



DRIVER SAFETY SYSTEMS



Table of Contents

The purpose of this paper	3
Terminology base.....	3
Driver safety systems (DSS).....	3
Automatic train protection (ATP) systems.....	4
SPAD.....	4
Overspeed	4
Deadman systems	5
Vigilance control devices (known in North America as alerter)	5
Physiological vigilance measurement devices	5
AWS-type systems	5
TPWS	6
Deadman systems.....	6
Disadvantages of deadman devices	9
Vigilance control (alerter) systems	11
Key difficulties.....	13
Protecting against automaton responses.....	16
Timing of the inactive part of the cycle.....	18
Pre-emption	20
Arrangements with two-person operation	21
What to do if a vigilance non-response occurs	22
Physiological vigilance measurement devices.....	23
AWS	26
TPWS.....	28
Devices which fall outside the DSS definition but relate to driver vigilance	29
Degraded operations.....	30
A general appraisal of driver safety systems.....	31



The purpose of this paper

This paper discusses usage, advantages, disadvantages and ways forward in respect of the systems as defined as “driver safety systems” in ITSRR’s broad information paper on driver safety systems and automatic train protection.

Unlike that information paper, this current paper includes some assessments which (although based on extensive factual information) are essentially value judgments. They should be viewed as professional assessments rather than absolute statements of fact.

Terminology base

Driver safety systems (DSS)

Systems directed to confirming the capacity of the driver to continue to operate the train, and to stop the train if this capacity is not confirmed. These directly include deadman systems, vigilance control devices and physiological vigilance measurement devices; various other systems such as AWS may also be substantially regarded as driver safety systems.

There are also various systems which may contribute to driver safety without having the potential to intervene, but such systems are excluded from this definition. They are however briefly alluded to in the text.



Automatic train protection (ATP) systems

Systems which provide security against inappropriate movements of a train if the driver is not managing the train appropriately, irrespective of whether the driver is currently capable of managing it. Notionally, this security may be anywhere from a very low level to a very high level, but the term “ATP” in isolation is normally applied to systems at the upper end of the range of control, in particular to what are described as “full” ATP systems below.

SPAD

Signal passed at danger, i.e. when a train proceeds (without authority) past a signal even though the signal instructed the train to stop.¹

Overspeed

Literally, any situation where the train exceeds the currently authorised speed. This logically includes the train moving when it should be stationary, but the term commonly refers to situations such as excessive speed through curves or turnouts, or non-compliance with track speed restrictions.

Intervention

Driver safety systems and ATP systems are said to “intervene” if they take over control of the train from the driver, typically by de-applying power and applying braking.

¹ The term “SPAD” excludes authorised passage past a signal at stop, e.g. where a signal is authorised to pass a signal at stop because there is a hardware failure preventing the signal from being cleared or for other reasons authorised under operating rules. SPADs may be separated into “genuine” SPADs, where the signalling system was performing entirely as appropriate, and “technical SPADs” or RIFODs (“return in face of driver”) where the signal was returned to stop in the face of the train by system failure or human action, leaving the train insufficient time or space to stop before reaching the signal. Hardware protections such as route locking and approach locking will usually (but not invariably) protect against technical SPADs creating accidents due to collision or derailment, but may generate adverse consequences in other respects (e.g. consequences of emergency braking).

³ It is a matter of debate whether a deadman device requiring intermittent release should be classified as a “vigilance control device”. In this paper, such a device is considered in both contexts.



Deadman systems

A form of driver safety system which detects a continuous input from the driver, e.g. by application of force to a pedal or handle. Some deadman systems require periodic release and reapplication of the pressure, to avoid continuous pressure being accepted as legitimate input from an inert person or from a physical object inserted (deliberately or otherwise) to provide continuous pressure.

Vigilance control devices (known in North America as alerters)

Typically require intermittent inputs from the driver, e.g. button pressing at fixed intervals or in response to flashing lights or sounds.³ Modern vigilance systems are often task-linked, i.e. they accept direct evidence that the driver is actively controlling the train without requiring a separate response to the vigilance device itself.

Physiological vigilance measurement devices

Devices which will assess physiological measures of alertness, and will take action if that alertness is determined to be inadequate. For instance, they may act to warn the driver if alertness is detected to be inadequate, and may intervene to stop the train if that warning is not acted upon within the specified timeframe.

AWS-type systems

These require a response from a driver if a warning is given, e.g. at a signal showing a restrictive indication or on the approach to a speed restriction. If the driver does not acknowledge the warning in a few seconds, the brakes are applied.



TPWS

Is a British system, combining AWS with what is effectively an electronic train stop system if a train passes a signal at stop, and with an intermediate train stop facility which will stop the train if it is going too fast at a designated point approaching a signal at stop or approaching some other hazard.⁴ Enhancements such as TPWS+ (for higher speeds) also exist but follow the same essential principles.

Deadman systems

The earliest driver safety devices (at least on a large scale) were deadman devices. Because trains in the 19th century were predominantly steam (scarcely requiring deadman devices since most locomotives were two-person operated), deadman devices were strongly perceived as relevant only with electric trains towards the end of that century.

From an early time, electric one-driver trains tended to have simple deadman devices, accompanied often by electro-mechanical train stop systems to “trip” a train that passes a signal at stop.

The initial deadman controls were spring-loaded devices which were kept activated by some force applied by the driver (typically downwards force on the master controller handle). If the driver died or collapsed, he would no longer apply the force, so the spring would return the device to the deactivated position and the train would stop (by some combination of brake application and switching off power).

Despite this being at least 110 years old, the system is still applied with very little change all around the world. There are literally dozens – and perhaps

⁴ Systems such as TPWS and conditionally-cleared train stops may be regarded as examples of “speed traps”, which respond to excessive speed of a train at specific locations as distinct from monitoring train speed continuously.



hundreds - of urban systems where at least some trains are fitted with hand-operated spring-pressure deadman controls, where the driver applies pressure against the spring to avoid the system intervening to stop the train.

The simplest form requires vertical pressure. For instance on Sydney single-deck electric trains, there was a master controller handle pivoted at one end, with rotation to determine the powering position and vertical movement to activate the deadman function. So the driver would keep the handle depressed at all times in addition to rotating it whenever required to change the power application.

A common more modern form (e.g. on tangaras) is a symmetric or asymmetric “T bar” handle, where the deadman pressure is applied by twisting the handle against the spring (with forwards/backwards movements then controlling the power levels).

Other forms of the same basic principle exist, e.g. the French favoured a “squeeze” device where the required action was to hold and lift a ring surrounding the master controller handle.

There are various ways in which release of the device will be translated into the required application of brakes and removal of power. Some systems work on the principle that release of the master-controller handle will physically spring the handle into the full-service or emergency braking position just as though it had been placed there manually. In others, release of the deadman device will release air from the air-brake system and/or deactivate electrical switches or relays.



Spring-loaded devices were less than optimal in at least five respects:

- Ergonomic factors, e.g. RSI
- Reliance on spring function, in that a “sluggish” spring or various forms of contamination could fail to respond to release of pressure
- Difficulties in determining or maintaining calibrated spring tensions so as to ensure that the system will reliably function for a full range of different drivers with different physical characteristics
- Ability for unintended circumvention – e.g. slumping over the handle when unconscious but still applying vertical pressure from body weight
- Ability for intended circumvention – e.g. hanging a weight from the controller handle, or tensioning the T-bar handle with a rubber band.

As a development away from deadman hand devices, deadman *pedals* were introduced in response to the above, but primarily to provide relief from requiring constant arm or hand pressure. Some other reasons also applied, for instance in some non-electric rolling stock, pedals were a more suitable engineering solution.

Some pedals require constant pressure in the same manner as deadman handles, although sometimes with an additional feature that excess pressure, as well as insufficient pressure, will trigger the device into stopping the train. Some require the pressure to be released intermittently (to avoid intentional or unintentional suppression by a jamming device), while some are normally undepressed but require periodic depression (see later re vigilance devices).

It is now common, but not universal, for drivers to be given the choice of hand or foot activation. For instance, a common behaviour is for drivers to use the handles at times when they are frequently varying the master controller settings, but to use the pedal when the train is running without frequent changes in power and braking levels.



Disadvantages of deadman devices

1. As demonstrated by the Waterfall accident, there is no guarantee that a dead or unconscious person will actually release the device. There is the possibility that the person will slump over a deadman handle, or will continue to have body weight applied to a pedal; and it is conceivable that in a crisis situation, a person could inadvertently clutch onto a handle or otherwise maintain pressure upon it.

It follows that to protect against this sort of problem, there must be some detection of *changes* in condition rather than just maintenance of the condition, e.g. a pedal device which has to be released periodically is more likely to avoid this problem than one which accepts continuous pressure.

2. Because of the physical inconvenience of applying continuous pressure, they invite circumvention, and it is very difficult to design something which cannot be circumvented. Such circumvention is not necessarily malicious; for instance, a driver may seek to leave the controls to deal with some in-cab contingency or even to look along the train for signs of a problem such as sticking brakes. Any constant-application device will have operational challenges of the latter nature.
3. Setting and maintaining the pressure parameters (e.g. spring and microswitch tension) may be difficult, especially since a suitable range of pressures for one person (e.g. a heavy male) may be very different from that for another (say) a light female. There is little evidence that any railway systems anywhere in the world have fully solved this problem; and indeed many have consciously headed away from solving it by bringing other devices into use or at least choosing the required-intermittent-release variants.



4. Failure conditions may not be readily apparent. For instance a “sticky” spring may not be revealed even under ordinary testing, only becoming apparent under particular forms of gradual release as might occur if a person gradually reduces their muscle tension rather than suddenly “letting go”.
5. It is not clear what forms of incapacitation are protected against (even by the intermittent-release varieties). Even apart from the fact that a dead person may not release a non-intermittent-release type, there appears to have been no research as to the capability of a person with other forms of lesser incapacitation to maintain pressure on the devices. This is made even more complex by some evidence that people become so habituated to holding the devices activated that they may continue to do so in a very involuntary way. Certainly sleep and other forms of inattention may not be detected, and indeed may be unlikely to be detected.⁵
6. Even if the deadman devices themselves were ideally effective in detecting incapacitation, there is a question of what negative side-effects might result from requiring a person to remain in the one place, with a very constant level of physical configuration, for long periods. For instance it may breed physical and/or mental fatigue, it may produce reduced ability to respond to a sudden event, it may mean that conditions are not attended to (e.g. the sticking brakes example above).

For further commentary on the intermittent-release varieties of deadman devices, see under “vigilance control” below, since such devices have much in common with other forms of vigilance control device.

⁵ The Victorian Footscray collision provides a prime example of this. The driver was inattentive due to some combination of health and associated conditions (medications), but none of these constituted “incapacitation” in a sense likely to be detected by a deadman device. Much complaint was directed to the alleged non-performance of the deadman device, yet there was no way in which it was designed or implemented so as to be relevant to such conditions.

Moreover it is commonly admitted by drivers that they enter into microsleep states, and yet these are rarely associated with deadman intervention.



Vigilance control (alerter) systems

The definition above stated that vigilance control devices “typically require intermittent inputs from the driver, e.g. button pressing at fixed intervals or in response to flashing lights or sounds. Modern vigilance systems are often task-linked, i.e. they accept direct evidence that the driver is actively controlling the train without requiring a separate response to the vigilance device itself.”

To some extent, therefore, vigilance control (VC) devices can be viewed as intermittent deadman devices – i.e. they test from time to time whether someone is capable of undertaking the specific muscular actions targeted by the device. Whether they test what the English language would normally mean by “vigilance” is another question altogether.

Typical VC configurations

Although there are innumerable variants in matters of fine detail, there is large commonality amongst the broad features. Typically, there is

- A timing mechanism, to count time in order to know when to initiate activation of the device, and to time the various action phases once initial activation has occurred
- Physical aural or visual devices (e.g. lights, bells, sirens) to request train crew to respond
- One or more devices (e.g. buttons) to enable train crew to acknowledge the system’s requests for response
- A mechanism to stop the train if the requests for response are not actually responded to within the specified time intervals, or if the responses are not in approved form (e.g. if a button is continuously depressed).

⁷ http://www.atsb.gov.au/publications/investigation_reports/1998/RAIR/rair1998001.aspx



There *may* also be:

- Mechanisms to interpret normal active driving functions (e.g. braking changes, powering changes, horn activation) as an indication that the crew is active, and therefore do not require further verification by the VC – called “task linking”; and
- Logic to vary the behaviour of the device according to the speed at which the train is moving – called “speed linking”.

The “classical” VC system – still used very widely around the world – typically functions in the following sequence:

1. The driver does something (e.g. puts the reverser out of neutral into either forward or reverse) to prepare the train for motion. The VC is then active until the train is again verified as stopped and inactive.
2. The VC starts a timer.
3. If a member of the train crew presses an acknowledgment button at any time, the timer restarts.
4. If after some specified period (typically 60-75 seconds in older installations, now often shorter), the acknowledgment button has not been pressed, a light flashes in the cab, and the timer is set for a further short period (5-10 seconds).
5. If acknowledgment is done in that period, the timer (and hence the whole action cycle) restarts back at step 2 above.
6. If there is no acknowledgment in that period, an audible warning sounds (with the light still flashing), with another short period (5-10 seconds) set on the timer.
7. If acknowledgment occurs, the timer restarts back at step 2 and the whole cycle restarts.



8. If there is still no acknowledgment in the permitted time, a brake application (referred to as a “penalty application”) is made in either full-service or emergency, and power is de-applied.

Some of the most common variants are:

- a) The option in step 3 – known as “pre-emption” – may not be available, i.e. the system always runs to the point where the light flashes and a response is then required.
- b) The time interval to the light-flashing stage is random (within some range) or is speed-linked (faster speed implying shorter intervals).
- c) Task-linking allows normal control inputs (braking, powering etc) to be accepted as equivalent to button-pressing in at least some of the stages.

Key difficulties

There are a number of key problems and trade-offs intrinsic to VC devices. These may be summarised as

1. Responses may only be “automaton” responses, i.e. they verify that the person is still alive but not that he/she is even awake, let alone “vigilant”. Habituation arises because of the many thousands of times when the button must be pressed (especially in non-task-linked designs) – for instance button pressing may occur several hundred times in a single shift.
2. Repetitive strain injuries may arise from repeated button-pressing.
3. Pre-emption encourages automaton behaviours, yet requiring a response in a narrow time interval in a non-pre-emptive system may cause distraction from other duties at a time of peak safety workload.
4. Task-linking, while reducing the need for distractive button-pressing, may encourage control inputs being made which are not optimal for train management.



5. It is difficult to set the timer intervals optimally, especially the interval from previous activation to onset of next activation. The longer this interval, the further the train can travel before a loss of driver capacity is detected; but the shorter the interval, the more likelihood of distraction or habituation.

Points 2 and 4 admit little further analysis. However the others deserve elaboration, and the issues can be seen more readily from three actual incidents – the Beresfield accident in NSW on 23 October 1997⁷, the NSW Waterfall accident, and a SPAD and potential collision at Rangitawa in New Zealand on 12 May 2005⁸.

At Beresfield, it was clear that both crew members were asleep or effectively so (arguably more than mere “microsleeps”), yet they had activated the vigilance control device successfully. However even if they had suddenly become incapacitated rather than becoming progressively sleepy, the VC would have provided no guaranteed protection against collision – the time of last acknowledgment could have been late enough for the VC not to have instituted braking until far too late to avoid the collision.

In the Waterfall accident, there was no VC fitted, and the Special Commission of Inquiry concluded that the driver had become incapacitated soon after leaving Waterfall. Calculations showed that depending on the precise behaviour of the driver and on the time cycles of a vigilance device if it had indeed been fitted, the outcome could have ranged from no effective intervention whatever (i.e. the accident would still have occurred much as it did) through to complete prevention.

At Rangitawa, a driver (because of extreme fatigue) passed two signals on approach to a single line crossing loop without being aware of either signal (showing respectively a restrictive and a stop indication). A train coming in the

⁸ <http://www.taic.org.nz/rail/05-117.pdf>



opposite direction was stopped by a signal reverting to stop and prompt response by its driver, and the train which had passed the signal at stop was brought to a stand short of collision when the driver awoke (evidently either because of traversing points or the sound of a radio call).

The sleep-affected driver had been activating the VC (with progressively slower response times) even though his failure to observe the first signal suggests that he was already effectively non-functional at that stage. The investigation report calculates various scenarios, and concludes that if the driver had not awoken and the VC had been relied upon to stop the train, it may not have activated soon enough to prevent a collision occurring. The report does not carry this to one further stage – i.e. if the driver had yet again activated the VC as he had done before (evidently effectively unconsciously) and had not otherwise awoken, there definitely would have been a collision of some severity.

New Zealand investigation reports reveal other examples of single-manned trains being involved in incidents in similar contexts, where fatigue has not been countered by the VC and unsatisfactory outcomes have occurred. However it should not be assumed that the problems are exclusive to single-manning; for instance the Beresfield accident involved double-manning, and double-manning runs an increased risk that each crew member will rely on the other remaining vigilant.⁹ Moreover although many North American severe accidents – typically overruns of limits of movement authority – occurred (even in recent times) without VC (“alerter”) devices even being fitted, others have occurred in multiple-manned trains with alerters.

⁹ This is of course most dangerous if the reliance is subconscious. However sometimes there is a deliberate “buddy system” where one crew member takes an agreed nap, thereby effectively reducing the operation to a one-person operation for that period.



Thus although VC cannot be relied upon to protect against sleep (nor other forms of inattention) nor against outcomes which arise before the VC has passed through its full time cycle, it is important at least to try to optimise the various design features and parameters for the particular context. Some aspects of this are discussed below.

The matter must also be viewed in the context of the occasions (typically unreported and possibly even unrecognised) in which the VC actually *does* serve the intended purpose. As the Rangitawa report succinctly states,

“The inability of the locomotive vigilance device to prevent short-duration microsleeps and the potential consequences raises doubts as to its suitability in its present form. However despite its limitations it does provide a warning and ultimate response system which must be weighed against the lack of defence if it were not present in the locomotive cab”.

Protecting against automaton responses

Theoretically, the best way is to have a response mechanism that does not permit automaton responses. The trouble is that anything satisfying that criterion is also likely to be a distraction to the driver.

For instance, it might be possible to require selection amongst several different buttons – e.g. the flashing button (randomly selected by the system) must be pressed, and pressing the wrong one will either have no effect or actually generate a penalty response. However at times of peak workload, is it safe to distract a driver from other actions to require him to make that sort of conscious effort?¹⁰

¹⁰ Of course task-linking goes a significant way to address this issue, since many – but not all – times of high workload involve frequent variations of control settings, which will be accepted as inputs to a task-linked VC. The button selection would therefore not be required at such times.



Speech recognition could be used, and the VC could select one of several words which the driver has to repeat. This might be slightly less distracting, using different parts of the brain and body to process, but presents obvious difficulties in ensuring that the system can reliably recognise different voices and might be industrially and operationally problematic. Such a system has apparently not actually been implemented anywhere¹¹.

However such conscious-selection paths have generally not been adopted. Instead, current thinking tends to be directed more to timing of the responses, and perhaps their sequence.

It is alleged, for instance, that setting random time intervals substantially reduces automaton behaviour. This does not appear very plausible, because it would appear that time-based habituation is much less likely than stimulus-based habituation – e.g. someone is more likely to respond automatically to a light or sound than merely to undertake an action repetitively at fixed time intervals like every 60 seconds. So the crew will likely respond to the light/sound irrespective of when it occurs.

Sequence-based systems may be somewhat more promising. For instance, mention was made earlier of deadman systems which require periodic release or increased pressure rather than accepting continuous pressure. Once such a feature is included, the deadman acquires some of the features of a VC.

¹¹ It should be noted, though, that systems have been suggested which would require a driver to vocalise a signal aspect as the signal is passed, to verify that the driver has correctly interpreted the signal display. Non-vocalised systems have physically existed which require a driver to confirm signal indications displayed in-cab, e.g. each time the display of an in-cab signalling device is changed, the driver has to register the newly-displayed item. These devices are not so much “vigilance control” but rather a development of AWS style systems, see later herein.

Moreover some railways require a driver to call each signal indication aloud, even if there is nobody else to hear; however there is no failsafe element in such a requirement. Some North American railroads require signal indications to be called over the radio, but (even apart from consequential saturation of radio communications in higher-traffic-density areas) this only detects non-vigilance if someone else observes that a signal has been passed without the corresponding radio call being correctly heard – and even then, it is not clear what short-term action is available.



Moreover it is more consciousness-demanding to (for instance) release a handle or pedal and then immediately re-apply the correct amount of pressure than it is to do either of those tasks alone. And if releasing/reapplying the deadman has to be followed by some other acceptable action, the probability of the complete sequence being done automatically is reduced further – but perhaps with increased risk of distraction at crucial times.¹²

Timing of the inactive part of the cycle

The light and sound phases have traditionally amounted to 10 to 20 seconds. There is not much room to move with these timings; they need to be long enough for reaction time and for non-immediate response due to other task demands, but must be short enough to provide effective action if non-response arises. Combined times as low as 8 seconds have apparently been used, but about 15 seems to be the norm.

This leaves the time between the previous activation and the first request for acknowledgment – i.e. the “inactive” part of the cycle - as the component most subject to determination. Periods as short as 25 seconds and as long as 90 seconds have been adopted in various contexts, although the shortest and perhaps the longest have typically been speed-linked, with the commonest range (especially for non-speed-linked) being 45 to 75 seconds. Non-task-linked intervals are typically longer than task-linked, in order to reduce task overload.

The key quantity is the time (and hence the distance) which can elapse between an incapacitation event and the train coming to a stop. This time consists of the total cycle time (i.e. inclusive of all phases including the inactive

¹² There is always a trade-off between reduction of automaton responses versus diversion of attention away from other safety or operational requirements. Increasing the conscious action required for vigilance-control response potentially distracts from other tasks, and in worst-case scenarios, may cause other tasks not to be done in a timely manner or to be done incorrectly.



stage) and the time for the train to come to a stand after braking is initiated. The worst-case distance is notionally based on the maximum permitted speed.

Thus for instance if the maximum speed is 100 km/h, a total cycle time of 90 seconds would entail the train travelling 2.5 km from the last vigilance acknowledgment to the time that braking is initiated, plus the distance involved in braking from that speed (which would depend on the train characteristics but could well exceed 1 km). Since many hazards can be encountered in 3.5 km, the VC will not prevent a large range of incidents.

But even if the inactive part of the cycle is reduced to 30 seconds, with a 15 second lights/sounds cycle, this still enables 1.25 km to be travelled in that time plus the braking distance. Even 2.25 km is still long enough for adverse consequences to occur.

Of course these are worst-case scenarios in the incapacitation context, in that typically the driver will not become totally incapacitated instantly after the last VC acknowledgment. On the other hand, it also emphasises the severity of consequences possible if an acknowledgment occurs while a person is effectively asleep or already in the throes of an incapacitation event.

Consequently it is always desirable to minimise the time of the inactive stage (that being the only available variable), but it cannot be reduced below a reasonable level from a workload viewpoint. Not only will over-frequent acknowledgment raise OH&S and driver distraction issues, but it will encourage the habituation/automaton risks.

Task-linking reduces the frequency with which the VC must be directly acknowledged, and hence potentially allows a shorter cycle time. For instance, it might be said that if only 50% of acknowledgments are “direct” by button-pressing (the remainder being effected by task-linking), then a 30 second true inactive interval will *appear* (to the driver) to be “on average” 60 seconds.



This is of course highly context-dependent, in that the frequency of control inputs in (say) a suburban application may be very different from that on long country runs. A common argument is that frequent acknowledgments are not a burden when there are no control inputs being required, so there is no harm in a short (say) 40 second “inactive” time under such circumstances; however a counter-argument is that if this increases the habituation/automaton risks at the very time when someone is most likely to be sleepy, the result could be counter-productive.

It follows that in the ultimate analysis, “whatever you do will be wrong”, in that any choice has significant disadvantages. Different jurisdictions will choose to weigh the conflicting criteria differently, and will hopefully reach a conclusion which is optimal given their weighting but which may not seem optimal to others.

Pre-emption

A closely allied issue is whether pre-emption is permitted, i.e. whether a driver can restart the inactive phase by pressing the button before he is requested to do so. It seems to be universally agreed that task-linking can be pre-emptive, but opinions differ regarding human pre-emption.

In favour of the practice are:

- the driver can avoid needing to activate the system at an anticipated high workload time by managing his acknowledgment actions appropriately;
- it reduces the occurrence of the distraction factors such as lights flashing and aural devices sounding;
- it provides the driver with a greater sense of “ownership”, of being in control with the VC as a backup rather than the driver being a slave to an impersonal device; and



- if the system is viewed as a sort of intermittent deadman device, it is beneficial to minimise the intermittent effect by having responses as frequently as convenient.

Against the practice are:

- the commonest manifestation of pre-emption is automated button-pressing, quite devoid of any conscious intent, and it may progressively increase habituation/automaton effects;
- it enables sleep to occur with no required external stimulus whatever to trigger pressing the button, whereas at least detecting a light flashing or something sounding involves some cognitive transfer; and
- if it leads to more frequent button pressing, it increases OH&S risks like RSI.

As with other aspects, decisions must be made by weighting the various criteria, and different people will reach different conclusions. Some jurisdictions are currently coming to favour an option of one pre-emption being allowed (to accommodate the first “in favour” dot point above) but a full cycle being required thereafter.

Arrangements with two-person operation

Although not a major issue, it should be mentioned that there are various views relating to how VC operations should be managed when two persons are both supposed to be “vigilant”. Options adopted within Australia have included:

- driver only to activate;
- driver and second-person to activate alternately; and
- driver or second-person may activate at any time.



Variants have even existed where driver and second-person had different acknowledgment arrangements (e.g. where the second-person actually had to lift himself out of his seat to reach the acknowledgment button, apparently on the theory that this made it more difficult for him to sleep). The arrangements may also be affected by whether means other than buttons and task-linking are permitted for acknowledgment purposes, e.g. a driver may have a deadman pedal on which pressure can be intermittently released/applied while the other person does not.

While the arrangements under two-person operation were once highly contentious, and may have required (for instance) that locomotives working between NSW and Victoria needed to be switched between systems, the topic no longer appears contentious, perhaps because it is realised that this also is a matter for trade-offs without conclusively “right” solutions.

What to do if a vigilance non-response occurs

Although the primary response to a non-response must be to stop the train, even then there are choices.

Is the brake application to be made in emergency or by a full service application? The former will usually (but not always) stop the train quicker; however it may be destructive of passengers or goods, may cause the train to break into two or more parts, may promote wheel skids, or may cause jackknifing or other derailment outcomes. Both emergency or non-emergency options are applied in different jurisdictions, although the former is probably more common (because, inter alia, it requires less sophisticated hardware).

Then when the train is stopped, what next? If there is a two-person crew, it is unlikely (but not impossible) that both are permanently incapacitated, but with one-person crewing, emergency assistance may be required. Consequently it is



becoming increasingly common for a vigilance non-response to be followed (either immediately or after a short period for over-ride) by automated radio communication to Network Control and other trains in the vicinity.

Another question – accommodated differently in different places – is whether immediate reset is possible by the train crew, or whether some time delay or other action is required. The former may be thought preferable to “keep the network” running in the case of inadvertent or otherwise non-problematical reasons for the original non-response; whereas the latter may be wise to prevent someone resetting and then relapsing immediately into the previous problem state if there is something really amiss. Where automated notification to Network Control is provided, the actions of the controller are also a key factor in this regard.

Physiological vigilance measurement devices

The road trucking industry has been concerned – for both commercial and safety reasons – at the incidence of fatigue and substance-induced inattention by truck drivers. Various studies have been conducted, both notionally and experimentally, to assess the physiological symptoms of such inattention states, and to try to design electronic means of triggering warnings or other responses if such symptoms are observed.¹³

Although such studies have found a number of measurable parameters which appear strongly correlated with inattentiveness, the correlations are typically not high enough to be conclusive (and hence anything based on them could not be failsafe). Moreover many possibly more reliable measures tend to be intrusive to measure (e.g. blood oxygen levels), while others such as eye movements or

¹³ Such studies date back to at least the 1980s; an informative example of a 1996 report, with some historical background, is at <http://www.fmcsa.dot.gov/facts-research/research-technology/publications/cmvmfatiguestudy.htm>



head nodding are difficult to measure reliably even irrespective of their reliability as indicators. Brain activity (as measured by EEG methods) is notionally an appropriate measure, but is highly variable even within one individual and different factors may be confounded. Parameters such as skin conductivity may be easy to measure but their relevance, reliability and calibration may be in question.

Except possibly for one ill-documented application in Russia and other countries in that vicinity, this current review has produced no evidence that there is any current application anywhere in the world where measurement of such parameters is used on a day-to-day (as distinct from experimental) safety-critical basis on either road or rail. Moreover, instances as have been partially applied (notably eye, head and facial movement detection devices in American trucks) appear generally not to have been persevered with, presumably because they were perceived as ineffective, insufficiently effective, or too difficult to calibrate.

In the road industry, there is interest in devices which monitor the characteristics of a driver's control inputs. For instance the tendency to make small steering adjustments may vary according to degree of sleepiness. However in view of the very different control behaviours required for trains, it is difficult to see application to the rail industry.

There is a serious risk that any ineffectiveness of an alertness-measurement device will be strongly counterproductive. As soon as the device is relied upon in any degree, it will generate a behavioural response along the lines of "I don't need to worry about becoming inattentive because the device will warn me and protect me". Even a low failure rate of the detection system can then result in a calamity.



Notwithstanding the rather negative summary above, some particular suppliers are promoting devices which are claimed to have adequate performance. ITSRR has not been in a position to evaluate these claims.

One is popularly known as the “Russian wrist watch” which is said to be associated with the Russian railway application mentioned above.¹⁴ It consists of a device physically very similar to (and indeed incorporating) a wrist watch, and is readily demonstrated to have effects such as waking people who go to sleep in meetings. In rail applications, the full system includes bringing the train to a stand if the system detects a non-responsive state. Whether its failsafe performance is adequate for safety-sensitive applications is not clear.¹⁵

It is claimed that a device called Optalert is ready for commercial application on roads, and satisfies the requirement of being substantially non-intrusive. It consists of spectacles akin to normal reading glasses, and an alarm system which responds to two levels of drowsiness by different alarm levels. However it does not appear that this has been merged with a system to do anything more than generate an alarm, which is less than useful as a response to more severe forms of incapacitation.

¹⁴ At <http://www.neurocom.ru/ru/about/australia.pdf>, it is claimed that the equipment known as DVTCS-L is or was in use in Russia, Ukraine and Latvia as long ago as 1997.

¹⁵ The British Rail Safety & Standards Board commissioned a report in 2002 from a company called Quintec. This report, together with the Board’s own commentary, is available at <http://www.rssb.co.uk/pdf/reports/research/Driver%20vigilance%20devices%20-%20systems%20review.pdf>

The consultants concluded at that time that the DVTCS was “mature in terms of design development and actual use, and would be the most suitable candidate to recommend for implementation on the railways”. However the RSSB in its response observed that considerably more testing and evaluation would be required. In the same response, the Board stated “RSSB do not intend to mandate the introduction of driver vigilance devices ... in the foreseeable future, if at all”. However such a conclusion must be viewed in the context of the British rail system having extensive use of AWS (as discussed below) and heading to use of TPWS.



Thus overall, while there may be hope that such devices are a forward path beyond deadman and vigilance devices of more traditional kinds, there is little evidence of their suitability at this point in time.

AWS

The systems are described above are essentially internal to the train. While they may take account of factors such as speed, they use no information about the track or anything else external to the train.

At the opposite extreme, devices under the broadest meaning of “automatic train protection” (ATP) take account of the external situation and make the train match the external requirements. (For fuller discussion, see the ITSRR discussion paper on ATP.)

In between, there are systems which take account of the external world and require some response from the driver, but do little or nothing to confirm that the driver’s response is appropriate. To a large degree, they may be viewed as an extension of deadman and vigilance devices in that they do little more than confirming that the driver is capable of making a response, but that they schedule that response based on some external condition.

Such systems exist, and have existed for a century, in a multitude of variants. It is conventional to use the British “Automatic Warning System” (AWS) as the standard example, although this convention overlooks the plurality of other similar systems, some of which are more sophisticated. AWS has itself been installed in a fairly direct form in Brisbane and Adelaide.

AWS philosophy derives from the British mechanical-signalling era, wherein there were individual signalboxes that each controlled a defined territory (typically at most a mile in length). On the approach to each signalbox from open line or from another signalled area, the train would first encounter a



“distant” signal. This signal could either give a caution indication – signifying that some restrictive condition existed in the area ahead – or it could give a clear indication that the train could proceed unhindered.

AWS (or its predecessors) was linked to each distant signal. When a train reached a distant signal, electromagnetic and/or mechanical devices would indicate to the AWS system on the locomotive whether the signal was displaying caution or clear. If the signal was clear, no response would be required from the driver; but if it was at caution, the driver would have to respond to an audible warning by activating an acknowledgment device (e.g a lever or button). If the driver failed to acknowledge the warning, the brakes would be applied.

The weakness was, of course, that there was no check that the driver actually did anything to respond in terms of train management. He could just carry on regardless of the warning – and sometimes did, with tragic consequences.

The system was gradually applied more broadly to other signals and to other locations such as approaches to heavily speed-restricted curves. However the more installations occurred, the more tendency there was for drivers to respond robot-style without actually absorbing or acting upon the warning – exacerbated yet further by the fact that for some trains, the warnings were actually superfluous, e.g. when a high-braking train received an AWS warning at a signal requiring immediate action only by a low-braking train.

Consequently AWS is sometimes viewed as little more than an intermittent externally-triggered deadman device – but equally, it can never be known how many accidents have actually been prevented by it. Certainly many examples exist in accident literature as to how it is known or reasonably presumed that a driver simply acknowledged an AWS warning (or even multiple AWS warnings) without acting thereon – but the contrary situations are typically not recorded.



Evidence allegedly exists that drivers often note the aspect of a signal for the first time when alerted to it by an AWS warning. This would appear to be somewhat problematic, not only in that wrong-side failure of the AWS device (however rare that may be) would likely result in the signal being missed, but more broadly because it breeds an over-reliance on what is intended to be a backup rather than primary resource.

In view of the disadvantages above and the existence of much more modern technologies, it seems unlikely that many fresh installations of AWS will be made nowadays.

More sophisticated versions of AWS do exist, notably versions where the actual aspect of a signal (or the nature of some other warning) is indicated to the driver rather than a simple warning or no-warning choice. However even if the driver is required to acknowledge the specific nature of the warning, there is still nothing to verify that the warning is acted upon – unless yet more sophistication is introduced (as is indeed sometimes done¹⁶), in which case the design is heading into ATP territory.

TPWS

TPWS is an example of an extension of AWS towards ATP. Although derived initially from AWS concepts, it contains “speed traps” which will cause a train to be brought under control if the AWS-style warnings are not responded to appropriately. Consequently it is not discussed further in this paper and the reader is referred to the ATP paper for further comment.

¹⁶ Various “cab signalling” systems evolved from two strands of thought: that signal information should be displayed in the driver’s cab rather than entirely externally, and that there should be on-train enforcement of what the internal and/or external signals specified. Consequently systems exist (or existed) which varied across the whole spectrum from mere in-cab display of signal aspects, through AWS-style warnings of restrictive aspects requiring various types of responses, through to partial or full enforcement. There may not be universal agreement as to where in this spectrum the transition is made between AWS-style and ATP-style systems.



Devices which fall outside the DSS definition but relate to driver vigilance

It should be mentioned for completeness that some reminder devices are sometimes included under the category of “driver safety systems”, although they do not fall within our definition. These devices are typically provided so that a driver can protect against forgetfulness (by himself or others) of some unsafe condition, such as a signal being at stop or someone working on the train.

A prime example is the British “Driver Reminder Appliance” (DRA) which is manually set by a driver when standing at a station with a signal at stop ahead. Setting the DRA prevents the train from powering until the appliance is reset, and protects against the common contingency that the train will start without further cognisance of the signal when platform duties are complete. It is intended basically as a SPAD-prevention initiative rather than a driver safety device in the context of this paper.

Even such a simple device does however have potential negative consequences which need to be included in an evaluation. Clearly if the driver does activate the device when intended to do so, it will substantially prevent start-away SPADs. However what if the driver neglects to activate the device? He/she may trust (consciously or subconsciously) that the device is protecting him/her from the possibility of a start-away SPAD, and hence may fail to undertake further review of the signal aspect - yet because the device had not in fact been activated, it cannot provide any protection and the SPAD will probably occur. Thus the device has then actually been counterproductive.

It might then be said that a device should be provided to activate the DRA automatically rather than requiring the driver to do so. But this would require yet more hardware and would typify the “tail chasing” syndrome where each new device requires yet further devices to protect against failure modes.



As a further example of this syndrome, the British have installed (in some small number of locations) a “SPAD indicator” to inform a driver visually that he has committed a SPAD at the previous signal. This appears to be a uniquely British development.

Degraded operations

Driver safety devices may fail

- right-side – stopping the train unnecessarily; and
- wrong-side – failing to respond correctly to an unsafe condition.

Failures of both kinds do occur, but wrong-side failures may not be apparent because the unsafe conditions do not arise to demonstrate whether they are correctly addressed. It follows that there needs to be a facility to respond to failures when they do occur, but also a testing program to verify that the devices will in fact operate as designed if called upon to do so.

The ways in which these situations are addressed appear highly variable around the world, being driven both by ideological or commercial issues (such as weighing service disruption against pure safety criteria) and by the availability of manpower. For instance, it may not be realistic to require a train to await the provision of an extra crew member if the train is in a remote location or if a qualified person is not likely to be available in reasonable time – the safety consequences (not to mention the service disruptions) arising from delaying a train under such circumstances may be worse than those arising from allowing the train to continue. This is exacerbated by the complexities of a multi-user network, wherein the effects of one operator’s actions may impinge on several other operators.

Consequently, unless the fault is immediately rectifiable, responses to such failure conditions vary amongst options (or combinations of options) such as



- providing override facilities which just “cut out” the offending device;
- providing override facilities which substitute for the offending device (e.g. having another person activating a push button);
- allowing the train to continue to the next defined suitable location;
- allowing the train to continue its present journey but be taken out of service on reaching its destination;
- taking the train out of service at “first reasonable opportunity”;
- obtaining an extra motive power unit to lead, or rearranging motive power units;
- restricting speed;
- obtaining an extra crew member in the cab (either from on-train staff or externally); and
- requiring extra oral communication.

In order to avoid a device being (deliberately or inadvertently) declared as defective when it is not, it is common to require a perceived failure situation to be discussed with a suitable authority off the train. However even this system is not reliable if not put into effect; for instance in Britain, a driver is required to consult a signaller for approval to proceed if an apparently false TPWS activation occurs, yet there have been serious incidents where a driver failed to consult with the signaller even though the activation had in fact been genuine.

A general appraisal of driver safety systems

As defined above, driver safety systems such as deadman devices and vigilance control devices have largely reached the limits of their potential, until such time as they can be designed to provide genuine assessments of *alertness/awareness* as distinct from mere robotic compliance.

The task-linked vigilance systems installed on RailCorp trains are reasonably state-of-the-art within their frames of reference, but should not be thought to



have qualities which they do not have (in particular, the use of the term “vigilance device” should not be interpreted to indicate that the driver is “vigilant” in the common usage of that word; the driver could be effectively asleep or otherwise unaware of his environment, and yet still be accepted as responding to the device).

It must be realised, however, that any assessment of the qualities of such systems involves an evaluation of trade-offs, of which frequent examples have been given above. For instance, there is not an objective basis for determining a unique optimum timing cycle for vigilance devices; by weighting the respective criteria, one may reach a considered decision, but it cannot be said that this decision has absolute truth nor that it would be applicable to another context.

A future path may exist with physiological vigilance assessment devices, especially since they may be more successful at detecting inattention states as distinct from physical incapacitation, but it would appear (despite the existence of some commercial products) that there is still some distance to go before a clearly established recommendation could be made.