

Security as a Safety Issue in Rail Communications

J. Smith

S. Russell

M. Looi

Information Security Research Centre
Queensland University of Technology,
GPO Box 2434, Brisbane, Queensland 4001,
Email: {j4.smith,s.russell,m.looi}@qut.edu.au

Abstract

Systems whose failure can lead to the damage of property or the environment, or loss of human life are regarded as safety-critical systems. It is no longer adequate to build safety-critical systems based on the control of errors and failures alone. Safety-critical systems must also deal with securing the data that is used in their operation. While safety and security engineering have evolved separately, there are a number of similarities. These similarities and efforts to integrate safety and security are identified. A project looking at securing safety-critical communications for the Australian rail network is also discussed.

Keywords: system safety, safety-critical systems, rail control, formal methods, security.

1 Introduction

Safety-critical systems are defined as systems whose failure could potentially result in loss of human life, damage to property, or damage to the environment (Knight 2002a). Modern railways are comprised of a range of systems that can be considered safety-critical according to this definition.

Safety results from a combination of system design and the operational environment in which that system is used. Changes in the operational environment of safety-critical systems can result in dramatic changes to the safety of such systems.

Safety-critical systems in general, and the railway industry specifically, are currently undergoing revolutionary changes. Mechanical and electromechanical devices are being replaced by solid state and programmable electronics that are often controlled remotely via communications networks; computing systems and associated software are being used in increasingly complex systems; and the environment in which these systems are deployed is becoming more overtly hostile. Designers and operators now not only have to contend with component failures and user errors, but also with the possibility that malicious entities are seeking to disrupt the services provided by their systems.

Security, now more than ever, is emerging as a safety issue for critical systems, and to date has failed to receive adequate attention.

Copyright ©2003, Australian Computer Society, Inc. This paper appeared at the *8th Australian Workshop on Safety Critical Systems and Software (SCS'03)*, Canberra. Conferences in Research and Practice in Information Technology, Vol. 33. P. Lindsay & T. Cant, Ed. Reproduction for academic, not-for profit purposes permitted provided this text is included.

This work has been conducted as part of an approved CRC for Railway Engineering and Technology (RailCRC) project, Theme 4, Project 19-54 <http://www.crcrail.com.au>.

This paper will introduce the concepts of dependability and safety in §2; approaches to safety in the railway industry, including application of formal methods, are discussed in §3. In §4 some of the interrelationships between safety and security are discussed, and the need for security in safety-critical systems is presented. The paper concludes in §5 with a description of a CRC for Railway Engineering and Technology project currently being undertaken to address the secure communications needs of emerging rail control networks.

2 System Safety

Historically safety was approached from a component reliability perspective in the belief that if each component of a system was safe that the system too would be safe. However, over time it was realised that many safety incidents were the result of complex interactions between components, each of which was reliable. This observation gave rise to the concept of system safety, where safety is considered in the context of not only the components that make up a system, but in the context of the interactions between the components, and between the system and its environment.

Safety can be considered in the context of the more general concept of dependability, which is discussed next.

2.1 Dependability

Dependability is described as a property of a computer system such that reliance can justifiably be placed on the service that such a system provides (Laprie 1996). The field of dependable computing is interested in the measures, impairments, and means of provisioning dependability of services offered by computing systems. Dependable computing provides a conceptual framework to discuss dependable systems in general and safety-critical systems in particular.

2.1.1 Measures

A number of measures of dependability have been put forward and include notions of: (1) reliability - a measure of continuous service accomplishment; (2) availability - a measure of service accomplishment with respect to service interruption; (3) maintainability - the ease with which the system is returned to service following an interruption; and (4) safety - a measure of the hazards present in the system. Collectively, reliability, availability, maintainability, and safety is known as RAMS and is referred to in numerous standards. Given these measures of dependability we can intuitively deduce that a dependable system is one which is reliable, available, maintainable and safe.




Figure 1: Failure Aetiology

2.1.2 Impairments

In the design and development of systems, a range of factors work to diminish dependability. These factors include: (1) faults - physical or human made, that result in a latent error being present in the system; (2) errors - initially latent, errors can be activated, having the potential to cause system failures; and (3) failures - the activation of an error that leads to a deviation in the delivery of a service (See Figure 1). An error can be considered the manifestation in the system of a fault, and a failure is the effect of an error, on a service (Laprie 1996).

2.1.3 Means

The techniques of fault prevention, fault tolerance, fault removal, and fault forecasting can be used to improve the dependability of a system, by reducing faults, errors and failures.

Fault prevention, also known as fault avoidance, is a strategy that makes use of techniques to prevent the introduction of faults into the system at design. Fault avoidance is referred to later in the paper when formal methods are discussed (See §3.2).

Fault tolerance, makes use of techniques such as redundancy, to ensure the correct operation of the system in spite of the presence of faults.

Fault removal, also known as error removal, aims to minimise the presence of latent errors through the process of verification ¹.

Fault forecasting, also known as error forecasting, estimates the presence, creation, and consequence of errors through evaluation.

Collectively, fault prevention and tolerance provide techniques for obtaining dependable systems, while fault removal and forecasting permit dependability validation.

A partial taxonomy of dependability concepts is shown in Figure 2 (Laprie 1996).

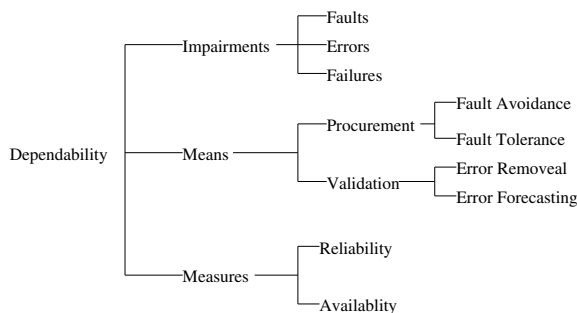


Figure 2: Dependability Taxonomy

¹Ensuring that the system is being built right (according to the specifications) as opposed to validation, which ensures that the right system is being built

2.2 Safety

As discussed in §2.1 safety is an attribute of a dependable system. Dependability is an important concept, but it is argued (Leveson 1995) that safety presents specific concerns² that warrant it to be more closely examined in this section. First, some definitions (Leveson 1995, Standards Australia (SA) 1999b):

Safety is freedom from accidents or losses.

Hazard a source of potential harm or a situation with a potential to cause loss.

Hazard level the combination of hazard severity and hazard likelihood.

Loss any negative consequence, financial or otherwise.

Risk the chance of something happening that will have an impact upon objectives. It is measured in terms of consequences and likelihood.

Safety is about managing risks through hazard elimination, reducing the likelihood of hazards that cannot be eliminated, and reducing the hazard exposure duration for those hazards that do occur.

Safety engineering is the process of applying scientific, management, and engineering principles through the lifecycle of a system (Leveson 1986) to ensure that the hazards, and therefore the risks, in the final system are reduced to an acceptable level. The goal of safety engineering is to identify and control hazards in a manner that permits a ‘proof of safety’ about a system to be made. While more complex and time consuming than traditional engineering approaches, the use of a safety engineering process is preferable to reactively improving the safety of a system following an accident, especially in safety-critical systems, or systems in which failures are potentially catastrophic.

There are many approaches that can be taken to safety engineering that are tailored to specific applications but in general, the safety engineering process will require a safety analysis, validation, and verification to be performed.

2.2.1 Safety Analysis

The initial safety analysis incorporates the following stages (Eames & Moffett 1999):

- Functional and technical analysis;
- Qualitative analysis;
- Quantitative analysis; and
- Synthesis.

Commencing with the functional and technical analysis provides the analyst with information on how the system works, the system structure, the environment it will operate in, and defines the system boundaries.

A qualitative analysis is then performed to explore the failure causes and potential hazards affecting the system. A range of hazard analysis techniques³ are

²Leveson argues that a system may be more dependable, but less safe and that the use of the generic term of dependability does not aid in the assessment of specific components such as safety.

³including: design reviews; checklists; fault tree analysis; event tree analysis; hazard and operability studies (HAZOP); random number simulation analysis (RNSA); failure modes and effects analysis (FMEA); and failure mode, effect, and criticality analysis (FMECA).

available to support this process and are described in the literature (Leveson 1986, Falla 1998).

The quantitative analysis phase then assigns values to the hazards identified and is used to rank the relative risk of particular hazards.

In the synthesis phase, the qualitative and quantitative outputs are combined to identify critical components and functions of the system that must be controlled to ensure the resulting system is acceptably safe.

2.2.2 Safety Verification and Validation

Safety verification and validation involves the provision of a proof of safety. This can be achieved through demonstration that faults cannot occur, or if faults can occur, that they are not hazardous. The technique of proof by contradiction may be used in providing safety assurance, as one of the goals of safety verification is to prove that something will not happen - the system entering an unsafe state for example.

Other techniques used in verification and validation can include: observing the system in operation; or performing static analysis via formal verification, fault tree analysis or independent verification and validation.

2.3 Software and Safety

Computer systems are increasingly being used in the control and operation of systems that have the potential to cause harm should those systems fail.

As newer control technologies have evolved the safety engineering process has also had to evolve from application to purely mechanical systems to electromechanical and more recently computer (software) controlled systems. The widespread replacement of dedicated controller hardware, with more general purpose programmable controllers and computers, local or remote, makes the safety analysis process more difficult.

2.3.1 Software Complexity

The use of general purpose computing devices introduces software complexities relating to the underlying operating system, the application program itself, and the runtime environment in which the application will operate. Each of these components needs to be adequately assessed, and more importantly, maintained in a manner that does not degrade the safety of the system. Few enterprises will have the in house expertise required to develop the operating systems, runtime environments and applications used in these general purpose systems, so the process of integration and evaluation becomes more complex.

The strategies and techniques employed by safety engineering in general need to be extended to address the additional safety issues that arise from the use of general purpose computers and software in the control of safety-critical systems. The evaluation of such devices is far more complex than the evaluation of dedicated, specialised devices traditionally used in control systems.

2.3.2 Software Errors

Over and above hardware faults and failures, computers may be exposed to errors in software logic that stem from incorrect requirements specifications, or incorrect implementation of the requirements. The following approaches are recommended for managing software errors (Leveson 1995):

1. Ensure the requirements and code are correct

2. Enhance software reliability through fault tolerance
3. Apply standard system safety techniques

Unfortunately, correctness and fault-tolerance are insufficient, as software related accidents can occur when the software is reliable and well specified. This can result as: (1) requirements specify unsafe behaviour; (2) requirements fail to specify a necessary safety behaviour; or (3) the software exhibits unintended behaviour, above that specified in the requirements.

Unintended behaviour may result from implementation issues, or from exposures to vulnerabilities in the underlying operating system or runtime environment that could potentially be exploited by viruses and other malicious code.

The combination of ensuring that requirements and code are correct, using reliability techniques such as fault tolerance, and the application of system safety techniques is currently the best method for minimising safety hazards that may result from the use of software in safety-critical systems.

2.4 Discussion

While safety improvements have the potential to be realised, they must be considered in the context in which such systems are used. The potential improvements in safety systems can be countered if the presence of such safety systems encourages users of the system to take greater risks, resulting in a reduction in overall safety.

The impact of complexity on safety must be considered. Systems used to enhance reliability may reduce safety by creating more points for failure or error. Additional complexity will also make the verification and validation less tractable.

Reliability has the goal of preventing failures and is relevant to preservation of safety in the presence of hazards caused by failures. However, while overlapping with safety, reliability is not synonymous with safety, as hazards can result from interactions between system components that do not fail. Failures do not always result in hazards, and hazards are not always the result of failures.

3 Safety in the Railway

Concepts associated with dependability and system safety have been introduced in §2. This section of the paper will identify safety strategies employed by the railway industry.

The rail industry has an excellent safety record based on the safe separation of trains and fail-safe protection of paths through junctions and crossings. The safe separation of trains is based on the assumption that lead trains are able to come to a complete stop instantaneously⁴ and that following trains exhibit suboptimal braking performance⁵.

Fail-safe protection of paths through crossings and junctions is provisioned through the use of interlockings that prevent conflicting signals and points being set.

3.1 Signalling

A number of signalling strategies can be used to support the safe separation of trains. These separation techniques are based on a train's current location, direction of travel, and speed in relation to other trains

⁴Known as a brick-wall stop

⁵Known as worst-case breaking

in the same area. The mechanisms used for determining the trains location can either be based on the track being subdivided into fixed-block lengths, or variable block lengths - known as moving-block systems (Transport Research Board 1995).

3.1.1 Fixed-Block Signalling

In fixed-block signalling the track is divided into segments. Each segment, or block of track, can only be occupied by a single train at a time. Segmentation of the track can be achieved with passive markers, line side signals, or track circuits. Track circuits are commonly used and are able to detect the presence of a train on a segment. Track circuits set the signals at the entry point to the block to stop, thereby preventing two trains from occupying the same block simultaneously⁶.

In addition to signals, known as aspects, being provided by wayside lights, it is possible to use in cab electronic signalling. With in cab signalling, instead of taking cues from wayside equipment, the train receives aspects directly from the track circuit in the form of code signals. These signals specify the maximum allowable speed for this section of track. This speed is displayed in the locomotive cabin, next to the current speed.

Fixed-block signalling based on track circuits is designed to be fail-safe, so that if track detection circuits fail, for example, the aspect guarding the entry to that segment will show stop.

Automatic Train Protection To ensure that trains observe the signals (wayside or in cab) governing which blocks can be occupied and maximum allowable speeds, automatic train stops and automatic train protection can be employed. Automatic train stops prevent a train passing through a red signal, through the application of the emergency brakes⁷. Automatic train protection (ATP) utilises the train's on board computer to compare maximum allowable and observed speeds. If the observed speed exceeds the maximum allowable speed, the brakes are applied.

ATP can be implemented as a warning or an enforcement mechanism. When operating as a warning mechanism, the driver of the train is warned when a signal or speed violation has occurred and has the opportunity to acknowledge the event and cancel the automatic application of the brakes. The implementation of ATP as a warning system does not prevent collisions, or overspeed derailments, that result from driver mistakes. An additional side-effect of ATP warning systems is that the driver may become reliant on the warning system to safely operate the train, assuming that in the absence of a warning that no hazard exists.

3.1.2 Moving-Block Signalling

While fixed-block signalling supports safe operation of the railway, it has some drawbacks including: a requirement to maintain a large amount of wayside and track circuit equipment; and that the block size determines the maximum track utilisation, even in

⁶This is a simplified description of fixed-block signalling. In order to efficiently manage rail assets it is usual for a range of aspects to be displayed in lead up blocks (distant signals). The website <http://www.trainweb.org/railwaytechnical/sighis.html> provides an excellent review of signalling systems and is recommended to the interested reader.

⁷A mechanical arm is raised along the trackside whenever the signal is showing red and will strike a trip cock on the train, activating the emergency brakes if a train ignores the signal.

the presence of trains with improved braking performance. To address these maintenance and efficiency costs, the concept of moving-block systems has emerged. In moving-block signalling systems, which replace fixed-block infrastructure with continuous two-way communications, a computer system calculates, in a fault-tolerant manner⁸, the safe braking distance on the basis of the trains current speed, direction of travel, track conditions, and surrounding traffic. The continuous nature of the safe separation distance calculations results in the ability to run trains safely with reduced headway, leading to greater track utilisation.

3.2 Application of Formal Methods

It is apparent from §3.1 that the rail industry makes use of a numerous safety techniques that involve monitoring and preventing the system from entering unsafe states through the use of supervisory systems such as ATP, interlockings at crossings and junctions and dealing with failures through the use of fail-safe designs and fault-tolerance. In addition to the use of these techniques, the rail industry has adopted the use of formal methods for modelling, requirements specification, design, and validation of safety-critical systems.

Formal methods permit the behaviour of a system to be mathematically verified, providing assurance that the system design and implementation satisfy system functional and safety properties. Formal methods can be considered a fault-avoidance technique, in dependability parlance, that prevent the entrance of errors to a system and have been applied to the following areas in safety-critical railway application:

- Signalling systems and train control(Bowen & Stavridou 1993, Dehbonei & Mejia 1994, Gnesi, Latella, Lenzi, Abbaneo, Amendola & Marmo 2000, Hussey 2000, Bohn, Damm, Wittke, Klose & Moik 2002);
- Automatic train protection systems(Dehbonei & Mejia 1994, Simpson 1994);
- Modelling of signalling rules(King 1994); and
- Verification of interlockings(Hansen 1994, Hartonas-Garmhausen, Campos, Cimatti, Clarke & Giunchiglia 2000)

3.2.1 Experience Summary

While it would not be practical to apply formal methods to highly complex systems in their entirety, the aforementioned applications were possible because the formal methods were applied to specific safety-critical systems, as identified by a system safety analysis.

The use of formal methods, as identified in §3.2 resulted in greater understanding of the systems being modelled (King 1994, Hansen 1994, Hartonas-Garmhausen et al. 2000), highlighted errors and ambiguities (Bowen & Stavridou 1993, Gnesi et al. 2000, Hartonas-Garmhausen et al. 2000), and assisted in conducting hazard analyses (Hussey 2000). Some of the projects reported on the value of combining the use of formal methods with simulation techniques to highlight errors in the underlying model used by formal methods (Hansen 1994).

⁸The computation may be made by three separate programs, each independently verified. The output selected is based on a threshold - eg 2 out of three agree, permitting fault-tolerant failure of one of the processes.

3.2.2 Limitations of Formal Methods

Formal methods only speak of correctness in regard to a specification (Bowen & Hinchey 1994). Formal methods cannot guarantee that the specification is correct, but as the experience reports above indicate, the adoption of a formal approach often aids in the detection of deficiencies or inconsistencies in models on which specifications will be based. It has been identified that while formal methods are useful in clarifying requirements, greater challenges may lie in teaching the proper modelling of systems (Gerhart, Craigen & Ralston 1994).

3.3 Standards

A range of standards relating to the rail industry and the use of electronics components and software in safety-critical systems. These standards aim to provide guidance on how the process of designing, procuring, and deploying safety-critical systems can be done in a manner that provides some assurance of the safety characteristics of such systems. Two standards are briefly discussed in this section.

3.3.1 AS 4292

Within Australia there is a Railway Safety Management Standard (AS 4292) that is comprised of six parts: (1) General and interstate requirements; (2) Track, civil, and electrical infrastructure; (3) Rolling stock; (4) Signalling and telecommunications systems and equipment; (5) Operational systems; and (6) Railway interface with other infrastructure. This standard specifies a process for railway safety management through the asset lifecycle, incorporating the phases of: design; construction and implementation; commissioning; monitoring and maintenance; modification; and decommissioning and disposal. The standard encourages an approach that incorporates hazard identification and risk analysis with specific regard to: safety and integrity, especially of signalling and telecommunications; reliability; compatibility; fail-safe operation; and safety. The standard requires that the design and construction and implementation phases be validated (right system) and verified (built right). Formal methods would be useful in meeting these aims, but the standard does not provide specific direction on how verification and validation are to be executed.

The similarities between the goals of this standard and the goals of dependable computing do not go unnoticed.

3.3.2 IEC 61508

IEC 61508 Functional Safety of Electrical / Electronic / Programmable Electronic (E/E/PE) Safety-Related Systems standard (International Electrotechnical Commission (IEC) 1998) has emerged in response to the recognition that E/E/PE components are increasingly being used in safety-critical systems and that the risks presented by such widespread use of these components has to be adequately managed. IEC 61508 promotes a design methodology framework that aims to prevent the presence of dangerous failures, or control them when they arise by providing guidance on each phase of the safety lifecycle. IEC 61508 has been adopted as an Australian standard (Standards Australia (SA) 1999a).

The safety requirements of a safety-related system (SRS) must be specified in terms of the functions to be performed by the SRS and the integrity required of each. This results in the specification of a safety integrity level (SIL) for each safety-related function.

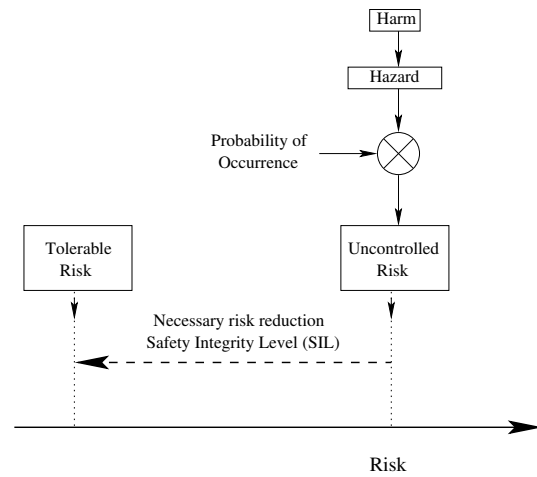


Figure 3: IEC65108 Risk/Requirements Model

The SIL is determined by the action necessary to reduce the risk associated with a function, from an uncontrolled risk, to a tolerable risk and is depicted in Figure 3 (Fowler & Bennett 2000). A SIL is a property of a safety function and corresponds to the required likelihood of failure probability. That is to say, a safety function with a requirement to have a probability of dangerous failure per operating hour of 10^{-8} to 10^{-9} will receive the highest SIL level of 4, while a safety function with a requirement to have a probability of dangerous failure per operating hour of 10^{-5} to 10^{-6} will receive the lowest SIL of 1. The SIL level is used to determine the requirements that must be met at each stage of the system lifecycle.

Once the safety requirements have been specified, evidence must be presented that demonstrates the risks presented by the system are tolerably low – this is known as providing safety-assurance. Safety-assurance, in IEC 61508, is demonstrated by presenting evidence: the system safety requirements are necessary, complete and correct; the safety requirements are fully realised in design and development; the installed system meets safety requirements; and the system is prepared for operational use.

Application of IEC 61508 to the certification of transport infrastructure systems is presented by Fowler and Bennett (Fowler & Bennett 2000).

3.4 Emerging Trends in Railway Communication

Signalling systems are becoming increasingly sophisticated and rely on advanced communications technologies, for location reporting⁹ and the issuing of movement authorities, to improve the efficiency and safe operation of the railway. Rail communications technologies in use include: GSM for Railways (ETSI 2001); Terrestrially Trunked Radio (ETSI 2003)¹⁰; Enhanced Position and Location Reporting System (Nishinaga, Evans & Mayhew 1994); Inductive Loop; Satellite; and a range of proprietary systems¹¹. Each of these: requires varying levels of infrastructure investment; supports different volumes of traffic¹²; and provides differing lev-

⁹Location reporting can be performed using a range of techniques including the use of Global Positioning Satellite services.

¹⁰Previously known as Trans-European Trunked Radio and before that MDTRS Mobile Digital Trunked Radio System

¹¹A summary of CBTC projects and associated communications technologies can be found at the Transportation Systems Design website <http://www.tsd.org/cbtc/index.htm>

¹²There is some concern as to GSM-R's ability to provide enough circuit switched channels at busy junctions (Jones, Porter & Wa-

els of assurance with regards to security of messaging, prioritisation of messaging, and quality of service guarantees.

Traditionally railway communications for safety-critical applications have been based on closed systems that are expensive to deploy, maintain and operate. Such systems are not appropriate for many parts of the Australian rail network, and the absence of clear alternatives for communications in these areas results in delays in the introduction of more flexible train control systems. The use of radio based communications technologies could help reduce the infrastructure costs associated with deploying alternative technologies, and obtaining such radio services from commercial providers may reduce these costs further. While such a solution is appealing from an economic view, the use of wireless and public communications systems introduces a security risk which could have a significant impact on safety.

The desire to migrate to remotely controlled trains is not new (Thorne-Booth 1968), but is gaining recent momentum as rail operators are increasingly having to operate in a highly competitive market place. This trend towards transmissions, or communications based train control (CBTC) presents unique safety and security challenges. Unfortunately the great expense associated with the deployment of such wide area communications cannot be justified by safety benefits alone and there is a strong desire to utilise the communications network that is deployed for CBTC for other less vital applications. This further compounds the safety and security challenges presented by this emerging trend.

4 Security as a Safety Issue

Computers and associated software are increasingly being used to control safety-critical systems, often these systems are controlled remotely over media such as fibre optic cable, microwave links, or other radio-based technologies. Critical systems connected to communications networks cannot be considered safe, unless the commands that are issued to those systems are adequately secured. As identified in §2.3.2 the use of operating systems, runtime environments, and application software exposes safety-systems to vulnerabilities (Thompson 1984, Davida, Desmedt & Matt 1989, Levy 2003) that may result in the exhibition of unintended behaviour unless each layer is adequately secured and maintained.

Section §3 summarised a range of strategies that are used to ensure the safe operation of critical systems in the railway. In this section of the paper, the importance of security and the potential impacts on the operation of safety-critical systems if security is not adequately provided is explored.

For the purposes of this discussion, security refers to computer and communications security. While physical security is an important consideration, it has generally been well considered by safety standards¹³.

4.1 Security and Safety

Security researchers have been interested in how security methodologies can benefit from the approaches used in safety-critical and dependable systems for some time (Meadows 1995, Meadows & McLean 1998, Brostoff & Sasse 2001), but it is only

boso 2003)

¹³The Australian Railway Safety Management Standard (Standards Australia (SA) 1997) addresses physical access control, but no direction is given with regard to communications or network security.

recently that the importance of security, and its potential impacts on safety have gained growing recognition (Eames & Moffett 1999, Winther, Johnsen & Gran 2001, Knight 2002a, Knight 2002b).

4.2 Security Analysis

Like the safety engineering community, the security community has established analysis methodologies that are used in determining the risks faced by a system and how those risks might be reduced to a tolerable level¹⁴. However, unlike the safety engineering community, security engineers are much more pessimistic about the disposition of entities considered in the system. Secure systems are engineered on the presumption that entities (usually referred to as attackers or adversaries) are intentionally trying to disrupt the system and that they have access to a wide range of resources to assist them in this pursuit. So while the safety community designs systems to accommodate errors and failures, the security community designs systems to accommodate malicious actions.

Intuitively this difference in assumptions makes sense, as users of safety systems would be unlikely to benefit from disrupting system safety efforts as this would only increase the possibility of themselves or their environment being harmed. Unfortunately, the environment in which safety-critical systems are being used is more overtly hostile, and the increasing use of communications networks in remotely controlling safety-critical processes presents an opportunity for attackers to disrupt these systems remotely, without having to bear the consequences of such a disruption.

There are numerous security risk analysis methodologies that can be employed but they all incorporate the following phases (Eames & Moffett 1999):

- Asset identification;
- Vulnerability analysis;
- Likelihood analysis; and
- Countermeasure evaluation.

Like safety engineers, security engineers have finite resources, so the security analysis process is designed to ensure that the right things are being protected.

Asset identification is crucial to the risk analysis process and will include tangible as well as the identification intangible assets (such as reputation).

Once the assets have been identified a vulnerability analysis is performed to determine all the possible ways in which those assets can be damaged and the consequences of such damage.

The likelihood analysis determines how often the system will be exposed to each of the identified vulnerabilities and the risks to the assets can be prioritised in a way that the most vulnerable and those with the biggest consequences are protected by adequate countermeasures.

4.3 Security Vulnerabilities

In the absence of adequate security countermeasures, safety-critical systems that rely on input via a communications link to operate safely are highly vulnerable. Current systems are designed to deal with errors and failures, not malicious and determined attempts to disrupt system operation.

Consider a safety-critical system, such as a train control application in which speed and braking profile information is transmitted to the locomotive via

¹⁴You could replace safety integrity level with security integrity level in Figure 3

some communications link. Existing safety standards and practices as discussed in §3 will assure to some degree that the communications end points have been well specified, verified, and implemented with respect to the assumed risk environment. Systems typically deal with transmission errors, resulting in data corruption, by requesting retransmission or through the use of error correcting codes. Communications link failure resulting from component failure or interference, intentional or otherwise, can be accommodated by the use of fail-safe designs. Of more concern, is how the system will handle the case where there is no error or failure, but where an attacker is intelligently manipulating messages, attempting to transmit forged speed and braking profile information for example?

As was stated in §1, safety is a combination of system design and the environment in which such a system is used. The environment in which rail control applications are likely to be deployed, is significantly different from the environment that current control infrastructure operates. In place of closed, point to point communications links, there is growing use of wireless transmission technologies. This is combined with a significant change in the assumed threat environment - it is no longer sufficient to design safety-critical systems on the assumption that the major hazards are due to errors or failures. The safety of such systems requires that the system operate safely in the presence of increasingly malicious and motivated attackers.

4.4 Security Countermeasures

In order to reduce the risk of an asset being damaged by an attacker to an acceptable level, security countermeasures can be deployed. In the context of system safety you can think of security being used to reduce the hazards that might result from misuse of system components. Before some security countermeasures are discussed some definitions are provided (Kaufman, Perlman & Speciner 2002):

Authentication The process of reliably determining the identity of a communicating party.

Confidentiality Protection against unauthorised disclosure of information.

Integrity Protection against the unauthorised modification of data.

The security services of authentication, confidentiality, and integrity can be used to protect information assets against the threats of: (1) masquerading - where one entity impersonates another; (2) eavesdropping - the unauthorised disclosure of information; and (3) unauthorised message manipulation.

Usually, authentication, confidentiality, and integrity are provided through the use of cryptography. Cryptography makes use of mathematical techniques, implemented as cryptographic primitives to provide the required security services¹⁵.

Generally, cryptographic primitives are based on an algorithm (the mathematical technique) and some secret information, known as the cryptographic key. The information to be protected is then transformed in some way by the algorithm and the key. For example, to provide a confidentiality service: an encryption function will take a message and a secret cryptographic key as input, and output ciphertext - an unintelligible message. The original message can only be recovered from the ciphertext by someone with

knowledge of the secret key. This process is shown in Figure 4. Cryptographic primitives can similarly be used to provide the services of authentication and message integrity protection.

While safety-critical systems may implement some form of error detection with regards to data integrity via the use of cyclic redundancy codes (CRC's), or use network address information for determining the origin of messages, these are very different from the notions of integrity protection and authentication in a security context. The use of CRC's and addressing information is considered weak, as it is trivial for an attacker to modify this information undetected. In contrast, the security context demands that strong authentication and integrity mechanisms are used. This strength is derived from the fact that the authentication and integrity protection relies on knowledge of the secret cryptographic key.

4.4.1 Applying Countermeasures

Previously in §4.3 a scenario was presented where an attacker might seek to send forged speed and braking profile information to a locomotive using a train control application. Assuming that some portion of the communication between the controller and the locomotive utilised wireless technology it would be trivial for the attacker to gain access to the communications medium. The attacker could then either inject bogus control data into that media with appropriately calculated CRC and addressing fields, or modify a message containing control data enroute, recalculating the CRC.

In order to safely execute control orders, a locomotive will need assurance: that the control message has come from an authorised controller; and that the message has not been modified enroute. These two goals can be met by the services of authentication and integrity protection.

The use of cryptographic authentication ensures that an attacker is not able to masquerade to locomotives as a controller, and the messages transmitted between the controller and locomotive are integrity protected, such that any message modification will be evident.

The specific cryptographic primitive chosen to provide a particular security service will depend on the required (Menezes, van Oorschot & Vanstone 1997): level of security; functionality; methods of operation; performance; and ease of implementation. An important, and often overlooked issue in cryptographic systems is the key management issue. It is not sufficient to have cryptographic primitives that meet the required security parameters, these must be complemented by an adequate infrastructure to support the generation, distribution, use, and destruction of the cryptographic keys used. The specific primitives used to protect the described communications must be judiciously chosen, as safety-critical systems often operate under stringent time, processing, and communications constraints. The application of security to safety-critical systems must ensure that system safety is enhanced and not degraded, otherwise hazards and risks are being traded against each other with no net improvement in safety.

4.5 Validation and Verification of Primitives

Like the safety community, the security community in general and cryptographers specifically recognise the need for the mechanisms that are used to reduce risk to be verifiable. This has resulted in great interest in provable security (Bellare 1998), in which formal methods are used to rigorously assess the correctness of cryptographic protocols.

¹⁵ Availability is an important component of security, that cannot be achieved through the application of cryptography.




Figure 4: Message Encryption

4.6 Harmonisation of Safety and Security

While the safety and security fields have evolved somewhat in isolation there are a number of observations that can be made including the following:

- Both safety and security are more effective when designed in to systems and not retrofitted.
- Complexity potentially diminishes both safety and security. It is unlikely that a highly complex system will be safe or secure, and providing assurance about either will be difficult.
- Both safety and security are concerned with managing risk.
- Both safety and security are significantly dependent on the human element - human factors can make or break well designed systems.
- Both safety and security consume resources that could be used in activities more directly related to production.
- Neither safety nor security are noticed when they are working, but get significant attention when they fail.
- Both safety and security incur certain costs, for what may appear to be uncertain benefits.

The first observation regarding designing safety and security into systems suggests that there may be benefits to greater coordination between the design efforts of safety and security engineers. Eames and Moffett (Eames & Moffett 1999) identify that there is a danger that conflicts can arise if safety and security requirements are developed in isolation. They consider the integration of safety and security requirements processes suggesting that this can be achieved through either unification or harmonisation of the requirements processes. They conclude that while there are significant similarities between the fields of safety and security, each has developed specific tools and skills, that if unified could result in compromises that would result in safety and security risks going unobserved. Harmonisation of the approaches on the other hand, while resulting in two sets of requirements, would ensure that safety and security, and their inter-relationships were considered at design time without the potential for requirements to go unobserved.

It appears that there is a significant opportunity for the safety and security communities to collaborate in the emerging rail control application environment to ensure that the systems deployed are both safe and secure.

5 Secure Communications for the Australian Rail Network

Recognising the importance of communications as a key enabler of emerging railway control and business applications the Cooperative Research Centre for Railway Engineering and Technology (RailCRC)

is jointly conducting a project with the Queensland University of Technology (QUT) to evaluate secure communications in the Australian rail network.

5.1 Project Background

The communications technologies used in the Australian rail network vary from state to state and between urban and rural areas, forcing rail operators who transit between states or regions to install and maintain communications equipment for each of the systems in use. Much of the equipment is unable to support data transmissions, so voice based messaging is prevalent. The issuing of movement authorities to a train can be a time consuming and error prone process, and could be greatly simplified by the judicious use of data communications.

Internationally, there is significant interest in improving track utilisation by replacing existing fixed-block signalling, with systems capable of providing a migration path to more efficient moving-block systems¹⁶ that utilise Communications Based Train Control (CBTC). CBTC is a system that uses two-way continuous communications to safely control a train's position, speed and other functions by exchanging control and telemetry data between the train and wayside equipment using either inductive loop (IL), or radio frequency (RF) transmissions.

In order to more efficiently manage Australian rail infrastructure assets, there is a desire to migrate the diverse voice based control systems currently in use, toward a standard communications infrastructure capable of supporting CBTC. Providing a consensus target for communications in the Australian rail industry will allow operators to make strategic investments in infrastructure that ensure interoperability with other systems in use around the country, and will provide a clear indication to manufacturers of communications equipment what is required in the rail control environment.

Traditionally, control systems such as CBTC are deployed using dedicated communications networks that are purpose built to meet the stringent safety and security requirements that such applications require. In Europe, the communications for CBTC are provided by a modified version of GSM, known as GSM for Railways or GSM-R (Hillenbrand 1999). In North America, CBTC is carried by a mixture of proprietary communications systems. The cost of deploying and maintaining such a dedicated infrastructure can be prohibitive and a more economical solution is desirable. By deploying a multipurpose communications infrastructure, capable of meeting the signalling and business application needs of rail stakeholders, the deployment and maintenance costs can be absorbed by a larger number of entities, making the infrastructure more economically sustainable.

While the motivation for the deployment of a network capable of supporting business and train control data is appealing, it is really only practical if there can be a high level of assurance that safety-critical

¹⁶ERTMS Level 3 permits moving block operation see <http://www.ertms.com/level03.html>

systems, such as CBTC, can be guaranteed a specific level of service from the network independent of other applications sharing the resources of that network. Additionally, the range of business data that is carried by the communications network should only be available to authorised entities.

5.1.1 Common Network Protocols

The Transmission Control Protocol / Internet Protocol (TCP/IP) suite (Postel 1981a, Postel 1981b) has emerged as the prominent networking protocol. Its ubiquity stimulates a desire, for many, to run all applications over a single TCP/IP communications infrastructure. Using a standardised and ubiquitous communications protocol such as TCP/IP has numerous benefits: it ensures a choice of vendors; interoperability; access to a large pool of expertise; and is more efficient to maintain and operate.

While IP may seem an alluring solution, it must be noted that IP assumes that all hosts connected to a network operate in a manner consistent with the defined protocols. Currently, IP networks are unable to provide the service assurance that would be required to support safety-critical applications. The provisioning of an IP network capable of performing robustly in the presence of misbehaving nodes is a challenging aspect of this project.

5.1.2 Integration of Rail Control Applications

To perform satisfactorily, critical applications such as rail control, require specific guarantees from a communications infrastructure. These guarantees may be as to the origin of messages (message authentication), that messages have not been deliberately or inadvertently changed in transit (message integrity), and that messages are delivered within strict time limits (message vitality). In addition to these requirements, network services must be accessible within a reasonable time (availability). Deliberate or unintentional actions that affect the availability of resources cause a denial of service (DoS). Networks that carry critical application data must have mechanisms in place to make them robust against denial of service activity.

Recognising that it is no longer economically feasible to build dedicated networks for critical applications we must consider the security and availability requirements of a network that will carry both critical and non-critical data. In meeting the security requirements of such a network, careful design decisions will need to be made to ensure that the security mechanisms implemented serve to enhance the overall safety of the system and not hinder the safe operation of the railway.

5.2 Project Goals

A major goal of this project is to provide the foundations for a dependable¹⁷ and secure communications infrastructure for the benefit of all rail stakeholders. This will require the evaluation and securing of train control orders, and location acquisition and reporting mechanisms over a communications network that is carrying this safety critical information, and other business application data.

Achieving these goals necessitates an understanding of the complex interrelationships between safety and security as described throughout this paper, so that both the security and the safety requirements of such a communications system can be adequately met.

¹⁷Reliable, available, maintainable, and safe

6 Acknowledgements

The authors would like to acknowledge the feedback provided by the reviewers in improving the quality of this paper.

References

- Bellare, M. (1998), Practice-oriented provable security, *in* I. Damgård, ed., ‘Lectures on data security: modern cryptology in theory and practise’, number 1561 *in* ‘Lecture Notes in Computer Science’, Springer-Verlag, Berlin Germany, pp. 1–15.
- Bohn, J., Damm, W., Wittke, H., Klose, J. & Moik, A. (2002), Modeling and Validating Train System Applications Using Statemate and Live Sequence Charts, *in* ‘Proceedings of Integrated Design and Process Technology 2002’.
- Bowen, J. P. & Hinchey, M. G. (1994), Seven More Myths of Formal Methods: Dispelling Industrial Prejudices, *in* Naftalin, Denvir & Bertran (1994), pp. 105 – 117.
- Bowen, J. P. & Stavridou, V. (1993), ‘Safety-Critical Systems, Formal Methods, and Standards’, *IEE/BCS Software Engineering Journal* 8(4), 189 – 209.
- Brostoff, S. & Sasse, M. . A. (2001), Safe and Sound: A Safety-Critical Approach to Security, *in* ‘Proceedings of New Security Paradigms Workshop, September 10–13, 2001’, pp. 41 – 50.
- Davida, G. I., Desmedt, Y. G. & Matt, B. J. (1989), Defending Systems Against Viruses through Cryptographic Authentication, *in* ‘Proceedings of the 1989 IEEE Symposium on Security and Privacy, Oakland, California, USA, May 1989’, IEEE Computer Society Press, pp. 312 – 318.
- Dehbonei, D. & Mejia, F. (1994), Formal Methods in the Railway Signalling Industry, *in* Naftalin et al. (1994), pp. 26 – 35.
- Eames, D. P. & Moffett, J. (1999), The Integration of Safety and Security Requirements, *in* M. Felici, K. Kanoun & A. Pasquini, eds, ‘Proceedings of 18th International Conference on Computer Safety, Reliability and Security, SAFE-COMP 1999, Toulouse, France.’, Vol. 1698 of *Lecture Notes in Computer Science*, Springer-Verlag, Berlin Germany, pp. 468 – 481.
- ETSI (2001), Global System for Mobile communication (GSM); Requirements for GSM operation on railways, Standard EN 301 515 V1.0.1 (2001), ETSI.
- ETSI (2003), Terrestrial Trunked Radio (TETRA); Voice plus Data (V+D), Standard EN 300 392-1 V1.2.1 (2003), ETSI.
- Falla, M. (1998), Advances in Safety Critical Systems: Results and Achievements from the DTI/EPSRC Programme in Safety Critical Systems, Technical report, Department of Trade and Industry.
- Fowler, D. & Bennett, P. (2000), IEC 61508 - A Suitable Basis for the Certification of Safety-critical Transport-Infrastructure Systems ??, *in* F. Koornneef & M. van der Meulen, eds, ‘Proceedings of 19th International Conference on Computer Safety, Reliability and Security,

- SAFECOMP 2000, Rotterdam, The Netherlands.', Vol. 1943 of *Lecture Notes in Computer Science*, Springer-Verlag, Berlin Germany, pp. 250 – 263.
- Gerhart, S., Craigen, D. & Ralston, T. (1994), 'Experience with Formal Methods in Critical Systems', *IEEE Software* **11**(1), 21 – 28.
- Gnesi, S., Latella, D., Lenzini, G., Abbaneo, C., Amendola, A. & Marmo, P. (2000), A Formal Specification and Verification of a Safety Critical Railway Control System, in 'Proceedings of the 5th International Workshop on Formal Methods for Industrial Critical Systems, Berlin, April 3–4, 2000'.
- Hansen, K. M. (1994), Validation of a Railway Interlocking Model, in Naftalin et al. (1994), pp. 582 – 601.
- Hartonas-Garmhausen, V., Campos, S., Cimatti, A., Clarke, E. & Giunchiglia, F. (2000), 'Verification of a Safety-Critical Railway Interlocking System with Real-time Constraints', *Science of Computer Programming* **36**(1), 53 – 64.
- Hillenbrand, W. (1999), GSM-R The Railways Integrated Mobile Communication System, Technical report, Siemens.
- Hussey, A. (2000), HAZOP Analysis of Formal Models of Safety-Critical Interactive Systems, in Fowler & Bennett (2000), pp. 250 – 263.
- International Electrotechnical Commission (IEC) (1998), Functional Safety of Electrical / Electronic / Programmable Electronic Safety-Related Systems, Standard IEC 61508, International Electrotechnical Commission (IEC).
- Jones, S., Porter, L. & Waboso, D. (2003), The ERTMS Programme Team Year 1 Progress Report, Technical report, Strategic Rail Authority.
- Kaufman, C., Perlman, R. & Speciner, M. (2002), *Network Security: Private Communication in a Public World*, 2nd edn, Prentice Hall PTR, Upper Saddle River, New Jersey. ISBN 0-13-046019-2.
- King, T. (1994), Formalising British Rail's Signalling Rules, in Naftalin et al. (1994), pp. 45 – 54.
- Knight, J. C. (2002a), Safety Critical Systems: Challenges and Directions, in 'Proceedings of the 24th International Conference on Software Engineering', pp. 547 – 550.
- Knight, J. C. (2002b), Software Challenges in Aviation Systems, in S. Anderson, S. Bologna & M. Felici, eds, 'Computer Safety, Reliability and Security, 21st International Conference, SAFECOMP 2002, Catania, Italy, September 10-13, 2002, Proceedings', Vol. 2434 of *Lecture Notes in Computer Science*, Springer-Verlag, Berlin Germany, pp. 106 – 112.
- Laprie, J. C. (1996), Dependable Computing and Fault Tolerance: Concepts and Terminology, in 'Proceedings of FTCS-25', Vol. 3, pp. 2 – 11. Reprinted from FTCS-15, 1985.
- Leveson, N. G. (1986), 'Software Safety: Why, What, and How', *ACM Computing Surveys* **18**(2), 125 – 163.
- Leveson, N. G. (1995), *Safeware: System Safety and Computers*, Addison-Wesley. ISBN 0-201-11972-2.
- Levy, E. (2003), 'Poisoning the Software Supply Chain', *IEEE Security & Privacy* **1**(3), 70 – 73.
- Meadows, C. (1995), Applying the Dependability Paradigm to Computer Security, in 'Proceedings of New Security Paradigms Workshop, 22 - 25 August, 1995', pp. 75 – 79.
- Meadows, C. & McLean, J. (1998), Security and Dependability: Then and Now, in 'Proceedings of Computer Security, Dependability, and Assurance', pp. 166 – 170.
- Menezes, A. J., van Oorschot, P. C. & Vanstone, S. A. (1997), *Handbook of Applied Cryptography*, CRC Press series on discrete mathematics and its applications, CRC Press. ISBN 0-8493-8523-7.
- Naftalin, M., Denvir, T. & Bertran, M., eds (1994), *FME'94 Industrial Benefit of Formal Methods, Second International Symposium of Formal Methods Europe, Barcelona, Spain*, Vol. 873 of *Lecture Notes in Computer Science*, Springer-Verlag, Berlin Germany.
- Nishinaga, E., Evans, J. A. & Mayhew, G. L. (1994), Wireless Advanced Automatic Train Control, in 'Proceedings of the 1994 ASME/IEEE Joint Railway Conference', pp. 31 – 46.
- Postel, J. (1981a), Internet Protocol, RFC 0791, IETF.
- Postel, J. (1981b), Transmission Control Protocol, RFC 0793, IETF.
- Simpson, A. (1994), A Formal Specification of an Automatic Train Protection System, in Naftalin et al. (1994).
- Standards Australia (SA) (1997), Railway Safety Management, Standard AS 4292, Standards Australia (SA).
- Standards Australia (SA) (1999a), Functional Safety of Electrical / Electronic / Programmable Electronic Safety-Related Systems, Standard AS 61508, Standards Australia (SA).
- Standards Australia (SA) (1999b), Risk Management, Standard AS/NZS 4360:1999, Standards Australia (SA).
- Thompson, K. (1984), 'Reflections on trusting trust', *Communications of the ACM* **27**(8), 761 – 763.
- Thorne-Booth, G. M. (1968), 'Signaling of Remotely Controlled Railway Trains', *IEEE Transactions on Communication Technology* **16**(3), 369 – 374.
- Transport Research Board (1995), Rail Transit Capacity, Report A-08, National Academies. Transit Cooperative Research Program Project 13.
- Winther, R., Johnsen, O. & Gran, B. A. (2001), Security Assessments of Safety Critical Systems Using HAZOPs, in U. Voges, ed., 'Proceedings of the 20th International Conference on Computer Safety, Reliability and Security, SAFECOMP 2001, Budapest, Hungary', Vol. 2187 of *Lecture Notes in Computer Science*, Springer-Verlag, Berlin Germany, pp. 14 – 24.