Compliance with Standards or Claim-based Justification? The Interplay and Complementarity of the Approaches for Nuclear Software-based Systems

Sofia Guerra and Dan Sheridan

Adelard

London, UK

Abstract The control and protection of nuclear power plants has become increasingly dependent on the use of computers. The UK nuclear regulatory regime requires that a safety case be developed to justify and communicate their safety. There are several ways of constructing such a safety case. In the past, safety justifications tended to be standards-based – compliance to accepted practice was deemed to imply adequate safety. Over the last 20 years, there has been a trend towards an explicit claim-based approach, where specific safety claims are supported by arguments and evidence at progressively more detailed levels. These approaches are not mutually exclusive, and a combination can be used to support a safety justification. In fact, for the most critical systems it can be argued that a safety case should consider both aspects. For less critical systems, one might believe that one approach would suffice. This paper discusses software-based systems with only a modest integrity requirement, and the interplay of the two approaches. It describes our experience with justifying such systems for the nuclear industry, and it claims that there are a number of benefits of taking both approaches together.

1 Introduction

Programmable components including personal computers (PCs) or smart instruments can offer considerable benefits in terms of usability, functionality and possibly price in a safety-related system. For example, the nuclear operators are increasingly replacing analogue instruments with their digital 'smart' counterparts, as they achieve greater accuracy, better noise filtering together with in-built linearisation, and provide better on-line calibration and diagnostics features. Similarly, there are usability and functionality advantages to using PCs for some applications

where the reliability requirement is fairly modest. These systems are typically not nuclear specific, or are based on components that are not nuclear specific (such as general purpose operating systems). Nevertheless, if the application is considered to have an impact on safety, the licensees are required to justify that level of safety achieved.

There are often difficulties with the safety justification of this type of system. This might be due to all or part of the system being sold as a 'black box', where the purchaser has no knowledge of the internal system structure or the development processes followed by the manufacturer. PCs and their operating systems are typical examples of this. Difficulties might also arise from the age of the component, since expectations of 'best practice' have changed over the years; even if a component was developed in accordance with best practice ten years ago it may not meet current expectations.

This paper discusses the justification of software-based systems with a modest reliability requirement (e.g., a probability of failure on demand of 10^{-1} or 10^{-2}, or SIL 1) in the nuclear industry in the UK. It builds on our experience of justifying a number of these systems, and discusses some of the issues encountered and how they have been addressed.

In Section 2, we summarise the UK nuclear regulatory regime. Section 3 describes the range of justification approaches and gives examples of how they are applied in the UK nuclear industry. In Section 4, we discuss the use of a combined compliance and claim-based justification approach, and give examples of how this approach has been beneficial to us. We draw conclusions from these case studies in Section 5.

2 UK nuclear regulatory regime

The UK has a clearly defined approach to how the assessment and licensing of command, control and protection systems should be carried out. There are still significant differences between the UK and other countries, despite the internationalization of the supply chain and efforts such as:

- effective collaboration with international agencies such as IAEA and OECD
- standards committees including the IEC
- working groups such as the Nuclear Regulators Working Group (NRWG)
- projects to encourage harmonisation such as CEMSIS (CEMSIS 2004) and HARMONICS (European Commission 2011).

In the UK, nuclear licensees are required to produce safety cases in order to comply with the licence conditions (HSE 2011). These are designed to be non-prescriptive, and instead set high-level goals for the nuclear operators. The Safety Assessment Principles (SAPs) (HSE 2006) are the primary principles that define the approach to assessment to be followed for nuclear installations in the UK. The SAPs were revised in 2006 and have been brought into line with IAEA guidance.

The SAPs have the following clauses on computer-based safety systems (our emphasis added).

> Where the system reliability is significantly dependent upon the performance of computer software, the establishment of and compliance with appropriate standards and practices throughout the software development life-cycle should be made, commensurate with the level of reliability required, by a demonstration of '*production excellence*' and '*confidence-building*' measures. [(HSE 2006) clause ESS.27]

> '*Production excellence*' requires a demonstration of excellence in all aspects of production, covering initial specification through to the finally commissioned system, comprising the following elements:

> a) Thorough application of technical design practice consistent with current accepted standards for the development of software for computer-based safety systems.
> b) Implementation of an adequate quality assurance programme and plan in accordance with appropriate quality assurance standards.
> c) Application of a comprehensive testing programme formulated to check every system function. [(HSE 2006) clause 360]

> Independent '*confidence-building*' should provide an independent and thorough assessment of a safety system's fitness for purpose. This comprises the following elements:

> a) Complete and preferably diverse checking of – the finally validated production software by a team that is independent of the systems suppliers, including:
>   – independent product checking providing a searching analysis of the product;
>   – independent checking of the design and production process, including activities needed to confirm the realisation of the design intention;
> b) Independent assessment of the test programme, covering the full scope of test activities. [(HSE 2006) clause 361]

> Should weaknesses be identified in the production process, *compensating measures* should be applied to address these. The type of compensating measures will depend on, and should be targeted at, the specific weaknesses found. [(HSE 2006) clause 362]

In the UK, we refer to this approach of demonstrating production excellence and confidence building measures as the *two-legged approach*. The expectation for both legs is that the depth of the assessment is 'commensurate with the level of reliability required'. For example, for systems with a modest reliability requirement, the production excellence leg is typically supported by:

- good commercial practice
- development process evidence
- compliance with international quality assurance standards.

The independent confidence building measures for a modest reliability system might include:

- commissioning tests to demonstrate that the instrument adequately performs its required functions
- data on prior use
- evidence of the manufacturer's pedigree.

In this paper we argue that, even for systems with modest integrity requirements, the production excellence leg needs to be developed using both a process and a product-based approach, in the sense described in Section 3 below.

3 Justification approaches

There are two principal ways of constructing a safety justification (Bishop et al. 2004). A *process-based* approach focuses on the development process and defined standards and practices, while a *product-based* approach focuses on the behaviour required of the system. Below, we describe these approaches in more detail and give examples from the UK nuclear industry.

3.1 Process-based approach

A widely used approach to justifying a Control and Instrumentation (C&I) system and its software is to provide evidence that they have been designed and verified following a well-structured development process and in accordance with the requirements and recommendations of rigorous standards. For example, when the justification is based on compliance with a standard such as IEC 61508, assessors argue that the system is acceptably safe by showing that the development process followed is consistent with that described in the standard and by applying a set of techniques and methods that the standards associate with a specific safety integrity level. This type of justification approach is also called a *rule-based* or *standards-based* approach, as it relies on the application of pre-defined rules on the development process and on compliance with standards.

3.1.1 Process-based approaches in the UK

There are a number of different standards or regulations that can be used as a basis for the justification of a software-based system. This varies with the type of system, reliability requirement and organization. For example:

- COTS products that have not been developed for the nuclear industry may be assessed against IEC 61508 (Stockham 2009).
- Nuclear specific standards may be used, such as IEC 61513 or IEC 62138.
- Organisations may have their internal standards, which will typically be based on an international standard adapted to the system and reliability requirements to which they will be used.

All of the standards and approaches cover a number of common areas including:

- configuration management
- the approach to development, including requirements, design and implementation
- design constraints or design principles
- the level of verification that may be considered adequate
- the documentation that is to be produced during development.

3.1.2 Criticism of the process-based approach

The process-based approach works well in stable environments where best practice is deemed to imply adequate safety. However, it is often criticised for being highly prescriptive and impeding the adoption of new and novel methods and techniques. The fundamental limitation of process-based safety justification lies in the observation that good tools, techniques and methods do not necessarily lead to the achievement of a specific level of integrity. The correlation between the prescribed techniques and the failure rate of the system is often infeasible to justify. Conversely, process-based approaches are also inadequate where high-quality systems were developed in accordance with older or different standards, or just meet industrial good practice but fall short of strict compliance with standards.

A purely process-based approach does not necessarily provide direct evidence that the C&I system and its software achieve the behaviour or the properties required to the desired level of reliability. However, despite these uncertainties, some might argue that process-based evidence is sufficient for systems of modest integrity requirements. In the next sections we argue that this not the case.

3.2 Product-based approach

To overcome the difficulties with the process-based approach, we can focus instead on directly justifying that the desired behaviour, property or reliability has been achieved, using product-specific and targeted evidence. This type of approach can be called the *product-based* or *property-based* approach.

A product-based approach allows the focus to be directly on the safety requirements for the system, making it applicable even when a standards compliance case cannot be made. This is often the case for off-the-shelf devices (such as smart instruments), where development follows industrial good practice and does not necessarily conform to a recognised safety lifecycle.

Alternative evidence can be presented to justify the safety properties depending on the characteristics of the device being justified and the process followed. For example, in applications with limited safety significance, extensive field experience for a device may provide alternative evidence of compliance to the required performance, or alternative arguments can be used to justify expected behaviour when process non-compliances might have been identified. Evidence can also be

related to a range of different safety standards by identifying how the require-ments in the standards support the various claims. This allows greater flexibility in making a justification while ensuring that all safety relevant attributes of the system are justified.

This approach is usually linked with specific claims about the product or system being justified, and may be *claim-based (*or *goal-based)*. This can follow a structured approach such as *Claims-Arguments-Evidence* (CAE) (Bishop and Bloomfield 1995) or *Goal Structuring Notation* (GSN) (Kelly 1999). The CAE approach was developed as part of a research project funded by the UK nuclear industry (Adelard 1998). It forms the basis for the generic design assessment currently being performed in the UK for new nuclear build, and it is used for justifying refurbishment projects in existing plants.

3.2.1 Product-based approaches in practice

In this section, we summarise two techniques used by Adelard in constructing product-based safety justifications. The main characteristic of these two tech-niques is that they focus on the system behaviour, rather than on the way it was developed.

Behavioural attributes. The justification of the system's behaviour may be done by arguing that specific high-level claims are met. These may be based on a set of *behavioural attributes* which allow us to provide separate arguments focusing on specific aspects of behaviour. Table 1 gives an example of behavioural attributes that have been used to justify an FPGA-based system (Guerra and Sheridan 2012).

Table 1. Example attributes

Category	Attributes	Discussion
Functionality	Functionality	The function performed by the system
Performance	Timing	Includes time response, permissible clock frequencies, propagation delays, etc.
	Accuracy	Affected by analogue/digital conversion, processing functions, IP cores, etc.
Availability	Availability	Readiness for correct service, a system-level attribute supported by component attributes such as reliability
Reliability	Absence of faults	This may be connected with a vulnerability analysis
	Fault detection and tolerance	Internal detection of faults
Robustness	Robustness	Tolerance to out-of-normal inputs and stressful conditions
Failure recovery	Failure recovery	The ability to recover from failures through error detection and reporting, such as sounding an alarm

LowSIL approach. An approach developed for the nuclear industry to justify systems that used *Non-Safety Assured Programmable Components* (NSPCs) that

cannot be directly justified, e.g., PC-based systems running Microsoft Windows (Bishop et al. 2010). This guidance is intended for use when:

- Failure of the system can affect nuclear safety, environmental protection or industrial safety, the integrity of plant actuations, safety-related information presented to operators, or the safety integrity of components or calibration data that will be used in the plant at some time in the future.
- The required integrity of the system safety function is at or below SIL 1. Typically the approach is used for an integrity of no more than 10^{-1} failures per demand or 10^{-4} dangerous failures per hour.
- The system contains one or more NSPCs (such as a PC, a programmable logic controller, or a configurable device such as a smart sensor), and there is insufficient assurance of the NSPC's safety integrity.

The main assessment strategy is to perform a hazard analysis to identify how the NSPC may fail, and how its failure could lead to a hazard for the system containing the NSPC. This analysis includes performance failures such as slow response. The impact of the failures is assessed together with the effectiveness of existing mitigations, and additional controls or mitigations may be recommended. These may result in design changes and additional requirements on the system surrounding the NSPC or, if technical mitigations are not possible, they may be procedural.

This approach has been used in a range of C&I systems in nuclear plants as a means of demonstrating adequate safety assurance when NSPCs are used.

3.3 Combining justification approaches

The process and product based approaches are not mutually exclusive, and a combination can be used to support a safety justification. This might be particularly important where the system consists of both off-the-shelf components and application-specific elements, for example. We have applied the integration of the approaches in a number of assessments conducted on behalf of the nuclear industry, including:

- a justification for smart instruments (Guerra et al. 2010), where the requirement for production excellence in the SAPs was met by demonstrating compliance to an acceptable standard, while the requirement for independent confidence-building measures was addressed by a range of claim-based assessments (e.g., demonstration of accuracy, reliability, etc.) together with assessments of potential vulnerabilities in the smart device implementation
- a justification of the first safety-related FPGA-based system deployed in the UK (Guerra and Sheridan 2012), based on a combination of approaches.

We have found that there are significant advantages in using both approaches together for systems with a modest integrity requirement, as they provide flexibility, understanding and documentation of the system behaviour that is commensurate

with the reliability requirement. In the next section, we review our experiences and discuss the advantages of the combination of approaches.

4 Discussion and experiences

To illustrate the effectiveness of using process and product based assessment techniques together, we describe below a selection of situations in conducting assessments. These situations are illustrations of the range of assessments we perform in practice; they have been modified and consolidated to show the type of issues encountered.

Assessment constructed during development. In this situation, the assessment work starts early in the lifecycle of the product development, and the work on the justification is fed into development process. The process-based approach establishes a clear need for documentation, which is especially valuable as the supplier might not be familiar with nuclear requirements. Having a clear set of documentation generated during the development process enables a stronger product-based justification to be constructed.

Like-for-like replacement. In this case, the stakeholders involved in the project may not be aware of the value of comprehensive requirements documentation, since the aim of the project is to replace old equipment with new technology following equivalent requirements. However, this type of replacement is seldom truly 'like-for-like', as new technologies, and the opportunity for operational improvements, are sources of new requirements that may interact with the existing plant in unforeseen ways.

In a replacement project, it is often the case that the documentation on the surrounding system is limited, and therefore tests or other types of analysis may be required to fully understand the impact that the system under development has on the overall plant; this can be performed using the LowSIL approach (see Section 3.2.1). By also following a process-based approach to justification, we are able to make the case for documenting the interfaces between the plant and the system, as well as any additional requirements and design decisions, all of which are crucial for the operation and maintenance of the system. These aspects may otherwise be overlooked, especially given the modest integrity target.

Smart instrument with extensive field experience. COTS products which have been in the market for many years are often seen as trustworthy items. We have found on a number of occasions that adopting a process-based approach to assessing older instruments is prone to difficulties. This can be because a development process which was considered good practice 20 years ago is no longer consistent with current standards, because documents have been archived or lost, or because staff members involved in the development are no longer available to support the assessment.

In some cases, the manufacturer may be able to supply some development process evidence, but not enough to provide sufficient confidence in the product. We could adopt an approach based partly on claims about compliance with specified device behaviour such as functionality, time response or robustness to abnormal inputs. Some of this evidence could be drawn from assessment techniques such as static analysis and black-box testing, while evidence of field experience and field-reported faults could also be analysed to demonstrate reliable operation.

Smart instrument with haphazard development process. At first glance, it might be difficult to distinguish this case from the earlier case of a smart instrument with lost or unavailable process evidence. Similar documentation may be missing for both instruments. However, given a more general understanding of the process followed, it is possible to determine whether a disciplined lifecycle process has been attempted, regardless of the available evidence. A lack of configuration management calls into question the validity of otherwise promising product-based evidence such as analysis and testing, since it is difficult to connect it up to the right part of the development process and the right version of the instrument. In this case, although we may attempt to complement a problematic process-based assessment with a product-based assessment, there may be sufficient doubt in the process case to jeopardise the product-based case.

This suggests some minimum process requirements needed to successfully complete a product-based assessment, despite their apparently different aims: an understanding of the development process followed, and confidence in the configuration management approach. The importance of configuration management in a claim-based approach is widely accepted as being needed to establish trust in the evidence (CAA 2013).

6 Conclusions

In the discussion above, we identified situations where we benefitted from using both a process and product-based assessment approach in conjunction, when developing a production excellence argument. This combination brings a number of advantages, described below.

Understanding hazards and properties in more detail. An approach based on understanding and mitigating the hazards of the system being developed on the overall system of which this is part of, such as the LowSIL approached described in Section 3.2.1, gives one an understanding of the impact of the failures that would be difficult to grasp if a pure compliance process-based approach was followed. In our experience, following this approach not only provides a way of dealing with NSPCs, but also a way of deriving system level requirements and mitigations that could easily be missed. Similarly, focusing on behavioural properties of the system gives you direct evidence of the product, which addresses some of the weaknesses of a pure compliance process-based approach.

Documenting process and requirements. For systems where the justification is developed as the system itself is being designed and implemented, the process-based approach establishes the need for specific documentation that will be useful both in the safety justification and for operation and maintenance. In a like-for-like replacement, the process-based approach can help to uncover requirements and design evidence, while a product-based approach can help by identifying hazards and the system-level impact of the replacement system.

The need for configuration management and process understanding. Clearly, for a purely behaviour based case, we would not need any claims about the process. Process evidence would be incorporated, if appropriate, through its impact on the claims about the product behaviour. Nevertheless, there are minimum process procedures and evidence that are essential. The most important of these is configuration management and other quality assurance activities that support the production of traceable and consistent evidence. A key aspect of a good safety case is that it should be supported by trustworthy evidence. In the cases where process evidence is limited, it is important to understand the reasons for these deficiencies; evidence may not exist because the component was developed a while ago, when best practice did not require the extensive documentation that is considered necessary today, or because an informal development process was followed, with limited planned verification and documentation. Understanding the process is important to be able to differentiate these cases. Demonstration that a sound development process and quality assurance principles have been followed increases our confidence in all the evidence generated to support the safety case, and consequently in the justification as a whole. This is particularly true when supporting evidence is limited.

In conclusion, our experience shows that even for systems with a modest reliability requirement, combining two styles of assessment, process-based and product-based, provides a number of advantages over either style on its own. These advantages are further reaching than simply increased confidence in the safety justification, extending to include a better understanding of the system in its context.

References

Adelard (1998) Adelard Safety Case Development Manual (ASCAD). http://www.adelard.com/resources/ascad/. Accessed 15 November 2013

Bishop P, Bloomfield RE (1995) The SHIP safety case – A combination of system and software methods. SRSS95, Proc 14th IFAC Conf on Safety and Reliability of Software-based Systems, Bruges, Belgium, 12–15 September

Bishop P, Bloomfield R, Guerra S (2004) The future of goal-based assurance cases. In Proceedings of Workshop on Assurance Cases, supplemental volume of the 2004 International Conference on Dependable Systems and Networks, pp. 390–395, Florence, Italy, June 2004

Bishop P, Tourlas K, Chozos N (2010) Programmable components in modest integrity systems. In: Erwin Schoitsch (ed.) Computer Safety, Reliability, and Security, 29th International Conference, SAFECOMP 2010, Vienna, Austria, September 14–17. Springer

CAA (2013) SW01 – Regulatory objectives for software safety assurance. CAP 670 Air traffic services safety requirements. CAA Safety Regulation Group

CEMSIS (2004) CEMSIS – Cost-Effective Modernisation of Systems Important to Safety. http://cemsis.org. Accessed 14 November 2013

European Commission (2011) HARMONICS – Harmonised Assessment of Reliability of Modern Nuclear I&C Software. http://cordis.europa.eu/projects/rcn/97431_en.html. Accessed 14 November 2013

Guerra S, Bishop P, Bloomfield R, Sheridan D (2010) Assessment and qualification of smart sensors. 7th American Nuclear Society International Topical Meeting on Nuclear Plant Instrumentation, Control and Human-Machine Interface Technologies (NPIC & HMIT)

Guerra S, Sheridan D (2012) Justification of an FPGA-based system performing a Category C function: development of the approach and application to a case study. 8th American Nuclear Society International Topical Meeting on Nuclear Plant Instrumentation, Control and Human-Machine Interface Technologies (NPIC & HMIT)

HSE (2006) Safety assessment principles for nuclear facilities, Health and Safety Executive, UK. http://www.hse.gov.uk/nuclear/saps/index.htm. Accessed 14 November 2013

HSE (2011) Licence condition handbook – Securing the protection of people and society from the hazards of the nuclear industry. Office for Nuclear Regulation, an agency of HSE. http://www.hse.gov.uk/nuclear/silicon.pdf. Accessed 14 November 2013

Kelly TP (1999) Arguing safety – a systematic approach to safety case management. PhD thesis, Department of Computer Science, University of York, UK

Stockham R (2009) Emphasis on safety. IET Magazine, Issue 02