, 37, 235–244.
- Lockyer, K. and Gordon, J., 1991. *Critical path analysis and other project network techniques*. London: Pitman.
- Moray, N., 1997. Human factors in process control. In: G. Salvendy, ed. *Handbook of human factors and ergonomics*. New York: Wiley, 1944–1971.

- Moray, N. and Rotenberg, I., 1989. Fault management in process control: eye movements and action. *Ergonomics*, 32, 1319–1342.
- Olsen, J.R. and Nielsen, E., 1988. The growth of cognitive modeling in human–computer interaction since GOMS. *Human–computer interaction*, 3, 309–350.
- Olsen, J.R. and Olsen, G.M., 1990. The growth of cognitive modeling in human–computer interaction since GOMS. *human–computer interaction*, 5, 221–265.
- Reason, J., 1990. *Human error*. Cambridge: Cambridge University Press.
- Riera, B. and Debernard, S., 2003. Basic cognitive principles applied to the design of advanced supervisory systems for process control. In: E. Hollnagel, ed. *Cognitive task design*. Mahwah, New Jersey: Lawrence Erlbaum Associates, 255–281.
- Roth, E., Woods, D.D., and Pople, H.E., 1992. Cognitive simulation as a tool for cognitive task analysis. *Ergonomics*, 35, 1163–1198.
- Shorrock, S.T. and Scaife, R., 2001. Evaluation of an alarm management system for an ATC centre. In: D. Harris, ed. *Engineering psychology and cognitive ergonomics – volume five: Aerospace and transportation systems*. Aldershot: Ashgate, 221–230.
- Singer, G. and Dekker, S.W.A., 2000. Pilot performance during multiple failures: an empirical study of different warning systems. *Transportation human factors*, 2, 63–76.
- Spence, C. and Driver, J., 1997. Audiovisual links in attention: implication for interface design. In: D. Harris, ed. *Engineering psychology and cognitive ergonomics*. Avebury: Ashgate.
- Stanton, N.A., 1994. *Human factors in alarm design*. London: Taylor and Francis.
- Stanton, N.A. and Baber, C., 1995. Alarm initiated activities: an analysis of alarm handling by operators using text-based alarm systems in supervisory control systems. *Ergonomics*, 38, 2414–2431.
- Stanton, N.A. and Baber, C., 1997. Comparing Verbal Media for Alarm Handling: speech versus textual displays. *Ergonomics*, 40, 1240–1254.
- Stanton, N.A. and Baber, C., 2006. The Ergonomics of command and control. *Ergonomics*, 49, 1131–1138.
- Stanton, N.A. and Edworthy, J., 1999. *Human factors in auditory warnings*. Aldershot: Ashgate.
- Stanton, N.A. and Stammers, R.B., 1998. Alarm initiated activities: matching formats to tasks. *International journal of cognitive ergonomics*, 2, 331–348.
- Stanton, N.A., Stewart, R., Harris, D., Houghton, R.J., Baber, C., McMaster, R., Salmon, P., Hoyle, G., Walker, G., Young, M.S., Linsell, M., Dymott, R., and Green, D., 2006. Distributed situation awareness in dynamic systems: theoretical development and application of an ergonomics methodology. *Ergonomics*, 49, 1288–1311.
- Schweickert, R., 1978. A critical path generalization of the additive factor method. *Journal of mathematical psychology*, 18, 105–139.
- Trujillo, A.C., 2001. Response times in correcting non-normal system events when collocating status, alerts, procedures and controls. In: *People in control: The international conference on human interfaces in control rooms, cockpits and command centres*, 19–21 June. UMIST, Manchester, UK. IEE conference publication number 481. London: IEE, 8–13.
- Tyler, R.R., 1997. *Response management timed to real-world alerts*. Orlando, FL: University of Central Florida.
- Vidulich, M.A. and Tsang, P.S., 2007. Methodological and theoretical concerns in multitask performance: a critique of Boles, Bursk, Phillips and Perdelwitz. *Human factors*, 49, 46–49.
- Walker, G.H., Gibson, H., Stanton, N.A., Baber, C., Salmon, P., and Green, D., 2006. Event Analysis of Systemic Teamwork (EAST): A novel integration of ergonomics methods to analyse C4i activity. *Ergonomics*, 49, 1345–1369.
- Wickens, C.D., 1984. Processing resources in attention. In: R. Parasuraman and D.R. Davies, eds. *Varieties of attention*. New York: Academic Press, 63–101.
- Wilson, J.R. and Rajan, J.A., 1995. Human-machine interfaces for systems control. In: J.R. Wilson and E.N. Corlett, eds. *Evaluation of human work*. London: Taylor and Francis, 357–405.