

The Role of the Evaluator in Australian Defence Standard Def (Aust) 5679

C.B.H. Edwards

AMW Pty Ltd

PO Box 468, Queanbeyan, NSW 2620

Chris.Edwards@amwps.com

Abstract

Issue 2 of Def (Aust) 5679 has evolved as a result of the application of Issue 1 to Defence projects. Experience with the use of Def (Aust) 5679 suggests that the standard provides the necessary guidance required to develop a high assurance safety critical system, but also suggests that the standard requires further strengthening in relation to the conduct of the system hazard analysis.

Keywords: Def (Aust) 5679, Safety Standard, Evaluation, Hazard Analysis.

1 Introduction

The objective of this note is to describe the experience of evaluating safety cases developed using the methodology of Defence Standard Def (Aust) 5679. Lessons learnt from this experience are discussed and further avenues for strengthening the standard are suggested.

The notation and terminology of Issue 2 of Def (Aust) 5679 is used in this paper to describe the safety process. However, reflections on the role of the Evaluator are largely based on experience gained from evaluating safety cases developed under the guidance of Issue 1 of the standard. There is no real conflict in this approach as the role of the Evaluator has not been fundamentally changed, but rather been strengthened in Issue 2.

2 Background

2.1 Evolution of the Def (Aust) 5679

The Australian Defence Standard Def (Aust) 5679 provides requirements and guidance for the procurement of computer-based safety critical systems. The standard was first published in 1998, is being used by the Australian Department of Defence and has achieved a degree of international recognition within the safety community. The standard has now undergone a major review and has appeared in revised form as Issue 2 (Draft Version 1.1: Issued for Comment), March 2007.

In order to support the use of Issue 2 the Defence Materiel Organisation (DMO) sought development of guidance material for users of the standard. DMO contracted Nova Defence to produce Implementation Guidance for Def (Aust) 5679 Issue 2. At the time of

writing the guidance material had not been released by the DMO. Importantly, while Issue 1 was primarily aimed at guiding and supporting Defence projects containing safety critical software intensive systems Issue 2 provides a broader system focus. Paragraph 1.1.1 of Issue 2 notes:

This STANDARD presents requirements and guidance for SAFETY ENGINEERING in the acquisition and sustainment of systems for the Australian Government Department of Defence (DOD).

2.2 Application of Def (Aust) 5679

Def (Aust) 5679 has only been applied to Defence related acquisition projects. By way of contrast, many assessments of system safety in Defence are based on alternative concepts, best characterised in MIL-STD-882. Comparison of the process and outcomes resulting from the use of these two standards throws the problems associated with the use of Def (Aust) 5679 into sharp relief.

3 Evaluation Concepts

3.1 Evaluation Objectives

Paragraphs 3.7.1 and 3.7.2 of Issue 2 of Def (Aust) 5679 notes:

A SYSTEM SAFETY EVALUATION is a formal review, by the EVALUATOR, of the technical validity of the SAFETY CASE.

The primary purpose of SAFETY EVALUATION is to provide an independent assessment of the safety of the System.

Again at paragraph 6.1.1 Issue 2 notes:

The aim of SAFETY EVALUATION is to increase the assurance of safety by providing a thorough independent and objective review of the SAFETY MANAGEMENT process and of the technical validity of the SAFETY CASE.

3.2 Role of the Evaluator

The role of an Evaluator is well described in Issue 2 of Def (Aust) 5679. It is not an easy one to fulfil. In principle the Evaluator acts as an adjudicator on the technical correctness of the Suppliers' safety argument. However, in fulfilling this function the Evaluator needs to understand the complexities of the Safety Case development and also needs to be sensitive to the capabilities of both the Acquirer and the Supplier. As such, issues raised by the Evaluator have the propensity to affect the cost and schedule of the safety program, and

as a result there can often be an uncomfortable, but natural, tension between the Evaluator, the Acquirer and the Supplier.

A more complete description of the role of the Evaluator is provided in Chapter 6 of Issue 2 of the standard. This section discusses appropriate access to documentation, the interaction between the Evaluator and the safety engineering process and the resolution of safety concerns raised by the Evaluator.

4 Application of the Def (Aust) 5679 Safety Process

4.1 Partitioning the Safety Process Discussion

Experience as an Evaluator suggests that a number of issues require particular attention if system safety is to be adequately assured. These issues are best described in terms of the sequence of activities involved in the development of the system safety case. Broadly, the discussion can be partitioned into three phases:

- a. Phase 1 – Safety and the system acquisition strategy;
- b. Phase 2 – The safety process as defined by Def (Aust) 5679; and
- c. Phase 3 - Acceptance of the outcomes of the safety process.

Within each of these three phases a number of qualitatively different and apparently intractable problems arise, varying from administrative process issues to fundamental mathematical problems requiring further research.

4.2 Safety and the System Acquisition Strategy

4.2.1 Funding of System Safety

Defence acquisition projects usually involve high technology and demanding schedules. Additionally, the Defence Materiel Organisation (DMO) increasingly leans toward the use of fixed price contracts. These factors force a situation that requires system safety to be included in the Work Breakdown Structure (WBS) and be adequately funded prior to the commencement of the project. Failure to meet these conditions will almost certainly ensure that development of system assurance will be inadequate.

Based on experience of implementing high assurance safety programs it is suggested that an initial estimate of the cost for development of a system safety case should range from 1% to 3% of total system acquisition cost (unpublished data). Some cost reductions are possible if the system is a Non Developmental Item that has a defensible safety pedigree. Section 3.9 of Issue 2 of Def (Aust) 5679 provides further commentary on the use of an extant Safety Case.

Experience also suggests that the early provision of adequate funding needed to ensure system safety is rarely met. Development of appropriate and validated cost models for Safety Programs would be of assistance in this regard.

4.2.2 System Safety Requirements

A common problem is the absence of meaningful system safety requirements embedded in the statement of requirements for the proposed system. This problem is addressed at section 2.2.2 of Issue 2 of Def (Aust) 5679 which notes:

Safety issues must be addressed early in the development lifecycle, and tracked throughout. This is of crucial importance. It is rarely possible for safety to be introduced as an afterthought. It is almost impossible to demonstrate the safety of a system for which safety requirements were not identified at the outset, and in which the design took no account of safety issues. Paying early attention to safety issues is the best way of ensuring that costly maintenance or re-engineering is not required later in the implementation and installation phases. Also, it has been shown that unsafe system behaviours can usually be traced to deficiencies in requirements specifications. Early consideration of safety issues may produce a design that entirely eliminates some sources of hazard. For these reasons, safety management must be applied stringently from the earliest stages of development, and addressed consistently throughout system development.

The development and integration of safety requirements with other functional requirements demands a conscious effort by system engineers responsible for the development of systems requirements. Wildman (1999) provides an example of the reformulation of a requirements document using a ‘Timed Interval Calculus’. This work illustrates that requirement analysis (including safety requirements) can be progressed from a purely English language expression to a more precise mathematical basis.

Unfortunately, the systems engineering community in Defence appears to rely on the use of tools such as DOORS and Core to provide requirements engineering support. While such tools assist in the collection of a diverse set of requirements they do not lead to a complete and internally consistent requirements set; let alone an adequate integration of appropriate safety requirements with other functional requirements. A further final step is required, i.e. the mathematical modelling of the requirements captured by the preceding informal processes. Such modelling will highlight internal inconsistencies in the informal statement of requirements. In summary, the need to develop a mathematically based requirements engineering process within Defence appears to be well overdue.

4.3 The Def (Aust) 5679 Safety Process

4.3.1 Structure of the Def (Aust) 5679 Safety Process

Issue 2 of DEF(AUST)5679 requires that a safety case be developed using a phased approach. These phases are:

- a. Hazard Analysis;
- b. Safety Architecture Assurance; and
- c. Design Assurance

This section provides a brief overview of the safety process provided by Issue 2 of Def (Aust) 5679 and leads into a discussion of issues observed during the application of the processes described in the standard.

4.3.2 The Hazard Analysis

Identification of system hazards and their potential to lead to accidents is a common first step in system safety standards. In Issue 2 of Def (Aust) 5679 the conduct of a Hazard Analysis (HA) is the first step in the development of a System Safety Case (SSC). The required methodology of the HA is described in Chapter 8 of the standard and is described briefly below.

The aim of the HA is stated in paragraph 8.1.1 to be:
to describe the SYSTEM, its OPERATIONAL CONTEXT and identify all possible ACCIDENT SCENARIOS (and their associated DANGER LEVELS) that may be caused by a combination of the states of the SYSTEM, environmental conditions and external events.

In order to allow the HA to proceed a system description is required. At paragraph 8.4.2 it is stated that:

The SYSTEM DESCRIPTION defines the scope of the SYSTEM, which includes the operators of the SYSTEM. Specifically, it draws the boundary defining what can be changed or designed by a SUPPLIER (including operator procedures). Factors that are beyond the control of the SYSTEM are called external and are determined through reference to the OPERATIONAL CONTEXT. Sufficient description of the SYSTEM environment should be provided to understand the role of the SYSTEM within the overall ACCIDENT SCENARIOS.

The establishment of the system description and system boundary with its associated interfaces can be a difficult task, as the processes required to provide the description and the boundary are interdependent. Also, there is a natural desire on the part of the Acquirer and Supplier to limit the scope of the boundary. This desire flows both from technical and financial and schedule considerations and has the propensity to lead to false conclusions about the level of assurance required for the system to be considered safe and suitable for service. Both these topics are discussed later.

System Hazards are defined at paragraph 8.6.1 to be:
SYSTEM HAZARDS are top-level states or events of the system from which an ACCIDENT, arising from a further chain of events external to the SYSTEM, could conceivably result. SYSTEM HAZARDS are expressed in terms of the SYSTEM interface and are the means by which the SYSTEM can play a role in ACCIDENTS.

Accident scenarios are specified in paragraph 8.7.1 to be:

ACCIDENT SCENARIOS are causally linked states or events, primarily internal SYSTEM HAZARDS and external COEFFECTORS that can result in an accident. ACCIDENT SCENARIOS can be enumerated by any suitably precise means (such as the use of event trees).

The development of accident scenarios resulting from the realisation of system hazards at the system boundary, in combination with external events is the most important activity in the development of the safety case. It is from these scenarios that all other safety activity follows. As such the development and description of the accident scenarios deserves the highest level of effort and discipline.

A simple English language description of the accident scenarios is not generally sufficient as semantic ambiguities have the potential to confuse interactions between Coeffectors and obscure the assignment of a Danger Level for that accident.

Section 4 of Issue 2 of Def (Aust) 5679 lists and outlines the roles of the Safety Management Agents (SMAs). A natural tension exists between the SMAs reflecting the need to provide a high level of system assurance within the constraints of a schedule and budget. Such tensions can be exploited to improve the outcomes from the HA process. Specifically, requiring a consensus agreement on the system description and system boundary would force a test of the completeness of the system description and increase the likelihood that a more complete identification of system hazards was achieved.

Similarly, the choice of methodology used to enumerate accident sequences would be more sharply focused if it were subject to scrutiny by the SMAs and a subsequent consensus on the applicability of the methodology required.

4.3.3 Safety Architecture Assurance

Section 9.1 of Issue 2 of Def (Aust) 5679 notes that:

The aim of the SAFETY ARCHITECTURE phase is to describe and analyse the SAFETY ARCHITECTURE and the proposed implementation of the SYSTEM in sufficient detail to be able to allow CRITICALITY ASSESSMENT.

The process required to complete this phase focuses on the 'components' of the system architecture and provides a criticality assessment of each component. At paragraph 9.11.1 it is noted that:

The aim of CRITICALITY ASSESSMENT is to assign a Safety Assurance Level (SAL) to each COMPONENT SAFETY REQUIREMENT that indicates the level of assurance effort to be expended on DESIGN analysis for each COMPONENT. The SAL is commensurate both with the individual DANGER LEVEL of the SYSTEM SAFETY REQUIREMENT from which it is derived and the SAFETY FEATURES reflected in the ARCHITECTURE.

The Safety Architecture Assurance process aims to ensure that the prerequisites of safety are incorporated into the more general system design process.

4.3.4 Design Assurance

Design Assurance has an ambitious but apparently simple objective. Section 10.1.1 of Issue 2 of Def (Aust) 5679 notes:

The aim of the DESIGN ASSURANCE phase is to analyse COMPONENT DESIGNS to provide sufficient assurance that, if the COMPONENT is implemented as designed, the COMPONENT SAFETY REQUIREMENTS are met.

In order to achieve the stated objective the standard calls for the application of a variety of technologies not commonly used by systems and software engineers. In the past these techniques were a product of academic research and were not readily available to safety practitioners. While many of those technologies, such as static analysis tools, are now available as commercial products, it should be noted that they can be expensive to learn and to apply to specific problems. Wildman et al (2008) provides some examples of the application of these tools and techniques.

4.4 Acceptance of the Outcomes of the Safety Process

4.4.1 The Technical Regulatory Authorities in Defence

There are a number of Technical Regulatory Authorities (TRAs) embedded within the fabric of the Department of Defence. The roles of these authorities vary in description, emphasis and basic function, but all claim 'safety' as part of their *raison d'être* for existence. So for example, the Defence Safety Management Agency will claim seniority in matters of Occupation Health and Safety (OH&S), whereas the Director General Technical Airworthiness claims ownership of air and ground systems safety within the Royal Australian Air Force (RAAF). The TRAs are supported by administrative proclamations issued by various levels within the Defence hierarchy.

There are a number of safety standards endorsed by the TRAs. These standards address issues ranging from the development of safety critical avionics software through the development and deployment of explosive ordnance to the methodology of reporting and investigating accidents within administrative areas.

4.4.2 The Use of Def (Aust) 5679 in Defence

A problem resulting from the diversity of TRAs and associated standards is the difficulty of determining acceptance criteria for a Def (Aust) 5679 Safety Case when arguing that a system is safe and suitable for service. This problem results from a fundamental philosophical difference between the approach taken by Def (Aust) 5679 and the various standards supported by the TRAs, and is characterized here as the 'Hazard Risk Index' (HRI) Problem. This problem is discussed below.

5 The HRI Problem

5.1 Background

Over the last two decades there has been a divergence in the theoretical basis of safety standards. In essence there are two lines of thought. The first approach is a qualitative approach based on a perceived severity of identified system hazards, which in turn dictates the level of rigour required to analyse and mitigate the hazard. The

second approach involves a process which attempts to identify and classify hazards according to some sort of acceptability criteria. Def (Aust) 5679 is an example of the first approach, where the standard asserts that if the severities of hazards are correctly identified and the defined approach is correctly implemented, then the necessary system assurance will be obtained. This is in contrast to other safety standards (e.g. the UK DEF STD 00-56 and the US MIL-STD-882) which are based on the At Least as Reasonably Practical (ALARP) approach to safety. The latter accepts the possibility of Residual Risk and attempts to quantify assurance through the use of the Hazard Risk Index (HRI). Note: The US MIL-STD-882 does not specifically call out ALARP but is instead based on a reasonability test similar to the ALARP approach.

The concept of a Hazard Risk Index (HRI) is designed to identify the risk of hazards and to guide the application of resources to minimize perceived risk. This concept is often applied to a wide range of situations, ranging from relatively simple OH&S problems, such as office safety, to the acquisition of complex weapons systems. After the derivation of the HRI for a particular hazard, an assessment of the application of resources required to mitigate or remove the risk is made. Often this assessment is based on the As Low as Reasonably Practicable (ALARP) principle. Notably, the ALARP method allows for a statement of Residual Risk, i.e. the risk remaining after the completion of the safety process.

It is important to note that while standards such as Def (Aust) 5679 eschew the concept of applying probabilistic and other frequency based concepts to the identification and mitigation of hazards, they still acknowledge that in certain circumstances such data can be usefully applied to a hazard analysis. By comparison, alternative standards such as MIL-STD-882 demand the development of arguments based on a two dimensional relationship between likelihood and consequence. Implicit in the concepts of likelihood and consequence is the availability of statistical data to support the arguments.

Interestingly, Issue 1 of Def (Aust) 5679 used probabilistic concepts to derive LOT and SIL values, while Section 8.8 of Issue 2 of Def (Aust) 5679 suggests a qualitative approach to the estimation of the likelihood of external mitigations, but suggests ranges of probability values that can be used to classify the strength of the mitigation. This approach conveniently skirts the issue of how such probabilities can be derived and the practical reality that the probabilities will have some sort of density function.

5.2 Derivation of the HRI

The derivation of a HRI for a particular hazard, i.e. a hazard derived from a Hazard Analysis process, is typically based on a tabulation of 'Likelihood' versus 'Consequence' as shown in Table 1. The acceptability of the HRI is then determined by a grouping of derived HRI. An example is shown in Table 2.

Likelihood	Consequence			
	Catastrophic	Critical	Major	Minor
Frequent	1	3	7	13
Probable	2	5	9	16
Occasional	4	6	11	18
Remote	8	10	14	19
Improbable	12	15	17	20

Table 1. Hazard Risk Index

HRI	Risk Level	Risk Acceptability
1 to 5	Extreme	Intolerable
6 to 9	High	Tolerable with continuous review
10 to 17	Medium	Tolerable with periodic review
18 to 20	Low	Acceptable with periodic review

Table 2. Acceptability of Risk

5.3 Application of the HRI

The Likelihood and Consequence scales on Table 1 are ordinal measures. Thus within a particular row or column of Table 1 entries are ranked and comparison of those rankings is valid. For example, it is reasonable to assert that within the Critical Consequence column an Occasional likelihood is a worse outcome than a Remote likelihood. Comparisons of rankings from different rows or columns are more problematic. To assert that Occasional but Catastrophic (HRI=4) is equivalent to Frequent and Critical (HRI=3) is difficult to justify in the absence of a quantified hazard context. Thus a grouping of HRI, for example as shown in Table 2, is difficult to justify. Groupings may have some meaning if the context of the Likelihood and Consequence assessments is known and done on a case-by-case basis, i.e. the system or issue under evaluation together with the operational environment are well understood. Importantly, because a general *a priori* statement of HRI groupings has no theoretical basis, groupings of HRI used to ascribe a level of risk acceptability must be done on a case by case basis prior to the safety analysis, in a manner that takes into account the system context. Stevens (1946) provides a useful discussion on the theory of scales of measurement.

A safety case developed under Def (Aust) 5679 will almost certainly provide enough context information to allow an informal grouping of HRI to be made. Thus the translation from the Def(Aust)5679 approach to a risk based approach should, in principle, be relatively easy. The point here is that while the process of moving from a severity based approach to a risk based approach is possible, the reverse is likely to be much more difficult.

5.4 HRI Limitations

The problem with a general application of the HRI as shown in Tables 1 and 2 lies in the fact that it is not always possible to apply context scaling (i.e. scaling appropriate to the system and its environment) to the Likelihood and Consequence classifiers. In particular, for complex systems, contexts can be difficult to define and as such the concept of acceptability of risks has no meaning. Any subsequent ALARP analysis is then based on a qualitative assessment of risk which may in turn lead to a sub optimal allocation of resources aimed at risk mitigation. Prasad and McDermid (1999) discuss the importance of the context of a system when attempting to identify emergent properties such as dependability.

The HRI based approach to safety has intuitive appeal to program managers because the ALARP concept appears to simplify the problem of resource allocation. This follows from the fact that the residual risk is qualitative, and once articulated provides a well defined end to a safety program.

5.5 The Way Ahead

Given the pervasive use of HRI there is a need to further develop the theoretical basis for the application of HRI in assessing system safety. In particular there is a need to better understand the relationship between this safety methodology, i.e. risk based, and those based on the concept of accident severity. The provision a method of mapping from the hazard severity based approach of Def (Aust) 5679 to a HRI description of the outcome of the safety process would assist in the acceptance of outcomes of the safety analysis.

6 Human Factors

6.1 Integration of Human Factors into the Safety Process

The inclusion of the consideration of Human Factors early in the Safety Process is important. Many accidents have resulted from the realization of hazards existing at the man-machine interface. Issue 2 of Def (Aust) 5679 recognises the importance of this issue, noting at section 10.13.3 and 10.13.4 that:

The SUPPLIER shall, in consultation with END USERS, ensure that human factors and task analysis are applied to design of OPERATOR ROLES, and

The DESIGN ASSURANCE REPORT must include the results of human factors and task analysis carried out.

Experience suggests that both the Supplier and the Acquirer are loath to include this consideration in the safety process because of time, cost and the almost inevitable controversial results of such analysis. The usual justification for the exclusion of such considerations is based on the argument that the system being acquired is a COTS product and that the presence of trained operators improves aggregate system safety and hence there is no need to waste time on this subject.

The involvement of the End User in considering human factor issues can be inhibited by the way the SMG is managed. The SMG is usually chaired by the Acquirer, who in turn is motivated to minimise discussion in the interest of cost and schedule. Experience suggests that the SMG must, to the maximum extent possible, include critical representation from the End User community.

Unfortunately, experience also suggests that the End User will often accept assurances from the supplier that the issue of human factors has been appropriately dealt with. This acceptance is reinforced by the common belief that the presence of trained operators will provide a higher level of assurance that the system is safe to use.

There is a large body of literature on human factors and it is suggested that the standard should provide further guidance on the use of this literature, specifically relating to techniques for the analysis of human factors in complex systems. For example, Sandom (2007) provides a guide to a different approach to managing operator error.

7 System Boundary and System Description

7.1 Abstraction of System Description

The description of a complex computer based system necessarily involves a process of abstraction in order to provide a description suitable for a specific audience. Issue 2 of Def (Aust) 5679 defines the audience and assumes that the supplier has the ability to provide a system description suitable for that audience, i.e. in a suitable format and at a suitable level of abstraction. This is a critical assumption and is a key to the successful development of the Safety Case.

Achieving an appropriate level of abstraction when describing a system can be a deceptively difficult task. This is particularly the case for computer based systems which are invariably described at different levels of abstraction during their life cycle. At the highest level of abstraction the system might be viewed as a single entity that interacts with an environment, generating and accepting data with no visibility of the internal workings of the system. At the lowest level all the computer code and hardware would be visible with data flows between software and hardware modules e.g. Field Programmable Gate Arrays (FPGAs) and with the external environment. It is probable that the latter description would be very large, complex and difficult to analyse.

At paragraph 8.4.2 Issue 2 of Def (Aust) 5679 notes:

The SYSTEM DESCRIPTION defines the scope of the SYSTEM, which includes the operators of the SYSTEM. Specifically, it draws the boundary defining what can be changed or designed by a SUPPLIER (including operator procedures). Factors that are beyond the control of the SYSTEM are called external and are determined through reference to the OPERATIONAL CONTEXT. Sufficient description of the SYSTEM environment should be provided to understand the role of the SYSTEM within the overall ACCIDENT SCENARIOS

In DRAFT MIL-STD-882E July 04 the system description requirements in the HA Task are specified as:

This will consist of summary descriptions of the physical and functional characteristics of the system and its components. Reference to more detailed system and component descriptions, including specifications and detailed review documentation shall be supplied when such documentation is available. The capabilities, limitations and interdependence of these components shall be expressed in terms relevant to safety. The system and components shall be addressed in relation to its mission and its operational environment. System block diagrams or functional flow diagrams may be used to clarify system descriptions. Software and its role(s) shall be included in this description.

Successful identification of system hazards will depend on the visibility of system functions and hence on the level of abstraction used to describe the system. Achieving an appropriate description of the system is thus critical to the success of the HA process and the standard needs to provide guidance on how this can be achieved.

7.2 Identification of System Boundary

The level of abstraction used in describing the system under investigation has the propensity to influence the definition of the system boundary. Thus while the simplicity of a high level of abstraction tempts the inclusion of parts of the environment within the boundary, a low level of abstraction will have the opposite effect. It follows that there is potential for correlation between the level of abstraction and the definition of the system boundary to develop. The extent of this correlation will vary from system to system and will be influenced by the experience and technical competence of the developer and personnel involved in the development of the Safety Case.

In general the level of rigour required to resolve hazards at the system boundary should be consistent with the degree of criticality of the system, irrespective of the level of abstraction of the system. This assumes that the boundary encompasses approximately the same number of system components and interfaces regardless of the level of abstraction.

The degree of coupling or integration with the external environment can affect the definition of the system boundary. In the case of a system tightly coupled to external sensors which provide data for safe operation there is a reasonable motivation to include those sensors within the system boundary. Yet such an inclusion might well increase the scope and complexity (and hence cost) of the HA. To argue that the system boundary should not include the sensors on the basis of cost or schedule might be to ignore a critical dependency between the sensors and the system presented for safety assessment.

Issue 2 of Def (Aust) 5679 notes that external factors (such as sensors) exist as Coeffectors and requires that they be included in the discovery of accident sequences. However, because they are often beyond the control of the Acquirer and Supplier there is strong motivation by those parties to minimize their consideration. In turn this

constrains the scope of the HA and has the potential to allow the various Safety Management Agents (SMAs) to selectively ignore broader system integration issues having safety implications. For example, a system might require wind vector data of a higher resolution than that provided by a currently installed wind sensor in order to operate safely. If the wind sensor is considered to be external to the system it is possible that the need to modify the wind sensor could be conveniently ignored (e.g. for cost and/or schedule reasons) during the development phase and the problem would only be discovered after the system has been accepted into service.

In DRAFT MIL-STD-882E, Jul 04 the system interface requirements in the HA Task is specified as:

Safety related interface considerations among various elements of the system (e.g., material compatibilities, electromagnetic interference, inadvertent activation, fire/explosive initiation and propagation, and hardware and software controls). This shall include consideration of the potential contribution by software (including software developed by other contractors/sources) to subsystem/system mishaps. Safety design criteria to control safety-critical software commands and responses (e.g., inadvertent command, failure to command, untimely command or responses, inappropriate magnitude, or MA-designated undesired events) shall be identified and appropriate action taken to incorporate them in the software (and related hardware) specifications.

Section 3.8.2 of DEF(AUST)5679 discusses the issue of large systems composed of a federation of subsystems and (inter alia) notes:

The SAFETY ENGINEERING process described in this STANDARD does not easily admit decomposition of the SAFETY CASE across subsystems. The strongest barrier to SAFETY CASE compositionality lies in the CRITICALITY ASSESSMENT of the SAFETY ARCHITECTURE. In particular, it is not always possible to claim as much assurance credit for safety features implemented external to the SYSTEM as for those implemented internal to the SYSTEM. Thus, the STANDARD may require more rigorous analysis activities when SAFETY ENGINEERING a subsystem separately, than it would require when treating the subsystem as part of a larger system.

And at Section 3.8.3 it notes:

This non-compositionality is unfortunate, but well justified on both philosophical and empirical grounds. Philosophically, it is reasonable to argue that the SUPPLIER should not be allowed to claim as much credit for, nor place as much trust in, safety functionality that is not under the SUPPLIER'S design control. Empirically, it seems well established that some of the most intractable engineering challenges occur at the interface between subsystems, so that is reasonable to require greater assurance rigour when integrating subsystems.

Complex systems are almost always composed of subsystems. The problem that faces the safety practitioner is that as the level of abstraction is reduced the number of interfaces between subsystems grows rapidly. This suggests that the lower the level of abstraction of the system, the higher the degree of effort required when analysing system hazards at the system boundary.

Sometimes a low level of abstraction is forced on the Supplier. For example, upgrading the function of an FPGA on a single circuit board in a missile system forces a focus on the FPGA and the associated functions of the circuit card. At this point the Acquirer and Supplier are tempted to constrain the scope of the system boundary to the single circuit board with the aim of containing the safety effort. By defining the system as a circuit board containing the FPGA, within the context of a missile system rather than the missile system within the context of its deployed environment it is possible that any change in the behaviour of the missile could be missed. The result would be that the HA would fail to fully enumerate the accident sequences required by the standard. Additionally, the focus on the circuit card could well demand such a high level of analysis effort that it would not be possible to provide the required level of system assurance.

It appears that there must be a 'sweet spot' where the level of abstraction in combination with the scope of the system boundary and its safety criticality offers the most cost effective basis for further safety analysis. This conclusion is not intuitive to program managers, who almost always see constraining the system boundary and raising the level of abstraction as a way to reduce the expenditure on safety assurance.

The establishment of an appropriate system boundary thus requires careful consideration, should include consideration of system integration issues, and should result from a consensus developed between the various SMAs. In developing such a consensus care needs to be taken to ensure that technical and not schedule and cost issues drive the process. The standard needs to provide guidance on a process for developing such a consensus.

7.3 Enumeration of Accident Sequences

At Section 8.7 of Issue 2 of Def (Aust) 5679 a conceptual process for the identification of accident sequences is described. The process takes into account the existence of external Coeffectors and provides some guidance on how they are to be integrated into the overall safety process.

A methodology for enumerating accident sequences is not spelt out in the standard. Anecdotal evidence suggest that the realization of multiple system hazards can be involved in a sequence of events leading to an accident and that more complex models than linear event trees are required to discover possible accident sequences. The assumption of the mutual independence of System Hazards is thus intuitively difficult to accept and suggests that the assumption is a weak one. However, mutual independence offers a manageable approach to enumerating accident sequences.

The enumeration of accident sequences must address the likelihood that more than one System Hazard may be

simultaneously realized. This requirement is addressed in Issue 2 of Def (Aust) 5679 but the approach is essentially qualitative in nature. While a variety of quantitative techniques have been discussed in the literature (e.g. Markov Chain models) the theoretical basis and associated assumptions of the chosen technique needs to be clearly understood. It is proposed that while the choice of technique is left to the developer, the developer should be required to demonstrate an appropriate understanding of the chosen technique to the Safety Management Group (SMG).

Experience suggests that guidance on techniques for enumerating accident sequences is urgently required. The applicability of specific techniques will vary depending on the system and its operational context.

8 Proposed Changes to Def (Aust) 5679

8.1 Hazard Analysis Process

The Hazard Analysis is the first of a number of common threads in most safety standards. The results of the following activities all depend on the quality of this first step. This dependence suggests that, for all the differences in approach to assessing system safety, this activity must be completed with a high level of rigour.

A consideration of the level of abstraction used to describe the system being analysed, together with the degree of the systems coupling with the external environment, needs to be part of the HA process. This suggests that the HA process outlined in Issue 2 of Def (Aust) 5679 requires strengthening. Changes proposed include:

- a. A more holistic consideration of the of the system description together with the associated system boundary and interfaces prior to further hazard analysis effort; and
- b. An improved definition of the process required to discover accident scenarios.

8.1.1 Development of the System Description and Function Definition

In order to facilitate the development of an appropriate system description and associated system boundary and interfaces it is proposed that Def (Aust) 5679 include an extra first step in the HA process. This step would require that:

- a. A consensus is developed between all the SMAs on the definitions of the system description (or model) to be used in the Safety Case, together with that of the system boundary and interfaces. This to be achieved prior to any further safety analysis being conducted. In achieving such a consensus the degree of system abstraction and identification of system integration issues (i.e. external factors) must be considered and the conclusions formally recorded and signed off by the SMAs.
- b. Further modification of the agreed system definition could only be made with the agreement of all the SMAs and such a

change would normally involve a reduction in the level of abstraction used in the model.

- c. In the event that a consensus on the system description and system boundary cannot be achieved between the SMAs, the development of the Safety Case must be halted and the problem referred to a higher, external management level for independent technical review.

8.1.2 Enumeration of Accident Scenarios

The enumeration of accident scenarios resulting from the realisation of system hazards at the system boundary, in combination with external coeffectors is the single most important activity in the development of a safety case. As such sections 8.7 and 8.8 in Issue 2 of Def (Aust) 5679 require strengthening. In particular, the process used to identify accident scenarios requires the provision of guidance on the use of appropriate techniques. Such guidance could take the form of a separate publication which, while referenced in the standard, could be readily updated.

In addition to the above changes it is also proposed that the Supplier be required to demonstrate an understanding of the methodology used to enumerate accident sequences. Such a demonstration would take the form of a documented presentation to the full SMG. Subsequent SMG endorsement of the presentation would be required prior to the developer implementing the chosen method.

9 Conclusions

9.1 Human Factors

As noted previously there is a large body of literature on human factors and it is suggested that Def (Aust) 5679 should provide guidance on the use of this literature, specifically relating to techniques for the analysis of human factors in complex systems.

9.2 Hazard Analysis

Experience with the use of Def (Aust) 5679 suggests that use of the standard provides the necessary guidance required to develop a high assurance safety critical system. However, the HA process outlined in the standard makes implicit assumptions about the system that requires the section to be strengthened. In particular, the process used to define the system boundary and its associated interfaces, together with the process of listing accident scenarios requires further development.

9.3 Further Research

In order to provide interoperability between safety standards it is suggested that research into the relationship between severity and risk based assessments of system safety be supported by Defence

10 References

Luke Wildman (2002) Requirements Reformulation using Formal Specifications: A Case Study, Software Verification Research Centre, University of Queensland.

S.S. Stevens (1946) On the Theory of Scales of Measurement, Science, Vol. 103, NO. 2684, June 7, 1946.

Luke Wildman et al, (2008) Guidance for Def(Aust) 5679 Issue 2, 13th Australian Conference on Safety Related Programmable Systems, Australian Computer Society, System Safety and Quality Engineering Pty Ltd.

Divya Prasad, John McDermid (1999) "Dependability Evaluation using a Multi-Criteria Decision Analysis Procedure," dcca, p. 339, Dependable Computing for Critical Applications (DCCA '99).

Carl Sandom (2007), Success and Failure: Human as Hero – Human as Hazard, 12th Australian Conference on Safety Related Programmable, Systems, Australian Computer Society.