Exploring the Psychological Factors involved in the Ladbroke Grove Rail Accident

Prof. Neville A. Stanton[1] *and Dr Guy H. Walker*[2]

[1]*School of Civil Engineering and the Environment, University of Southampton, Southampton [UK]SO17 1BJ*

[2]*School of the Built Environment, Heriot-Watt University, Edinburgh [UK]EH14 4AS.*

Abstract

Ten years after the event and the question as to exactly why a driver passed a signal at danger to cause the Ladbroke Grove rail disaster is still an open one. This paper uses the literature on human error and cognition, combined with critical path analysis, to provide further insight. Five aspects of train operation are drawn out of the known facts surrounding the incident: custom and practice in the use of the Driver's Reminder Appliance, operation and use of the Automatic Warning System, the sequence of signalling information, methods of supplying route information, and speed restrictions. Associated with each are several important human factors issues which, combined, give rise to five potential explanations. Critical path analysis is used to map these explanations onto the known facts of the situation. It is suggested that the proximal cause of the Ladbroke Grove rail crash was a combination of an association-activation error and a mode error (leading the driver to mistakenly assume he had activated the Reminder Appliance) together with a loss-of-activation error (the driver failing to remember that a previous signal was showing caution) and a data-driven-activation error (by associating an in-cab warning to the wrong external source). The findings support the original inquiry recommendations, but also go further into predictive methods of detecting problems at the human/transport system interface.

Keywords: SPAD; Schema; Error; Event Analysis; Critical Path Analysis

Contact: N.Stanton@soton.ac.uk

Introduction

Ladbroke Grove has become shorthand for Lord Cullen's inquiry report whose recommendations had a wide ranging impact on the UK railway industry. Ladbroke Grove itself is a suburban district located in West London through which the London to West Country main line runs. It is approximately two miles from the line's terminus, London Paddington. It was at this location, at 08:09 on the 5th October 1999, that an outbound three car commuter train collided with an eight coach high speed train at a combined speed of 130mph, injuring over 400 people and causing the deaths of 31, including both drivers. Although accidents like this are extremely rare on the UK rail network, this incident was the result of a long-standing human factors problem – a Signal Passed At Danger (SPAD).

A SPAD is when a train driver passes a signal displaying a stop aspect. To the outside observer, problems like this appear deeply perplexing. How is it possible that a driver can acknowledge up to three in cab warnings, override three automatic applications of the emergency brake, but still proceed past a point at which he or she is required to stop? Within the scientific community, and despite a considerable body of research, this category of problem is far from fully resolved: SPADs have occurred for well over a hundred years and still occur to this day, with many examples remaining stubbornly resistant to a wide range of well intentioned safety measures. SPADs are more than specific incidents, they also represent a perplexing 'category' of human factors problem, a category that is to be found in virtually all transport domains. Road transport, for example, has 'unintended acceleration' (e.g. Schmidt, 1989) and 'highway hypnosis' (e.g. May & Gale, 1998). Aviation has 'controlled flight into terrain' (e.g. Shappell & Wiegmann, 2003), 'mode errors' (e.g. Endsley & Kiris, 1995) and 'automation surprises' (e.g. Sarter & Woods, 1997). What all these cases share in common are functioning safety systems which are defeated or ignored, leading to perfectly serviceable vehicles being placed in highly dangerous conditions.

The problem with SPADs (and SPAD-like phenomenon) is that they appear both paradoxical yet simple. After all, the task of the driver is to spot a signal displaying a stop indication and apply the brakes. But this apparent simplicity is misleading. In the case of the Ladbroke Grove rail crash a clue to the presence of deeper, more psychological issues is revealed by the fact that the inquiry team had no doubt that the driver *believed* he had a proceed aspect showing at the last signal he passed. The report, furthermore, alludes to the fact that "SN109 [the code given to the signal that was passed at danger] was a multi-SPAD signal" (Cullen, 2001, pg 2) meaning that the driver in question was not alone in passing this particular signal at danger.

Although the Ladbroke Grove accident happened ten years ago SPADs still occur. An opportunity exists to re-visit Ladbroke Grove and go beyond the official enquiry report into the misleadingly simple question of 'why' the driver in question passed a signal he was not meant to. The inquiry report invites such analysis. In the words of Lord Cullen, the inquiry chairman, "...it is not possible for me to arrive at a full explanation of [the driver's] actions" (Cullen, 2001, para 5.111, p. 80). The question of specifically 'why' the driver passed the signal is therefore an open one.

Several authors have responded to this opportunity previously and have used the great depth of information provided by the formal investigation to try and go further. Lawton and Ward (2005) and Santos-Reyes and Beard (2006) both adopt a systems approach, shifting the level of analysis to encompass the combined effects of the operational and organisational environment. Both conclude, quite rightly, that train crashes like this cannot be distilled into a single causal factor and that a systemic failure took place which enabled direct, indirect, latent and active factors to propagate through the interacting system elements. A paper by Evans (2005) presents an interesting counterpoint. Despite the severity of the Ladbroke Grove accident, and the public perceptions which surrounded it, rail safety continued its improving trend regardless. This brings into sharp relief a particular facet of SPADs as a category of human factors problem; their low probability but extremely high cost.

This unfortunate probability/cost trade-off, requires a particularly fine grained analysis. The present paper aims to make a contribution of this sort by shifting the systemic level of analysis closer to the accident itself. It does this by looking at the interacting psychological mechanisms underpinning the actions and behaviour of the driver who committed the SPAD. The paper combines the known facts of the Ladbroke Grove incident with the latest knowledge of human factors in order to develop a number of plausible explanations for the accident, which can then be evaluated. This probing is intended to go further than the inquiry report and respond to the tacit invitation therein: simply, why did the driver believe SN109 was displaying a proceed aspect? Given that a similar class of problem occurs in most other high reliability domains, additional insight into the underlying psychological mechanisms is as valuable as it is transferrable. The remainder of the introduction deals with the events leading up to the crash and wider aspects of rail operations which have a bearing on understanding its cause.

Timeline of Events

In order to understand the situation and context surrounding the Ladbroke Grove rail accident it is necessary to re-play the train journey in question.

It is 08:05 on the 5th October 1999. A 'Thames Turbo' multiple unit commuter train, identifying number 1K20, is waiting on platform 9 of London's Paddington station. On board are the driver and 147 passengers.

At 08:06 the On-Train Monitoring and Recording (OTMR) equipment and records from the computerised signalling centre show that the train passed the signal (code number SN17) at the end of the platform. SN17 was showing a green aspect (proceed) and an illuminated number four (indicating that the train would be routed to line four). So called 'starting signals' at the ends of platforms in terminus stations like Paddington do not have the Automatic Warning System (AWS) fitted to them, thus no audible or visual indication of SN17's signal aspect would be provided in the cab.

Between SN17 and SN43 (the next signal in sequence) the driver gradually increased speed to 34 mph. He did this by moving the speed control successively up through seven speed notches in the following order: 1, 2, 3, 4, 7. Two speed restriction signs were passed, one reading 40mph and later one reading 60mph.

The next signal, code number SN43, was showing a green aspect (proceed) and an illuminated number four, indicating that the train would be continuing along line four. 200 yards from SN43 the Automatic Warning System (AWS) triggered an audible 'chime' in the cab, confirming to the driver that SN43 was showing a proceed aspect and that he need not perform any actions additional to those already being performed. From this point, speed increased further to 44 mph.

The next signal had the code number SN63 and it was showing a double yellow aspect. This means 'preliminary caution' and that two signalised sections of track ahead are clear but that the third is occupied. This signal was not presenting any route information meaning that the driver could assume he was staying on the same track. Approximately 200 yards before the signal, the AWS triggered a horn sound in the cab and gave the driver six seconds to press the AWS cancel button in order to prevent an automatic application of the train's brakes. The driver correctly pressed the button causing the horn sound to stop and the in-cab AWS display to change colour. The display now provided an ongoing reminder that the last AWS activation was a warning. Shortly after SN63, the driver put the speed control in the neutral position and applied the brake at level one for about seven seconds, enough to reduce speed from 44 to 39 mph as it approached the next signal in sequence, SN87.

SN87 was displaying a single yellow aspect meaning 'caution', only the next signalised section of track is clear and that the next signal is likely to be at danger. Route information was also being displayed on an associated 'position light junction indicator' where a row of white coloured lights was pointing left. This informed the driver that the train was going to be routed onto the adjacent left-hand line after the next signal. Again, 200 yards from SN87 the AWS horn sounded in the cab and a request for an automatic application of the brakes was cancelled. The in-cab AWS display remained in its previous state and provided an ongoing reminder of the last AWS indication received (i.e. a warning). At this point, it would have been normal practice for the driver to start slowing down in order to stop in advance of the next signal. In this case, however, the train continued to coast for another 30 seconds.

The next signal, SN109, was displaying a red aspect meaning stop. Approximately 200 yards before SN109 the train passed over the associated Automatic Warning System (AWS) equipment on the track. A warning horn was received in the cab and a request for an automatic brake application was once again cancelled. The in-cab AWS display remained in the state that had been triggered by SN43 (two signals previously) and the driver moved the speed controller from notch zero (coasting) to notch 5. This kept the train cruising at around 38mph.

Just ahead of SN109 were two speed restriction signs located at the trackside. The first was a round sign warning of an 85mph restriction, the second was a triangular warning sign referring to a 70mph restriction on the crossover to the right. Concurrent with the signs, the driver moved the speed controller from notch five to notch seven (maximum). The train accelerated past SN109 at danger, the only in-cab warning for this event having been cancelled by the driver approximately two seconds previously. The only alarms raised as the train entered a section of track it was not meant to were those in the remote signalling centre 15 miles away in Slough. The signaller in charge of monitoring the situation made an attempt to halt an approaching train by turning a signal with the code number SN120 from green (proceed) to red (stop). The driver of the train travelling in the opposite direction, a high speed train comprised of eight coaches with a power-car at each end, immediately applied the brakes.

Thirty six seconds after selecting speed control notch seven, with the Thames

Turbo's speed now approaching 50mph, the On-Train Monitoring and Recording (OTMR) equipment showed that the driver moved the control from notch seven straight to notch zero (the neutral position) followed immediately by an application of the emergency brake.

Five seconds later the Thames Turbo collided with the High Speed Train near to signal SN120. The combined speed of impact was 130mph.

The events and surrounding context have been assembled into Figure 1 which is based on data from the On-Train Monitoring and Recording (OTMR) equipment and the known layout of the rail infrastructure.

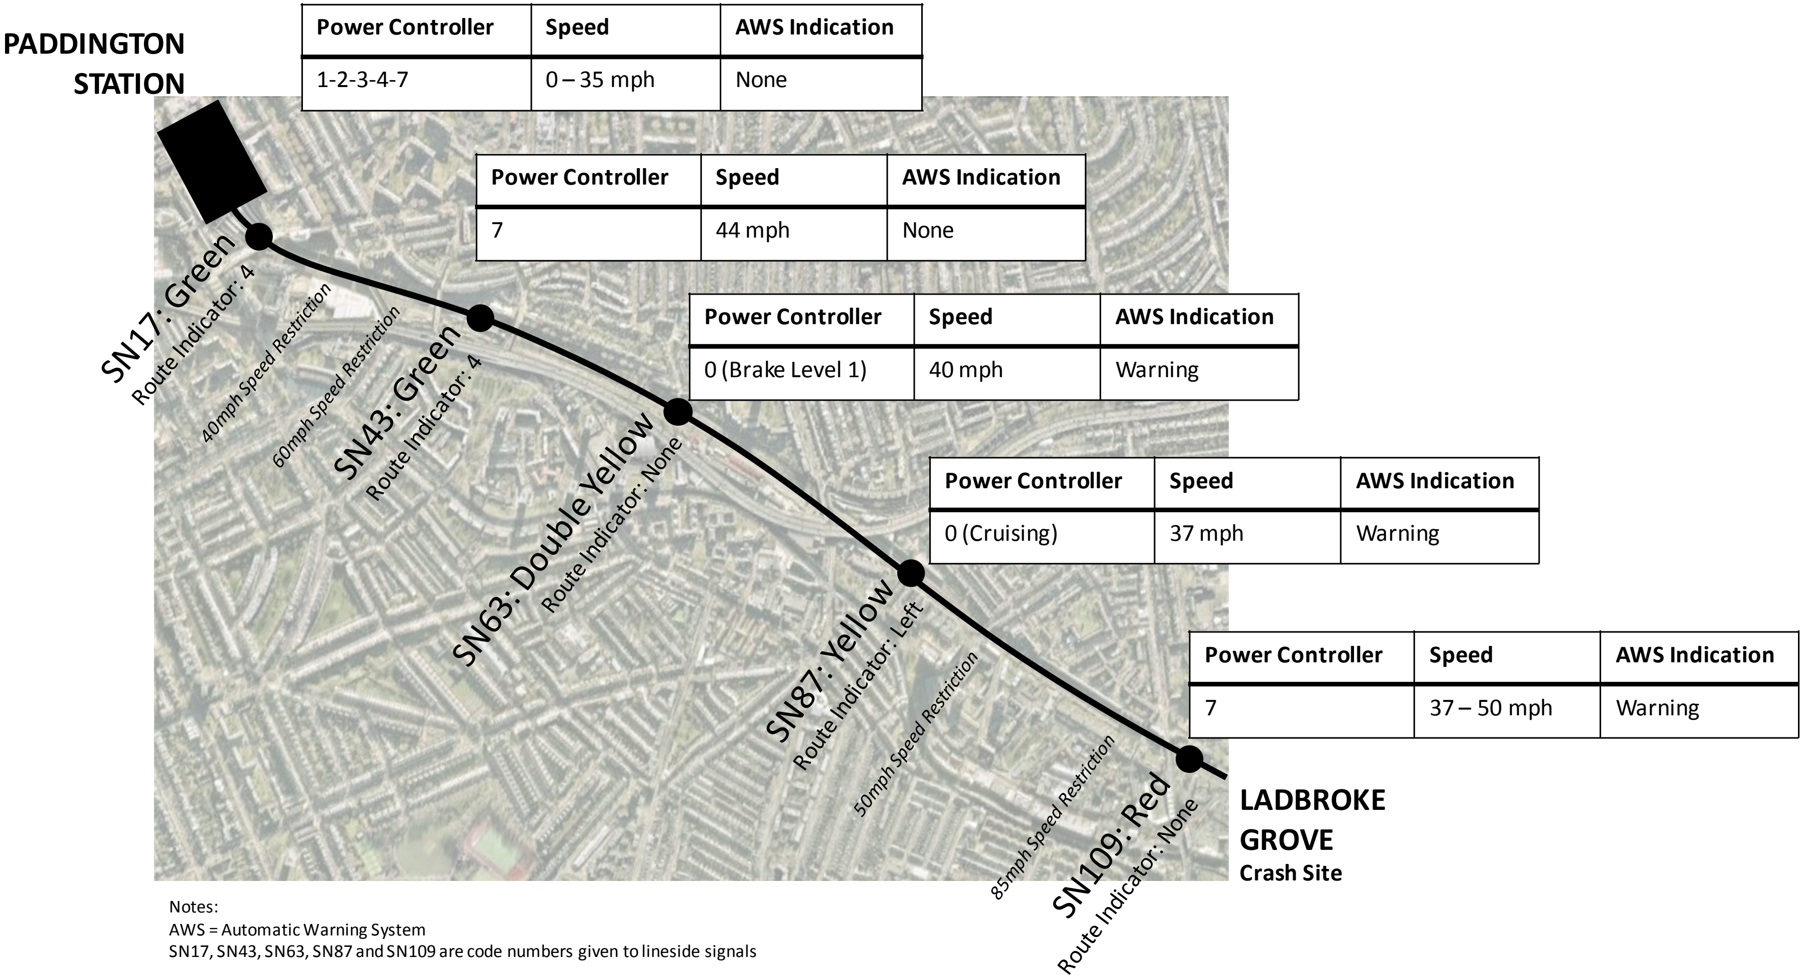

Figure 1 – Schematic of Ladbroke Grove event timeline

Defences in Depth

The railway system, like all safety-critical systems, has multiple layers of defences each with the potential to prevent an accident (Reason, 1990, 1988). Multiple defences have the benefit that should any one defence fail, there will be other defences in place to prevent anything untoward occurring. Reason argues that in reality these defences tend to interact dynamically with local demands and human error. As shown in Table 1, there were at least seven such defences in place on the day of the incident, with the driver managing to defeat the last five. Specifically, had the driver not selected the speed control notch 5, or had applied the DRA after SN87, or correctly identified the automatic warning at SN109, or not selected the speed control notch 7, it is probable that the accident either would not have occurred or not been as devastating as it was.

Table 1 – Analysis of defences in depth

Defences-in-depth	Driver response	Train response
1. SN63 AWS horn	Cancel AWS	AWS horn cancelled
2. SN63 at double yellow	Select speed control notch 0 and brake level 1 applied and released	Train coasting at 45mph then braking to 41mph and coasting to 39mph
3. SN87 AWS horn	Cancel AWS	AWS horn cancelled
4. SN87 at single yellow	Allow train to continue coasting for 30 seconds, then select speed control notch 5	Train coasting at 39mph then cruising at 38mph
5. Driver Reminder Appliance	Not applied	Train coasting but driver is still able to draw power
6. SN109 AWS horn	Cancel AWS Select speed control notch 7	AWS horn cancelled Train accelerating from 38mph to 50mph
7. SN109 at red	No response	Train accelerating from 38mph to 50mph

The Ladbroke Grove accident clearly follows a trajectory through each layer of defence which, for a host of interacting local and global reasons, had aligned in such a way as to permit the crash to take place. The multiple defences in the system offer an explanation as to why these events are so rare and why problems like SPADS appear, on the surface, to be so unique and perplexing.

Background to Driver Operations

At first sight, it seems inconceivable that a driver can acknowledge the receipt of no less than three warnings, each one requiring the cancellation of an automatic emergency brake request, yet still fail to stop at a red signal. But as numerous studies have shown, it does happen (e.g. Hall, 1999; Andersen, 1999; Williams, 1977; Van der Flier & Schoonman, 1988; Gilchrist, 1990; Gibson, 1999; Downes & Robinson, 1999; May, Horberry & Gale, 1996; Wright, Ross & Davies, 2000; Hill, 1995). All of these sources allude to the broader context of train driving operations as an explanation for the counter-intuitive nature of SPADS, a context which can now be explored in more detail.

We begin with a lay-man's description of train driving provided by a driver in a focus group:

"The best way to describe [driving a train] is a bit like having to drive your car on ice – at 100 mph – with only the handbrake to slow you down." (McLeod, Walker & Moray, 2003: 2005).

It takes approximately two miles for a train travelling at 125mph to be brought to a complete halt. As a result, it is comparatively rare for a driver to begin braking from high speed for an event that can, at that moment, be clearly seen ahead. Rather like pilots who make extensive use of their instruments, train drivers rely heavily on information conveyed to them by line-side signals. This combines with in-depth route knowledge to generate a response (like braking or acceleration) which is appropriate for the circumstances. What this means is that the kind of feedback control encountered in car driving, where a driver will see a red traffic light and apply the brakes in order to stop before it, does not often apply. Neither do the simple Input (i.e. red traffic light) –Processing (i.e. what is the braking distance) -Output (i.e. apply the brakes) models of cognition frequently encountered in human factors apply particularly well to SPADs. This is because 'the input' is derived indirectly and constructed from an expectation of what is ahead. Train driving is characterized by large amounts of this predictive, feed-forward control, with drivers being actively required to *believe* that certain system states are imminent based on what they know currently: they could not control a train at speed otherwise. The key in railway operations is to ensure those beliefs and predictions, and the corresponding behaviours, are correct. Put simply, because drivers have to prepare to stop for up to two miles, and several signals in advance, the problem becomes less to do with signals being passed at danger and more to do with other events, principally the succession of previous signals passed at caution.

Psychological Mechanisms in Error Production

In order to relate train driving operations to the system defences that were defeated in the run up to the Ladbroke Grove crash, it is necessary to consult the psychological literature on human error, working memory and attention.

Interactive Schemata

Norman (1981) reports research on the categorisation of action slips based on an analysis of 1,000 incidents. Underpinning the analysis was a psychological theory of schema activation. He argued that action sequences are triggered by knowledge structures organised as memory units and called schemas. The mind comprises a hierarchy of schemas that are invoked (or triggered) if particular conditions are satisfied or events occur. Related to this is Neisser's (1976) seminal work on 'Cognition and Reality', where he puts forward a view of how human thought is closely coupled with a person's interaction with the world. He argues that knowledge of how the world works (e.g. schemas) leads to certain kinds of information being expected, which in turn directs behaviour to seek it out and provide a ready means of interpreting the situation. During the course of events, as the environment is sampled, the information serves to update and modify the internal cognitive schema of the world, which will again direct further search.

The perceptual cycle can be used to explain human information processing in train driving. Assuming that the individual has the correct route knowledge for operating the train, their schema will enable them to anticipate events (such as the signals and signs they expect to see and routes they expect to take), search for confirmatory evidence (such as looking at the signal aspect, trackside objects, routing information, speed indicators and notes), direct a course of action (such as braking and accelerating) and continually check that the outcome is as expected (such as the slowing down or the speeding up of the train). If they uncover some data they do not expect (such as an unexpected warning or track routing) they are required to source a wider knowledge of the world to consider possible explanations that will direct future search activities. The completeness of this model is in the description of process (the cyclical nature of sampling the world) and product (the updating of the world model at any point in time).

Error Taxonomy

The interactive schema model works well for explaining how we act in the world and what can go wrong. As Norman's (1981) research shows, it may also explain why errors occur as they do. If, as schema theory predicts, action is directed by schemata, then faulty schemata, or faulty activation, will lead to erroneous performance. As Table 2 shows, this can occur in at least three main ways. First, we can select the wrong schema due to misinterpretation of the situation. Second, we can activate the wrong schema because of similarities in the trigger conditions. Third, we can activate schemas too early or too late.

Table 2 - Taxonomy of errors with examples

Taxonomy of errors	Examples of error types
Errors that result from the formation of intention	Mode errors: erroneous classification of the situation
	Description errors: ambiguous or incomplete specification of intention
Errors that result from faulty activation of schemas	Loss-of-activation errors: schemas that lose activation after they have been activated
	Data-driven-activation errors: external events that cause the activation of schemas
	Association-activation errors: currently active schemas that activate other schemas with which they are associated
	Capture errors: similar sequences of action, where stronger sequence takes control
Errors that result from faulty triggering of active schemas	Blend errors: combination of components from competing schemas
	Premature activation errors: schemas that are activated too early
	Failure to activate errors: failure of the trigger condition or event to activate the schema

Reference to the known facts of the Ladbroke Grove crash described above, and to the error types shown in Table 2, shows only weak evidence for a description error, association-activation error, capture errors or blend/premature/failure to activate errors. The driver demonstrated no apparent ambiguity in his intention and made no attempt to correct his actions. This, in turn, supports a view that the actions being performed were consistent with individual schemas. The known facts (particularly the OTMR data) also support the view that actions were carried out in response to, and apparently triggered by, external stimuli. What the known facts do communicate, however, is stronger evidence for a loss-of-activation error, followed by a mode error, followed again by a data-driven-activation error. These error types also fit the timeline of events well: for example, the driver fails to remember that he has passed a single yellow at SN87, perhaps confuses the coasting of the train with application of the DRA and then, possibly, mistakes the speed restriction as the cause of the activation of SN109's in-cab automatic warning. Specific explanations will be developed shortly, for the time being it seems reasonable to suspend further consideration of the former error types (description, association-activation, capture, blend, premature, failure to activate) in favour of the latter three (loss of activation, mode and data-driven). These are taken forward in the next sections.

Loss of Activation Error

There are two main ways in which information is lost from working memory: displacement and decay. In displacement, older information is displaced from working memory by newer information. Working memory has upper limits on capacity of (see Miller's classic 1956 work). This means that as the journey progresses, information about signals and signs passed will be displaced by more recent signs and signals. In decay, information is forgotten as time elapses. The more information held in working memory the faster it decays. Individual chunks (such as remembering a single yellow aspect at a previous signal) have a half-life in working memory of approximately 70 seconds (the delay after which recall is reduced by one half: see Card, Moran & Newell, 1986). Whilst this is well within the actual time taken by the driver to travel from SN87 to SN109 (the journey took approximately 30 seconds) the half-life is reduced by a factor of ten when the numbers of chunks are tripled (such as remembering route information and speed restrictions in addition to remembering a single yellow aspect at the signal just passed). The loss of information from working memory is thus plausible in these circumstances.

Mode Error

Norman (1981) singles out mode errors as requiring special attention in the design of technological systems. He pointed out that the misclassification of the mode that the system was in could lead to input errors which may have serious effects. Woods et al (1994) state that mode errors occur when there is a breakdown in situation assessment. They argue that keeping track of the mode that a system is in increases the demands placed on memory and situation assessment. In the present case, the loss of activation error seems to have compounded the problem by giving the driver a mistaken belief as to the current mode. The selection of notch 5 on the speed control just before SN109 is consistent with a mode error, and the driver's mistaken belief that he had not passed a previous signal displaying caution and therefore would not be expecting to see a signal displaying danger.

Data Driven Activation Error

The third error, the attribution of an in-cab warning at referring to SN109 incorrectly to a nearby speed restriction sign, is consistent with the driver's almost immediate selection of speed control notch 7 (the data-driven-activation error). The concept of limited pools of attentional resources (Wickens, 1992) is relevant here. The basic premise is that allocation of attentional resources to one task will result in fewer resources available for another. For example, attentional resources focused on detecting hazards or checking the line ahead for crossovers will mean that there are fewer resources available for determining alternative courses of action and executing manual responses.

Attention has often been described using the metaphor of a searchlight (Barber, 1988). The direction of the driver's attention is like the beam of the searchlight and everything that falls within the beam of the searchlight is processed. The limits of human attention can cause problems in the train driving task in cases where the driver is required to focus on some stimuli rather than others, is distracted and/or too thinly divided across too many concurrent tasks. There are two main factors to be considered: attentional resources and direction of visual focus. These factors offer an explanation for the demands placed upon the driver at SN109. On hearing and cancelling the in-cab warning referring to SN109's signal aspect (i.e. red) the driver's visual attention may have focused on the speed restriction sign below the bridge. The information perceived and processed within the attentional "searchlight" is highly detailed, whilst the surrounding information tends to be perceived and processed in a peripheral manner. Whilst the driver may have been very aware of the speed restriction signs, he may only have been vaguely aware of the overhead lines and signal gantry. Of course, this goes against all driver training and common experience to the point of implausibility. Yet, after cancelling the AWS the driver behaved in a way consistent with this by selecting speed control notch 7. To do this he may have switched his attention to the controls and speed indicator inside the cab. By the time his attention was focused back on the line the gantry at SN109 might not have been within the attentional "searchlight", and the opportunity to fully appreciate the consequence of his actions lost. Prior experience at SN109's gantry, combined with artefacts of the driver's training and experience are also likely to have played a role, as discussed in the official inquiry report (Cullen, 2000).

So far we have provided details of the crash, the background to driver operations, the defences present in the system and the wider psychological literature on human error. Combined, these provide the basis for the next stage of analysis.

The next section presents five detailed explanations related to specific aspects of rail operations which contribute directly towards understanding why the driver believed SN109 was displaying an aspect which would have allowed him to proceed safely.

Explanations for Signal SN109 Passed at Danger

Explanation 1: The driver misread the aspect of signal SN109.

In the UK, railway signals communicate information about the upcoming route. In the context of the Ladbroke Grove incident (and most of the UK rail network) this information is communicated using coloured lights. Different signal aspects can be displayed (e.g. red, yellow, double yellow, green) which inform the driver of route occupancy, that is, how many signal-controlled sections of track ahead are clear of conflicting traffic/traffic arrangements. The task of the driver is to combine this information with in-depth route knowledge to decide what action is appropriate. For example, with certain types of rolling stock and in certain conditions a ‘double yellow’ signal aspect requires no action from the driver at all. In other circumstances, however, it could mean an immediate reduction in speed is required. This practice differs from many European countries which use speed-based signalling whereby a fixed behavioural response is required for any given signal aspect (e.g. a single yellow signal requires the driver to reduce speed to 40mph regardless of circumstances).

The signalling regime applicable to the Ladbroke Grove incident is four aspect colour light signalling. The four aspects (red, yellow, double yellow and green) are displayed by individual lens assemblies mounted in a common signal head. The sequence of information that can be presented, and the possible interpretations by the driver, are presented in Table 3.

Table 3 – Definition of the four main signal aspects from the Master Rule Book (RS&SD, 1999)

Colour of signal	Definition	Driver interpretation
Green	Proceed	Three signalised sections of track ahead are clear and it is safe to maintain or accelerate to the line speed.
Double yellow	Proceed (but either start slowing down or expect to have to start slowing down at the next signal)	Two signalised sections of track ahead are clear. Depending on the rolling stock and current speed, either begin braking or else maintain speed in anticipation of information conveyed by the next signal.
Single yellow	Proceed (but start slowing down in anticipation of having to stop at the next signal)	Only one signalised section of track ahead is clear. Continue or start to brake in order to be at an appropriate speed at which to stop when the next signal becomes visible.
Red	Stop	The next signalised section of track is not clear. Bring the train to a complete stop in advance of the signal.

An example of a conventional four lamp signal head is presented in Figure 2. Signal SN109 is also reproduced where it is important to note its unusual design, being shaped in the fashion of a mirrored 'L' with the red aspect on the horizontal leg of the 'L'.

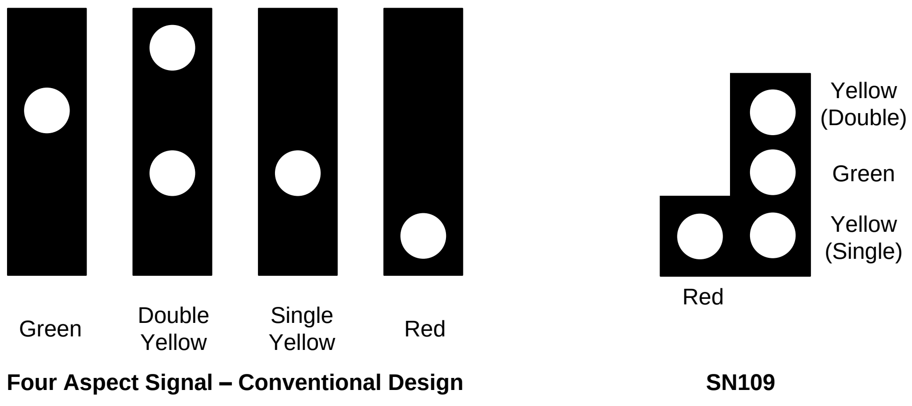

Figure 2 – Signal aspects for a four aspect colour light signal

The Ladbroke Grove inquiry concludes that the reason why the driver believed he had a proceed aspect at SN109 was due to "the poor sighting of SN109, both in itself and in comparison with the other signals on and at gantry 8, allied to the effect of bright sunlight at a low angle, […] which led him to believe that he had a proceed aspect and so that it was appropriate for him to accelerate as he did after passing SN87 [the previous signal]" (Cullen, 2001, pg. 2). Clearly, the critical human performance requirement for trackside signals is that they are not only visible but understandable. Railway Group Standards go into extensive detail regarding signal sighting and conspicuity, including the oft-quoted statement that "Signals shall normally be positioned to give drivers an approach view for a minimum of seven seconds and an uninterrupted view for at least four seconds" (GK/RT 0031, 2002). SN109 was visible to the driver for eight seconds, even with the presence of possible obstructions, suggesting a number of further factors relevant to the role of signalling in causing the Ladbroke Grove crash.

1. It is important to note that apart from signals at danger, there is no fixed behavioural response required of the driver in respect to the route signalling philosophy. The required and appropriate behaviour depends on many contextual factors and on the ability of the driver to keep track of them.
2. The driver's experience was a factor that the formal inquiry examined in detail. Research has shown that drivers in their first year of driving were most at risk from SPADs and that between 30 and 47% of SPADs were accounted for by drivers with less than five years experience (Downes & Robinson, 1999). The research goes on to show that the

problems are rarely to do with failures in mastering the control of the train, rather, the problem lies in misinterpretation of the information supplied to them. With just over a month of qualified driving experience the driver in question clearly fell into this 'at risk' group. His training and subsequent supervision were also criticized by the inquiry report.

3. The driver's experience of the signals on Gantry 8 (on which SN109 is mounted) is also important. It was stated in the inquiry report that the driver had only driven trains out of Paddington 20 times. No less than 19 of those times the signals the driver had encountered on Gantry 8 were showing a green/proceed aspect. The only other occasion was when the driver had encountered SN109 on one previous occasion, three days prior to the crash, in which SN109 cleared from red to green as the train approached. Therefore, on no previous occasion had the driver been required to stop at any of the signals mounted to Gantry 8. In addition, the dynamics of signal aspect presentation were such as to favour proceed aspects. Phenomenon like Driving Without Attention Mode (DWAM; May & Gale, 1998) show that it is possible for expectancy of a particular signal aspect to override the external cues to the contrary.

Whilst signal sighting and conspicuity play a key role in the immediate approach to SN109, inconsistencies in the driver's behaviour had started to occur before the signal was visible (i.e., when compared to a notional ideal driver who would have been preparing to stop the train from SN63 and SN87). This suggests that the driver may have started forming his belief about the upcoming aspect of SN109 long before sighting and conspicuity issues became a factor. Indeed, not only was the driver's behaviour towards the previous signal (SN87) inconsistent (i.e. he did not begin to slow down in the manner normally expected when in receipt of the AWS horn and double yellow at SN63, and again at the single yellow of SN87), so was that of a number of other drivers who had also experienced SPADs at SN109. When combined, it seems that although the proximal cause of the crash was a Signal Passed at Danger, the distal reason was due to a previous 'Signal Passed At Caution'.

Explanation 2: The driver mistook the position line junction indicator at position one for an indication that he was to be routed onto line one.

Where routes diverge, the route to be taken is communicated by an attendant line route information display. Two principle methods are used. 1) the driver is supplied with alphanumeric information (i.e. a letter or number that refers to particular lines the driver should take), or 2) the position line junction indicator with rows of lights that illuminate and point to the relevant line and/or direction the train should take. These two methods provide different types of information. The alphanumeric displays present the absolute route in line number(s) or letter(s) whereas the position line junction indicator displays the relative direction. Illustrations and mappings are provided in Figure 3:

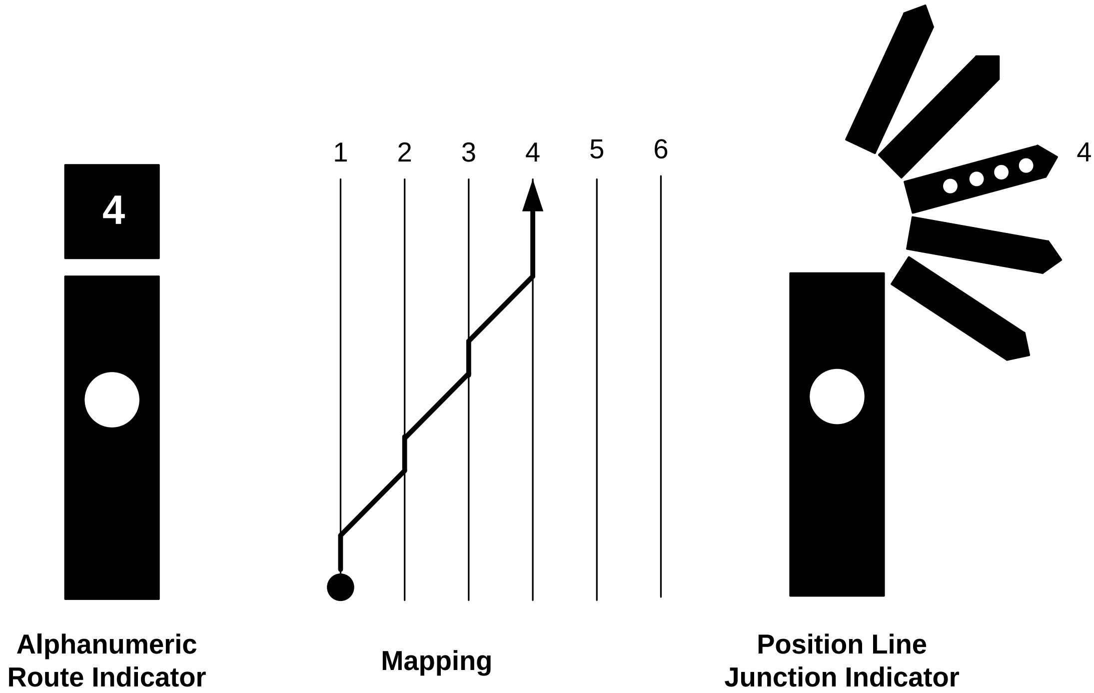

Figure 3 – Alpha numeric and position line indicators with accompanying mappings onto track layouts.

There are two important differences between the position line junction indicator displays and the alphanumeric displays. Firstly, no line information is supplied by the position line junction indicator display if the line remains unchanged, whereas the alphanumeric display has the potential to display the current as well as diverging line. Second, the position line junction indicator display provides relative line routing information, whereas the alphanumeric display provides ‘absolute’ information about the line to be taken. The driver may encounter both types of indicator display on a route, as was the case at Ladbroke Grove. As shown in Figure 2, the driver was presented with two alphanumeric line indications, one prior to a change of route (after SN17) and one serving as a reminder (SN43). This was followed by no line information at the next signal (SN63), because the route remained unchanged, followed by the position line junction indicator at position one at SN87.

If the driver formed a belief prior to SN109 coming into view that he had a proceed aspect, and subsequently did not see, or else misinterpreted the position line indication at SN109, then it is conceivable that an expectation could be formed that the train would not be diverging but proceeding on the straight ahead route. It is difficult to say, based on the evidence available, whether this happened or not. The facts of the situation are, however, that two different methods of position line display were used with the potential to cause confusion, and that the position line indicator at SN109 could in fact represent one of the defences in the system, albeit a rather weak one. If this is the case, then it would offer a partial explanation for the driver selecting notch 7 (full power) on his speed controller prior to SN109. In other words, the sequence of information presented to the driver may have activated schema which would not have prompted him to search for relevant cues in the environment.

Explanation 3: The driver forgot to apply the DRA after signal SN87.

The Driver Reminder Appliance (DRA) is an additional aid to help reduce the likelihood of a driver starting to move past a red signal, or a so-called Starting Against Signal (SAS) SPAD. Prior to performing station duties such as operation and monitoring of the train doors, interacting with the guard and so forth, the driver pushes the DRA button in the cab as a reminder that the signal ahead is red. The action of pushing the button causes a prominent red light to illuminate in the cab and prevents power from being drawn from the engine until it has been disengaged.

The DRA is designed to be used in stations and can only be applied after the driver has selected the speed control at notch zero, the neutral position. However, it is possible to activate it while the train is coasting past a signal displaying caution. It can also be activated during the braking manoeuvre for a red signal itself. In both cases, the speed controller is in the neutral position. Within several train operating companies the practice of using the DRA on-the-move had developed, and in fact was seen as best practice (see Cullen, 2001 pgs 65 & 78). Indeed, the driver who passed SN109 was trained by one of his on-the-job mentors to apply the DRA at signals displaying single yellow signal aspects. This he did consistently on the journey into Paddington Station on the morning of the accident, just ten minutes before the outward journey began. From this it seems reasonable to suggest that the driver would have normally applied the DRA at single yellow signals. Subsequent research into the use of DRA on the move reveals a number of problems:

1. The use of the DRA is not consistently applied at every caution signal. This lack of consistency could lead the driver to generate false expectations as to what signal is ahead. For example, if the DRA is usually applied at a single yellow signal, and if the driver does not apply the DRA at a single yellow, they may assume that the next signal will not be showing a red aspect.
2. The operation of engaging the DRA could be a significant distraction under certain operational scenarios. For example, the driver may have to attend to other in-cab warnings simultaneously, some of which are more safety critical than the DRA.
3. The DRA could be set too early in the sequence, meaning that there is not sufficient power to get the train to the next signal.
4. The setting and cancelling of the DRA in a succession of caution signals could lead the driver to get out of sequence, resetting the DRA rather than setting it (McCorquodale et al, 2000).

One explanation related to the Ladbroke Grove crash is that the driver forgot to apply the DRA at SN87. This certainly seems to be the case based on the available OTMR data. For drivers who used DRA 'on the move', application of the DRA would be strongly associated with the train coasting, and as the train was already coasting in the region of SN87, the driver could have mistakenly associated this with the DRA having been applied. Upon discovering that he could in fact draw power it may have changed his belief about the state of SN87 and the upcoming state of SN109: a classic mode error. This belief could have been further modified by previous experience of signal aspects displayed by SN109.

Explanation 4: The driver was cancelling the Automatic Warning System without proper reference to the signal aspects triggering the system

McLeod, Walker and Moray (2003) describe the purpose of the Automatic Warning System (AWS) thus:

> *"AWS serves two functions. The first function is to provide an audible alert to direct the driver's attention to an imminent event (such as a signal or a sign). The second function, linked to the first, is to provide an ongoing visual reminder to the driver about the last warning. [AWS] is there to help provide advance notice about the nature of the route ahead, and thus communicate to the driver the need to slow down or stop" (p.4).*

AWS alerts and reminders are triggered by an electro-magnetic device placed between the tracks approximately 200 yards prior to the signal, sign or other event to which it refers. Sensors underneath the train detect the presence of a magnetic field and activate AWS accordingly. AWS has two system states, but as Table 4 shows, multiple referents:

Table 4 – Summary of AWS states/referents when used in connection with four aspect colour light signalling as encountered in the region of Ladbroke Grove

STATE OF THE ROUTE AHEAD	EXTERNAL VISUAL INFORMATION AVAILABLE TO THE DRIVER	AWS AUDITORY ALERT	AWS VISUAL INDICATION	IMMEDIATE DRIVER ACTION	DRIVER BEHAVIOUR
Three signalised sections of track ahead are clear of other traffic and the route is configured for the train	Single green signal aspect	Bell or simulated chime at 1200 Hz	Display shows solid black	No additional action required	Proceed
Two signalised sections of track ahead are clear, but the third is occupied and/or not configured for the train	Double yellow signal aspect	Horn sound or steady alarm at 800 Hz	Display shows black and yellow	Cancel audible alarm and begin to slow down (subject to driving strategy, train type and/or other local circumstances)	Proceed with caution and expect to slow down further at the next signal
Only the next signalised section of track ahead is clear. The one after is occupied and/or not configured for the train	Single yellow signal aspect			Cancel audible alarm. Start (or continue) to brake	Proceed with caution and anticipate having to stop at the next signal
The next section of track is occupied and/or the route is not configured for the train	Single red signal aspect			Cancel audible alarm and bring train to a complete stop 20 metres before signal	Stop and wait for further signal aspects to be displayed

Note: AWS warnings are also triggered for flashing double and single yellow aspects, which occur prior to a train diverging from its current route (not applicable to the current situation).

AWS is a legacy system that dates back to the 1930's and was originally conceived as a means to prevent SPADS, demonstrating that SPADs do indeed have a long history. Several major accidents have seen the use of AWS, and the number of events it now refers to, being extended. AWS now provides warnings in six circumstances:

1. (Certain types of) permanent speed restriction (see next section),
2. All temporary speed restrictions,
3. (Some) level crossings,
4. SPAD indicators,
5. Cancelling boards,
6. And other locations (such as unsuppressed track magnets, depot test magnets etc.).

Unfortunately, the simple two state warning (bell/horn) and visual reminder (a black/yellow disc) are unable to discriminate between these six different events. The confusion that this could cause for drivers was cited in the Ladbroke Grove inquiry report (Cullen, 2001).

An unwanted AWS brake demand is a highly inconvenient event so there is great incentive to cancel AWS warnings quickly. This combines with (a) the visibility of the AWS equipment mounted on the track, (b) the driver's high level of route knowledge and (c) the predictive manner in which trains must be driven, to explain the extremely quick 'reaction times' drivers exhibit in response to AWS warnings. Walker and McLeod (2003) provide some insight into this based on a small random sample of OTMR data. It was observed that the mean reaction time to an AWS warning (the time taken between the horn being activated and the cancellation button being pressed) is just 0.6 seconds (min = 0.49, max = 0.89). This response is so fast that the term 'reaction time' may not be entirely appropriate. Rather than 'reacting', drivers seem to be 'pro-acting'. In-cab observations of drivers show them to be covering the AWS button with their hand in expectation of an AWS warning. OTMR data also shows numerous occasions where driver's cancelled the AWS horn before it had even started to sound (suggestive of the driver pressing the button a number of times on the approach to the on-track AWS equipment, which is clearly visible from the train cab). The frequency of AWS events in the Paddington and Ladbroke Grove area would be compatible with this behaviour, with four AWS events occurring in the space of 2.5 miles or just three minutes.

The actions following an AWS warning are of further interest. In McLeod, Walker and Moray's sample of OTMR data there were 21 AWS warnings. Three of those events were followed by the brake being applied but a greater number (5) were followed (quite appropriately for the circumstances) by the train being accelerated. There is, therefore; *"no single, fixed behavioural response expected of a driver when in receipt of an AWS warning. Many factors specific to the driver, the class of rolling stock involved, the nature of the movement, and the situation at the time the warning occurs will determine how and when an individual driver reacts"* (2003, p.9). The task of the driver, therefore, is to keep track of these factors and reach a safe and effective decision based on them.

On the morning of the Ladbroke Grove crash, the first signal which required the driver of IK20 to respond to an AWS warning was SN63. The reaction time between hearing the warning and pressing the AWS cancel button was 1.15 seconds. At the next signal, SN87, this reaction time shortened by 0.5 seconds to 0.65 seconds.

The same fast reaction time occurred in respect to the AWS warning at the final signal at danger, SN109. In both cases, 650ms is an extremely fast response. It is suggestive of the action being performed in a predictive, feed-forward manner. What this response time data seems to show is a switch from deliberate feedback control (i.e. perceive – decide – act) to predictive feed forward control (i.e. decide – perceive/act) between SN63 and SN109.

Explanation 5: The driver mistook the cause of the AWS activation prior to signal SN109 for a triangular speed restriction.

As a train leaves Paddington Station the line speed is initially set at 25 mph and a reminder of this speed restriction is placed close to the starting signal SN17, the first signal the driver will pass on leaving platform nine. Subsequent speed restriction signs allow the driver to gradually increase their speed as they move further down the line. There were five trackside speed restriction signs applicable to IK20's route, and a further five applicable to adjacent tracks. Various group standards define where, and the manner in which, speed limits should be posted. Warning signs are sometimes provided in advance of speed restriction signs (warning indicators are triangular with a yellow border whereas permanent speed limits are posted with a round sign and red border). AWS is attributed to warning indicators *"where the permissible speed on the approach is 60mph or greater and the required reduction in speed is one third or more of that permissible speed* [..]*"*. (GK/RT0038, 2000 p.14). Although this is the case at around 1000 sites on the UK rail network (e.g. Walker & McLeod, 2003) this was not the case in the vicinity of SN109 even though a warning indicator was placed under the bridge before signal SN109.

It is possible that from the driver's position in the cab that upon cancelling the AWS he saw the triangular speed restriction under the bridge in front of signal SN109 and took that as the AWS trigger, rather than the signal gantry. If the driver thought that the reason for the AWS was the triangular speed restriction he may not have searched any further for possible triggers.

Critical Path Analysis

The explanations above describe a discrete set of events which individually, and combined, could have played a role in the driver believing he could proceed past signal SN109 despite it showing a red aspect. In this section, multi-model critical path analysis will be used to map these discrete events onto the accident timeline, and from this, attempt to draw some conclusions about why the driver behaved as he did.

Critical Path Analysis (CPA) is a planning technique used to estimate the duration of projects in which some tasks can be performed in parallel. The assumption is that a given task cannot start until all preceding tasks that contribute to it are complete. This means that some tasks might be completed and the process is waiting for other tasks before it is possible to proceed. The tasks which are completed but waiting for others are said to be 'floating', i.e., they can shift their start times with little impact on the overall process. Tasks that the others wait for are said to lie on the critical path, and any change to these tasks will have an impact on the overall process time. It is possible to apply these ideas to any time-based activity, including human performance and the situation at Ladbroke Grove.

As a method of modelling human performance, CPA is based on the idea that if two or more tasks occupied the same modality they must be performed in series, but if they occupied different modalities they may be performed in parallel (Olsen & Olsen, 1990; Baber & Mellor, 2001). The allocation of tasks to modalities is shown in Table 5, along with generic task timings derived from the Key-Stroke-Level Method (KLM) and similar literature (Graham, 1999; Baber & Mellor, 2001; Olsen & Olsen; 1990).

Table 5 – Allocation of tasks to modalities

Modalities	Tasks	Timing	Source
Central processing	Identify intervention strategy		
Visual	Look at track Identify signal Identify button	340ms	Baber & Mellor, 2001
Auditory	Hear warning	300ms	Graham, 1999
Manual	Reach for button Press button	320ms 230ms	Baber & Mellor, 2001 Olsen & Olsen, 1990

The tasks are put into the order of occurrence, checking the logic for parallel and serial tasks. For serial tasks, the logical sequence is determined by a task analysis. For parallel tasks, the sensory modality determines their placement in the representation. To make the CPA easier to view, tasks of the same modality are always positioned in the same row in the diagram. The time that the task may be performed can be found by tracing through the CPA using the longest node-to-node values. Figure 4 shows the CPA model for the driver's responses to the AWS at signal SN63 compared to signals SN87 and SN109, where a switch from feedback to feed forward control is visually apparent.

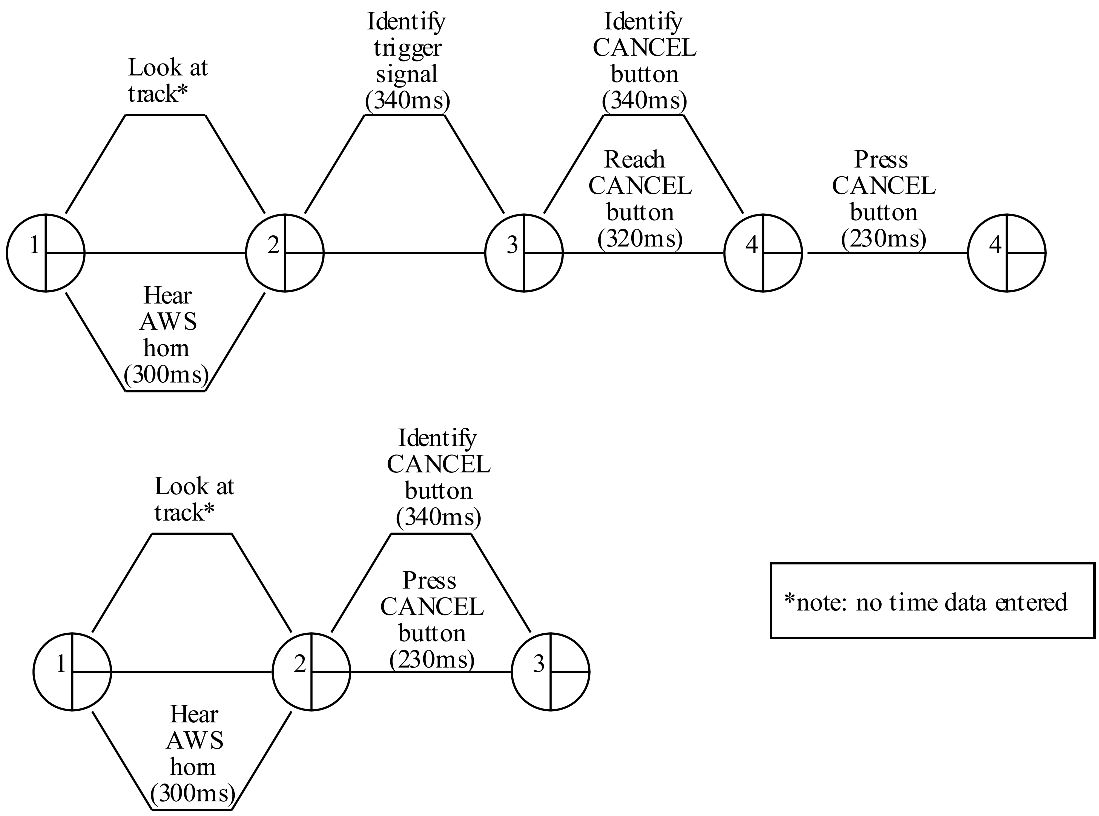

Figure 4 – Critical path analysis of driver responses to AWS at SN63 (top) and SN87/SN109 (bottom) where a change of cognitive mode is visually apparent

The CPA diagram is constructed from the known facts of the journey as described above, with detailed timings and sequences of action provided from the on-train monitoring and recording equipment. This analysis shows a total activity time of 1.2 seconds for signal SN63, which is very close indeed to the actual response time of 1.15 seconds. Likewise, a total activity time of 0.64 seconds is found for the response time to signals SN87 and SN109, virtually identical once more to the observed times. (Total activity times are based on tracing the path of the longest response times through Figure 4). Figure 4 shows that the driver's first response to SN109, after cancelling the AWS horn, was to increase speed. It is estimated that from the first presentation of the AWS horn the driver would have been engaged in the tasks of identifying the AWS trigger (i.e., mistakenly identifying the triangular speed restriction) then focusing in the cab to increase the speed, then confirming that speed was increasing. If the driver was expecting to cross over to the down main or down relief lines, he may well have been focusing on the line for the points leading to the crossover.

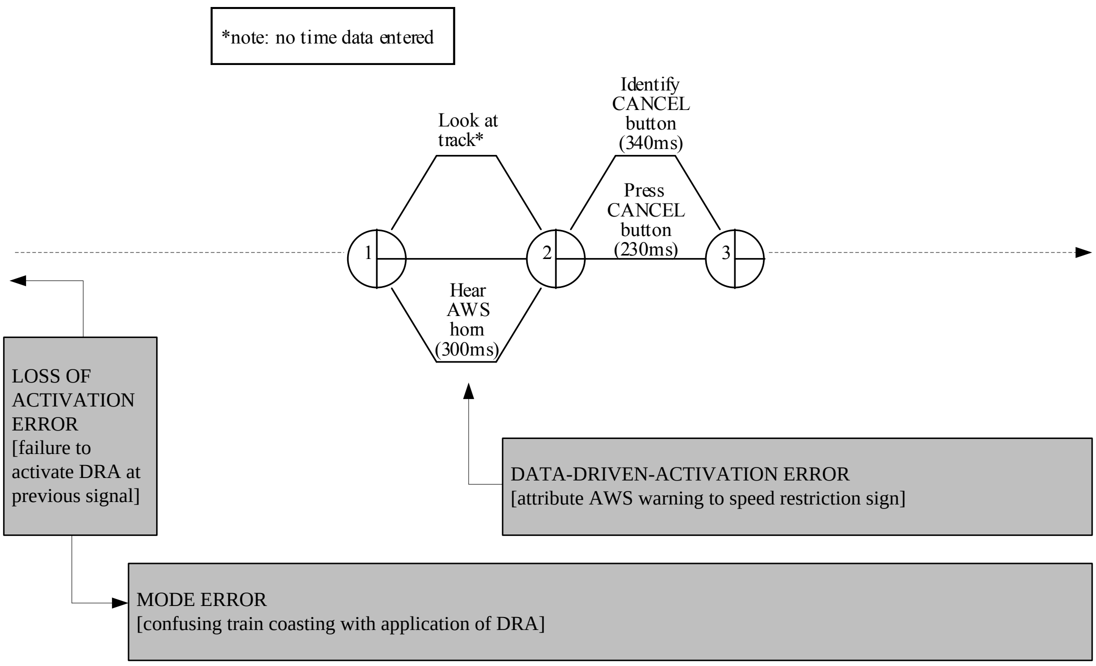

Figure 5 – Error types mapped on to the critical path analysis of activities in the region of SN109

The preceding sections have provided a detailed account of the situation, context and plausible explanations based on these and the psychological literature. From this it becomes possible to extend the critical path analysis presented in Figure 5 further back in time to show what error conditions become active prior to SN109, when they become active, and how they combine to defeat the last defence in the system (i.e. signal SN109). From Figure 5 it can be argued that a combination of a loss-of-activation (i.e. the failure to use the DRA after passing signal SN87 leading the driver to forget its aspect) created the conditions for a mode error (i.e. the realisation that power could be drawn and a consequent change in expectation as to future events) coupled with a data-driven-activation error (i.e. interpreting the AWS activation trigger as the triangular speed restriction sign under the bridge before signal SN109). These conditions would be sufficient to cause the driver to proceed past SN109 even though it was showing danger. Moreover, this sequence of events fits the timeline of actual events and would render such behaviour as 'locally rational' on the part of the driver.

Conclusions

Despite the Ladbroke Grove accident taking place over ten years ago, and despite a wide range of interventions and changes to the industry as a direct result, SPADs still occur. Ladbroke Grove is now something of a classic case study, yet regardless of a seemingly exhaustive and in-depth analysis, the inquiry itself was not able to reach a definitive answer as to why the driver in question believed he could proceed past SN109 despite it showing danger. The inquiry recognised this as an open question with its 19th formal recommendation (out of 89) being: "Further research should be carried out to develop the understanding of human factors as they relate to train driving" (Cullen, 2000, p. 229).

The simplest answer as to why the driver in question passed SN109 at danger is that he simply misread the aspect. As such, the issue becomes one of signal sighting and conspicuity, an issue which rightly informs many human factors and engineering interventions. But this simplistic answer is unsatisfactory in this case. The clue to deeper, more psychologically orientated factors lies not only in the fact that many SPADs occur in conditions of perfect sighting and conspicuity, but in this case, that a number of subtle behaviour changes were detected well in advance of the signal passed at danger. To probe these deeper issues this paper considers the actions of the driver, the behaviour of the train, the technology used, the context of rail operations, and the underpinning psychological and technological issues. Compatible with Lawton and Ward's previous work in this area we take a broader systems view but this time we take that view in conjunction with the literature on human error, distributed and situated cognition. Via this approach it is possible to arrive at a specific multi-causal explanation of why the driver in question passed signal SN109 at danger and, moreover, why it may have made sense to him to do so in those circumstances. It is proposed that the following error types could have combined to create a form of local rationality for the driver: a combination of loss-of-activation (i.e. the failure to use the DRA after passing signal SN87 ultimately led the driver to forgetting its aspect) created the conditions for a mode error (i.e. the realisation that power could be drawn and a consequent change in expectation as to future events) coupled with a data-driven-activation error (i.e. attributing the AWS activation trigger to the triangular speed restriction sign under the bridge before signal SN109). Taken together, these factors offer a plausible explanation for a seemingly implausible set of behaviours.

The next question, simply stated, is what can be done. There are several aspects of system design relevant to reducing the likelihood of these errors, many of which have already been actioned as a result of Ladbroke Grove. This includes eliminating the use of DRA on-the-move, integrating the line-side signals into the drivers' cab interfaces, and separating the warning sounds and visual indicators for different stimuli. Following the accident, the use of DRA on-the-move was indeed halted (McCorquodale et al, 2002). The rail industry's long term ambition to place signals in the cab, and provide better discrimination between warnings, is subsumed under the European Rail Traffic Management (ERTMS) system. These outcomes are consistent with the current findings as well as those of the official inquiry and in themselves are not especially novel.

Whilst specific error types can be attributed to specific items of infrastructure, and that infrastructure changed in order to prevent those errors occurring, it is evident from this work that such a deterministic approach may still leave gaps in system defences. What the present findings show, above all, is how errors combine over time, multiply rather than add together, and emerge in forms that are not only extremely difficult to predict but which cause a form of local rationality for those on the operational front line, preventing them from providing significant insight into their own behaviour. So whilst the present paper provides support for the many interventions that have flowed from the Ladbroke Grove crash it is possible to go further.

The problem with high reliability transport modes, such as rail, is that quite often major accidents occur and interventions are put in place that are not necessarily well evolved to the situation they are being placed in. Evolution requires time whereas accident investigations require decisive action, and usually a deterministic approach to the management of errors. This is clearly evident in several important aspects of rail operations. For example, the problem of the simple two state AWS system with multiple warning referents arises from previous accident investigations which recommended the use of AWS being extended beyond its original design. A further irony is that a system designed to overcome certain types of SPAD, the DRA, may have played a part in causing the SPAD at Ladbroke Grove. The use of lagging indicators to inform the design of infrastructure, therefore, tends to deal with ‘known unknowns' simply because they have revealed themselves in the course of a major incident. The real challenge for high reliability transport domains lies elsewhere, in the extended pre-accident timeline and the conditions which exist in the system which enable as yet unknown error types to emerge in the future.

This paper supports the concept of Train Data Monitoring. Instead of using the data which informed the inquiry and analysis into Ladbroke Grove 'after the event', an opportunity now exists (with on train monitoring and recording now mandatory) to use it 'before the event'. That is, to use it in a manner similar to the mature Flight Data Monitoring programmes within the aviation industry, whereby data on routine journeys is collected and assembled into a central source (CAA, 2008). By these means, very subtle out of specification events of the sort discussed in this paper can be detected before they manifest themselves as a serious safety concern. Furthermore, it has become apparent through this paper that it is eminently possible to drive human factors methods such as CPA using this type of data. For example, it would be possible to identify particular locations where sequences of actions deviate from accepted good practice. It would also be possible to identify changes in mode, from feedback to feed forward control using aggregated industry data on reaction times. This data could then be used to inform engineering, training and/or procedural interventions which, in turn, could be monitored to test their effect. In conclusion, the work presented in this paper motivates an ongoing programme of research into the concept of Rail Data Monitoring and the ability to use on train monitoring and recording outputs to drive human factors methods.

References

Andersen, T. (1999). Human reliability and railway safety. In: Safety and reliability in transport. 16th ESReDA seminar, Oslo.

Baber, C. & Mellor, B. (2001) Using critical path analysis to model multimodal human-computer interaction. International Journal of Human-Computer Studies, 54, 613-636.

Barber, P. (1988) Applied Cognitive Psychology. Routledge: London.

CAA (2003). Flight data monitoring: a guide to good practice. Gatwick: CAA Safety Regulation Group.

Card, S. K.; Moran, T. P. & Newell, A. (1983) The Psychology of Human-Computer Interaction. Lawrence Erlbaum Associates: Hillsdale, New Jersey.

Cullen, W. D. (2001). The Ladbroke Grove rail enquiry: part 1 report. London: HMSO.

Downes, M. & Robinson, W. W. (1999). Driver's route experience and SPAD propensity. Special Topic Report for Railtrack S&SD.

Endesley, M. R., & Kiris, E. O. (1995). The out-of-the-loop performance problem and level of control in automation. Human Factors, 37, (2), 381-394.

Evans, A.W.,(2007). Rail safety and rail privatisation in Britain. Accident Analysis & Prevention, 39(3). 510-523.

Gibson, J. J. (1999). Human factors review of category A SPADs. Special topic report rtspad/REP1 DOC produced for Railtrack S&SD.

Gilchrist, A. (1990). An investigation into the causation of signals passed at danger. Review and analysis produced for British Rail Research.

Graham, R. (1999) Use of auditory icons as emergency warnings: evaluation within a vehicle collision avoidance application. Ergonomics, 42 (9), 1233-1248.

GK/RT0031 (Feb, 2002). Railway Group Standard; Line-side Signals and Indicators (4). London: Railtrack.

GK/RT8035 (Oct, 2001). Railway Group Standard; Automatic Warning System (1). London: Railtrack.

GK/RT0038 (Dec, 2000). Railway Group Standard; Signing of Permissible Speeds and Speed Restrictions (2). London: Railtrack.

Hall, S. (1999). Hidden dangers: railway safety in the era of privatisation. Shepperton, Surrey: Ian Allen.

Hill, D. N. (1995). SPAD reduction and mitigation: summary report. Report produced for British Rail Research.

Lawton, R. & Ward, N. J. (2005). A systems analysis of the Ladbroke Grove rail crash. Accident Analysis and Prevention, 37, 235-244.

May, J. L., & Gale, A. G. (1998). How did I get here? – Driving without attention mode. In M. Hanson (Ed.). Contemporary ergonomics 1998. London: Taylor & Francis.

May, J. L., Horberry, T. & Gale, A. G. (1996). Railway signals passed at danger – the prevention of human error. In S. Robertson (Ed), Contemporary ergonomics. London: Taylor & Francis.

McCorquodale, B., Chissick, C., McGuffog, A., Rowley, I. & Bunting, A. (Sept. 2002). Drivers reminder appliance: effectiveness study. Farnborough, UK: QinetiQ.

McLeod, R. W., Walker, G. H., & Moray, N. (2005). Analysing and modelling train driver performance. Applied Ergonomics, 36, 671-680.

McLeod, R. W., Walker, G. H. & Moray, N. (2003). Extended AWS study: Review of the knowledge base. London: RSSB.

Miller (1956) The magical number seven plus or minus two: some limits on our capacity for processing information. Psychological Review, 63, 81-97

Neisser, U. (1976) Cognition and Reality: principles and implications of cognitive psychology. Freeman: San Francisco.

Norman, D. A. (1981) Categorisation of action slips. Psychological Review, 88, 1-15.

Olsen, J. R. & Olsen, G. M. (1990) The growth of cognitive modelling in human-computer interaction since GOMS. Human-Computer Interaction, 3, 309-350.

Reason, J. (1990). Human error. Cambridge: Cambridge University Press.

Reason, J. (1988) Managing the Risks of Organisational Accidents. Ashgate: Aldershot.

RS&SD (1999). Master rule book. Newark, Notts: Willsons.

Santos-Reyes, J. & Beard, A. N. (2006). A systemic analysis of the Paddington railway accident. Proceedings of the Institution of Mechanical Engineers Part F-Journal of Rail and Rapid Transit, 220(2).. 121-151.

Sarter, N. B. & Woods, D. D. (1997). Team play with a powerful and independent agent: operational experiences and automation surprises on the Airbus A-320. Human Factors, 39(4), 553-569.

Schmidt, R. A. (1989). Unintended acceleration: A review of human factors contributions. Human Factors, 31, (3), 345-364.

Shappell, S. A. & Wiegmann, D. A. (2003). A human error analysis of general aviation controlled flight into terrain accidents occurring between 1990-1998. Washington, DC: U.S. Department of Transportation/FAA.

Stanton N.A., Baber C., (2008). Modelling of human alarm handling responses times: A case of the Ladbroke Grove rail accident in the UK. Ergonomics, 51(4), 423-440

Van der Flier, H. & Schoonman, W. (1988). Railway signals passed at danger – situational and personal factors underlying stop signal abuse. Applied Ergonomics, 19(2), 135-141.

Walker, G. H. & McLeod, R. W. (2003). Report on Industry Data. London: RSSB.

Wickens, C. D. (1992) Engineering Psychology and Human Performance. Merrill: Columbus, Ohio.

Williams, J. C. (1977). Railway signals passed at danger – some further research. Ergonomics Abstracts, 10(3), 74661.

Woods, D. D.; Johannesen, L. J.; Cook, R. I. & Sarter, N. B. (1994) Behind human error: cognitive systems, computers and hindsight. CSERIAC: Wright-Patterson Air Force Base, Ohio, USA.

Wright, L., Ross, A. & Davies, J. (2000). SPAD risk factors: results of a focus group study. Report produced by ScotRail Railways Ltd.

Exploring the Psychological Factors involved in the Ladbroke Grove Rail Accident

Prof. Neville A. Stanton[1] *and Dr Guy H. Walker*[2]

[1]*School of Civil Engineering and the Environment, University of Southampton, Southampton [UK]SO17 1BJ*

[2]*School of the Built Environment, Heriot-Watt University, Edinburgh [UK]EH14 4AS.*

Abstract

Ten years after the event and the question as to exactly why a driver passed a signal at danger to cause the Ladbroke Grove rail disaster is still an open one. This paper uses the literature on human error and cognition, combined with critical path analysis, to provide further insight. Five aspects of train operation are drawn out of the known facts surrounding the incident: custom and practice in the use of the Driver's Reminder Appliance, operation and use of the Automatic Warning System, the sequence of signalling information, methods of supplying route information, and speed restrictions. Associated with each are several important human factors issues which, combined, give rise to five potential explanations. Critical path analysis is used to map these explanations onto the known facts of the situation. It is suggested that the proximal cause of the Ladbroke Grove rail crash was a combination of an association-activation error and a mode error (leading the driver to mistakenly assume he had activated the Reminder Appliance) together with a loss-of-activation error (the driver failing to remember that a previous signal was showing caution) and a data-driven-activation error (by associating an in-cab warning to the wrong external source). The findings support the original inquiry recommendations, but also go further into predictive methods of detecting problems at the human/transport system interface.

Keywords: SPAD; Schema; Error; Event Analysis; Critical Path Analysis

Contact: N.Stanton@soton.ac.uk

Introduction

Ladbroke Grove has become shorthand for Lord Cullen's inquiry report whose recommendations had a wide ranging impact on the UK railway industry. Ladbroke Grove itself is a suburban district located in West London through which the London to West Country main line runs. It is approximately two miles from the line's terminus, London Paddington. It was at this location, at 08:09 on the 5th October 1999, that an outbound three car commuter train collided with an eight coach high speed train at a combined speed of 130mph, injuring over 400 people and causing the deaths of 31, including both drivers. Although accidents like this are extremely rare on the UK rail network, this incident was the result of a long-standing human factors problem – a Signal Passed At Danger (SPAD).

A SPAD is when a train driver passes a signal displaying a stop aspect. To the outside observer, problems like this appear deeply perplexing. How is it possible that a driver can acknowledge up to three in cab warnings, override three automatic applications of the emergency brake, but still proceed past a point at which he or she is required to stop? Within the scientific community, and despite a considerable body of research, this category of problem is far from fully resolved: SPADs have occurred for well over a hundred years and still occur to this day, with many examples remaining stubbornly resistant to a wide range of well intentioned safety measures. SPADs are more than specific incidents, they also represent a perplexing 'category' of human factors problem, a category that is to be found in virtually all transport domains. Road transport, for example, has 'unintended acceleration' (e.g. Schmidt, 1989) and 'highway hypnosis' (e.g. May & Gale, 1998). Aviation has 'controlled flight into terrain' (e.g. Shappell & Wiegmann, 2003), 'mode errors' (e.g. Endsley & Kiris, 1995) and 'automation surprises' (e.g. Sarter & Woods, 1997). What all these cases share in common are functioning safety systems which are defeated or ignored, leading to perfectly serviceable vehicles being placed in highly dangerous conditions.

The problem with SPADs (and SPAD-like phenomenon) is that they appear both paradoxical yet simple. After all, the task of the driver is to spot a signal displaying a stop indication and apply the brakes. But this apparent simplicity is misleading. In the case of the Ladbroke Grove rail crash a clue to the presence of deeper, more psychological issues is revealed by the fact that the inquiry team had no doubt that the driver *believed* he had a proceed aspect showing at the last signal he passed. The report, furthermore, alludes to the fact that "SN109 [the code given to the signal that was passed at danger] was a multi-SPAD signal" (Cullen, 2001, pg 2) meaning that the driver in question was not alone in passing this particular signal at danger.

Although the Ladbroke Grove accident happened ten years ago SPADs still occur. An opportunity exists to re-visit Ladbroke Grove and go beyond the official enquiry report into the misleadingly simple question of ‘why’ the driver in question passed a signal he was not meant to. The inquiry report invites such analysis. In the words of Lord Cullen, the inquiry chairman, “...it is not possible for me to arrive at a full explanation of [the driver’s] actions” (Cullen, 2001, para 5.111, p. 80). The question of specifically ‘why’ the driver passed the signal is therefore an open one.

Several authors have responded to this opportunity previously and have used the great depth of information provided by the formal investigation to try and go further. Lawton and Ward (2005) and Santos-Reyes and Beard (2006) both adopt a systems approach, shifting the level of analysis to encompass the combined effects of the operational and organisational environment. Both conclude, quite rightly, that train crashes like this cannot be distilled into a single causal factor and that a systemic failure took place which enabled direct, indirect, latent and active factors to propagate through the interacting system elements. A paper by Evans (2005) presents an interesting counterpoint. Despite the severity of the Ladbroke Grove accident, and the public perceptions which surrounded it, rail safety continued its improving trend regardless. This brings into sharp relief a particular facet of SPADs as a category of human factors problem; their low probability but extremely high cost.

This unfortunate probability/cost trade-off, requires a particularly fine grained analysis. The present paper aims to make a contribution of this sort by shifting the systemic level of analysis closer to the accident itself. It does this by looking at the interacting psychological mechanisms underpinning the actions and behaviour of the driver who committed the SPAD. The paper combines the known facts of the Ladbroke Grove incident with the latest knowledge of human factors in order to develop a number of plausible explanations for the accident, which can then be evaluated. This probing is intended to go further than the inquiry report and respond to the tacit invitation therein: simply, why did the driver believe SN109 was displaying a proceed aspect? Given that a similar class of problem occurs in most other high reliability domains, additional insight into the underlying psychological mechanisms is as valuable as it is transferrable. The remainder of the introduction deals with the events leading up to the crash and wider aspects of rail operations which have a bearing on understanding its cause.

Timeline of Events

In order to understand the situation and context surrounding the Ladbroke Grove rail accident it is necessary to re-play the train journey in question.

It is 08:05 on the 5th October 1999. A ‘Thames Turbo’ multiple unit commuter train, identifying number 1K20, is waiting on platform 9 of London’s Paddington station. On board are the driver and 147 passengers.

At 08:06 the On-Train Monitoring and Recording (OTMR) equipment and records from the computerised signalling centre show that the train passed the signal (code number SN17) at the end of the platform. SN17 was showing a green aspect (proceed) and an illuminated number four (indicating that the train would be routed to line four). So called ‘starting signals’ at the ends of platforms in terminus stations like Paddington do not have the Automatic Warning System (AWS) fitted to them, thus no audible or visual indication of SN17’s signal aspect would be provided in the cab.

Between SN17 and SN43 (the next signal in sequence) the driver gradually increased speed to 34 mph. He did this by moving the speed control successively up through seven speed notches in the following order: 1, 2, 3, 4, 7. Two speed restriction signs were passed, one reading 40mph and later one reading 60mph.

The next signal, code number SN43, was showing a green aspect (proceed) and an illuminated number four, indicating that the train would be continuing along line four. 200 yards from SN43 the Automatic Warning System (AWS) triggered an audible 'chime' in the cab, confirming to the driver that SN43 was showing a proceed aspect and that he need not perform any actions additional to those already being performed. From this point, speed increased further to 44 mph.

The next signal had the code number SN63 and it was showing a double yellow aspect. This means 'preliminary caution' and that two signalised sections of track ahead are clear but that the third is occupied. This signal was not presenting any route information meaning that the driver could assume he was staying on the same track. Approximately 200 yards before the signal, the AWS triggered a horn sound in the cab and gave the driver six seconds to press the AWS cancel button in order to prevent an automatic application of the train's brakes. The driver correctly pressed the button causing the horn sound to stop and the in-cab AWS display to change colour. The display now provided an ongoing reminder that the last AWS activation was a warning. Shortly after SN63, the driver put the speed control in the neutral position and applied the brake at level one for about seven seconds, enough to reduce speed from 44 to 39 mph as it approached the next signal in sequence, SN87.

SN87 was displaying a single yellow aspect meaning 'caution', only the next signalised section of track is clear and that the next signal is likely to be at danger. Route information was also being displayed on an associated 'position light junction indicator' where a row of white coloured lights was pointing left. This informed the driver that the train was going to be routed onto the adjacent left-hand line after the next signal. Again, 200 yards from SN87 the AWS horn sounded in the cab and a request for an automatic application of the brakes was cancelled. The in-cab AWS display remained in its previous state and provided an ongoing reminder of the last AWS indication received (i.e. a warning). At this point, it would have been normal practice for the driver to start slowing down in order to stop in advance of the next signal. In this case, however, the train continued to coast for another 30 seconds.

The next signal, SN109, was displaying a red aspect meaning stop. Approximately 200 yards before SN109 the train passed over the associated Automatic Warning System (AWS) equipment on the track. A warning horn was received in the cab and a request for an automatic brake application was once again cancelled. The in-cab AWS display remained in the state that had been triggered by SN43 (two signals previously) and the driver moved the speed controller from notch zero (coasting) to notch 5. This kept the train cruising at around 38mph.

Just ahead of SN109 were two speed restriction signs located at the trackside. The first was a round sign warning of an 85mph restriction, the second was a triangular warning sign referring to a 70mph restriction on the crossover to the right. Concurrent with the signs, the driver moved the speed controller from notch five to notch seven (maximum). The train accelerated past SN109 at danger, the only in-cab warning for this event having been cancelled by the driver approximately two seconds previously. The only alarms raised as the train entered a section of track it was not meant to were those in the remote signalling centre 15 miles away in Slough. The signaller in charge of monitoring the situation made an attempt to halt an approaching train by turning a signal with the code number SN120 from green (proceed) to red (stop). The driver of the train travelling in the opposite direction, a high speed train comprised of eight coaches with a power-car at each end, immediately applied the brakes.

Thirty six seconds after selecting speed control notch seven, with the Thames

Turbo's speed now approaching 50mph, the On-Train Monitoring and Recording (OTMR) equipment showed that the driver moved the control from notch seven straight to notch zero (the neutral position) followed immediately by an application of the emergency brake.

Five seconds later the Thames Turbo collided with the High Speed Train near to signal SN120. The combined speed of impact was 130mph.

The events and surrounding context have been assembled into Figure 1 which is based on data from the On-Train Monitoring and Recording (OTMR) equipment and the known layout of the rail infrastructure.

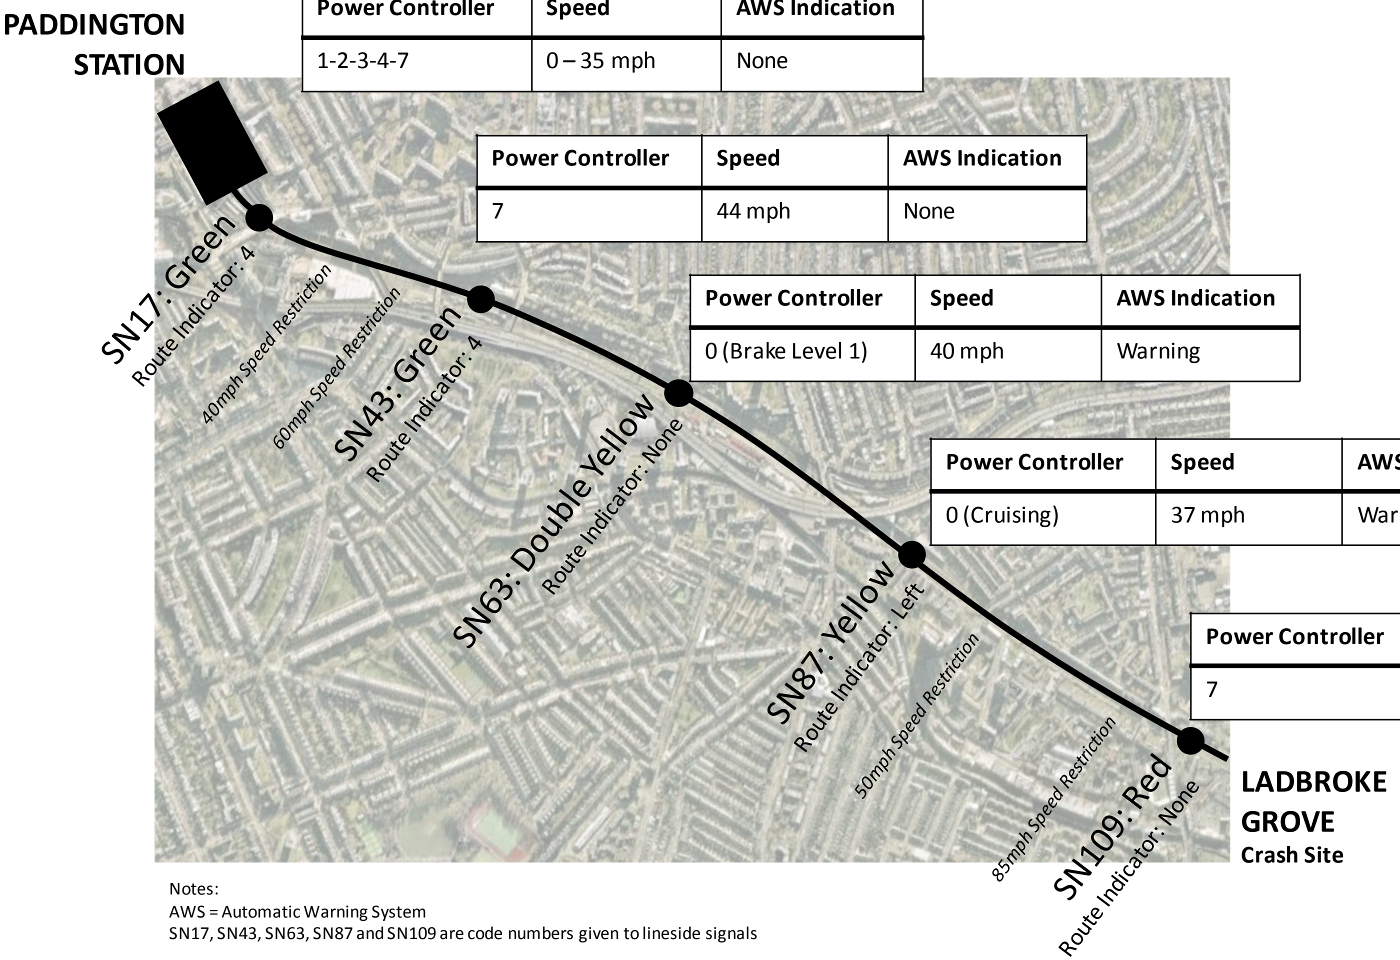

Power Controller	Speed	AWS Indication
1-2-3-4-7	0 – 35 mph	None

Power Controller	Speed	AWS Indication
7	44 mph	None

Power Controller	Speed	AWS Indication
0 (Brake Level 1)	40 mph	Warning

Power Controller	Speed	AWS Indication
0 (Cruising)	37 mph	Warning

Power Controller	Speed	AWS Indication
7	37 – 50 mph	Warning

Notes:
AWS = Automatic Warning System
SN17, SN43, SN63, SN87 and SN109 are code numbers given to lineside signals

Figure 1 – Schematic of Ladbroke Grove event timeline

Defences in Depth

The railway system, like all safety-critical systems, has multiple layers of defences each with the potential to prevent an accident (Reason, 1990, 1988). Multiple defences have the benefit that should any one defence fail, there will be other defences in place to prevent anything untoward occurring. Reason argues that in reality these defences tend to interact dynamically with local demands and human error. As shown in Table 1, there were at least seven such defences in place on the day of the incident, with the driver managing to defeat the last five. Specifically, had the driver not selected the speed control notch 5, or had applied the DRA after SN87, or correctly identified the automatic warning at SN109, or not selected the speed control notch 7, it is probable that the accident either would not have occurred or not been as devastating as it was.

Table 1 – Analysis of defences in depth

Defences-in-depth	Driver response	Train response
1. SN63 AWS horn	Cancel AWS	AWS horn cancelled
2. SN63 at double yellow	Select speed control notch 0 and brake level 1 applied and released	Train coasting at 45mph then braking to 41mph and coasting to 39mph
3. SN87 AWS horn	Cancel AWS	AWS horn cancelled
4. SN87 at single yellow	Allow train to continue coasting for 30 seconds, then select speed control notch 5	Train coasting at 39mph then cruising at 38mph
5. Driver Reminder Appliance	Not applied	Train coasting but driver is still able to draw power
6. SN109 AWS horn	Cancel AWS Select speed control notch 7	AWS horn cancelled Train accelerating from 38mph to 50mph
7. SN109 at red	No response	Train accelerating from 38mph to 50mph

The Ladbroke Grove accident clearly follows a trajectory through each layer of defence which, for a host of interacting local and global reasons, had aligned in such a way as to permit the crash to take place. The multiple defences in the system offer an explanation as to why these events are so rare and why problems like SPADS appear, on the surface, to be so unique and perplexing.

Background to Driver Operations

At first sight, it seems inconceivable that a driver can acknowledge the receipt of no less than three warnings, each one requiring the cancellation of an automatic emergency brake request, yet still fail to stop at a red signal. But as numerous studies have shown, it does happen (e.g. Hall, 1999; Andersen, 1999; Williams, 1977; Van der Flier & Schoonman, 1988; Gilchrist, 1990; Gibson, 1999; Downes & Robinson, 1999; May, Horberry & Gale, 1996; Wright, Ross & Davies, 2000; Hill, 1995). All of these sources allude to the broader context of train driving operations as an explanation for the counter-intuitive nature of SPADS, a context which can now be explored in more detail.

We begin with a lay-man's description of train driving provided by a driver in a focus group:

"The best way to describe [driving a train] is a bit like having to drive your car on ice – at 100 mph – with only the handbrake to slow you down." (McLeod, Walker & Moray, 2003: 2005).

It takes approximately two miles for a train travelling at 125mph to be brought to a complete halt. As a result, it is comparatively rare for a driver to begin braking from high speed for an event that can, at that moment, be clearly seen ahead. Rather like pilots who make extensive use of their instruments, train drivers rely heavily on information conveyed to them by line-side signals. This combines with in-depth route knowledge to generate a response (like braking or acceleration) which is appropriate for the circumstances. What this means is that the kind of feedback control encountered in car driving, where a driver will see a red traffic light and apply the brakes in order to stop before it, does not often apply. Neither do the simple Input (i.e. red traffic light) –Processing (i.e. what is the braking distance) -Output (i.e. apply the brakes) models of cognition frequently encountered in human factors apply particularly well to SPADs. This is because 'the input' is derived indirectly and constructed from an expectation of what is ahead. Train driving is characterized by large amounts of this predictive, feed-forward control, with drivers being actively required to *believe* that certain system states are imminent based on what they know currently: they could not control a train at speed otherwise. The key in railway operations is to ensure those beliefs and predictions, and the corresponding behaviours, are correct. Put simply, because drivers have to prepare to stop for up to two miles, and several signals in advance, the problem becomes less to do with signals being passed at danger and more to do with other events, principally the succession of previous signals passed at caution.

Psychological Mechanisms in Error Production

In order to relate train driving operations to the system defences that were defeated in the run up to the Ladbroke Grove crash, it is necessary to consult the psychological literature on human error, working memory and attention.

Interactive Schemata

Norman (1981) reports research on the categorisation of action slips based on an analysis of 1,000 incidents. Underpinning the analysis was a psychological theory of schema activation. He argued that action sequences are triggered by knowledge structures organised as memory units and called schemas. The mind comprises a hierarchy of schemas that are invoked (or triggered) if particular conditions are satisfied or events occur. Related to this is Neisser's (1976) seminal work on 'Cognition and Reality', where he puts forward a view of how human thought is closely coupled with a person's interaction with the world. He argues that knowledge of how the world works (e.g. schemas) leads to certain kinds of information being expected, which in turn directs behaviour to seek it out and provide a ready means of interpreting the situation. During the course of events, as the environment is sampled, the information serves to update and modify the internal cognitive schema of the world, which will again direct further search.

The perceptual cycle can be used to explain human information processing in train driving. Assuming that the individual has the correct route knowledge for operating the train, their schema will enable them to anticipate events (such as the signals and signs they expect to see and routes they expect to take), search for confirmatory evidence (such as looking at the signal aspect, trackside objects, routing information, speed indicators and notes), direct a course of action (such as braking and accelerating) and continually check that the outcome is as expected (such as the slowing down or the speeding up of the train). If they uncover some data they do not expect (such as an unexpected warning or track routing) they are required to source a wider knowledge of the world to consider possible explanations that will direct future search activities. The completeness of this model is in the description of process (the cyclical nature of sampling the world) and product (the updating of the world model at any point in time).

Error Taxonomy

The interactive schema model works well for explaining how we act in the world and what can go wrong. As Norman's (1981) research shows, it may also explain why errors occur as they do. If, as schema theory predicts, action is directed by schemata, then faulty schemata, or faulty activation, will lead to erroneous performance. As Table 2 shows, this can occur in at least three main ways. First, we can select the wrong schema due to misinterpretation of the situation. Second, we can activate the wrong schema because of similarities in the trigger conditions. Third, we can activate schemas too early or too late.

Table 2 - Taxonomy of errors with examples

Taxonomy of errors	Examples of error types
Errors that result from the formation of intention	Mode errors: erroneous classification of the situation
	Description errors: ambiguous or incomplete specification of intention
Errors that result from faulty activation of schemas	Loss-of-activation errors: schemas that lose activation after they have been activated
	Data-driven-activation errors: external events that cause the activation of schemas
	Association-activation errors: currently active schemas that activate other schemas with which they are associated
	Capture errors: similar sequences of action, where stronger sequence takes control
Errors that result from faulty triggering of active schemas	Blend errors: combination of components from competing schemas
	Premature activation errors: schemas that are activated too early
	Failure to activate errors: failure of the trigger condition or event to activate the schema

Reference to the known facts of the Ladbroke Grove crash described above, and to the error types shown in Table 2, shows only weak evidence for a description error, association-activation error, capture errors or blend/premature/failure to activate errors. The driver demonstrated no apparent ambiguity in his intention and made no attempt to correct his actions. This, in turn, supports a view that the actions being performed were consistent with individual schemas. The known facts (particularly the OTMR data) also support the view that actions were carried out in response to, and apparently triggered by, external stimuli. What the known facts do communicate, however, is stronger evidence for a loss-of-activation error, followed by a mode error, followed again by a data-driven-activation error. These error types also fit the timeline of events well: for example, the driver fails to remember that he has passed a single yellow at SN87, perhaps confuses the coasting of the train with application of the DRA and then, possibly, mistakes the speed restriction as the cause of the activation of SN109's in-cab automatic warning. Specific explanations will be developed shortly, for the time being it seems reasonable to suspend further consideration of the former error types (description, association-activation, capture, blend, premature, failure to activate) in favour of the latter three (loss of activation, mode and data-driven). These are taken forward in the next sections.

Loss of Activation Error

There are two main ways in which information is lost from working memory: displacement and decay. In displacement, older information is displaced from working memory by newer information. Working memory has upper limits on capacity of (see Miller's classic 1956 work). This means that as the journey progresses, information about signals and signs passed will be displaced by more recent signs and signals. In decay, information is forgotten as time elapses. The more information held in working memory the faster it decays. Individual chunks (such as remembering a single yellow aspect at a previous signal) have a half-life in working memory of approximately 70 seconds (the delay after which recall is reduced by one half: see Card, Moran & Newell, 1986). Whilst this is well within the actual time taken by the driver to travel from SN87 to SN109 (the journey took approximately 30 seconds) the half-life is reduced by a factor of ten when the numbers of chunks are tripled (such as remembering route information and speed restrictions in addition to remembering a single yellow aspect at the signal just passed). The loss of information from working memory is thus plausible in these circumstances.

Mode Error

Norman (1981) singles out mode errors as requiring special attention in the design of technological systems. He pointed out that the misclassification of the mode that the system was in could lead to input errors which may have serious effects. Woods et al (1994) state that mode errors occur when there is a breakdown in situation assessment. They argue that keeping track of the mode that a system is in increases the demands placed on memory and situation assessment. In the present case, the loss of activation error seems to have compounded the problem by giving the driver a mistaken belief as to the current mode. The selection of notch 5 on the speed control just before SN109 is consistent with a mode error, and the driver's mistaken belief that he had not passed a previous signal displaying caution and therefore would not be expecting to see a signal displaying danger.

Data Driven Activation Error

The third error, the attribution of an in-cab warning at referring to SN109 incorrectly to a nearby speed restriction sign, is consistent with the driver's almost immediate selection of speed control notch 7 (the data-driven-activation error). The concept of limited pools of attentional resources (Wickens, 1992) is relevant here. The basic premise is that allocation of attentional resources to one task will result in fewer resources available for another. For example, attentional resources focused on detecting hazards or checking the line ahead for crossovers will mean that there are fewer resources available for determining alternative courses of action and executing manual responses.

Attention has often been described using the metaphor of a searchlight (Barber, 1988). The direction of the driver's attention is like the beam of the searchlight and everything that falls within the beam of the searchlight is processed. The limits of human attention can cause problems in the train driving task in cases where the driver is required to focus on some stimuli rather than others, is distracted and/or too thinly divided across too many concurrent tasks. There are two main factors to be considered: attentional resources and direction of visual focus. These factors offer an explanation for the demands placed upon the driver at SN109. On hearing and cancelling the in-cab warning referring to SN109's signal aspect (i.e. red) the driver's visual attention may have focused on the speed restriction sign below the bridge. The information perceived and processed within the attentional "searchlight" is highly detailed, whilst the surrounding information tends to be perceived and processed in a peripheral manner. Whilst the driver may have been very aware of the speed restriction signs, he may only have been vaguely aware of the overhead lines and signal gantry. Of course, this goes against all driver training and common experience to the point of implausibility. Yet, after cancelling the AWS the driver behaved in a way consistent with this by selecting speed control notch 7. To do this he may have switched his attention to the controls and speed indicator inside the cab. By the time his attention was focused back on the line the gantry at SN109 might not have been within the attentional "searchlight", and the opportunity to fully appreciate the consequence of his actions lost. Prior experience at SN109's gantry, combined with artefacts of the driver's training and experience are also likely to have played a role, as discussed in the official inquiry report (Cullen, 2000).

So far we have provided details of the crash, the background to driver operations, the defences present in the system and the wider psychological literature on human error. Combined, these provide the basis for the next stage of analysis.

The next section presents five detailed explanations related to specific aspects of rail operations which contribute directly towards understanding why the driver believed SN109 was displaying an aspect which would have allowed him to proceed safely.

Explanations for Signal SN109 Passed at Danger

Explanation 1: The driver misread the aspect of signal SN109.

In the UK, railway signals communicate information about the upcoming route. In the context of the Ladbroke Grove incident (and most of the UK rail network) this information is communicated using coloured lights. Different signal aspects can be displayed (e.g. red, yellow, double yellow, green) which inform the driver of route occupancy, that is, how many signal-controlled sections of track ahead are clear of conflicting traffic/traffic arrangements. The task of the driver is to combine this information with in-depth route knowledge to decide what action is appropriate. For example, with certain types of rolling stock and in certain conditions a ‘double yellow’ signal aspect requires no action from the driver at all. In other circumstances, however, it could mean an immediate reduction in speed is required. This practice differs from many European countries which use speed-based signalling whereby a fixed behavioural response is required for any given signal aspect (e.g. a single yellow signal requires the driver to reduce speed to 40mph regardless of circumstances).

The signalling regime applicable to the Ladbroke Grove incident is four aspect colour light signalling. The four aspects (red, yellow, double yellow and green) are displayed by individual lens assemblies mounted in a common signal head. The sequence of information that can be presented, and the possible interpretations by the driver, are presented in Table 3.

Table 3 – Definition of the four main signal aspects from the Master Rule Book (RS&SD, 1999)

Colour of signal	Definition	Driver interpretation
Green	Proceed	Three signalised sections of track ahead are clear and it is safe to maintain or accelerate to the line speed.
Double yellow	Proceed (but either start slowing down or expect to have to start slowing down at the next signal)	Two signalised sections of track ahead are clear. Depending on the rolling stock and current speed, either begin braking or else maintain speed in anticipation of information conveyed by the next signal.
Single yellow	Proceed (but start slowing down in anticipation of having to stop at the next signal)	Only one signalised section of track ahead is clear. Continue or start to brake in order to be at an appropriate speed at which to stop when the next signal becomes visible.
Red	Stop	The next signalised section of track is not clear. Bring the train to a complete stop in advance of the signal.

An example of a conventional four lamp signal head is presented in Figure 2. Signal SN109 is also reproduced where it is important to note its unusual design, being shaped in the fashion of a mirrored ‘L’ with the red aspect on the horizontal leg of the ‘L’.

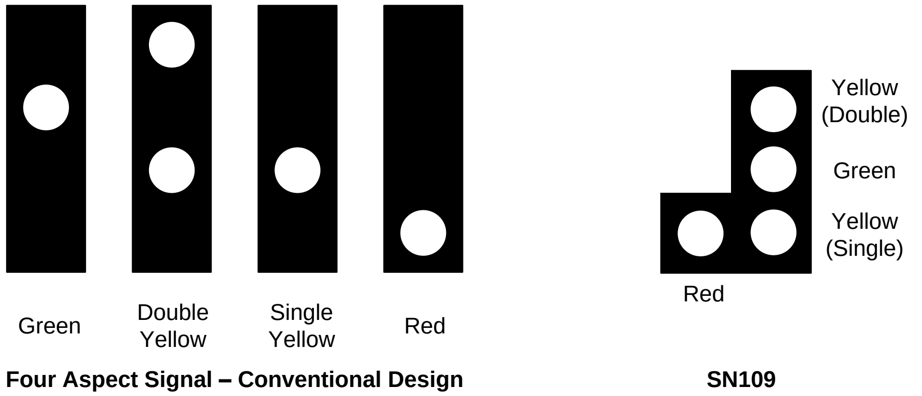

Figure 2 – Signal aspects for a four aspect colour light signal

The Ladbroke Grove inquiry concludes that the reason why the driver believed he had a proceed aspect at SN109 was due to "the poor sighting of SN109, both in itself and in comparison with the other signals on and at gantry 8, allied to the effect of bright sunlight at a low angle, […] which led him to believe that he had a proceed aspect and so that it was appropriate for him to accelerate as he did after passing SN87 [the previous signal]" (Cullen, 2001, pg. 2). Clearly, the critical human performance requirement for trackside signals is that they are not only visible but understandable. Railway Group Standards go into extensive detail regarding signal sighting and conspicuity, including the oft-quoted statement that "Signals shall normally be positioned to give drivers an approach view for a minimum of seven seconds and an uninterrupted view for at least four seconds" (GK/RT 0031, 2002). SN109 was visible to the driver for eight seconds, even with the presence of possible obstructions, suggesting a number of further factors relevant to the role of signalling in causing the Ladbroke Grove crash.

1. It is important to note that apart from signals at danger, there is no fixed behavioural response required of the driver in respect to the route signalling philosophy. The required and appropriate behaviour depends on many contextual factors and on the ability of the driver to keep track of them.
2. The driver's experience was a factor that the formal inquiry examined in detail. Research has shown that drivers in their first year of driving were most at risk from SPADs and that between 30 and 47% of SPADs were accounted for by drivers with less than five years experience (Downes & Robinson, 1999). The research goes on to show that the

problems are rarely to do with failures in mastering the control of the train, rather, the problem lies in misinterpretation of the information supplied to them. With just over a month of qualified driving experience the driver in question clearly fell into this 'at risk' group. His training and subsequent supervision were also criticized by the inquiry report.

3. The driver's experience of the signals on Gantry 8 (on which SN109 is mounted) is also important. It was stated in the inquiry report that the driver had only driven trains out of Paddington 20 times. No less than 19 of those times the signals the driver had encountered on Gantry 8 were showing a green/proceed aspect. The only other occasion was when the driver had encountered SN109 on one previous occasion, three days prior to the crash, in which SN109 cleared from red to green as the train approached. Therefore, on no previous occasion had the driver been required to stop at any of the signals mounted to Gantry 8. In addition, the dynamics of signal aspect presentation were such as to favour proceed aspects. Phenomenon like Driving Without Attention Mode (DWAM; May & Gale, 1998) show that it is possible for expectancy of a particular signal aspect to override the external cues to the contrary.

Whilst signal sighting and conspicuity play a key role in the immediate approach to SN109, inconsistencies in the driver's behaviour had started to occur before the signal was visible (i.e., when compared to a notional ideal driver who would have been preparing to stop the train from SN63 and SN87). This suggests that the driver may have started forming his belief about the upcoming aspect of SN109 long before sighting and conspicuity issues became a factor. Indeed, not only was the driver's behaviour towards the previous signal (SN87) inconsistent (i.e. he did not begin to slow down in the manner normally expected when in receipt of the AWS horn and double yellow at SN63, and again at the single yellow of SN87), so was that of a number of other drivers who had also experienced SPADs at SN109. When combined, it seems that although the proximal cause of the crash was a Signal Passed at Danger, the distal reason was due to a previous 'Signal Passed At Caution'.

Explanation 2: The driver mistook the position line junction indicator at position one for an indication that he was to be routed onto line one.

Where routes diverge, the route to be taken is communicated by an attendant line route information display. Two principle methods are used. 1) the driver is supplied with alphanumeric information (i.e. a letter or number that refers to particular lines the driver should take), or 2) the position line junction indicator with rows of lights that illuminate and point to the relevant line and/or direction the train should take. These two methods provide different types of information. The alphanumeric displays present the absolute route in line number(s) or letter(s) whereas the position line junction indicator displays the relative direction. Illustrations and mappings are provided in Figure 3:

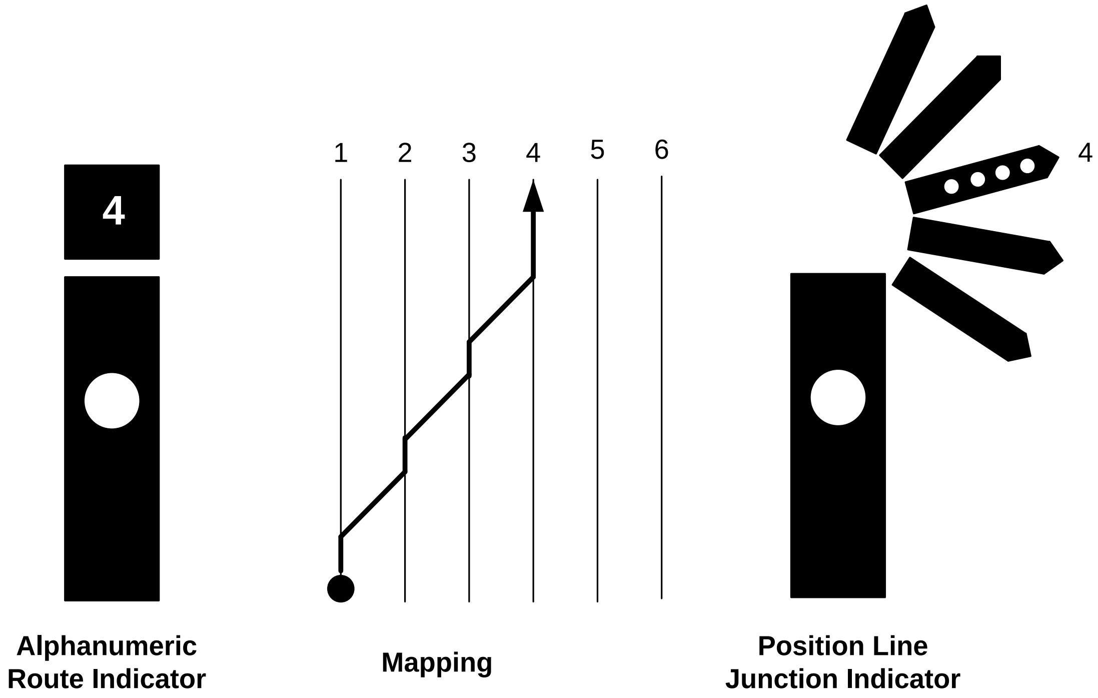

Figure 3 – Alpha numeric and position line indicators with accompanying mappings onto track layouts.

There are two important differences between the position line junction indicator displays and the alphanumeric displays. Firstly, no line information is supplied by the position line junction indicator display if the line remains unchanged, whereas the alphanumeric display has the potential to display the current as well as diverging line. Second, the position line junction indicator display provides relative line routing information, whereas the alphanumeric display provides ‘absolute’ information about the line to be taken. The driver may encounter both types of indicator display on a route, as was the case at Ladbroke Grove. As shown in Figure 2, the driver was presented with two alphanumeric line indications, one prior to a change of route (after SN17) and one serving as a reminder (SN43). This was followed by no line information at the next signal (SN63), because the route remained unchanged, followed by the position line junction indicator at position one at SN87.

If the driver formed a belief prior to SN109 coming into view that he had a proceed aspect, and subsequently did not see, or else misinterpreted the position line indication at SN109, then it is conceivable that an expectation could be formed that the train would not be diverging but proceeding on the straight ahead route. It is difficult to say, based on the evidence available, whether this happened or not. The facts of the situation are, however, that two different methods of position line display were used with the potential to cause confusion, and that the position line indicator at SN109 could in fact represent one of the defences in the system, albeit a rather weak one. If this is the case, then it would offer a partial explanation for the driver selecting notch 7 (full power) on his speed controller prior to SN109. In other words, the sequence of information presented to the driver may have activated schema which would not have prompted him to search for relevant cues in the environment.

Explanation 3: The driver forgot to apply the DRA after signal SN87.

The Driver Reminder Appliance (DRA) is an additional aid to help reduce the likelihood of a driver starting to move past a red signal, or a so-called Starting Against Signal (SAS) SPAD. Prior to performing station duties such as operation and monitoring of the train doors, interacting with the guard and so forth, the driver pushes the DRA button in the cab as a reminder that the signal ahead is red. The action of pushing the button causes a prominent red light to illuminate in the cab and prevents power from being drawn from the engine until it has been disengaged.

The DRA is designed to be used in stations and can only be applied after the driver has selected the speed control at notch zero, the neutral position. However, it is possible to activate it while the train is coasting past a signal displaying caution. It can also be activated during the braking manoeuvre for a red signal itself. In both cases, the speed controller is in the neutral position. Within several train operating companies the practice of using the DRA on-the-move had developed, and in fact was seen as best practice (see Cullen, 2001 pgs 65 & 78). Indeed, the driver who passed SN109 was trained by one of his on-the-job mentors to apply the DRA at signals displaying single yellow signal aspects. This he did consistently on the journey into Paddington Station on the morning of the accident, just ten minutes before the outward journey began. From this it seems reasonable to suggest that the driver would have normally applied the DRA at single yellow signals. Subsequent research into the use of DRA on the move reveals a number of problems:

1. The use of the DRA is not consistently applied at every caution signal. This lack of consistency could lead the driver to generate false expectations as to what signal is ahead. For example, if the DRA is usually applied at a single yellow signal, and if the driver does not apply the DRA at a single yellow, they may assume that the next signal will not be showing a red aspect.
2. The operation of engaging the DRA could be a significant distraction under certain operational scenarios. For example, the driver may have to attend to other in-cab warnings simultaneously, some of which are more safety critical than the DRA.
3. The DRA could be set too early in the sequence, meaning that there is not sufficient power to get the train to the next signal.
4. The setting and cancelling of the DRA in a succession of caution signals could lead the driver to get out of sequence, resetting the DRA rather than setting it (McCorquodale et al, 2000).

One explanation related to the Ladbroke Grove crash is that the driver forgot to apply the DRA at SN87. This certainly seems to be the case based on the available OTMR data. For drivers who used DRA 'on the move', application of the DRA would be strongly associated with the train coasting, and as the train was already coasting in the region of SN87, the driver could have mistakenly associated this with the DRA having been applied. Upon discovering that he could in fact draw power it may have changed his belief about the state of SN87 and the upcoming state of SN109: a classic mode error. This belief could have been further modified by previous experience of signal aspects displayed by SN109.

Explanation 4: The driver was cancelling the Automatic Warning System without proper reference to the signal aspects triggering the system

McLeod, Walker and Moray (2003) describe the purpose of the Automatic Warning System (AWS) thus:

> *"AWS serves two functions. The first function is to provide an audible alert to direct the driver's attention to an imminent event (such as a signal or a sign). The second function, linked to the first, is to provide an ongoing visual reminder to the driver about the last warning. [AWS] is there to help provide advance notice about the nature of the route ahead, and thus communicate to the driver the need to slow down or stop" (p.4).*

AWS alerts and reminders are triggered by an electro-magnetic device placed between the tracks approximately 200 yards prior to the signal, sign or other event to which it refers. Sensors underneath the train detect the presence of a magnetic field and activate AWS accordingly. AWS has two system states, but as Table 4 shows, multiple referents:

Table 4 – Summary of AWS states/referents when used in connection with four aspect colour light signalling as encountered in the region of Ladbroke Grove

STATE OF THE ROUTE AHEAD	EXTERNAL VISUAL INFORMATION AVAILABLE TO THE DRIVER	AWS AUDITORY ALERT	AWS VISUAL INDICATION	IMMEDIATE DRIVER ACTION	DRIVER BEHAVIOUR
Three signalised sections of track ahead are clear of other traffic and the route is configured for the train	Single green signal aspect	Bell or simulated chime at 1200 Hz	Display shows solid black	No additional action required	Proceed
Two signalised sections of track ahead are clear, but the third is occupied and/or not configured for the train	Double yellow signal aspect	Horn sound or steady alarm at 800 Hz	Display shows black and yellow	Cancel audible alarm and begin to slow down (subject to driving strategy, train type and/or other local circumstances)	Proceed with caution and expect to slow down further at the next signal
Only the next signalised section of track ahead is clear. The one after is occupied and/or not configured for the train	Single yellow signal aspect			Cancel audible alarm. Start (or continue) to brake	Proceed with caution and anticipate having to stop at the next signal
The next section of track is occupied and/or the route is not configured for the train	Single red signal aspect			Cancel audible alarm and bring train to a complete stop 20 metres before signal	Stop and wait for further signal aspects to be displayed

Note: AWS warnings are also triggered for flashing double and single yellow aspects, which occur prior to a train diverging from its current route (not applicable to the current situation).

AWS is a legacy system that dates back to the 1930's and was originally conceived as a means to prevent SPADS, demonstrating that SPADs do indeed have a long history. Several major accidents have seen the use of AWS, and the number of events it now refers to, being extended. AWS now provides warnings in six circumstances:

1. (Certain types of) permanent speed restriction (see next section),
2. All temporary speed restrictions,
3. (Some) level crossings,
4. SPAD indicators,
5. Cancelling boards,
6. And other locations (such as unsuppressed track magnets, depot test magnets etc.).

Unfortunately, the simple two state warning (bell/horn) and visual reminder (a black/yellow disc) are unable to discriminate between these six different events. The confusion that this could cause for drivers was cited in the Ladbroke Grove inquiry report (Cullen, 2001).

An unwanted AWS brake demand is a highly inconvenient event so there is great incentive to cancel AWS warnings quickly. This combines with (a) the visibility of the AWS equipment mounted on the track, (b) the driver's high level of route knowledge and (c) the predictive manner in which trains must be driven, to explain the extremely quick 'reaction times' drivers exhibit in response to AWS warnings. Walker and McLeod (2003) provide some insight into this based on a small random sample of OTMR data. It was observed that the mean reaction time to an AWS warning (the time taken between the horn being activated and the cancellation button being pressed) is just 0.6 seconds (min = 0.49, max = 0.89). This response is so fast that the term 'reaction time' may not be entirely appropriate. Rather than 'reacting', drivers seem to be 'pro-acting'. In-cab observations of drivers show them to be covering the AWS button with their hand in expectation of an AWS warning. OTMR data also shows numerous occasions where driver's cancelled the AWS horn before it had even started to sound (suggestive of the driver pressing the button a number of times on the approach to the on-track AWS equipment, which is clearly visible from the train cab). The frequency of AWS events in the Paddington and Ladbroke Grove area would be compatible with this behaviour, with four AWS events occurring in the space of 2.5 miles or just three minutes.

The actions following an AWS warning are of further interest. In McLeod, Walker and Moray's sample of OTMR data there were 21 AWS warnings. Three of those events were followed by the brake being applied but a greater number (5) were followed (quite appropriately for the circumstances) by the train being accelerated. There is, therefore; *"no single, fixed behavioural response expected of a driver when in receipt of an AWS warning. Many factors specific to the driver, the class of rolling stock involved, the nature of the movement, and the situation at the time the warning occurs will determine how and when an individual driver reacts"* (2003, p.9). The task of the driver, therefore, is to keep track of these factors and reach a safe and effective decision based on them.

On the morning of the Ladbroke Grove crash, the first signal which required the driver of IK20 to respond to an AWS warning was SN63. The reaction time between hearing the warning and pressing the AWS cancel button was 1.15 seconds. At the next signal, SN87, this reaction time shortened by 0.5 seconds to 0.65 seconds.

The same fast reaction time occurred in respect to the AWS warning at the final signal at danger, SN109. In both cases, 650ms is an extremely fast response. It is suggestive of the action being performed in a predictive, feed-forward manner. What this response time data seems to show is a switch from deliberate feedback control (i.e. perceive – decide – act) to predictive feed forward control (i.e. decide – perceive/act) between SN63 and SN109.

Explanation 5: The driver mistook the cause of the AWS activation prior to signal SN109 for a triangular speed restriction.

As a train leaves Paddington Station the line speed is initially set at 25 mph and a reminder of this speed restriction is placed close to the starting signal SN17, the first signal the driver will pass on leaving platform nine. Subsequent speed restriction signs allow the driver to gradually increase their speed as they move further down the line. There were five trackside speed restriction signs applicable to IK20's route, and a further five applicable to adjacent tracks. Various group standards define where, and the manner in which, speed limits should be posted. Warning signs are sometimes provided in advance of speed restriction signs (warning indicators are triangular with a yellow border whereas permanent speed limits are posted with a round sign and red border). AWS is attributed to warning indicators *"where the permissible speed on the approach is 60mph or greater and the required reduction in speed is one third or more of that permissible speed* [..]*"*. (GK/RT0038, 2000 p.14). Although this is the case at around 1000 sites on the UK rail network (e.g. Walker & McLeod, 2003) this was not the case in the vicinity of SN109 even though a warning indicator was placed under the bridge before signal SN109.

It is possible that from the driver's position in the cab that upon cancelling the AWS he saw the triangular speed restriction under the bridge in front of signal SN109 and took that as the AWS trigger, rather than the signal gantry. If the driver thought that the reason for the AWS was the triangular speed restriction he may not have searched any further for possible triggers.

Critical Path Analysis

The explanations above describe a discrete set of events which individually, and combined, could have played a role in the driver believing he could proceed past signal SN109 despite it showing a red aspect. In this section, multi-model critical path analysis will be used to map these discrete events onto the accident timeline, and from this, attempt to draw some conclusions about why the driver behaved as he did.

Critical Path Analysis (CPA) is a planning technique used to estimate the duration of projects in which some tasks can be performed in parallel. The assumption is that a given task cannot start until all preceding tasks that contribute to it are complete. This means that some tasks might be completed and the process is waiting for other tasks before it is possible to proceed. The tasks which are completed but waiting for others are said to be 'floating', i.e., they can shift their start times with little impact on the overall process. Tasks that the others wait for are said to lie on the critical path, and any change to these tasks will have an impact on the overall process time. It is possible to apply these ideas to any time-based activity, including human performance and the situation at Ladbroke Grove.

As a method of modelling human performance, CPA is based on the idea that if two or more tasks occupied the same modality they must be performed in series, but if they occupied different modalities they may be performed in parallel (Olsen & Olsen, 1990; Baber & Mellor, 2001). The allocation of tasks to modalities is shown in Table 5, along with generic task timings derived from the Key-Stroke-Level Method (KLM) and similar literature (Graham, 1999; Baber & Mellor, 2001; Olsen & Olsen; 1990).

Table 5 – Allocation of tasks to modalities

Modalities	Tasks	Timing	Source
Central processing	Identify intervention strategy		
Visual	Look at track Identify signal Identify button	340ms	Baber & Mellor, 2001
Auditory	Hear warning	300ms	Graham, 1999
Manual	Reach for button Press button	320ms 230ms	Baber & Mellor, 2001 Olsen & Olsen, 1990

The tasks are put into the order of occurrence, checking the logic for parallel and serial tasks. For serial tasks, the logical sequence is determined by a task analysis. For parallel tasks, the sensory modality determines their placement in the representation. To make the CPA easier to view, tasks of the same modality are always positioned in the same row in the diagram. The time that the task may be performed can be found by tracing through the CPA using the longest node-to-node values. Figure 4 shows the CPA model for the driver's responses to the AWS at signal SN63 compared to signals SN87 and SN109, where a switch from feedback to feed forward control is visually apparent.

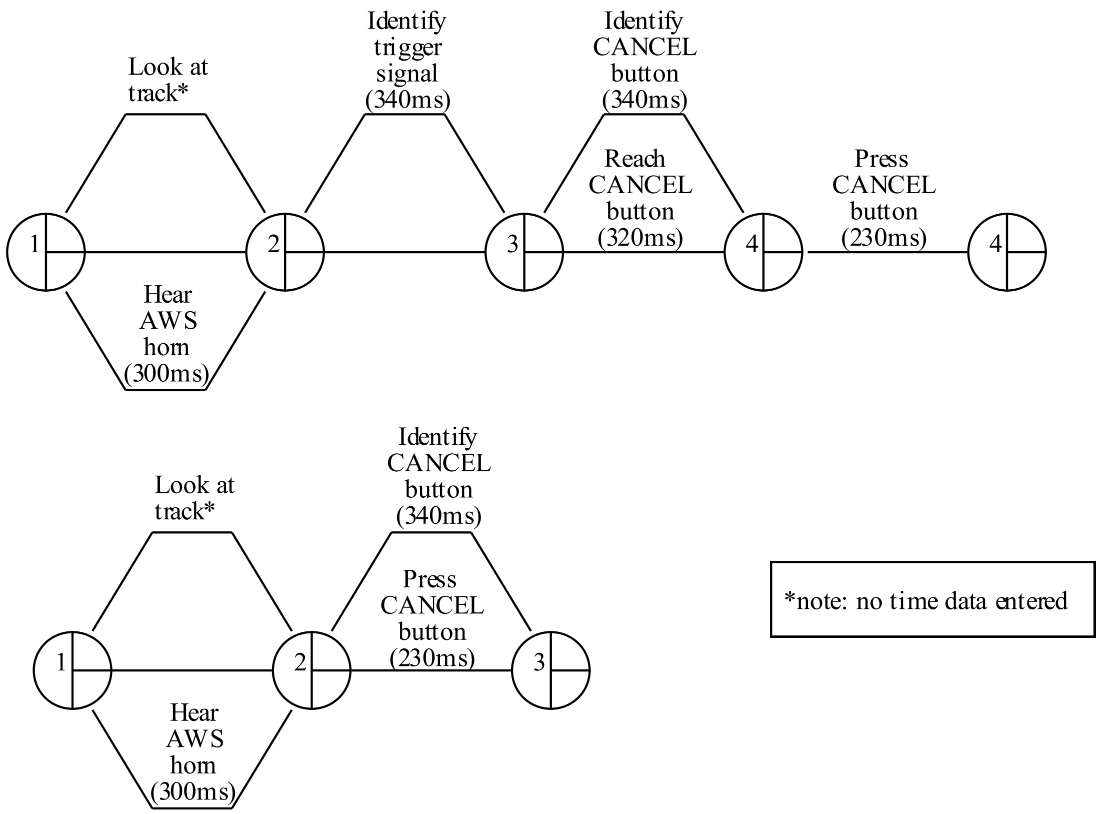

Figure 4 – Critical path analysis of driver responses to AWS at SN63 (top) and SN87/SN109 (bottom) where a change of cognitive mode is visually apparent

The CPA diagram is constructed from the known facts of the journey as described above, with detailed timings and sequences of action provided from the on-train monitoring and recording equipment. This analysis shows a total activity time of 1.2 seconds for signal SN63, which is very close indeed to the actual response time of 1.15 seconds. Likewise, a total activity time of 0.64 seconds is found for the response time to signals SN87 and SN109, virtually identical once more to the observed times. (Total activity times are based on tracing the path of the longest response times through Figure 4). Figure 4 shows that the driver's first response to SN109, after cancelling the AWS horn, was to increase speed. It is estimated that from the first presentation of the AWS horn the driver would have been engaged in the tasks of identifying the AWS trigger (i.e., mistakenly identifying the triangular speed restriction) then focusing in the cab to increase the speed, then confirming that speed was increasing. If the driver was expecting to cross over to the down main or down relief lines, he may well have been focusing on the line for the points leading to the crossover.

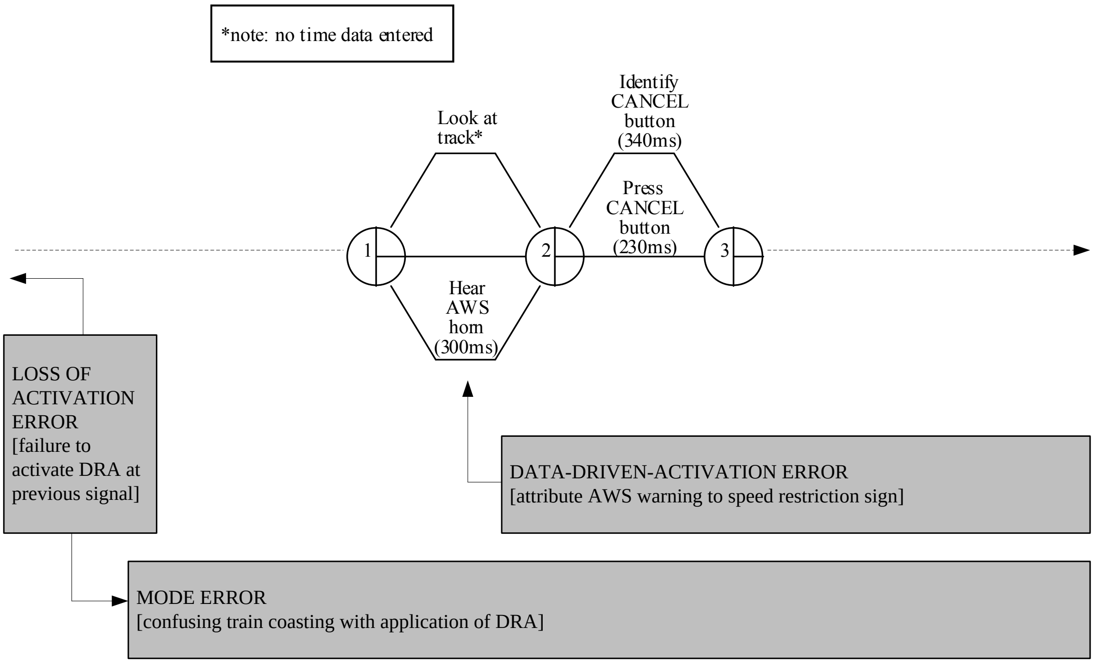

Figure 5 – Error types mapped on to the critical path analysis of activities in the region of SN109

The preceding sections have provided a detailed account of the situation, context and plausible explanations based on these and the psychological literature. From this it becomes possible to extend the critical path analysis presented in Figure 5 further back in time to show what error conditions become active prior to SN109, when they become active, and how they combine to defeat the last defence in the system (i.e. signal SN109). From Figure 5 it can be argued that a combination of a loss-of-activation (i.e. the failure to use the DRA after passing signal SN87 leading the driver to forget its aspect) created the conditions for a mode error (i.e. the realisation that power could be drawn and a consequent change in expectation as to future events) coupled with a data-driven-activation error (i.e. interpreting the AWS activation trigger as the triangular speed restriction sign under the bridge before signal SN109). These conditions would be sufficient to cause the driver to proceed past SN109 even though it was showing danger. Moreover, this sequence of events fits the timeline of actual events and would render such behaviour as 'locally rational' on the part of the driver.

Conclusions

Despite the Ladbroke Grove accident taking place over ten years ago, and despite a wide range of interventions and changes to the industry as a direct result, SPADs still occur. Ladbroke Grove is now something of a classic case study, yet regardless of a seemingly exhaustive and in-depth analysis, the inquiry itself was not able to reach a definitive answer as to why the driver in question believed he could proceed past SN109 despite it showing danger. The inquiry recognised this as an open question with its 19th formal recommendation (out of 89) being: "Further research should be carried out to develop the understanding of human factors as they relate to train driving" (Cullen, 2000, p. 229).

The simplest answer as to why the driver in question passed SN109 at danger is that he simply misread the aspect. As such, the issue becomes one of signal sighting and conspicuity, an issue which rightly informs many human factors and engineering interventions. But this simplistic answer is unsatisfactory in this case. The clue to deeper, more psychologically orientated factors lies not only in the fact that many SPADs occur in conditions of perfect sighting and conspicuity, but in this case, that a number of subtle behaviour changes were detected well in advance of the signal passed at danger. To probe these deeper issues this paper considers the actions of the driver, the behaviour of the train, the technology used, the context of rail operations, and the underpinning psychological and technological issues. Compatible with Lawton and Ward's previous work in this area we take a broader systems view but this time we take that view in conjunction with the literature on human error, distributed and situated cognition. Via this approach it is possible to arrive at a specific multi-causal explanation of why the driver in question passed signal SN109 at danger and, moreover, why it may have made sense to him to do so in those circumstances. It is proposed that the following error types could have combined to create a form of local rationality for the driver: a combination of loss-of-activation (i.e. the failure to use the DRA after passing signal SN87 ultimately led the driver to forgetting its aspect) created the conditions for a mode error (i.e. the realisation that power could be drawn and a consequent change in expectation as to future events) coupled with a data-driven-activation error (i.e. attributing the AWS activation trigger to the triangular speed restriction sign under the bridge before signal SN109). Taken together, these factors offer a plausible explanation for a seemingly implausible set of behaviours.

The next question, simply stated, is what can be done. There are several aspects of system design relevant to reducing the likelihood of these errors, many of which have already been actioned as a result of Ladbroke Grove. This includes eliminating the use of DRA on-the-move, integrating the line-side signals into the drivers' cab interfaces, and separating the warning sounds and visual indicators for different stimuli. Following the accident, the use of DRA on-the-move was indeed halted (McCorquodale et al, 2002). The rail industry's long term ambition to place signals in the cab, and provide better discrimination between warnings, is subsumed under the European Rail Traffic Management (ERTMS) system. These outcomes are consistent with the current findings as well as those of the official inquiry and in themselves are not especially novel.

Whilst specific error types can be attributed to specific items of infrastructure, and that infrastructure changed in order to prevent those errors occurring, it is evident from this work that such a deterministic approach may still leave gaps in system defences. What the present findings show, above all, is how errors combine over time, multiply rather than add together, and emerge in forms that are not only extremely difficult to predict but which cause a form of local rationality for those on the operational front line, preventing them from providing significant insight into their own behaviour. So whilst the present paper provides support for the many interventions that have flowed from the Ladbroke Grove crash it is possible to go further.

The problem with high reliability transport modes, such as rail, is that quite often major accidents occur and interventions are put in place that are not necessarily well evolved to the situation they are being placed in. Evolution requires time whereas accident investigations require decisive action, and usually a deterministic approach to the management of errors. This is clearly evident in several important aspects of rail operations. For example, the problem of the simple two state AWS system with multiple warning referents arises from previous accident investigations which recommended the use of AWS being extended beyond its original design. A further irony is that a system designed to overcome certain types of SPAD, the DRA, may have played a part in causing the SPAD at Ladbroke Grove. The use of lagging indicators to inform the design of infrastructure, therefore, tends to deal with 'known unknowns' simply because they have revealed themselves in the course of a major incident. The real challenge for high reliability transport domains lies elsewhere, in the extended pre-accident timeline and the conditions which exist in the system which enable as yet unknown error types to emerge in the future.

This paper supports the concept of Train Data Monitoring. Instead of using the data which informed the inquiry and analysis into Ladbroke Grove 'after the event', an opportunity now exists (with on train monitoring and recording now mandatory) to use it 'before the event'. That is, to use it in a manner similar to the mature Flight Data Monitoring programmes within the aviation industry, whereby data on routine journeys is collected and assembled into a central source (CAA, 2008). By these means, very subtle out of specification events of the sort discussed in this paper can be detected before they manifest themselves as a serious safety concern. Furthermore, it has become apparent through this paper that it is eminently possible to drive human factors methods such as CPA using this type of data. For example, it would be possible to identify particular locations where sequences of actions deviate from accepted good practice. It would also be possible to identify changes in mode, from feedback to feed forward control using aggregated industry data on reaction times. This data could then be used to inform engineering, training and/or procedural interventions which, in turn, could be monitored to test their effect. In conclusion, the work presented in this paper motivates an ongoing programme of research into the concept of Rail Data Monitoring and the ability to use on train monitoring and recording outputs to drive human factors methods.

References

Andersen, T. (1999). Human reliability and railway safety. In: Safety and reliability in transport. 16th ESReDA seminar, Oslo.

Baber, C. & Mellor, B. (2001) Using critical path analysis to model multimodal human-computer interaction. International Journal of Human-Computer Studies, 54, 613-636.

Barber, P. (1988) Applied Cognitive Psychology. Routledge: London.

CAA (2003). Flight data monitoring: a guide to good practice. Gatwick: CAA Safety Regulation Group.

Card, S. K.; Moran, T. P. & Newell, A. (1983) The Psychology of Human-Computer Interaction. Lawrence Erlbaum Associates: Hillsdale, New Jersey.

Cullen, W. D. (2001). The Ladbroke Grove rail enquiry: part 1 report. London: HMSO.

Downes, M. & Robinson, W. W. (1999). Driver's route experience and SPAD propensity. Special Topic Report for Railtrack S&SD.

Endesley, M. R., & Kiris, E. O. (1995). The out-of-the-loop performance problem and level of control in automation. Human Factors, 37, (2), 381-394.

Evans, A.W.,(2007). Rail safety and rail privatisation in Britain. Accident Analysis & Prevention, 39(3). 510-523.

Gibson, J. J. (1999). Human factors review of category A SPADs. Special topic report rtspad/REP1 DOC produced for Railtrack S&SD.

Gilchrist, A. (1990). An investigation into the causation of signals passed at danger. Review and analysis produced for British Rail Research.

Graham, R. (1999) Use of auditory icons as emergency warnings: evaluation within a vehicle collision avoidance application. Ergonomics, 42 (9), 1233-1248.

GK/RT0031 (Feb, 2002). Railway Group Standard; Line-side Signals and Indicators (4). London: Railtrack.

GK/RT8035 (Oct, 2001). Railway Group Standard; Automatic Warning System (1). London: Railtrack.

GK/RT0038 (Dec, 2000). Railway Group Standard; Signing of Permissible Speeds and Speed Restrictions (2). London: Railtrack.

Hall, S. (1999). Hidden dangers: railway safety in the era of privatisation. Shepperton, Surrey: Ian Allen.

Hill, D. N. (1995). SPAD reduction and mitigation: summary report. Report produced for British Rail Research.

Lawton, R. & Ward, N. J. (2005). A systems analysis of the Ladbroke Grove rail crash. Accident Analysis and Prevention, 37, 235-244.

May, J. L., & Gale, A. G. (1998). How did I get here? – Driving without attention mode. In M. Hanson (Ed.). Contemporary ergonomics 1998. London: Taylor & Francis.

May, J. L., Horberry, T. & Gale, A. G. (1996). Railway signals passed at danger – the prevention of human error. In S. Robertson (Ed), Contemporary ergonomics. London: Taylor & Francis.

McCorquodale, B., Chissick, C., McGuffog, A., Rowley, I. & Bunting, A. (Sept. 2002). Drivers reminder appliance: effectiveness study. Farnborough, UK: QinetiQ.

McLeod, R. W., Walker, G. H., & Moray, N. (2005). Analysing and modelling train driver performance. Applied Ergonomics, 36, 671-680.

McLeod, R. W., Walker, G. H. & Moray, N. (2003). Extended AWS study: Review of the knowledge base. London: RSSB.

Miller (1956) The magical number seven plus or minus two: some limits on our capacity for processing information. Psychological Review, 63, 81-97

Neisser, U. (1976) Cognition and Reality: principles and implications of cognitive psychology. Freeman: San Francisco.

Norman, D. A. (1981) Categorisation of action slips. Psychological Review, 88, 1-15.

Olsen, J. R. & Olsen, G. M. (1990) The growth of cognitive modelling in human-computer interaction since GOMS. Human-Computer Interaction, 3, 309-350.

Reason, J. (1990). Human error. Cambridge: Cambridge University Press.

Reason, J. (1988) Managing the Risks of Organisational Accidents. Ashgate: Aldershot.

RS&SD (1999). Master rule book. Newark, Notts: Willsons.

Santos-Reyes, J. & Beard, A. N. (2006). A systemic analysis of the Paddington railway accident. Proceedings of the Institution of Mechanical Engineers Part F-Journal of Rail and Rapid Transit, 220(2).. 121-151.

Sarter, N. B. & Woods, D. D. (1997). Team play with a powerful and independent agent: operational experiences and automation surprises on the Airbus A-320. Human Factors, 39(4), 553-569.

Schmidt, R. A. (1989). Unintended acceleration: A review of human factors contributions. Human Factors, 31, (3), 345-364.

Shappell, S. A. & Wiegmann, D. A. (2003). A human error analysis of general aviation controlled flight into terrain accidents occurring between 1990-1998. Washington, DC: U.S. Department of Transportation/FAA.

Stanton N.A., Baber C., (2008). Modelling of human alarm handling responses times: A case of the Ladbroke Grove rail accident in the UK. Ergonomics, 51(4), 423-440

Van der Flier, H. & Schoonman, W. (1988). Railway signals passed at danger – situational and personal factors underlying stop signal abuse. Applied Ergonomics, 19(2), 135-141.

Walker, G. H. & McLeod, R. W. (2003). Report on Industry Data. London: RSSB.

Wickens, C. D. (1992) Engineering Psychology and Human Performance. Merrill: Columbus, Ohio.

Williams, J. C. (1977). Railway signals passed at danger – some further research. Ergonomics Abstracts, 10(3), 74661.

Woods, D. D.; Johannesen, L. J.; Cook, R. I. & Sarter, N. B. (1994) Behind human error: cognitive systems, computers and hindsight. CSERIAC: Wright-Patterson Air Force Base, Ohio, USA.

Wright, L., Ross, A. & Davies, J. (2000). SPAD risk factors: results of a focus group study. Report produced by ScotRail Railways Ltd.