Complex Reactive Real Time Systems and the Safety Case

Gordon R. Stone
Compucat Research Pty Limited
14 Wales Street, Belconnen 2617, Australian Capital Territory
gordon.stone@defence.gov.au

Abstract

Complex reactive real time systems are systems of systems that interact with the external world to perform selections of tasks in real time. They may be required to carry out selected tasks when some components of the system cannot be considered acceptably safe. They may be required to operate continuously.

The requirements of a safety case for such systems can be considered by use of a "response to tasking model" based on a representative system. Consideration of the response to tasking model indicates that what we are used to considering as a safety case for more straightforward systems may be inadequate for such systems.

The application of safety cases to existing complex reactive real time systems may benefit from review. As semi-autonomous and autonomous systems become more prevalent, there will be more urgency for safety cases to take account of a sophisticated, automated decision-making process that is flexible, responsive to changing circumstances and intelligent. This may require re-evaluation of the format, content and presentation of safety cases. A navy warship is used as an example complex reactive real time system to illustrate the issues in demonstrating achievement of acceptable safety.

Keywords: Safety case, complex systems.

1 Introduction

This paper is a practitioner's view of safety cases and complex reactive real time systems. It is largely based on experience gained in the field.

For the sake of brevity, throughout this paper the terms 'acceptably safe' and 'safely' are used in place of 'have an acceptable safety risk' and 'with an acceptable safety risk' respectively. This issue is discussed in more detail in Section 3.1.

For purposes of discussion, the term 'complex reactive real time system' is used to describe systems that interact with the external world to perform a range of tasks in real time. On occasion, such systems may be required to carry out all tasks concurrently or a selection of tasks concurrently. Frequently, suspension of the tasks currently being performed is not an option. These systems are responsive to their environments and modify their environments. They may be required to operate for long periods without assistance. When tasked with one or more concurrent tasks, they need to determine their ability to perform the tasks, either wholly or completely, and then set about performing the tasks. As circumstances change during performance of the tasks, the systems may need to modify their behaviour based on external conditions or in response to the systems' internal status. The systems may also need to modify the number and scope of the tasks that they carry out in response to the same factors.

To enable them to perform their assigned tasks, the systems must:

a. recognise the elements comprising the task,

b. synthesise a view of their environment based on information provided by a range of sensor systems,

c. determine the status of the system components based on information from internal sensors,

d. decide whether the task can be safely carried out in whole or part, and

e. assess the consequences of not doing the task if the task is deemed not safe to carry out in whole or part.

When deciding whether tasks can be carried out safely in whole or part, these systems may need to consider whether there are alternative ways of performing parts of the task when one or more components of the system are not fully functional or one or more safety controls are unavailable.

In many such systems, safe behaviour is often the responsibility of the 'human in the loop'. As the requirement to deal with complex situations more rapidly increases, more autonomous behaviour will be assigned to the non-human components of systems. This is already evident in naval warfare where the human in the loop is sometimes unable to act as a safety circuit breaker due to the requirement for rapid response and the almost total reliance on the view of the external environment synthesised by the system.

Kelly, Bate, McDermid and Burns (1997) state: "The safety case is the document, or set of documents, presenting the argument that a system is acceptably safe to operate in a given context. For safety-critical and related systems, an acceptable safety case must typically be presented to the appropriate regulatory authority prior to a system being allowed to enter service". This paper is

based on the above definition, with the qualification that prescribing the given context for the types of systems under discussion is not simple, and variations in contextual factors at different times can have significant consequences for system safety.

A safety practitioner's view of the safety case is that it provides assurance (in the form of a defensible argument) that a given system is acceptably safe within prescribed conditions of use. Such assurance is valuable because it gives comfort to the organisation acquiring and using the system that the organisation's employees will not be subject to unnecessary safety risk. It is also useful in the event of the occurrence of safety incidents; it demonstrates duty of care and thus helps to limit the effect of legal action under Occupational Heath and Safety Acts.

The safety case is also informs the owner's processes that define the usage patterns of the system by clearly articulating the conditions of use under which the system is acceptably safe. The result ought to be policies and procedures for employment of the system such that it remains acceptably safe at all times.

The argument in the safety case relies on:

a. a description of the system as known at the time, including its capabilities and limitations;

b. the intended conditions of use of the system as known at the time;

c. definitions of key terms such as 'acceptably safe';

d. a set of assumptions; and

e. evidence gathered from the design, analysis and testing of the system.

Current complex reactive real time systems rely on humans to ensure that they carry out their tasks safely. However, difficulties in retaining labour and current trends in reducing labour costs through personnel reductions are not only causing the workloads of personnel managing systems to increase, they are also reducing the pool of experienced system managers and reducing the numbers of personnel to whom system designers and managers can pass on their experience. As these trends continue, pressure for adoption of semi-autonomous and autonomous systems will increase.

The means by which the safe operation of such semi-autonomous and autonomous vehicles can be achieved, and the means by which that safety can be assured, must be significant challenges. The advent of such systems may require re-evaluation of the format, content and presentation of safety cases.

2 Complex Reactive Real Time Systems – Response to Tasks

Complex reactive real time systems respond to tasks. To perform tasks, such systems require sensors, a decision-making capability and effectors.

To determine whether the system can respond to tasks, and how it can respond, the system needs to be able to determine, with sufficient granularity for decision-making, the:

a. tasks to be performed and the elements of the tasks;

b. capabilities of each system component, where a component may be a system, equipment or item of software;

c. system components that are required to perform each element of each task;

d. conditions under which each of those system components can perform its required elements of the task safely;

e. status of each component; and

f. environment in which the tasks are to be performed.

The tasks need to be decomposed into collections of task elements with sufficient granularity to enable the following:

a. Identification of task elements that are shared among tasks.

b. Assignment of system components to the task elements.

The assignment of system components to task elements needs to be at a level that allows decision-making based on task assignment, environmental conditions and system component status. For example, a task to navigate a warship through a shipping lane should require, amongst other things, that the position and relative velocity of all objects of interest be determined and maintained and that collisions be avoided. Task decomposition should enable the ship kinematics and representative kinematics of other ships and watercraft to be taken into account so that collision avoidance can be effected. In addition, the task should provide for the inclusion of environmental constraints such as local navigation rules and the effect of weather on the sensor systems.

The capabilities of each system component need to be identified in terms that enable the decision-making component to assess its contribution to the performance of one or more elements of the tasks. The factors that influence the system component's capabilities also need to be described in terms that allow them to be associated with the environmental conditions during the performance of tasks. Finally, the residual capabilities of the system component in the presence of partial failures need to be identified so that the effects of component status can be assessed. For example, the capabilities of the warship's navigation radar must include its range performance against representative types of watercraft under ideal conditions, plus functions that relate its performance degradation to environmental conditions such as rain. If the radar is able to function with reduced performance when it has some failures, the degradation in performance related to each failure must be known. Further, if failures in the radar limit the time for which

the radar can be operated, the failures and the restrictions on time need to be described in a way that will support decision-making.

Where a system component is not completely serviceable, the system needs to determine whether the serviceable parts of the component are sufficient to support safe performance of the required elements of the task. Where the serviceable parts are insufficient to support safe performance of the required elements of the task, the system needs to determine whether a suitable substitute for the unserviceable component is available or whether the unserviceable component can be supplemented to support safe performance of the task. Where the available serviceable system components cannot assure safe performance of the task using normal task processes, the system needs to consider whether changes to the task processes can be made such that the task can be performed safely. If none of the above measures will result in safe performance of the task, the system must consider de-scoping the task or rejecting the task entirely. Part of this consideration is whether de-scoping or rejecting the task may have consequences that are unacceptable for safety in the context broader than the system.

Complex reactive real time systems have a range of sensors to support all of the required tasks. Some sensors support decision-making by providing information about objects of interest in the external environment, some sensors provide information about the environment itself and some sensors provide information about the status of the system itself.

From the safety perspective, sensors that provide information about the objects of interest in the external environment provide information essential to prevent the system coming to harm, such as colliding with geographical features. They also provide information essential to prevent the system harming entities that it should not, such as running down a pleasure craft.

To provide flexibility and adaptability, the capabilities of each sensor overlap with capabilities of one or more other sensors. The overlap may be in the same sensor space or a complementary sensor space. As an example of the first case, a radar used to detect surface vehicles can be used for collision avoidance when the navigation radar is unavailable. As an example of the second case, an electro-optical system can be used in certain environmental conditions to supplement or be substituted for a radar used to detect surface vehicles when that radar is temporarily below its full capability, albeit with reduced overall system capability.

Sensors that provide information about the environment itself are essential to enable the system to determine the actual capabilities of the system when one or more tasks are to be performed. For example, weather affects the mobility of a warship, rain affects the ability of some sensors to detect objects of interest, and some types of weather affect the ability of the system to use its effectors. Only when the actual capabilities of the system components are known can the safety of performing the required tasks be assessed.

Information from the sensors is not very useful until it is synthesised into a coherent view of the external environment. Once a coherent view of the external environment is available, the problems inherent in carrying out the required tasks must be determined and strategies applied to overcome the problems. As each task is being performed, the problems must be continually reassessed and modified as required. Where the system is performing several tasks concurrently, conflicts in the use of system resources to overcome problems must be resolved. For example, if a warship is required to engage surface targets and air targets, there may be insufficient resources to do both at the same time. Resources need to be scheduled to maximize the probability of successful task completion.

3 Some Issues With Safety Cases for Complex Real Time Systems

Application of a safety case to systems that have a limited, unvarying task set is relatively straightforward, with the possible exception of the software component. However, the application of a safety case to complex reactive real time systems may not be straightforward for the following inter-related reasons:

a. The term ‘acceptably safe’ is vexed.

b. The system description may not be complete nor completely accurate.

c. The set of intended system behaviours may not be complete and/or completely accurate.

d. The system may exhibit emergent behaviour.

e. The system may be required to operate in circumstances not envisaged during its design.

f. The system may not be required to exercise all of its capability all of the time.

g. The system may be required to carry out tasks when it contains failures.

h. The volume of work required to develop and maintain the safety case is considerable.

3.1 Acceptably Safe

The terms ‘safe’ and ‘acceptably safe’ are used in this paper to mean that the safety risk of the system is acceptable to the owner of the system within current applicable laws and regulations. This usually means that the system is acceptably safe when it is used under the conditions of use prescribed during its acquisition. However, in some circumstances the system may need to be operated outside the prescribed conditions of use for which the safety case is valid. In such cases, the owner of the system accepts an increased safety risk based on a safety risk assessment made at the time.

In many cases, the acquirer prescribes the broad conditions of use; the developer specifies additional conditions of use and, in consultation with the acquirer, refines the acquirer’s conditions of use.

The means by which the owner is convinced that the safety risk is sufficiently low for the system to be acceptable is the safety program run during acquisition of the system. One output of that program is the safety case. The safety case report submitted when the system is offered for acceptance is a primary consideration.

The terms ‘safe’, ‘acceptably safe’ and ‘acceptable safety risk’ are vexed where complex reactive real time systems are concerned for the following reasons:

a. The nature of some systems, such as the warship, is inherently dangerous and the tasks that they are required to carry often involve significant safety risk. Not all safety risks can be reduced to levels considered acceptable for the general community.

b. Some components of the system may not be able to be made safe without making them unable to contribute effectively to performance of the required tasks. In such cases, the owner of the system accepts the residual safety risk for prescribed conditions of use.

c. The ‘safety’ of the system needs to be established at the instance that a mishap occurred. To do so requires consideration of the tasks, material state and environment of the system at the time.

The application of the law, in particular the Occupational Health and Safety Acts, to complex reactive real time systems is not entirely clear. For example, if a system needs to be operated outside its prescribed conditions of use, or is required to carry out inherently dangerous tasks, establishing that duty of care has been exercised may not be easy. In such circumstances the time available to make decisions may be insufficient to make and record a comprehensive safety risk assessment.

3.2 System Description

The system description may not be complete nor completely accurate. This can arise where the design of the system is very complex and/or involves new techniques. The challenge of understanding thousands of pages of design documents is not easily surmountable, even with design decomposition. It seems axiomatic that interfaces between systems with complex behaviour are also complex, and that systems that are not well understood are likely to have interfaces that are not well understood.

Incomplete system design can also arise through incomplete definitions of the tasks that the system is to carry out, including the conditions and constraints under which the system is to carry them out. Incomplete task definitions flow through into incomplete or inaccurate system specifications.

Incompleteness and inaccuracy in design description can be exacerbated by changes to the system requirements during and after system development. Systems that need to respond to an evolving environment, such as fighting vehicles, are particularly susceptible to this effect.

Systems that include components developed previously for other purposes may not obtain a sufficiently detailed description of each such component to ensure that the description of the entire system will be valid. Design disclosure may be limited due to intellectual property concerns, such as occur with certain commercial software applications. Establishing the conditions of use under which components of the systems are safe for inclusion in a complex reactive real time system is an essential part of the system description required to support a safety case. However, the conditions of use under which components previously developed for other purposes are safe may not be articulated or may be different from the conditions of use intended for the component when included in the system. In particular, the conditions of use for commercial items may not be sufficiently well articulated to support a sound safety argument.

3.3 System Behaviours

The set of intended system behaviours in the safety case may not be complete nor completely accurate. To some extent, this reflects inadequacies of the requirements analysis and design. However, determining the entire set of behaviours of a complex reactive real time system is not trivial. In addition to defining the tasks that the system must perform and the elements of those tasks, the relationships among those tasks and their elements must also be defined. The design must consider all ways that the system can carry out its tasks. For example, the order in which part tasks can be carried out can affect the system behaviour.

3.4 Emergent Behaviour

The system may exhibit emergent behaviour in response to particular circumstances not subjected to analysis or testing. Emergent behaviour may arise due to incompleteness in component specifications, particularly software. It may also arise from exploitation of incompleteness, ambiguities and errors in inter-component specifications. It may also arise from software logic errors, software coding errors and inadequate handling of exceptions within software components. In addition, operators may introduce emergent behaviour through carrying out actions in ways not anticipated, particularly if the operators are tired. Although emergent behaviour is avoided, limited or eliminated by robust design and construction of the system, eliminating emergent behaviour in complex reactive real time systems is a significant challenge.

Avoiding emergent behaviour through complete software specifications is not easy. As an example, the order in which part tasks are performed can be important. In a warship command and control system, several operators manipulate common data objects in real time. The objects and attributes under control of each operator vary in time, as do the sequences in which each operator manipulates the attributes. The resulting non-deterministic behaviour of the system results in many possible combinations of software paths, with a consequential increase in the likelihood of a behaviour that is only partially described.

To illustrate the complexity of eliminating emergent behaviour, consider the case of exception handling in a warship command and control system. Suppose that an exception occurs due to temporary memory corruption, a distinct possibility in a warship. Exception handling that will result in the safest behaviour depends on what the warship is doing at the time that the exception occurs. It also depends on where the exception occurs. For example, if it occurs at the beginning of a safety-related thread, it may need to be handled differently than if it occurs near the end of the thread. If it occurs in a method that is invoked by a method that is invoked by a method, propagating the exception may or may not be a safe strategy depending on the construction of the software component in which it occurs. The safety of the outcome of handling some exceptions will also depend on exactly what the complex reactive real time system is doing when the exception occurs.

Operators embedded in complex reactive real time systems can also be sources of emergent behaviour. Incomplete training, insufficient practice and inadequate understanding of the system can contribute to a failure to respond appropriately to novel situations. Deficient delineation of duties and responsibilities can also lead to emergent behaviour when operators attempt to perform tasks for which they do not have the appropriate competencies, or assume responsibility for tasks additional to their expected role to the detriment of the tasks included in their assigned roles. Finally, through experience as part of the system over an extended period operators may gain undue trust in the system and fail to recognise emergent behaviour in automated capabilities that provide recommendations or initiate actions that will complete without operator veto.

3.5 Circumstances Not Envisaged During System Design

The system may be required to operate in circumstances not envisaged during its design, construction and test. Even where best endeavours were applied during requirements definition to identify all tasks and environmental conditions in which those tasks are to be carried out, unanticipated circumstances can arise through permanent changes to the external environment or through temporary combinations of circumstances. These unanticipated circumstances can stimulate emergent behaviour. Modifications to the system as a result of unanticipated circumstances may introduce further scope for emergent behaviour. For example, when sea skimming missiles were first introduced, radical changes were required to warship combat systems. Whether emergent behaviour resulted from the modifications depended on how well the original system was designed and constructed, the design and construction was described to the personnel modifying the system, and the robustness of the program that developed the modifications.

3.6 Not All Capability Exercised All The Time

Some complex reactive real time systems are not required to exercise all of their capability all of the time. This means that system safety must be assessable for subsets of tasks. A simplistic approach asserts that assurance of safety when all capabilities are being exercised concurrently in conditions of greatest system stress will cover the assurance of safety when subsets of capability are exercised; the availability and capability of resources for performance of subsets of tasks should be no less than when the system is under its greatest stress. However, the approach may mask the ability of operators to carry out subsets of tasks safely under less stress. The approach is also unhelpful in considering safety when the performance of some system components is degraded.

3.7 Carrying Out Tasks In The Presence Of Failures

The system may be required to carry out tasks when it contains failures. There will be times at which some safety controls are ineffective but the overall system remains safe because the components of the system for which the safety controls are ineffective are not required for the tasks in force at the time. There will also be times where supplementation or substitution of system components that contain failures can permit safe performance of tasks that would otherwise not be safe.

For complex reactive real time systems that are self-sustaining for long periods, some system components may be unavailable due to planned maintenance activities. Although not failures as such, the effect is the same as when failures occur, except that the ‘failure’ can be scheduled at a time when it is expected to have the least adverse effect on required system capabilities.

3.8 Volume of Work

The above considerations indicate that a safety case for a complex reactive real time system is in itself complex. Part of the complexity is the scale of the safety case. As stated by Kelly (2001), a promising approach to arguing the safety of modular systems is the use of compositional safety cases. In addition to minimising the effects on the safety argument of changes in system components, they assist with intellectual manageability of safety arguments for large systems. The effort required to develop a safety case that adequately addresses all combinations of task subsets, environmental conditions and system states is considerable.

4 The Warship as an Example of a Complex Reactive Real Time System

An additional qualification on the term ‘acceptably safe’ is required for warships. The Occupational Health and Safety (Commonwealth Employees) Act applies to the crews of Australian warships. Even in peacetime, warships are sometimes required to operate outside their prescribed conditions of use. In times of war, application of the Act in relation to warships is not simple. The Act excludes detriment to national security. However, death or injuries to crew due to actions of the ship are still covered by the Act. Death or injury to friendly units and neutral units due to actions of the ship may also be subject to the Act. The Act may also apply to actions of

hostile units directed against the warship or units it is protecting at the time.

The application of 'acceptably safe' to a warship that is part of a task force under command of a foreign power is unclear.

A warship is a system of systems that operates within a certain context. Figure 1, which is meant to be illustrative rather than exhaustive, illustrates a partial warship context. Some contextual items are omitted, but they do not affect the argument.

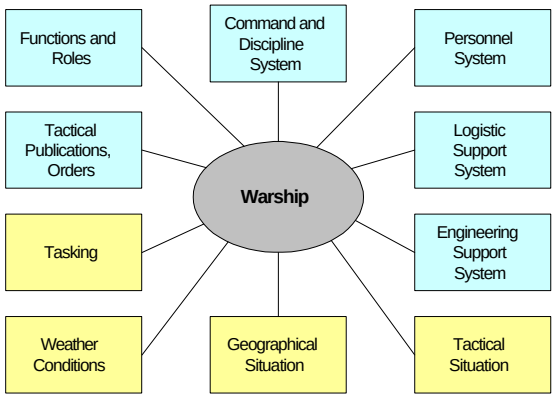

Figure 1: Partial Context of a Warship

The context of a warship includes factors that define what the warship is required to do, how it does what it is required to do and how it is supported while it does what it is required to do. Some of these factors are shown in Figure 1 in light shading.

In addition, the context of the warship contains environmental factors. Some environmental factors are shown in Figure 1 in darker shading.

Functions and Roles define what the warship is required to do. Examples of peacetime functions include maritime surveillance, fisheries protection and border protection. An example of a peacetime role is 'patrol vessel'. Examples of wartime functions include protection of shipping, self-defence and maritime strike. An example of a wartime role is 'escort'.

The functions translate to tasks. Tasks can share common elements. For example, maritime surveillance, fisheries protection and border protection all require the warship to detect and track other vessels. Boarding other vessels can be a task element of fisheries protection and border protection, but is not a task element of maritime surveillance.

The Command and Discipline System is responsible for ensuring that the chain of command is well formed, personnel are aware of their duties and responsibilities and that orders are carried out in a timely fashion.

The Personnel System, Logistic Support System and Engineering Support System are examples of external systems that provide support to the warship. Some of these systems have a bearing on whether the warship can carry out its assigned tasks safely. The navy's personnel system is responsible for ensuring that sufficient personnel with suitable competencies are provided to crew the warship. Unless this occurs, the warship may not be operated safely and its systems may not be maintained safely and/or maintained to be safe. The logistic support system is responsible for ensuring that the warship is provided with the spare parts for its systems. Unless the spare parts are kept up to the warship, some of its systems will become unsafe in themselves and may lead to unsafe behaviour by other systems. Amongst other things, the engineering support system is responsible for ensuring that all modifications made to the warship during its life are acceptably safe in themselves and do not otherwise cause the warship to become unsafe.

Tactical publications determine how the warship is used to carry out tactical tasks. They have a bearing on the safety of the warship as it carries out its tasking.

The tasking of the warship also has a bearing on its safety, especially when not all systems comprising the warship are able to exercise their full capabilities at the time. For example, where an assigned task causes the warship to operate in conditions for which the crew does not have sufficient operators with the right competencies, the tasking may put the warship at risk.

Figure 2 illustrates the types of systems comprising a warship. It is illustrative rather than exhaustive and rigorous.

The systems illustrated in Figure 2 are of varying complexity and capability. They also have varying contributions to the safe performance of the warship's tasks.

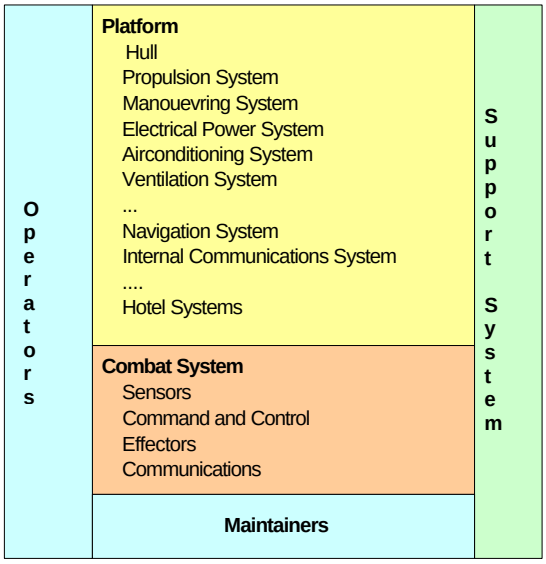

Figure2: Some Components of a Warship

Each of the systems identified in Figure 2 consists of lower-level configuration items. Some of the lower-level configuration items are themselves complex real time systems. For example, a fire control system consists of one or more computer-controlled radars, each with a control console and a radar director used to point the radar antenna at the object to be tracked, a computer for calculating where to aim the weapon and a means of

controlling the launch of the weapon. Missile fire control systems typically have additional subsystems used to provide information that the missiles can use to acquire the target against which they are fired.

Some systems are used whenever the warship is performing any of its tasks. Since these systems are essential, they have built-in redundancy and/or backup systems. The propulsion system is an example: a warship cannot afford to have a single point of failure in its propulsion system.

Some of the systems are not essential to all of the warship's tasks. For example, the under water sensors are not required for fisheries protection tasks.

Some tasks can be carried out when some systems are unable to make their expected contribution to safe conduct of a task provided that a substitution of capability can be made. For example, it the radar used to detect other vessels is unavailable, another radar can be used or, if the weather conditions are suitable, an electro optical system can be used.

Establishing a safety case for the platform is relatively easy. Classification societies establish the safety of merchant shipping, and a similar approach can be used for warships. There are some restrictions on the use of standards and criteria due to the differences between warships and merchant ships.

Establishing a safety case for the combat system is a more difficult task. Hazards that may arise due to the way that components of the combat system are built are relatively easy to identify and manage. Most of the classes of workplace hazards due to the physical design and construction of combat system equipment have been established or have commercial equivalents. Examples of hazard classes are radiation hazards, acoustic hazards, hazardous materials and noxious gas hazards. Most of the hazard classes can be treated by requiring that the equipment satisfies relevant standards. The engineering process can then be used to ensure that the standards are met and residual risk is reduced to an acceptable level.

However, establishing a safety case for the behaviour of the combat system is much more difficult and expensive. It may also be an endeavour that does not improve the safety of the warship to any significant extent.

The difficulty in establishing that a warship is acceptably safe when performing the tasks that it is intended to do is illustrated by consideration of the number of subsets of tasks and the number of subsets of the system capabilities due to environmental degradation and equipment failure. If there are n subsets of tasks and m subsets of system capabilities, potentially there are n factorial m partial safety cases to consider. The actual number will be fewer than n factorial m because not all combinations of task elements are disjoint.

Even if a safety case is developed to prove that the warship is acceptably safe when performing any subset of its tasks with certain allowable combinations of degraded capabilities, the worth of the effort is questionable unless it can affect crew behaviour at sea. Warships in the Royal Australian Navy have traditionally had system safety programs that established that the warship is acceptably safe under some nominal conditions of use. The responsibility for handling safety risk in other conditions of use rests with the crew. This presupposes that the crew is equipped with adequate information and processes to make such decisions.

The current way of handling the complexity inherent in operating and maintaining warships is analogous to that used for large transport vehicles: that is the operators and maintainers are licensed in some way. The primary difference is that the navy is self-regulating. The operators and maintainers of warships attain qualifications relevant to the tasks that they perform. The competencies required of each operator and maintainer are established, as far as reasonably practicable, and training courses are set up to instill the competencies. Continuation training is provided to avoid skill degradation. A continuous improvement program is conducted on the training courses, and circumstances not covered by the initial training can be included as they are exposed. On-the-job training is used to supplement formal training and ensure that competencies are retained.

The previous section outlines some reasons why arguing that a complex reactive real time system is acceptably safe is a difficult proposition. The following section uses a response to tasking model to investigate what is required to achieve proof of acceptably safe behaviour.

5 Response To Tasking Model

A response to tasking model is used to illustrate what is required to support safe operation of a complex reactive real time system. The model is based on a simplified version of decision making on a warship. The model takes account of the tasks required, the external environment and the material state of the system.

A tasking model was chosen as the basis for discussion since the tasks that the system must perform are the source of the requirements that define the:

a. system, including its components,

b. capabilities of the components,

c. interfaces among the components, and

d. processes that individual components and aggregations of components must perform.

An advantage of the response to tasking model is that it complements the development of compositional safety cases as described in Kelly (2001). The modules of the safety argument correspond to the system components.

Although the model does not explicitly address safety, it can be used to determine whether tasks can be performed safely. To do this, it is necessary that the:

a. safety-relevant behaviour contained in each task is identified, and

b. contribution of the capabilities of each system component to safe performance of tasks is also identified.

The foundation of the response to tasking model consists of a set of tasks, decomposed into element tasks, and an hierarchical decomposition of the system into its components. Although represented as a tree structure in the figure, the set of tasks is not a true tree structure because task elements may have more than one parent.

Each of the task elements can be associated with one or more system components and each system component can be associated with one or more task elements.

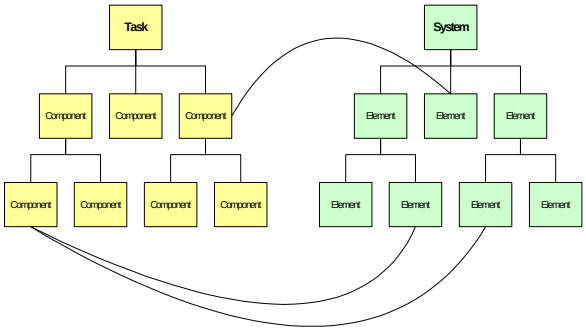

Figure 3: Foundations of Decision Making

Each task element has the following attributes:

a. An identifier.

b. Linkages to higher-level tasks.

c. Potential hazardous behaviour involved in the task.

d. Procedures for carrying out the task.

e. Safety rules relevant to the task.

f. Association with capabilities required to carry out the task element.

Each system component has the following attributes:

a. An identifier.

b. A linkage to the parent component.

c. A list of capabilities.

d. For each capability, the performance in ideal environmental conditions.

e. For each capability, the environmental factors that affect performance.

f. A list of failure modes.

g. For each failure mode, whether the safety risk of the component per se is increased and the extent of any increase.

h. For each failure mode, the affected capability and the extent of performance degradation of that capability.

Although the system is a tree structure of components, some of the components contribute capability to functions associated with more than one part of the tree. For example, a radar used to detect objects on or near the surface of the sea may be used for navigation as well as surface contact detection in the combat system. In addition, some system components are used to provide secondary capabilities if the component providing the primary capability fails.

The decision-making process (Figure 4) is supported by information repositories. Some of these repositories contain data alone and some contain data, models and rules. The data is along the lines outlined above.

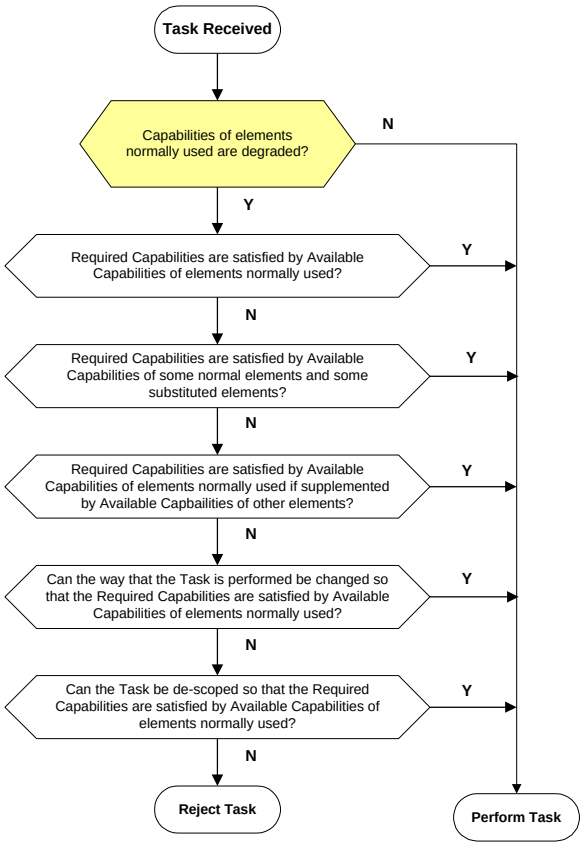

Figure 4: Decision Making

In conjunction with the task information and the system information outlined above, the decision-making process uses rules for preserving as much capability as possible. The rules are supported by strategies such as substituting components for components with degraded performance and/or supplementing components with degraded performance with components that have complementary capabilities. The decision-making process also has rules that use the task element information about potentially hazardous behaviour together with safety rules to determine whether the task is carried out in full, de-scoped or rejected. The safety rules are supported by strategies for keeping the behaviour safe.

Figure 5 illustrates that the process of establishing the safety of a new system begins at concept definition. The activities represented in Figure 5 are carried out as a matter of course at the beginning of system acquisition.

The tasks are the basis of the operational requirements of the system. The only unusual step is the distillation of tasks and task elements into a consolidated list of task elements. The consolidated list of task elements is then used to ensure that all of the operational hazards are

identified and all the associated safety factors are documented. This is part of preliminary hazard assessment and results in a Preliminary Hazard List (PHL) or equivalent. The PHL is used to develop a set of safety requirements associated with, and under, the operational requirements. Some of the safety requirements relate to behaviours required for the tasks to be performed safely and some of the safety requirements relate to behaviours that should not occur if the tasks are to be conducted safely. The operational requirements and safety requirements are used to generate system requirements. The operational requirements are needed for top-level requirements verification as part of the system engineering process and the safety requirements are needed for the safety argument.

A useful approach to developing safety-related requirements from consideration of system tasks is described in Allenby and Kelly (2001).

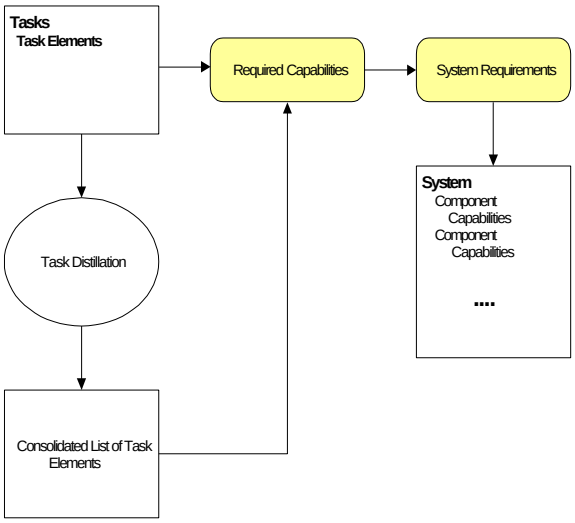

Figure 5: System Definition

The system requirements are decomposed into requirements for System Level Configuration Items (SLCIs), Hardware Configuration Items (HWCI) and Computer Software Configuration Items (CSCI) as necessary to encapsulate the construction and behavioural attributes of the overall system. The complex system acquired to satisfy these requirements consists of an hierarchy of system components that correspond to configuration items.

The process illustrated by Figure 5 results in the hierarchies represented by Figure 3. The linkages between tasks and system components are available through requirements traceability generated as part of the normal system engineering process.

The process illustrated by Figure 5 requires robust record keeping. Use of databases for task definition, also known as operational concept definition, and requirements definition, design management and test management is essential.

Extracting the capabilities of the system and its components and associating them with tasks is an activity that is required to support safety of complex reactive real time systems. This is a process that begins at acquisition, but continues throughout the acquisition and support of the system. As the system is developed, the capabilities delivered by the developer may vary from those required, and during its life the system may be modified to add capabilities. Not only should the association of capabilities with systems and system components be revised when changes occur, but the task and task element information should also be revised.

During the development of the system, failure mode analysis is carried out. This is an essential part of establishing that the system is safe. In addition to determining whether system components will be safe under certain failure modes, the results of failure mode analysis can be used by the decision-making process to determine the effects on the availability of the system's capabilities to support safe performance of tasks.

Hazard identification is also carried out during system development. Hazards that relate to system behaviour are of special interest to this model. The results of the hazard analysis and treatment are used in the engineering process to reduce the safety risk to an acceptable level. Results of hazard analysis can also be used by the decision-making process, particularly in relation to the effects of safety controls that are degraded or unavailable.

Figure 6 depicts a decision support system for a complex reactive real time system. The figure is illustrative only. It simplifies the decision-making process used, for example, by the crew of a warship. The decision-making process identified in Figure 6 implements the flowchart shown in Figure 4.

A decision support system that assists with the safe performance of tasks merits further consideration, as does the contribution to such a system that the safety case and the evidence it contains may make. However, both topics are beyond the scope of this paper.

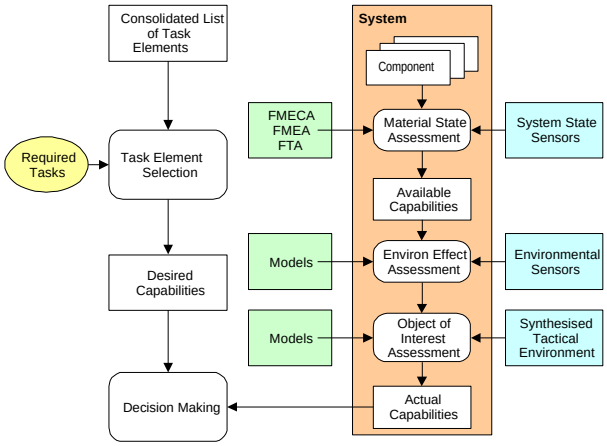

Figure 6: Decision Support

Currently, the decision-making process is carried out on warships by trained and experienced crew. Much of the information used by the crew is based on training and exercises, although some is documented in publications carried on board. The information is generally not

formalised in information libraries. The use of models to determine environmental effects is also limited. An example of model use is determining the propagation of sound through the water. It is one of a few. Most of the environmental effects are assessed by the crew based on informal or semi-formal rules. Formal advice is available in some cases.

Currently, the ability of the warship to perform its tasks safely largely rests on the crew. As automation increases and crew numbers decrease, this will become less tenable because:

a. the crew will have less exposure to the information on which decisions are made,

b. there will be fewer people who have the requisite skills, and

c. communication of the knowledge and skills will become less certain.

There will also be a secondary effect. As the number of people with the knowledge and skills to make the safety-related decisions decreases, there will be fewer people with the requisite domain experience to advise the developers of the automated systems that will replace the people.

It is quite clear that the there is a need for a comprehensive and robust understanding of how complex reactive real time systems can be managed to perform their tasks safely at all times. It is not a simple problem, and it needs attention whilst people with the right domain knowledge and skills can still be consulted.

The contribution of the safety case to this understanding also merits consideration. However, such consideration is beyond the scope of this paper, except to say that parts of the argument and parts of the evidence from the safety case appear to be useful in the safety decision support system.

6 Conclusion

The physical design and construction of complex reactive real time systems can be established as acceptably safe through current practices. The ability of this class of systems to perform its tasks safely requires that its behaviour is safe in all circumstances. This is a difficult proposition even when only a representative set of conditions is taken into account. When combinations of tasks and performance degradation are taken into account, the complexity is increased significantly. For warships, this complexity handled by employing qualified operators and maintainers.

For safety cases to be useful as a means of determining the safety of using such systems in the differing conditions under which they may be used, safety cases need to take account of the flexibility and adaptiveness inherent in these systems. If they don't, they will not adequately contribute to supporting decision-making for the systems' entire conditions of use. They may also lead to a false sense of security that the system is safe in all circumstances. For warships this is particularly important because the safety case relies on the crew making appropriate decisions about safety as they carry out their assigned tasks.

The cost-benefit of compiling a complete safety case for a warship needs examination. A safety argument for a warship for a nominal set of conditions is of benefit, provided that the nominal conditions are suitably representative. The structure of such a safety argument ought to be the hierarchy that results from the engineering process. Safety cases for components of the system should be prepared, provided that the suitable information/evidence is available to support them. Where the evidence is insufficient to support robust proof of acceptable safety risk, a lesser safety argument may need to be used, provided that the overall safety risk is at least bounded and the bounds fall within the acceptable limits.

To move to a safety case supported by an on-line safety management system that supports the safe operation of complex reactive real time systems with humans in the loop will require a great deal of rigor in the definition of the tasks that the system is to perform, including the safety factors of those tasks. It will also require a great deal of rigor in the safety analysis of the components of these systems, including the effect on safety of system degradation. It will further require an increased level of sophistication in the arguments that are used, as well as a change to the presentation methods so that the relevant information from the safety case can be used by the personnel carrying out the task. Unless this is done, there will be cynics who regard the safety case as a curiosity of the procurement process and a burden for the ongoing maintenance of the systems.

As autonomous complex reactive real time systems become more prevalent, safety management within these systems will require an approach similar to that outlined within this paper. For such systems to be proved to be safe, the safety case will need to ensure that the information is accurate and complete and that the decision-making component of the system can perform its functions correctly and without error.

The format, content and presentation of safety cases need to be reviewed in the context of complex real time systems.

7 References

Allenby, K and Kelly, T. Deriving Safety Requirements Using Scenarios. *Presented at the 5th IEEE International Symposium on Requirements Engineering (RE'01), proceedings published by IEEE Computer Society Press.*

Kelly, T. (2001): Concepts and Principles of Compositional Safety Case Construction. *(Contract Research Report for QinetiQ COMSA/2001/1/1),* Department of Computer Science, University of York (available from www.cs.york.ac.uk/~tpk/pubs.htm)

Kelly T, Bate I, McDermid, J and Burns A. (1997): Building a Preliminary Safety Case: An Example from Aerospace. *Proceedings of the 1997 Australian Workshop on Industrial Experience with Safety Critical Systems and Software, Australian Computer Society, Sydney, Australia, October 1997.*