
Proceedings Paper:
Fenn, Jane, Hawkins, Richard David and Nicholson, Mark (2024) A New Approach to Creating Clear Operational Safety Arguments. In: SAFECOMP 2024 (43rd International Conference on Computer Safety, Reliability and Security):11th International Workshop on Next Generation of System Assurance Approaches for Critical Systems (SASSUR 2024). , pp. 227-238.

https://doi.org/10.1007/978-3-031-68738-9_17

A New Approach to Creating Clear Operational Safety Arguments

Jane Fenn[0000−0002−6239−0177], Richard Hawkins[0000−0001−7347−3413], and Mark Nicholson[0000−0002−0954−4448]

Department of Computer Science, University of York, Deramore Lane, York, England, YO10 5GH
jlf541@york.ac.uk, richard.hawkins@york.ac.uk, mark.nicholson@york.ac.uk

Abstract. The use of Safety Cases has become relatively commonplace, particularly for high hazard industries. Safety cases should provide a compelling argument and evidence to demonstrate that a system is sufficiently safe both in design and in operation. Much of the guidance for developing safety cases has focussed on creating safety cases at design time to support the deployment of a system. Operational safety is significantly less well-handled in current safety case practice. In this paper, to start addressing the challenges of operational safety cases, we propose to extend the ideas of splitting complex safety cases into risk, confidence and compliance arguments to also consider operational safety arguments. We propose that the operational safety arguments should be separate but explicitly connected to the design–time risk argument through the use of ***operational claim points (OCPs)*** to ensure clarity in both the design–time risk argument and the operational argument, whilst still ensuring an explicitly defined relationship exists. We describe how this approach can bring a number of benefits by: 1) ensuring that system operators are able to focus on just the operational aspects of the safety case that are relevant to them (hiding irrelevant and potentially confusing design details); 2) making sure that, at the same time, the crucial relationship between the operational safety case and the design-time risk argument is explicitly documented and maintained (helping operators to better understand the safety impact of their work); 3) allowing design-time safety engineers to specify, in the risk argument, safety claims relating to system operation.

Keywords: Operational safety · Safety case · Safety arguments

1 Introduction

The introduction of Safety Cases for engineered systems is commonly traced back to the findings of the Piper Alpha Public Inquiry in 1990 [4]. This inquiry considered the need for operational safety cases that would provide argument and evidence for the safety of offshore drilling operations. Subsequent guidance around safety cases has been largely written from the perspective of the design

and pre-operation of systems, rather than the operational phase. Indeed, operational safety cases are significantly less well-handled than design-time safety cases and we explore why this might be the case in section 2.

In this paper, to start to address the challenges of operational safety cases, we propose to extend the existing concept of distinct risk, confidence and conformance arguments to also consider operational safety arguments and their supporting operational evidence, by which we mean, any evidence that relates to operation of the system, e.g. operator training, where the evidence could be generated before or after deployment. The Nimrod Review in 2009 [6], strongly criticised safety case practice for being too focused on 'box ticking' rather than risk management. As a partial response to this criticism, in [2] it was identified that safety cases should normally contain arguments relating to how risks are addressed, and arguments relating to confidence measures (such as the soundness of the processes used) as well as arguments relating to compliance with regulations and standards. Further, historically, projects would normally develop a single, unified safety argument that did not distinguish the arguments of safety and confidence. This merged what are essentially different but interrelated arguments. The proposal in [2] was therefore to split safety case arguments into three distinct 'legs' of argument, namely :

- A ***risk argument*** that records the arguments and evidence used to establish direct claims on the acceptability of safety risk
- A ***confidence argument*** that justifies the sufficiency of confidence in the safety risk argument
- A ***conformance argument*** that justifies belief in conformance with the requirements of a standard or regulation

The advantage of splitting the safety case into separate arguments in this way is that it improves clarity by simplifying the core risk argument and ensures that the role of all claims in the argument is clear; everything in the risk argument must have a direct role as part of the causal chain to a hazard and confidence arguments can only refer to elements of the risk argument. Careful attention to linking the separate arguments provides a mechanism for guiding analysis of the interrelationship between risk, confidence and compliance without overly complex and bloated safety arguments.

To represent these relationships between the risk and confidence arguments, the concept of Assurance Claim Points (ACP) was developed [7]. These represent specific points in the risk argument where a confidence argument is required. For example, if a safety argument considers the adequate mitigation of all hazards, a confidence argument is needed to demonstrate why it is believed that all hazards are identified. An ACP can be used to indicate where the confidence argument is used to support the risk argument, thus enabling the creation of a confidence argument that is separate but explicitly connected to the risk argument.

We propose that the operational safety arguments should (in a similar manner as for confidence arguments) be separate but explicitly connected to the design-time risk argument (through the use of our proposed *operational claim*

points (OCPs)) to ensure clarity in both the risk argument and operational argument, whilst still ensuring an explicitly defined relationship exists. This approach brings the following benefits:

- It ensures that system operators are able to focus on just the operational aspects of the safety case that are relevant to them (hiding irrelevant and potentially confusing design details).
- It makes sure that, at the same time, the crucial relationship between the operational safety case and the design-time risk argument is explicitly documented and maintained (helping operators to better understand the safety impact of their work).
- It allows design-time safety engineers to specify, in the risk argument, safety claims relating to system operation, as discussed in Section 4.

The paper is structured as follows. Section 2 explains our rationale and motivation for this proposal. Section 3 describes our proposal and Section 4 uses examples to illustrate our proposed approach to creating operational safety arguments. Section 5 provides concluding remarks and planned further work.

2 Related Work

The transition between design time safety case and the operational safety case is not well-described in literature, standards, or guidance. Logical progression from early design safety cases, such as 'Preliminary' to 'Interim' are described as a maturation and expansion as the design matures, [9], but the shift to 'Operational' safety case is not well explored. It is not clearly articulated whether this is a further maturation of the design case or something quite different. A number of publications actually indicate that it is a change in viewpoint and paradigm that is actually occurring at this boundary.

Hayes, [8], describes that *"Operational managers rarely use risk-based concepts as a way of thinking about specific situations or deciding on actions ... For them, safety is an active concept. Actions focus on two aspects: Compliance with rules; Ensuring sufficient integrity of the barriers that prevent a specific hazard from becoming a reality."* . Likewise, Green [5] contextualises operational decision making as *"Exact solutions are information intensive and require specialist knowledge or machinery beyond the capabilities of generalist engineering managers, often compelling decision-makers to use their subjective judgement in an unsupported way."*

Acfield and Weaver [1] also recognise this change in paradigm and go so far as to conclude that the safety case is not the correct approach for risk assessment and decision making in the operational phase, favouring instead the use of Bow Ties and analysis of the barriers in the causal and outcome chains. This reduction of risk evaluation in the operational phase to an assessment of the effectiveness of barriers is recognised as an option for simplifying risk presentation to non-safety experts. Unfortunately, this inherently brings a lack of flexibility due to a lack of understanding of exactly how the barriers contribute to safety or mitigate

hazards. Further, confidence requirements in the barriers may not be clearly communicated. Often, the argument becomes closer in nature to the conformance type argument.

The rationale cited in many of these papers suggest that the design safety argument is too complicated to be used in the operational domain. We contend that, what is needed is a simpler way of presenting the contribution of operational safety measures to the overall safety case and risk evaluation that is needed. We believe that a distinct operational safety case is useful, but it should explicitly link to the argument created in the design phase in a much more transparent and traceable way. We believe the design safety case is enhanced by the clear articulation of dependencies on operational mitigations, such as operating procedures, rather than a loose reference to 'standard operating procedures' which may be otherwise noted as an assumption or context of the design safety case. We believe that the operational safety case should address the dependencies placed upon it by the design safety case, allowing clarity as to how overall system risk is affected by operational mitigations and confidence in their effectiveness. Whilst these may seem like small, semantic changes, bringing together and providing much clearer interfaces between the design and operational safety case will benefit current systems and is also a gateway to facilitate opportunities for more dynamic and potentially real-time risk assessment in the future.

3 Proposed Approach

We propose that the operational safety arguments should be separate but explicitly connected to the design-time risk argument (through the use of our proposed ***Operational Claim Points (OCPs)***) to ensure clarity in both the risk argument and operational argument, whilst still ensuring an explicitly defined relationship exists.

Our proposed approach is that the operational safety arguments should be separate but explicitly connected to the design-time risk argument through the use of OCPs. An OCP represents an interface between the design safety case and the operational safety case. Whereas ACPs are used where arguments of confidence are required, OCPs represent points in the design-time risk argument where operational aspects are required to be considered. OCPs are therefore addressed by arguments with supporting evidence that exists, or will exist, in the operational domain. The operational safety case will then provide argument and evidence for each OCP which is separate from, but directly traceable to the risk argument. For each OCP, the argument and evidence in the operational safety case must relate specifically to the operational aspects associated with the point in the risk argument to which the OCP relates. In this way we can ensure that operational arguments specifically address aspects of risk mitigation in the design safety case (rather than appealing to general claims of "good operation").

Likely uses of the OCP would most commonly be around asserted context or evidence. Asserted context, (which may include assumptions), in the design safety case are declarations within which the system is argued to be safe, so

might, for example, include operating limits of the system such as temperature or pressure. An OCP associated with this asserted context would need to demonstrate how adherence to the system limitations would be achieved through life. The operational safety case should decompose this OCP by arguing and providing evidence of how these operating limits would be observed during operation, through life, typically perhaps through operating procedures that operators are trained to enact. An OCP on evidence in the design safety case indicates that the evidence may be affected by the operation of the system, e.g. by change in personnel and hence operator competency, or system servicing history, so the operational safety case must argue the sufficiency of that evidence throughout the operational phase.

In the next section we illustrate the use of our approach through some examples.

4 Illustrative Examples

It is recognised that evidence necessary to support operational hazard mitigations may not be available in the design phase and only produced in the operational phase itself. We propose that the requirement for operational argument and evidence should be identified in the design phase and captured as part of the design safety case. For example, an Air System Safety Case, is required by regulation [11] to include evidence about the presence of collision avoidance equipment, such as the Traffic Collision Avoidance System (TCAS), fitted to military aircraft. Some evidence to support this conformance requirement might be available pre-operation such as the inclusion of TCAS in the design standard for the aircraft. Other relevant evidence may not be available at design time. For example, manufacturing evidence records that the particular aircraft has the equipment fitted and there are, for example, no reported equipment failures. There may also be maintenance requirements to keep the equipment serviceable that are only available from in-service records. There could also be pre-flight and post take-off tests that are relevant to this equipment and must be carried out for every flight. All these post-design evidence items are typically not transparent in the design argument itself. Our proposed approach will make this operation-time evidence explicit in the design safety case.

Figure 1, uses an example safety argument structure adapted from [7] to illustrate how ACPs are represented graphically as part of an argument structure using Goal Structuring Notation (GSN) [3] (a key to the GSN symbology is provided in Figure 2). We propose that the same GSN 'decorator' be used to represent where OCPs are present in a safety argument (labelled OCP instead of ACP).

Figure 1 shows three ACPs representing different parts of the risk argument for an insulin pump where a design confidence argument is provided. A confidence argument, supported by suitable evidence is provided for each of the ACPs to demonstrate that sufficient confidence exists in that element of the argument. We have also introduced an OCP (OCP20) on the asserted context

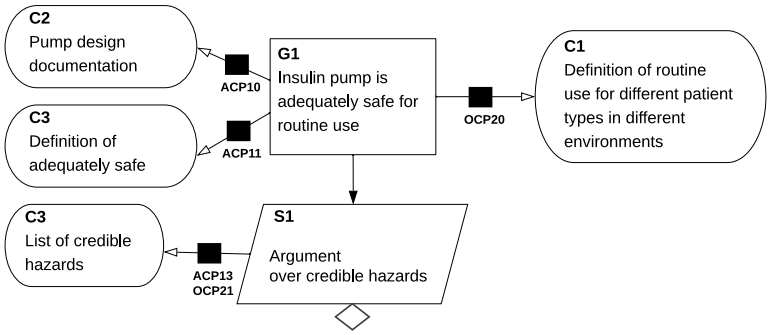

Fig. 1. Extract from design safety case for an insulin pump showing use of OCPs

C1, indicating that an operational argument is required to be developed to show that the context will remain valid throughout operation. This operational argument is shown in Figure 3. We also recognise that at any point in the risk argument both a design confidence argument and an operational argument may be required. This can be seen to be the case in Figure 1, where the context of a list of credible hazards (C3) requires both a confidence argument to be provided (to demonstrate sufficient completeness of the hazards identified at design time) as well as an operational argument (to demonstrate that new hazards arising in operation will be identified if they occur). Both an ACP (ACP13) and an OCP (OCP21) are created to capture the required confidence and operational arguments respectively. This example illustrates how the use of ACPs and OCPs allows complex confidence and operational arguments to be included in the safety case in a traceable manner without distracting from the core risk argument.

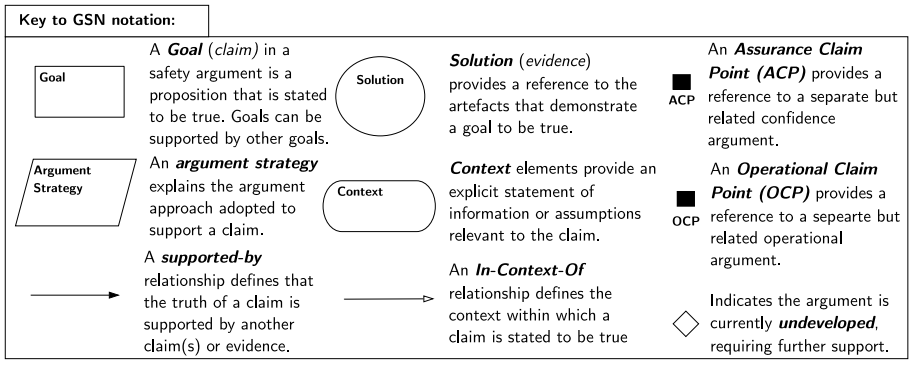

Fig. 2. A key to the GSN symbols used this paper

In Figure 3 we show the operational safety case relating to OCP20 which provides an argument regrading the operational activities undertaken in order to ensure the continued validity of context **C1** from the design risk argument (Figure 1). **C1** defines the expected routine use of the pump for different patient types in different environments. The claim **G1** in the risk argument is only valid within that context. It is therefore important to ensure that during operation, the actual patient types and operational environment are as expected. We add an OCP (OCP20) to show where the operational argument is used to support the risk case. We then provide that operational argument to show how this is demonstrated during operation in Figure 3. Here we split the argument over medically supervised and medically un-supervised environments as the arguments are very different in these cases. In medical settings, competent medical staff enact or supervise insulin pump usage so the confidence in the risk argument comes from the availability and adherence to the hospital's medical procedures. In the medically unsupervised setting, the general approach to building confidence is to provide training to patients and periodic follow-up monitoring to check and, if necessary, correct patient's practice to conform to the intended 'routine use' expected. Training type and follow up schedule are expected to vary by patient type.

Figure 4 shows another example of how OCPs can be used to link an operational argument to the risk argument using the door interlock on commercial aircraft as an example. The design intent is that, if doors are opened in an emergency situation, emergency evacuation chutes should automatically deploy. The system is intended to be 'fail safe' so this emergency chute deployment needs to be de-activated for normal, non-emergency door opening at the end of each flight. Here, we have split the argument over the emergency situation, where the chutes should deploy, and the normal operations where the system needs to be deactivated. The correct functioning of the emergency chute deployment is part of the design safety case so the evidence asserted is argued as sufficient through an ACP. However, this argument and evidence is made in the context of the Emergency Chute System Design (**C31**), whose design includes a requirement for a routine maintenance activity, which is a facet that only applies in the operational phase. So, we add OCP 1 to this context. On the normal operation leg or argument, switching the aircraft doors to manual is achieved, through life, by use of an operational procedure so we label this solution as OCP 2.

We expand on the two operational claims for OCP 1 (asserted context) in Figure 5 and OCP 2 (asserted solution) in Figure 6. As both OCPs have a very important role in the risk argument, it seems appropriate to bring to bear strong arguments for both OCPs. For OCP 1, we note that some maintenance checks are considered so important it is decided that an independent inspector must check the work of the maintainer to provide confidence in their work. We reflect this arrangement, as well as a 'backstop' that aircraft maintenance logs are routinely reviewed to provide confidence that all scheduled work has been successfully completed before the aircraft is declared fit to fly. In terms of the inferred solution, OCP 2, additional confidence is again achieved through

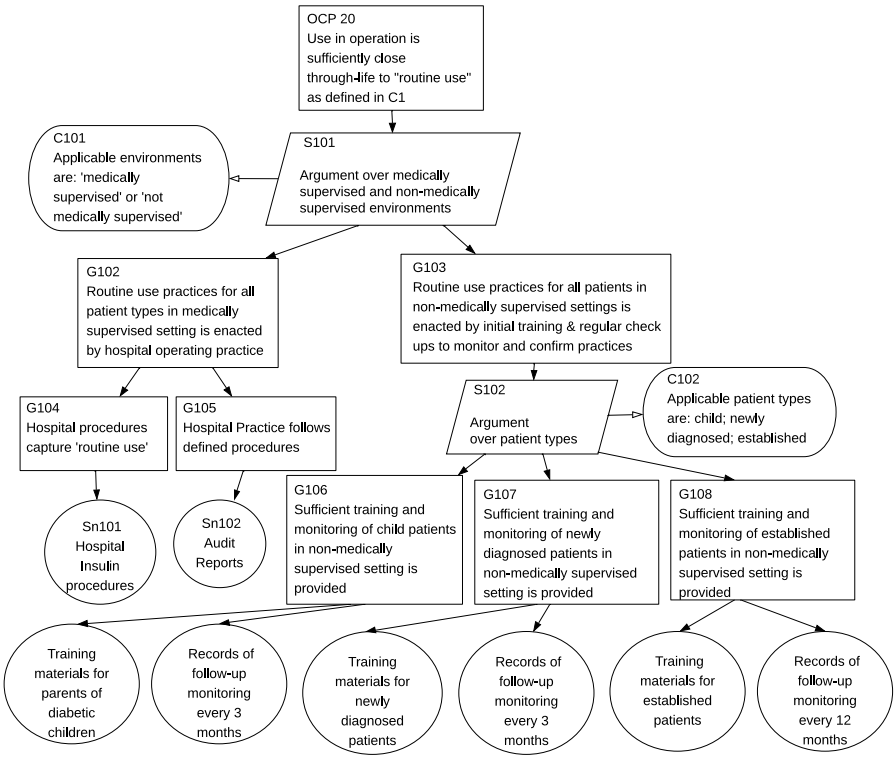

Fig. 3. Operational safety case for insulin pump showing satisfaction of OCP.

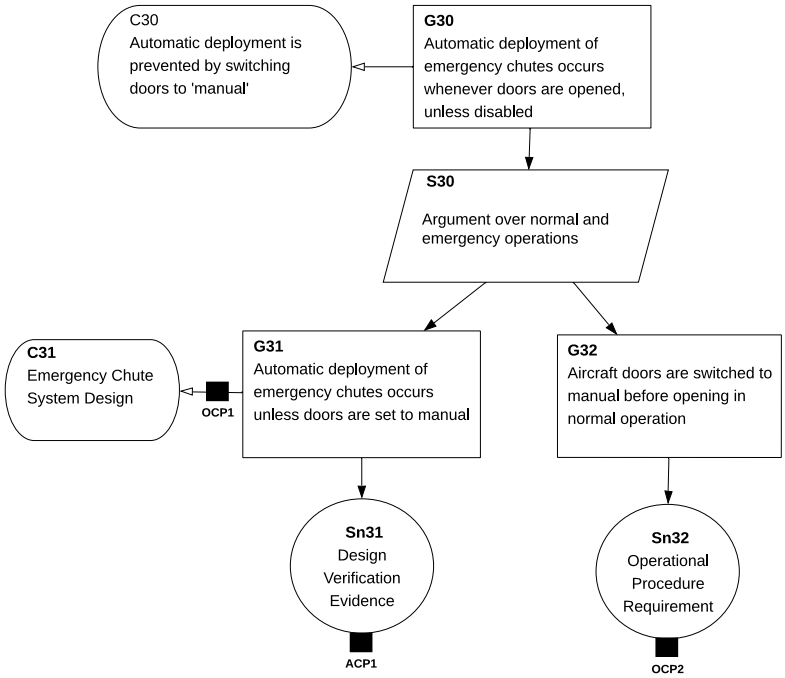

Fig. 4. Design safety case for emergency chute system showing use of OCPs.

independent human verification of the procedure for every flight, which is part of cabin crew training with further confidence provided by requiring cabin crew to practice this procedure during simulated rehearsal training.

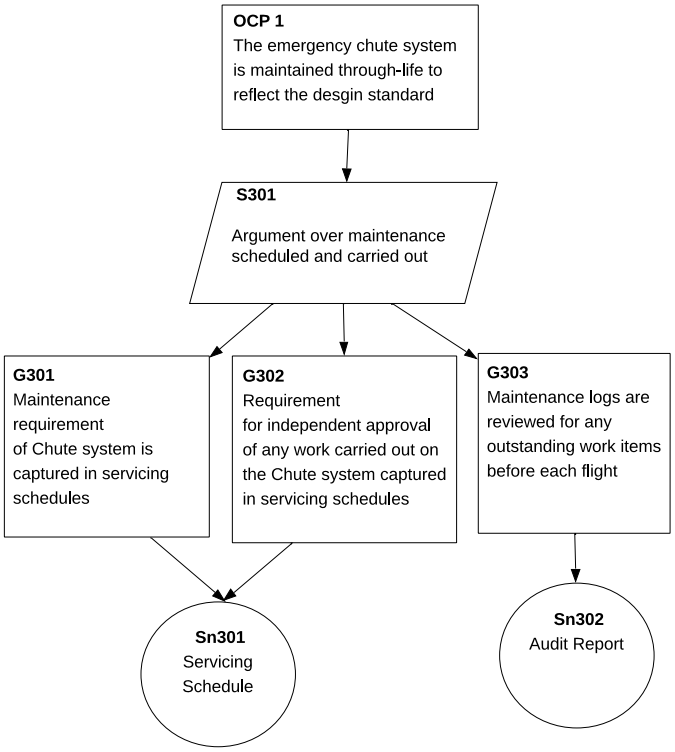

Fig. 5. Operational safety case for emergency chute system showing satisfaction of OCP1.

We believe that this transparent visualisation of operational claim points may provide us with a better opportunity to undertake more explicit trade-offs between mitigation by design and mitigation by operation throughout the life-cycle. It allows us, for instance, to consider trade-offs during trialling of systems, ahead of full deployment. We also envisage it to be useful, once tool-supported, to show different viewpoints to different stakeholders in the risk management process.

5 Conclusions and Future Work

The currently prevalent paradigm in industrial practice of shifting between risk assurance and evaluation in the design phase of systems, to a simplified barrier

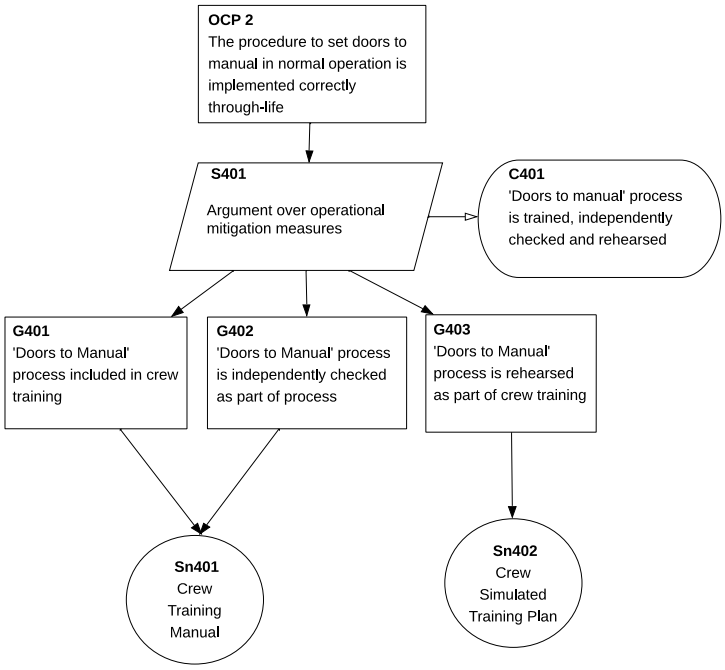

Fig. 6. Operational safety case for emergency chute system showing satisfaction of OCP2.

conformance evaluation in the operational phase of the lifecycle of a product is undesirable. We recognise why this has been necessary in the past, but note that future objectives to measure risk through the life of a system, required by approaches such as DevOps [10] and the potential use of learning algorithms, require better continuity of approach. We believe that the introduction of OCPs into safety cases provides:

- a mechanism to clearly articulate in the design phase any assumptions about processes or contextual (e.g. operating) limits that will be applied in the operational phase and/or set requirements on operational mitigations such as operating limits, operating procedures or training.
- greater clarity for responsible persons, at the handover of responsibility between a design organisation and an operating organisation, as to what interfaces and dependencies have been agreed between the two phases and potentially two different responsible persons
- a clearer interface that can be used to assess day to day impact of operational decisions on risk, and uncertainties associated with risk levels, rather than reverting to less well supported barrier effectiveness assessment

In future, we plan to further explore and test the utility of the OCP concept and notation, including through evaluation with stakeholders in both the design and operational domains. We also plan to extend this work through further research in the following areas:

- We have so far mainly focused on use of OCPs for asserted context and evidence. We believe there may also be utility in considering OCPs for other types of argument assertion and plan to explore this further.
- We recognise the need for processes for the consolidation of OCPs, to check for overlap and to assess for common mode failures of operational mitigations.
- There may be more than one operational safety case associated with any design safety case, so we will explore how the safety case interfaces can be configured for multiple operators, perhaps by the use of safety case contracts.
- We will consider how this approach may be beneficial in articulating and facilitating through life safety consideration in early design concepts
- We will consider the potential role of OCPs in confidence and conformance arguments as well as in the risk argument

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Acfield, A.P., Weaver, R.A.: Integrating safety management through the bowtie concept a move away from the safety case focus. In: Proceedings of the Australian System Safety Conference-Volume 145. pp. 3–12 (2012)

2. ACWG: Assurance Case Guidance. Tech. Rep. SCSC-159 v1, Safety Critical Systems Club (2021), https://scsc.uk/r159:1
3. ACWG: Goal Structuring Notation Community Standard. Tech. Rep. SCSC-141C v3.0, Safety Critical Systems Club (2021), https://scsc.uk/scsc-141C
4. Cullen, L.: The public inquiry into the Piper Alpha disaster. Department of Energy (1993)
5. Green, Richard: Supporting operational decision making concerning aircraft structural integrity damage identified during maintenance. Tech. rep., Cranfield University MSc thesis (2021)
6. Haddon-Cave, C.: The Nimrod Review. An Independent Review into the Broader Issues Surrounding the Loss of the RAF Nimrod MR2 Aircraft XV230 in Afghanistan in 2006. London: The Stationery Office (2009)
7. Hawkins, R., Kelly, T., Knight, J., Graydon, P.: A new approach to creating clear safety arguments. In: Advances in Systems Safety: Proceedings of the Nineteenth Safety-Critical Systems Symposium, Southampton, UK, 8-10th February 2011. pp. 3–23. Springer (2011)
8. Hayes, J.: Use of safety barriers in operational safety decision making. Safety Science **50**, 424–432 (2012)
9. Kelly, T.P.: Arguing safety - a systematic approach to managing safety cases. Tech. rep., University of York DPhil thesis (1998)
10. Kim, G., Humble, J., Debois, P., Willis, J., Forsgren, N.: The DevOps handbook: How to create world-class agility, reliability, & security in technology organizations. IT Revolution (2021)
11. MAA: The Air System Safety Case and Air System Safety Case Report(s). Regulatory Article 1205 Issue 8, UK MAA (2022), https://assets.publishing.service.gov.uk/media/6527b49b2548ca0014ddf18a/RA1205_Issue_8.pdf