

Urgent Operational Requirements: Impact on the Safety Case

Tony Cant and Brendan Mahony

Command, Control, Communications and Intelligence Division
 Defence Science and Technology Organisation
 PO Box 1500, Edinburgh, South Australia 5111
 Email: Tony.Cant@dsto.defence.gov.au, Brendan.Mahony@dsto.defence.gov.au

Abstract

Modern Defence systems are complex and software-intensive. In response to the technical challenges posed by such systems Defence has developed a capability lifecycle with suitably rigorous quality control measures. Unfortunately, in today's rapidly evolving Defence environment, unforeseen threats can lead to capability gaps that require rapidly fielded solutions. Such *Urgent Operational Requirements* (UOR) can accelerate (and perhaps curtail) the normal capability lifecycle.

Defence systems are often safety-critical: they have the potential to cause death or injury as a result of accidents arising from unintended system behaviour. For such systems an effective safety engineering process (along with choice of the appropriate safety standards) must be established at an early stage of the capability lifecycle, and reflected in contract documents. This process culminates in a safety case, which is a structured argument, supported by a body of evidence, that provides a compelling, comprehensible valid case that a system is safe for a given application in a given environment.

In this paper we discuss the impact of Urgent Operational Requirements and the above lifecycle issues on the Safety Case. We use the processes and terminology of the recently published standard DEF(AUST)5679 Issue 2. In discussing the impact of UORs on the safety case, we find it useful to distinguish three cases: *Greenfield Acquisition*, *In-Service Modification* and *Modified Operational Context*.

Keywords: Safety case, safety assurance, rapid acquisition, urgent operational requirements.

1 Introduction

Modern Defence systems (such as combat systems, avionics systems, command support systems, precision weapons systems etc) are complex and software-intensive systems. In response to the technical challenges posed by such systems Defence has developed a capability lifecycle with suitably rigorous quality control measures.

1.1 The Capability Lifecycle

In the Australian context the *capability lifecycle* for such systems is divided into the following phases:

Copyright © 2011, Commonwealth of Australia. This paper appeared at the Australian System Safety Conference (ASSC 2011), Melbourne, Australia. Conferences in Research and Practice in Information Technology (CRPIT), Vol. 133. Tony Cant, Ed. Reproduction for academic, not-for profit purposes permitted provided this text is included.

1. *The Needs Phase* — involves the articulation of capability goals in the context of the current and planned force.
2. *The Requirements Phase* — involves the detailed planning required for converting capability needs into an integrated set of changes in each of the Fundamental Inputs to Capability (FIC). The Requirements Phase also incorporates a *Decision Making Process*, consisting of:
 - *First Pass Approval* — at which Government considers alternatives and approves capability development options; and
 - *Second Pass Approval* — at which Government agrees to fund the acquisition and through-life costs of a specific capability.
3. *The Acquisition Phase* — the process by which Defence acquires a specific capability via requests for tender, risk reduction activities and management of system procurement via an Australian Defence contract. At the end of the Acquisition Phase, an *Acceptance into Service* decision is made in light of an assessment made by the relevant service's technical regulator.
4. *The In-Service Phase* — the normal operating life of the system in service.
5. *The Disposal Phase* — controlled and managed decommissioning of the system.

1.2 Urgent Operational Requirements

The Capability Lifecycle is a measured and well-instrumented process, designed to make well-informed decisions about the acquisition and deployment of capabilities with in-service lifetimes of up to a quarter of a century.

Unfortunately, in today's rapidly evolving Defence environment, unforeseen threats can lead to capability gaps that require rapidly fielded solutions. This is often referred to as an *Urgent Operational Requirement* (UOR). In some countries, the term Rapid Acquisition is used instead of UOR. We essentially use them as synonyms in this paper.

The general tendency of UORs is to accelerate (and perhaps abbreviate) the normal capability lifecycle. The Capability Lifecycle is intended to mitigate the risks inherent to capability development. Capability development risk include the following classes: engineering risks (project failure, capability gap); economic risks (budget overrun); security risks (release of classified information) and safety risk (death or injury of personnel or the public). Accelerating the process necessarily reduces the level of risk mitigation, but this is balanced against the mission risk that gives rise to the UOR. The evaluation of these competing

risks is fundamentally different for the various risk classes.

The primary motivation for mitigation of engineering risk lies in the potential for leaving a pressing capability gap, but this is precisely the problem which leads to the UOR, so that it is highly likely that a timely effort is better than a low engineering risk effort.

A high level of mitigation of economic risk is inherent to the time pressures posed by the UOR. By definition, a rapid acquisition will be of relatively fixed duration and cost. It is at low risk of overruns, provided that the rate of spend is well contained. Thus, it is fairly straightforward to make a rational decision in balancing economic and mission risks.

The primary motivation for mitigation of security risk lies in the potential for breaches to lead to mission failure. Again, the very existence of the UOR means that there is already a high risk of mission failure, so that it is highly likely the a timely effort is better than a low security risk effort.

The primary motivation for the mitigation of safety risk is to protect defence personnel and the general public from death or injury. In order to rationally balance safety risk against the mission risk motivating the UOR, it is necessary to properly identify the level of safety risk posed by the system. In contrast to the other risk classes discussed, there is no natural tendency for the UOR to limit the level of safety risk. Making rational decisions about the safety risk associated with a system requires the existence of an appropriately rigorous safety case.

1.3 Safety Cases

In Australia, the Occupational Health and Safety (OH&S) Act¹ requires that parties involved in the acquisition and sustainment of systems for Defence have a duty of care arising from their legal obligation to take “reasonably practicable steps to avert harm to members of the public, as well as their own employees.” A breach of this duty could make them liable in the case of an accident.

Defence systems often have the potential to cause death or injury as a result of accidents arising from unintended system behaviour. For such systems an effective safety engineering process (along with choice of the appropriate safety standards) must be established at an early stage of the acquisition lifecycle, and reflected in contract documents. This process culminates in a *safety case* that is presented to safety evaluators and certifiers for assessment. A safety case has been defined to be (Ministry of Defence 2007):

... a structured argument, supported by a body of evidence, that provides a compelling, comprehensible valid case that a system is safe for a given application in a given environment.

The safety case is the natural vehicle for the assessment and communication of the safety risk that is potentially introduced by use of the system — not least in the case of UORs. In fact, the accelerated nature of Rapid Acquisition requires a corresponding increase in the rate of safety effort to ensure a timely assessment of safety risk. In practice, there are known to be cases in which a UOR system has not been accepted into service due to high levels of safety risk.

In discussing the impact of UORs on the provision of safety cases, we find it useful to distinguish three system acquisition classes.

- *Greenfield Acquisition*: a new capability is acquired from scratch.
- *In-Service Modification*: a system is modified during its operational life.
- *Modified Operational Context*: a system is used in situations for which it was not originally intended.

Each of these classes occur quite naturally in capability development, but each provides different insights into the challenges posed by UORs.

1.4 Outline

In this paper we are interested in the implications that UORs can have on the safety case. First of all, in Section 2 we provide general background on the issue of Urgent Operational Requirements. In Section 3 we discuss the Nimrod Review. In Section 4 we discuss the structure of the safety case using the terminology of the recently released standard DEF(AUST)5679 Issue 2 (Department of Defence 2008c). Section 5 summarises the issues involved in the procurement of Non-Development Systems. In Section 6, we discuss the impact of UOR on the safety case; while in Section 7, we consider the three class of System Acquisition so as to identify situations that are highly favourable to Rapid Acquisition. Finally, Section 8 presents some conclusions.

2 Urgent Operational Requirements

A key driver for Defence organisations in, for example, the US, UK and Australia is the need to support peacekeeping or military operations across a range of environments. Such operations (for example the USA’s Operation Iraqi Freedom or the UK’s Operation HERRICK in Afghanistan) present huge challenges owing to the nature of the terrain, the political landscape and the threat posed by asymmetric warfare. Current Australian Defence operations are: CATALYST (Iraq); SLIPPER (focused on Afghanistan); ASTUTE (East Timor) and ANODE (Solomon Islands). These are smaller in scale than the corresponding UK or US operations, but present a similar range of challenges.

UOR is a complex area: in the following we highlight some aspects of UOR in the UK, USA and Australia that are especially relevant for our later discussions on safety.

2.1 UK

In the UK special Treasury funding is used to support UORs, for example the Ridgback and Mastiff Protected Patrol Vehicles used in Iraq and Afghanistan. The definition of UOR used in the UK is as follows (Ministry of Defence 2011):

UORs arise from the identification of previously un-provisioned and emerging capability gaps as a result of current or imminent operations or where deliveries under existing contracts for equipment or services require accelerating due to an increased urgency to bring the capability they provided into service. These capability shortfalls are addressed by the urgent procurement of either new or additional equipment, enhancing existing capability, within a timescale that cannot be met by the normal acquisition cycle.

¹Occupational Health and Safety (Commonwealth Employment) Act, 1991

In a recent speech entitled “Performance under pressure: the reality of acquisition in the world’s most complex environment”, Andrew Tyler (Tyler 2009) (the UK MOD’s Defence Equipment and Support (DE&S) Chief Operating Officer) points out that the MOD is “an organisation that is on a war footing.” DE&S has around 850 staff involved in Urgent Operational Requirements (UOR); over recent times they have responded to about 1600 urgent requirements, resulting in 700 items of new equipment being delivered into theatres of operations (often in less than six months). Tyler also comments that: “as much leading-edge technology is being brought to bear on the incredibly complex problem of counter Improvised Explosive Devices (IEDs) as is going into low observability on the Joint Strike Fighter”.

Tyler draws the distinction between UOR processes and the “normal” acquisition process:

Applying UOR processes to the purchase of a nuclear submarine is an absolute non-starter. UORs are about meeting an immediate military need, using rapidly modified off-the-shelf equipment where possible, which may be discarded quickly when the immediate requirement is removed. No enduring support solution is required and integration with wider systems is often minimised for expediency. Furthermore the degree of scrutiny of public spending is balanced against the rapid delivery times required to support crucial operations. None of this applies to nuclear submarines and fighter aircraft which take many years to design and build, usually succeed complex equipment already in service and are designed to meet the long-term military capabilities required in future decades.

Highly skilled, versatile and diverse teams tend to be involved in the problem-solving that is necessary to meet UORs.

2.2 USA

In the USA, the Report of the Defense Science Board Task Force on the Fulfilment of Urgent Operational Needs recommends that Rapid Acquisitions be acknowledged as processes that are formally different from (and incompatible with) deliberate (i.e. normal) acquisition processes. It also recommends that a separate funding stream and organisation be established to handle Rapid Acquisitions.

The Office of the Director, Defense Research and Engineering (DDR&E) has commissioned a study of tools suitable for Rapid Acquisition. This study has highlighted the need to focus on the “front-end” of the capability lifecycle by creating a strategic effort in “accelerated concept engineering” (from anticipated or emerging need to initial design). There is heavy emphasis on exploiting gaming technologies for need and concept exploration; explicit accounting of potential threat evolution and vulnerabilities (“red teaming”); modelling and simulation tools to support concept engineering; and agile and adaptive systems engineering.

During the study, it was observed that most Rapid Acquisitions are not new: they start with some existing capability, and their objective is to build on, adapt, or integrate.

2.3 Australia

In Australia, the recently published Defence Instruction (General) DI(G) LOG 4-1-008 (Department of

Defence 2008b) recognises the challenges posed by asymmetric warfare and provides an overall policy framework for Rapid Acquisition of Capability. It makes the Prime Minister the approving authority for Rapid Acquisitions, and includes the following policy statements relating to safety aspects of rapidly procured equipment:

- Procurement via Rapid Acquisition must not be used to circumvent or over-ride extant Government or Departmental policy.
- Capabilities acquired through Rapid Acquisition shall be certified as fit for service, safe and, where appropriate, comply with regulations for the protection of the environment.

However, the document also allows for the acceptance of risks at higher levels of authority:

10. Capability Managers must identify any risks associated with equipment procured under the Rapid Acquisition process. Risks identified under this process must only be waived at the correct level. Only the Government, CDF and Service Chiefs have the authority to accept the risks associated with the use of items acquired under Rapid Acquisition where, due to time critical requirements, normal due process cannot be followed.

It also allows (in Annex E) for some dilution in the degree of technical regulation:

2. Regulation. In developing the Rapid Acquisition proposal, Capability Managers are to refer to Defence Instruction (General) LOG 0815 — Regulation of Technical Integrity of Australian Defence Force Materiel. TRAs are to be mindful of the timeframes by which Rapid Acquisition capabilities may need to be deployed, which will necessitate risk assessments and judgements to be made concerning the degree to which regulation of the materiel is to be applied. Risks in the areas of safety, performance and environmental compliance are to be documented, reported and managed as part of the Rapid Acquisition process.

The recently published Strategic Reform Program (Delivering Force 2030) (Department of Defence 2009) outlines a program of savings within Defence that will deliver gross savings of \$20 billion. This money is to be reinvested in key areas of Defence to deliver stronger military capabilities; to remediate poorly funded areas; and to modernise the Defence enterprise backbone. The following reference is made to safety:

This program is not about compromising capability to save costs; it is about delivering improved levels of capability at less cost by improving productivity and eliminating waste. While efficiencies can be found in support areas, quality and safety will not be compromised.

In the Technical Regulatory framework for the Australian Army, policy has been developed to address issues arising from Rapid Acquisition (RA). The RA process considers: (1) risks to fitness for service (i.e. mission risk); (2) safety risks to the personnel or public; and (3) environmental risks.

In a normal acquisition, these three aspects will be articulated in a User Requirement and subsequent

Functional & Performance Specification (FPS). Then tendered options are assessed against the Statement of Work (a document that includes the FPS), and a preferred solution is selected. The aim is for the matériel to go into service with a residual risk baseline that is LOW, or at a level of risk that is assessed to be As Low As Reasonably Practicable (ALARP).

The policy recognises that, in a Rapid Acquisition, there is sometimes insufficient time to do this work, and that equipment has the potential to enter service with a significant level of residual risk. The aim of a Rapid Acquisition is to minimise risk as low as reasonably practicable in the time frame available, that is, as much of the above process should be followed as time permits.

The Defence Materiel Organisation (DMO), in carrying out a Rapid Acquisition, often cannot assure risk free operation to the user. Rather, the emphasis is on the DMO to understand the technical risks and inform the user accordingly so that they are able to make informed decisions regarding the equipment's use and operational impact.

3 The Nimrod Review

The recently released Nimrod Review (Haddon-Cave 2009) is an independent review by Charles Haddon-Cave QC into the broader issues surrounding the loss of the RAF Nimrod MR2 Aircraft XV230 in Afghanistan in 2006. It is an example of the issues that can arise with UORs.

3.1 Background

The Falklands War in April 1982 gave rise to an Urgent Operational Requirement (UOR) to equip the Nimrod MR2 with an Air-to-Air Refuelling (AAR) capability, thereby extending the Nimrod's endurance to 20 hours in the air so that they could better support British operations during the war (a new operational context). An in-service modification was made to the aircraft to provide the required AAR capability.

The initial UOR design was modified in 1989 to meet the requirements of Def-Stan 00-970 (Ministry of Defence 1983). The AAR modification changed the function of refuel pipes within No. 7 Tank Dry Bay (previously they had not been used in flight). The review states that: "In making these pipes 'live', the AAR modification introduced a significant new element to the risk of fire because of their close proximity to the hot Cross-Feed/SCP duct".

The review concludes that the accident most likely resulted from ignition (via the Cross-Feed/SCP duct) of fuel in the No 7 Tank Dry Bay that had accumulated as a result of AAR. The review further states that design flaws introduced over the life of the aircraft played a crucial part in the loss of the aircraft.

The review also claims that organisational factors also played a major role in the loss of XV230, and is critical of the Military Airworthiness System. Following the 1998 Strategic Defence Review, financial pressures and the shift in culture towards business and financial targets led to a "dilution of the airworthiness regime and culture within the MOD, and distraction from safety and airworthiness issues as the top priority".

The loss of the XV230 aircraft is illustrative of the consequences of extending the lifecycle beyond its intended end-point. The Nimrod Review points to "an inadequate appreciation of the needs of Aged Aircraft" and goes on to state: "But for the delays in the Nimrod MRA4 replacement programme, XV230 would probably have no longer have been flying in

September 2006, because it would have reached its Out-of-Service Date and already been scrapped or stripped for conversion."

3.2 Safety Case Criticisms

The Nimrod Review is especially critical of the inadequacy of the Nimrod Safety Case. For example, the safety case had a number of open or not properly assessed hazards, including the catastrophic fire hazard relating to the Cross-Feed/SCP duct that was the ignition source in the accident.

The Nimrod Review is likely to have a significant impact on the UK MOD procurement policy for safety-critical systems. We will not reflect on all of these in this paper, but concentrate on the comments and recommendations relevant for safety cases that are made in the report. The Nimrod Review (in Chapter 22) says: "The safety case regime has lost its way. It has led to a culture of 'paper safety' at the expense of *real safety*." Safety cases are too lengthy and complex; use obscure language; lack operator input; tend to be compliance-only exercises; involve audits of process only; and make prior assumptions of safety of 'shelf-ware' (another term for non-development items).

The Review makes the point that the definition of safety case given earlier tends to encourage a "laborious, discursive, document-heavy argument ('a structured argument', 'a body of evidence') aimed at justifying a self-fulfilling prophecy ('system is safe')."

It is recommended by the Nimrod Review that safety cases be re-named *risk cases*, ("to focus attention on the fact that they are about managing risk, not assuming safety"). The risk case is intended to provide "reasonable confirmation that risks are managed to ALARP." As used in the Review, the term 'risk case' implies the need to focus attention on the most significant hazards and the ways that they can lead to dangerous situations.² It must conform to six principles (abbreviated as *SHAPED*): Succinct; Home-grown; Accessible; Proportionate; Easy to understand; and Document-lite.

The Review also comments that "care should be taken when utilising techniques such as Goal Structured Notation or Claims-Arguments-Evidence to avoid falling into the trap of assuming the conclusion (the platform is safe), or looking for supporting evidence for the conclusion instead of carrying out a proper analysis of risk."

The Review states that "care should be taken when using quantitative probabilities ... Such figures and their associated nomenclature give the illusion and comfort of accuracy and a well-honed scientific approach. Outside the world of structures, numbers are far from exact. Quantitative Risk Assessment is an art not a science. There is no substitute for engineering judgement."

4 Structure of the Safety Case

The exact structure of a safety case depends on the application domain and the relevant standard(s); however, all safety cases have a number of features in common. The safety case structure described in

²While the term 'safety case' is a shorthand for 'the (evidence-based) case for system safety', the term 'risk case' means something like: 'the (evidence-based and streamlined) case for system safety in which the system hazards and safety risks are clearly stated, understood and accepted'. The term 'risk case' does *not* imply, as some might think, a focus on consideration of system risks other than safety. The authors do not believe that 'risk case' is a helpful concept and it will not be used in the rest of this paper. Having said that, we recognize that technically unsound terms can nevertheless be effective in a management or political context.

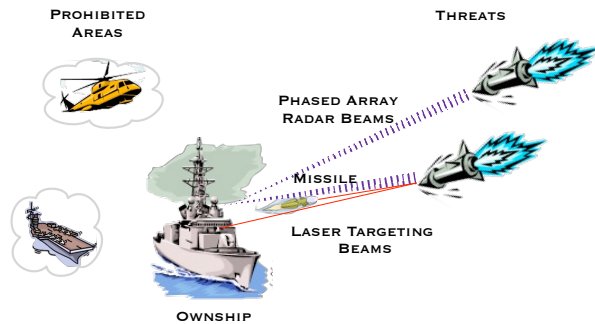


Figure 1: PARTI System Overview

this paper is taken from DEF(AUST)5679. There are three key reports in the DEF(AUST)5679 safety case:

- *Hazard Analysis* – identify the potential hazards posed by the system;
- *Safety Architecture* – demonstrate that the system is architected to be safe; and
- *Design Assurance* – demonstrate that the components are designed to be safe.

We illustrate the structure of the safety case using a case study from DEF(AUST)10679 (Department of Defence 2008a, Mahony & Cant 2008). The PARTI (Phased Array Radar and Target Illumination) System is a ship-borne Surface to Air Missile (SAM) targeting support system. It uses a Phased Array Radar (PAR) to direct laser illumination of hostile missiles and aircraft. The laser illumination provides targeting information to an existing ownship SAM capability. The main items of interest in the PARTI and environment are depicted in Figure 1.

4.1 Hazard Analysis

The first report of the safety case is called the *hazard analysis*. It provides an assessment of the danger (or threat to safety) that is potentially presented by the system. The hazard analysis must describe the system, its operational context and how the two interface from a safety viewpoint. Potential hazards posed by the system are then identified through a series of thought experiments about possible ways in which the system and its environment may interact to cause harm.

An *accident* is an external event that could directly lead to death or injury. An *accident scenario* describes a causally related mixture of system behaviours (*hazards*) and environment behaviours (*co-effectors*) that may culminate in an accident. The *severity* of an accident is a measure of the degree of its seriousness in terms of the extent of injury or death that may result. The *external mitigation level* associated with a hazard is a (qualitative) measure of the likelihood that an accident will result, given that the hazard is raised. The combination of severity and external mitigation level determine the *danger level* posed by each of the hazards individually and thereby the system in aggregate.

The need for a comprehensive identification of the relevant system hazards is probably self evident, but of particular interest to the discussion of UORs is the need for a complete description of the operational context. In the case of the PARTI system this involves such factors as: the ownship CMS (Combat Management System); the SAM capability for which the PARTI is providing a targeting service; ship support systems that provided power and physical security;

ship sensors that provided situation awareness; ship helicopters and other ordnance; personnel placements and procedures; friendly ships and aircraft; weather and sea conditions etc.

While such factors have immediate and obvious implications for the level of danger posed by a system, there are also more subtle implications for the suitability and effectiveness of a system's safety architecture or even on the nature of the system level hazards. The two primary hazards of the PARTI system are derived from the emission of radar and laser beams. These beams are both inherently hazardous (when directed at friendly assets) and necessary to the function of the system (when directed at missile threats), so the system can only be safely operated in a context that is aware of the hazard and is regulated to mitigate the hazard. In this case a protocol of prohibited areas is introduced into which the PARTI does not radiate and that vulnerable assets in the environment do not leave. The system hazards associated with the beams then become *radiating into a prohibited area*, rather than the unavoidable *emitting hazardous radiation*.

4.2 Safety Architecture

The aim of the *safety architecture* report is to describe and analyse the broad structure of the system from a safety viewpoint.

The first step is the development of a collection of *system safety requirements*, which collectively assert that the system hazards do not occur. The next step is to decompose the system into *components* and to describe how they combine to carry out the safety functions of the system. The interaction between components is described in terms of component *interfaces*, both between components and with the environment. Finally, the effectiveness of the safety architecture is shown by proposing *component safety requirements* and providing a *correctness* argument that shows how these component safety requirements ensure satisfaction of the system safety requirements (this is called *architecture verification*).

The architecture verification shows that the system will operate safely in its intended or *nominal* mode of operation. The safety architecture will, in general, also include *internal mitigations* that serve to make the system robust to unintended or *failure* modes of operation. Internal mitigations generally serve either to contain hazards (*partitioning*) or to distribute risk (*redundancy*). An argument must be made that the robustness of the internal mitigations and of the individual components is adequate to the dangers posed by the system.

The safety architecture of the PARTI system is depicted in Figure 2. The system's heavy reliance on situational awareness provided by the ownship's CMS is explicit in the diagram, but the safety argument may also make use of assumed properties of the operational context such as the deck placement of components, personnel placement at combat stations, sea state limitations etc. In fact, the system's two most prominent safety features, namely the components *PAR Filter* and *Interlock* provide redundancy in the safety functions of not radiating into protected zones. As described above, the effectiveness of these safety functions derive directly from the presence of mitigating factors in the operational context.

4.3 Design Assurance

The aim of the *design assurance* report is to provide evidence that components are designed and implemented so as to satisfy their component safety requirements.

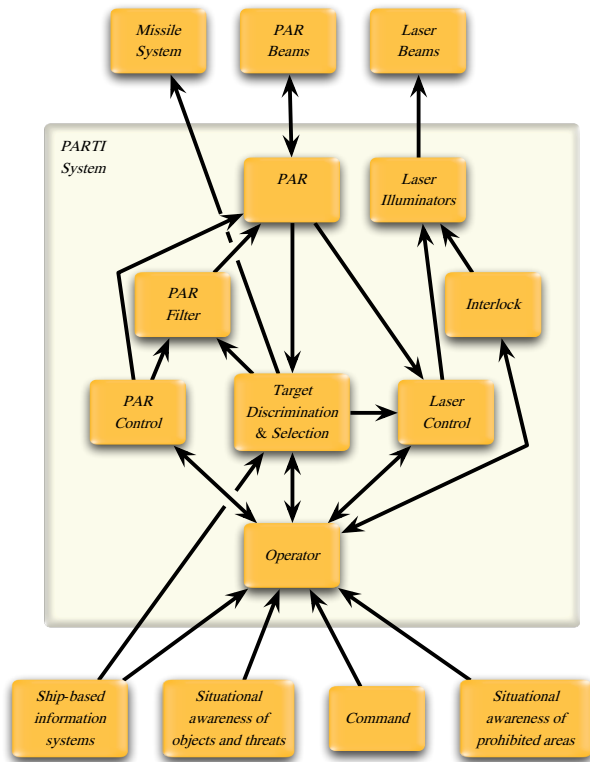


Figure 2: PARTI Architecture

The first step is to describe and justify the *implementation technology* for each component. This encompasses the design techniques and equipment used in the component. The choice of implementation technology must be justified as appropriate. In particular, the equipment must be shown to be suitably robust in consideration of the dangers posed and the assumed operating conditions.

Implementation technologies fall roughly into one of four classes, depending on whether the functions of the component are carried out using *analogue hardware*; *digital hardware*; *software* or *operator*. Specific assurance activities are prescribed, depending on the given implementation technology class.

Design analysis proceeds by re-expressing the component safety requirements in a form appropriate to the chosen implementation technology: this is called the *component safety specification*. A *design model* is developed for each component and a correctness argument developed that this model meets the component safety specification. *Design testing* is also carried out. Both design verification and testing may depend on assumptions about the behaviour of other components and even of the environment.

4.4 Safety Case Summary

The *safety case summary* is an overall narrative (or high-level argument) that is convincing to a third-party and pulls the results of the above phases together.

4.5 Observations on the Nimrod Review

Having discussed the structure of the safety case as provided in DEF(AUST)5679, and its application to the PARTI system, we now return to the conclusions of the Nimrod review and ask how DEF(AUST)5679 measures up against its recommendations.

First of all, consider the Nimrod review's recommendation to rename "safety cases" to "risk cases".

Although we do not agree with the change of name, we do agree with the intent: safety risks must be clearly identified. For this reason, we strongly believe that the hazard analysis phase remains central to *any* safety case. The main reason for this is that hazard analysis *identifies* the potential risks to human safety: without it, no sensible decisions can be made about whether sufficient effort has been made to eliminate or reduce these risks. It is the starting point for effective safety assessment of *any* system. This view is in direct accordance with the conclusions of the Nimrod Review. Roughly speaking it presents "evidence for unsafety" rather than "evidence for safety".

We strongly concur with the Nimrod Review's criticism of safety cases that are too process-focused. A primary aim of DEF(AUST)5679 is to focus attention on the safety of the actual system (product assurance). To this end a small, targeted collection of documents are mandated, each addressing a crucial aspect of system safety. There is a minimal number of purely process requirements. We call this approach *document-focused*. Adopting a document focus helps direct attention back to the system itself and its safety properties. It promotes a more general ownership of the safety case by de-emphasising the agents and processes involved in safety management and engineering. Similarly, it promotes reuse of safety case artefacts in subsequent maintenance and re-development phases.

We also concur with the need to properly involve operators (called *End Users* in the Standard) in all aspects of safety engineering. A number of the process-focused requirements of DEF(AUST)5679 are specifically designed to ensure appropriate levels of End User input to the safety case.

The danger of safety engineering devolving to a "compliance only exercise" is necessarily a concern regardless of the safety standard adopted. DEF(AUST)5679 addresses this through the Evaluator, an independent agent whose primary responsibility is to assess the technical safety of the system itself, focussing his/her attention on the quality and completeness of the arguments presented in the safety case. By bringing an independent set of experiences and biases to the Safety Case, the Evaluator serves as a second line of defence against "safety as a self-fulfilling prophecy."

Example: for a software-controlled explosive round, it is claimed in the safety case that there are no safety issues once the weapon has successfully been fired, and consequently there is no analysis of hazards relating to impact of the weapon in areas other than the intended target area. The hazard analysis (and subsequent safety case phases) may appear to address the requirements of the standard and the evaluator may acknowledge that the safety case is process compliant with the standard. However, in the absence of proper treatment of post-firing safety the evaluator will find the hazard analysis to be incomplete and thus unacceptable.

The abuse of quantitative risk assessment techniques has long been a concern of the authors. Numbers are often used to hide qualitative assessments on the basis that it helps them to fit into the overall risk management process. However, hiding qualitative assessments behind hard numbers can give them an unjustified level of technical authority — "You can't argue with the numbers." Often the underlying safety argument has little technical merit, safety becoming essentially a "self-fulfilling prophecy."

DEF(AUST)5679 strongly downplays the role of quantitative risks in safety management. There is no explicit requirement for quantifying risks; qualitative safety arguments are allowed (and usually preferred) at every level. This position derives from the soft-

ware focus of the standard and technical inadequacy of quantitative risk assessment for software based systems.

We note in passing that the most widely used military safety standard, being MIL-STD 882C (Department of Defence 1993), is both strongly process-focused and driven by quantitative risk assessment methodologies.

Our comments against the Nimrod Review's six principles for risk cases (*SHAPED*) are as follows:

- *Succinct* — this principle is reflected in the process described in DEF(AUST)5679. It focuses on system safety requirements and their decomposition into component safety requirements. It does not require elaborate flowing down of hazard analysis into subsystems. It uses diagrams to provide a clear picture of accident sequences.
- *Home-grown* — DEF(AUST)5679 stresses the need for End User participation in Safety Case activities and in particular the vital Hazard Analysis. This serves to promote End User awareness of system hazards and ownership of the Safety Case.
- *Accessible* — the safety case summary provides an overview of the safety case, and the safety case documentation should allow for easy searching and viewing of information.
- *Proportionate* — we believe that the process described by DEF(AUST)5679 represents an approach to safety case development that is proportionate to the level of danger presented by the system.
- *Easy to understand* — we agree with this in principle, although we consider the more basic principle to be that simple architectures promote safe systems. That said, the safety case should reflect the actual system. A simple easy to understand safety case for a complex hard to understand system will almost certainly be a wrong safety case. Furthermore, for any system, some of the assurance artefacts will, by their nature, only be understood by experts. The Evaluator's role is to provide independent judgement of the validity and strength of these artefacts in such a way as to be understood by the general reader. The safety case summary should also as a rule be simple and easy to understand.
- *Document-lite* — this is reasonable if the system is not too complex. In accordance with our remarks above, we would prefer to say 'document-focused'.

5 Non-Development Items

Defence procurements often involve what are called *non-development items* (NDIs). These are essentially items over which the supplier of the system has no design control. The use of NDIs, and their role in safety-critical systems, presents a number of complex issues that also arise for Urgent Operational Requirements.

Issue 2 of DEF(AUST)5679 views the use of NDIs as a necessary part of System Development. However, it makes no provision for tailoring or modification of Safety Case requirements for NDIs. The Safety Case is intended to discharge the responsibility under the OH&S Act to "take reasonably practicable steps to avert harm" which is not diminished by a decision to make use of a third party's development effort. It is not acceptable to make prior assumptions of the

safety of NDIs. Similarly, this OH&S responsibility is not diminished by Urgent Operational Requirements (see Section 6).

DEF(AUST)10679 — which provides Guidance Material for DEF(AUST)5679 — includes an Issues Paper on the use of NDIs (Department of Defence 2008a, IGP-004). This Issues Paper discusses the implications of NDIs for safety, with special reference to DEF(AUST)5679. The Issues Paper highlights three cases where Non-Development Items may appear.

- In general, *all* systems will normally make use of *non-development equipment* as part of the implementation technology of a specific component. Non-trivial examples include software components built in the framework of a commercial operating system; disk drives used for logging data etc.
- A specific *non-development component* may be used as part of the overall system design. For example, the PARTI system includes an already developed laser illuminator component.
- The system itself may be a *non-development system*. Examples include:
 - a commercial or military "off-the-shelf" response to a capability gap;
 - a system that was developed for another military context and is to be customised for use in a new operational environment; or
 - an upgrade of an existing or 'legacy' system (this could involve replacement of obsolescent hardware or a modification to software).

The Issues Paper stresses the importance of the hazard analysis phase — no matter what kind of NDIs are used in the system. This theme will be taken up again in the next section. Even for (perhaps especially for) a non-development system, a full hazard analysis must to be carried out, identifying and analysing the proposed operational context. The paper then goes on to discuss in detail NDI issues for the safety architecture and design assurance phases.

Notable examples illustrating the key roles played by NDIs in the Australian context are the Collins Class Submarine and the Air Warfare Destroyer (AWD).

The six Collins class submarines are the largest conventionally powered submarines in the world. They are based on the Västergötland class design built by Kockums Marine AG of Sweden. Long-standing issues with the originally envisaged combat system are being addressed by a replacement program using an "off-the-shelf" system (AN/BYG-1) from the US.

The Air Warfare Destroyer exemplifies a modern sophisticated defence platform that incorporates a number of capabilities. It will provide air defence for accompanying ships (as well as land forces and nearby coastal infrastructure), and offers self-protection against attacking missiles and aircraft. The AWD will make use of a special-purpose Aegis Weapon System incorporating long range anti-ship missiles. The AWDs can conduct undersea warfare via modern sonar systems, decoys and surface-launched torpedoes. The existing Spanish Navantia designed F100 class destroyer has been selected as the basis for the Hobart Class AWDs.

Each of these examples illustrates the use of existing designs, significant off-the-shelf subsystems and major system modifications in a complex defence platform.

Systems that involve NDIs present special challenges for the safety case. In particular:

- the safety case for the system may be non-existent, inadequate or developed in accordance with a different safety standard;
- the system may not have been designed and built with a rigorous safety engineering process; or
- there may be limited access to system development artefacts (including assurance evidence).

Nevertheless a safety case must be developed that properly addresses the hazards that arise from introducing the system to its intended operational context.

6 Impact on Safety Case Phases

Having described – at least in the terminology of DEF(AUST)5679 – the structure of the safety case, we consider how Urgent Operational Requirements can or should have an impact on it.

When there is an Urgent Operational Requirement, there might be political or schedule pressure to streamline — or even circumvent — normal safety case activities. However, there is no reduction in the duty of care required by the OH&S Act, so there is an equal need to be able to argue that the system is suitably safe when accepted into service. A safety case must be produced and it must be adequate to make a rational determination of system safety risk.

Beginning from this premise, we consider the vital question: *what reductions in safety case scope might be acceptable in the context of a UOR?* Such considerations are, of course, meta-level ones that are outside the scope of the standard itself. They need to be addressed by the technical regulatory and policy framework of which the standard is part and are finally a matter for the individual acceptance authority. In any case, it is clear that such non-compliance would have to be explicitly highlighted, acknowledged and accepted by the various parties in the safety case.

In the following, we address the implications of various conceivable reductions in the scope of the safety case.

6.1 Hazard Analysis

As discussed above, the hazard analysis phase of the safety case identifies potentially dangerous system behaviour. Of critical interest to those assessing the safety case is the list of accidents (and their severities). The accident list lays out in stark detail how dangerous the system might be. The accident sequences show how these accidents could actually arise from certain system states (hazards).

For the safety case to be adequate to the task of assessing system safety risk, it is absolutely necessary to determine (correctly) the potential hazards that can arise and the severity of the accidents they may cause. No UOR can be sufficiently pressing to justify the acceptance of a safety risk that is unknown.

Hazard analysis is not as onerous as might be thought. It involves a thought experiment by a diverse group of people with sufficient knowledge of potentially dangerous flows from the system, across its boundary and out into the environment. It does need to be done in a systematic way to ensure complete coverage of hazardous interfaces.

Once the severities of the system hazards have been determined, they provide an initial upper bound on the system safety risk. It might be tempting to use this to determine acceptability of the system; however, a determination of safety risk should not be

based entirely on accident severity. An assessment of accident likelihood is required to properly assess system safety risk. A minor accident that occurs with high frequency may be of more concern than a catastrophic accident that occurs with negligible frequency.

In order to address this aspect the operational context must be properly described, allowing the analysis to be further refined by consideration of external mitigations. Once danger levels have been properly assigned to hazards and to the overall system, they reasonably be thought to serve as an upper bound to the system safety risk in its intended operating context. However, this upper bound is likely to significantly over-estimate the system safety risk as the quality and robustness of the system itself have not been assessed and must therefore be assumed to be at the lowest of levels.

Even if this over-estimated system safety risk is assessed as being sufficiently low when balanced against the mission risk posed by the UOR, it may be difficult to argue that this level of safety analysis is sufficient to constitute taking “reasonably practicable steps to avert harm ...” Generally the acceptance authority will prefer to see argument that steps had been taken to ensure that the system possesses safety qualities and functional robustness commensurate with the identified system danger level.

6.2 Safety Architecture

The safety architecture phase has essentially two components: a safety correctness argument and a safety robustness argument. In response to a UOR, the developer may consider providing a safety case that omits one or the other of these arguments.

First suppose that only the robustness argument is made. This would allow the safety case to identify the internal mitigations present in the system, thus demonstrating that reasonable steps had been taken to make the system safe. This would also allow the strength of these internal mitigations to be used to provide a tighter bound on the system safety risk. However, in the absence of the safety correctness argument it will be hard to defend the technical validity of the robustness argument. Recall that the safety correctness argument demonstrates that the system is architected to be safe to operate when free of internal equipment failure. If the system is not safe in the *absence* of equipment failure, the robustness of system function in the *presence* of failure is cold comfort.

Conversely, suppose that only the correctness argument is made. This provides assurance that the system is architected to be safe to operate in the absence of equipment failure, but it will not be possible to confidently argue a reduced bound on the system safety risk. An understanding of system failure modes and their potential to realise system hazards is critical to assessing the system safety risk.

In summary, the robustness argument is essential to a proper assessment of system safety risk, but it cannot be trusted in the absence of a safety correctness argument. They are complementary activities, mutually informing each other, and both are required to provide a credible assessment of the risk posed by the system safety architecture.

As above, by making worst case assumptions about the quality and robustness of system components, the architecture assessment can be used to determine an upper bound on the system safety risk. Again, at best, this bound remains a significant over-estimate of system safety risk and it is questionable whether the developer can be said to have taken “reasonable steps *etc*” if appropriate analysis of component design is not undertaken.

6.3 Design Assurance

Having established that the system is architected for safety with an appropriate level of robustness to equipment failure, the design phase of the safety case turns attention to the fitness of individual components for the purpose assigned them by the architecture.

Again, design assurance consists of highly complementary correctness and robustness arguments; so as argued above it is hard to make use of one in the absence of the other.

6.4 Conclusion

We claim (perhaps unsurprisingly) that the body of evidence required in the DEF(AUST)5679 safety case is the minimum needed to provide a credible argument that safety risk has been properly assessed and that the developer has taken “reasonably practicable steps to avert harm to members of the public, as well as their own employees.”

The levels of rigour dictated by DEF(AUST)5679 are perhaps more open to debate and we do not consider them here in any detail. They simply represent a reasonable attempt to provide a mapping between the current range of commercially feasible levels of rigour and system danger levels.

Even if it is considered that the UOR makes lower levels of rigour acceptable, timely safety case development will require the application of a significantly higher safety analysis effort as a proportion of overall development effort. The safety case needs to provide essentially the same body of evidence as for standard acquisition, but over a compressed time period.

7 The Impact of Acquisition Class

Recall the three acquisition classes described earlier: Greenfield Acquisition, In-Service Modifications and Modified Operational Context. Each of these provide different advantages and disadvantages for any attempt to shorten the duration of safety case development. Generally speaking, timely response to UORs is most favoured in circumstances in which significant reuse of existing safety analyses is possible. We consider each class briefly for potential reuse, illustrating our discussion with simple example systems.

7.1 Greenfield Acquisition

For this acquisition class, there is no existing system for addressing the desired capability and hence no existing safety case. Both the system and its accompanying safety case must be developed to meet a pressing deadline.

Clearly, in most cases it will be very challenging to develop a completely new solution to meet a UOR in a timely fashion. For this to be contemplated with significant chance of success, it is likely either that a very simple solution system is envisaged or else that some existing third-party system is known to address the UOR.

In the former case, the simplicity of the system is likely to favour timely safety case development as much as it does general system development. It has often been observed that simplicity is a great friend of safety.

Example: A UOR results in a proposal to develop a new flak jacket based on a novel material. A hazard analysis of the new jackets is likely to focus primarily on the chemical properties of the new material (toxicity, heat resistance, etc) and the ergonomic hazards of the jacket design and it is likely that the safety

case will be relatively small in scope. Such systems as these also benefit from being developed in a highly mature discipline. The science of combat clothing is well studied, with well documented history of use on military operations.³

The latter case essentially devolves to the use of an NDI system. As observed in Section 5, this presents a considerable challenge in the absence of an existing safety case. Producing a safety case for an NDI can require more time and effort than for a bespoke system, even if commercial secrecy does not render it infeasible. The most favourable situation would follow from the NDI being a common consumer level device with few safety hazards or at else produced by an industry with a mature safety culture.

Example: A UOR results in a proposal to make use of commercial tablet devices to gather and communicate military intelligence. Hazard analysis may show that the equipment itself presents few safety hazards. However, depending on the nature of the intelligence and the purpose it is used for, there may be significant safety concerns requiring extensive safety engineering effort to address.

If the NDI system is provided with an extant safety case, the main concerns will revolve around the degree to which the Operational Context of the UOR matches that used in the safety case. The situation is essentially the same as for a modified operational context acquisition as discussed in Section 7.3.

Example: A UOR results in a proposal to procure a commercial bus. Hazard analysis will concentrate on the ways in which the envisaged military operational context may differ from the typical civilian operational context for the bus. If the operational context is essentially unchanged, the safety analysis will be able to depend largely on the civilian safety certification of the bus and may be concluded quickly.

7.2 In-Service Modifications

In this situation we have an existing system, with a safety case that has been accepted, and we intend to modify the system. The safety case must be updated to reflect the modification.

Firstly, we observe that this is a most favourable situation for rapid safety case development. For the contemplated modification to be feasible in a short time frame, it is likely that the scope of the proposed modification is small and much of the existing architecture and design is to be re-used. Often this will also be true of the safety architecture and design, so that much of the safety case is also re-usable.

It is also of considerable advantage if the existing operational context is maintained (we deal with the situation where this is not so in Section 7.3). In this case, it is likely that much, if not all, of the existing hazard analysis remains valid. Even so, it is necessary to reconsider the hazard analysis in a careful manner.

The simplest kind of modification that might be proposed would be the substitution of one piece of equipment with another as it may be possible to reuse the existing safety case almost totally. If hazard analysis does not reveal hazardous properties of the new equipment itself and the modified functionality is not related to component safety functions, then the safety architecture remains unchanged and the component design is changed only in as much as the equipment list changes. For once, the distinction between mission and safety functions may work in favour of speedy safety case development.

³The sinking of the HMS Sheffield by an Exocet Missile during the Falklands War resulted in changes to protective clothing; the synthetic materials worn by sailors were found to melt onto skin, increasing the severity of burn injuries in the victims.

Example: a UOR results in a proposal to swap an armoured vehicle's existing illuminator for night operations with one that provides better performance in harsh conditions. The intent is to improve on performance and reliability. Quite possibly the original illuminator had no direct bearing on the safety case (since it was always regarded as non-development equipment anyway). Thus updating the safety case simply involves re-visiting the hazard analysis (to ensure the higher performance illuminator is not itself dangerous) and noting the change of equipment in the design assurance.

The next step up the design hierarchy is a modification that replaces an existing component in total. Again, although the replacement component may be expected to provide different mission functionality, it very well may retain the same safety functionality. If hazard analysis reveals no new hazards associated the replacement component, it may be possible that the update to the existing safety case can concentrate largely on design assurance for the new component.

Example: A UOR results in a proposal to improve the ability of the PARTI system (see Figure 2) to illuminate multiple incoming missiles. The existing design makes use of a two laser illuminator component, which has been superseded by a three laser unit. Provided that the individual lasers of the new unit are functionally equivalent, producing the modified safety case may require little more than a re-evaluation of the hazard analysis. A complicating feature of this modification is the fact that laser illuminator is an NDI. The original safety case made use of a DefStan 00-56 (Ministry of Defence 2007) based component safety audit. If the new unit is not provided with similar design assurance data, the re-development of the safety case may be considerably more difficult.

Finally, we note that modifications that involve significant changes to the system safety interface or the safety architecture of a system may require the re-development of significant parts of the original safety case.

7.3 Modified Operational Context

In this situation, we have an existing system, with corresponding safety case, and we intend to make use of it in a different operational context. This can easily be an unfavourable situation, as any change in the operational context has the potential to cause major revision to the safety case and even modification of the safety architecture and component designs. All phases of the safety case could make use of assumptions about the operational context.

Clearly, the operational context is a critical factor in hazard analysis and the accident scenarios have a direct dependence on this context. It follows that if the operational context is modified, then the hazard analysis must be thoroughly reviewed and may need extensive redevelopment. Not only is it possible that new accident scenarios may arise, but existing ones may involve co-effectors that no longer exist or external mitigations that have been weakened or strengthened.

The operational context is also a critical factor in safety architecture analysis. If the architecture correctness argument makes use of properties of the original context that are not present in the new context, it may be necessary to completely re-develop the safety architecture. Of course, the safety architecture may use context properties that are not required by any mission function, so that the need to re-architect for safety may not be immediately apparent when the change in context is first considered. This is especially so where there is little or no safety analysis in early planning.

The operational context may even be a factor in component design assurance. All in all there is considerable potential for a change in operational context to result in significant re-engineering of the system safety case.

Example: A UOR results in a proposal to deploy a PARTI system in a land-based operational context. This project will face considerable challenges due the tight integration of the PARTI system with the ship's CMS, but even if this can be overcome, the heavy reliance of the existing safety case on the ability of the context to define and enforce protected zones for friendly assets will cause significant headaches in producing a modified safety case. Land combat environments are considerably more crowded and less structured than sea environments.

8 Final Remarks

In this paper, we have discussed Urgent Operational Requirements and the impact that they may have on the safety case. We believe that the threats to the safety of personnel and civilians arising from the installation and use of Defence equipment remain of paramount consideration. No UOR can be so urgent as to warrant ignoring safety issues or failing to properly assess them. Those accepting systems into service need to properly understand the safety risks associated with a system if they are to properly weigh them against the UOR and mission goals. We do recognise that, in a combat situation, commanders frequently put their personnel at risk in order to achieve mission goals; in particular, a commander in the field can make a command decision to override or ignore a safety-related issue or procedure.

While a system need not be safe in some absolute sense when adopted into service, a clear and accurate assessment of safety risk is a critical input to good decision making. Many of the properties of the system that have bearing on safety (reliability, robustness, operation context) are also critical to the operational success of the system. Thus, the effort put into understanding the system from a safety viewpoint may even serve to improve the operational outcome more generally.

Armed forces are likely to become more and more required to deploy rapidly in different regions of the world and to adapt quickly to the conditions that they face. Thus, there will be increased pressure for rapid acquisition of capabilities. Those responsible for developing safety cases need to have robust and efficient processes for carrying out hazard analysis and other safety case phases. They must be steadfast in analysing and highlighting safety risk to decision makers — especially when “reasonably practicable steps” have not been taken to avert harm.

References

- Department of Defence (1993), System Safety Program Requirements, Military Standard MIL-STD-882C, United States of America.
- Department of Defence (2008*a*), Guidance Material for DEF(AUST)5679/Issue 2, Australian Defence Handbook DEF(AUST)10679/Issue 1, Australian Government.
- Department of Defence (2008*b*), Rapid Acquisition of Capability, DI(G) LOG 4-1-008, Australian Government.
- Department of Defence (2008*c*), Safety Engineering for Defence Systems, Australian Defence Standard DEF(AUST)5679/Issue 2, Australian Government.
- Department of Defence (2009), *The Strategic Reform Program 2009, Delivering Force 2030*, Australian Government.
- Haddon-Cave, C. (2009), *The Nimrod Review*, The Stationery Office Limited UK.
- Mahony, B. & Cant, A. (2008), The PARTI architecture assurance, in 'Proceedings of the 13th Australian Conference on Safety-Related Programmable Systems', ACS.
- Ministry of Defence (1983), Design and Airworthiness Requirements for Service Aircraft, Volume 1 - Aeroplanes, Defence Standard 00-970, United Kingdom.
- Ministry of Defence (2007), Safety Management Requirements for Defence Systems, Part 1 Requirements, Defence Standard 00-56, United Kingdom.
- Ministry of Defence (2011), *Commercial Guidance for the UK MOD Defence Acquisition Community: Urgent Operational Requirements*, United Kingdom. <http://www.aof.mod.uk/aofcontent/tactical/toolkit/>.
- Tyler, A. (2009), 'Performance under pressure', *The RUSI Journal* **154**(5), 30-33.