Safety Cases and their role in ISO 26262 Functional Safety Assessment

John Birch[1], Roger Rivett[2], Ibrahim Habli[3], Ben Bradshaw[4], John Botham[5], Dave Higham[6], Peter Jesty[7], Helen Monkhouse[8], Robert Palin[9]

[1] AVL Powertrain UK Ltd, Basildon, UK
[2] Jaguar Land Rover, Coventry, UK
[3] University of York, York, UK
[4] TRW Conekt, Solihull, UK
[5] Ricardo UK Ltd, Cambridge, UK
[6] Delphi Diesel Systems
[7] Peter Jesty Consulting Ltd, Tadcaster, UK
[8] Protean Electric Ltd, Surrey, UK
[9] MIRA Ltd, Nuneaton, UK

Abstract. Compliance with the automotive standard ISO 26262 requires the development of a safety case for electrical and/or electronic (E/E) systems whose malfunction has the potential to lead to an unreasonable level of risk. In order to justify freedom from unreasonable risk, a safety argument should be developed in which the safety requirements are shown to be complete and satisfied by the evidence generated from the ISO 26262 work products. However, the standard does not provide practical guidelines for how it should be developed and reviewed. More importantly, the standard does not describe how the safety argument should be evaluated in the functional safety assessment process. In this paper, we categorise and analyse the main argument structures required of a safety case and specify the relationships that exist between these structures. Particular emphasis is placed on the importance of the product-based safety rationale within the argument and the role this rationale should play in assessing functional safety. The approach is evaluated in an industrial case study. The paper concludes with a discussion of the potential benefits and challenges of structured safety arguments for evaluating the rationale, assumptions and evidence put forward when claiming compliance with ISO 26262.

Keywords. Safety cases, safety arguments, ISO 26262, automotive safety.

1 Introduction

Critical functions in road vehicles are increasingly being implemented using electrical and/or electronic (E/E) systems. The malfunctioning behaviour of these systems can contribute to the safety risk to the vehicle occupants and/or other road users. As such, it is necessary to provide assurance that any unreasonable residual risks have been avoided. The safety standard ISO 26262 has been developed to address this necessity

by providing guidance, in the form of requirements and processes, for avoiding unreasonable residual risk caused by the malfunctioning behaviour of E/E systems [1]. Like many safety standards that cover complex software-based systems, ISO 26262 defines requirements for the creation of work products i.e. outputs from the safety lifecycle, and leaves it to the developers to interpret these requirements in the context of their products [2]. In order to provide a product-specific justification, compliance with the ISO 26262 standard requires the development and evaluation of a safety case for the safety-related items. The standard defines an item as a "*system or array of systems to implement a function at the vehicle level*" [1]. In order to justify freedom from unreasonable risk, a safety case argument should be developed in which the safety requirements are shown to be complete and satisfied by the evidence generated from the ISO 26262 work products. However, the standard does not provide practical guidance on the development and review of the safety argument, nor does it describe how the safety argument should be evaluated in the functional safety assessment process.

In this paper, we build on the experience of the authors in developing and evaluating safety cases in the context of ISO 26262. We examine the significance and nature of the product-based safety rationale within the argument and the role this rationale should play in assessing functional safety. The paper also builds on existing work on safety cases across different domains [3-5], and in the automotive industry in particular [6], [7], taking into account issues related to product-based and process-based assurance [8], the process of compliance [9] and assessment of confidence [10], [11].

The paper is organised as follows. In Section 2, we categorise and analyse the main argument structures of a safety case and the relationships that exist between the safety case and the ISO 26262 functional safety assessment. The approach is evaluated in an industrial case study in Section 3. In Section 4, we discuss the potential benefits and challenges of structured safety arguments for evaluating the rationale, assumptions and evidence put forward when claiming compliance with the ISO 26262 standard.

2 Safety Argument Categories in ISO 26262

ISO 26262 defines a safety case as an "*argument that the safety requirements for an item are complete and satisfied by evidence compiled from work products of the safety activities during development*" [1]. That is, the argument should play a central role in justifying why the available evidence, in the form of work products (e.g. design and analysis artefacts), has achieved a set of safety requirements and, therefore, why an acceptable level of safety has been achieved. Compliance with ISO 26262, based on the normative parts of the standard, mandates the satisfaction of a specific set of objectives by the generation of a concrete set of work products. As a result, all E/E systems that are compliant with the standard share a common safety argument structure linking the top-level safety requirements to the available evidence. Unfortunately, this common argument structure is implicit and is not documented in the standard.

2.1 Implicit Safety Argument in ISO 26262

The implicit safety argument in ISO 26262 is centred on the following chain of reasoning (Fig. 1). A sufficient and an acceptable level of safety of an E/E system is achieved by demonstrating absence of unreasonable risk associated with each hazardous event caused by the malfunctioning behaviour of the item (other hazard causes are outside the scope of the standard). This is achieved by defining safety goals to avoid unreasonable risk through the prevention or mitigation of the identified hazardous events. A hazardous event is the occurrence of a hazard in particular operational situations. Each hazardous event is assigned an Automotive Safety Integrity Level (ASIL), based on the combination of three parameters: severity (extent of human harm), probability of exposure (to operational situations) and controllability (ability for persons at risk to take action to avoid harm). Claims are then asserted that each safety goal is satisfied by the development of a functional safety concept. The functional safety concept specifies safety measures within the context of the vehicle architecture, including fault detection and failure mitigation mechanisms, to satisfy the safety goals. Two further hierarchies of claim are defined for asserting how the functional safety concept is adequately refined and satisfied by a technical safety concept and hardware and software components (again to the required ASIL). As a result, the implicit argument follows a hierarchy of claims that can be grouped as follows:

- Safety Goals (hierarchy 1) – the vehicle in its environment;
- Functional Safety Requirements (hierarchy 2) – the vehicle and its systems;
- Technical Safety Requirements (hierarchy 3) – the E/E system; and
- Hardware and software requirements (hierarchy 4) – component and part level.

For each hierarchy, ISO 26262 prescribes evidence, in the form of work products, for substantiating these claims. Additionally, the standard identifies methods for generating these work products in accordance with the required ASIL. For example, in order to substantiate a claim that the technical safety requirements have been correctly implemented at the hardware-software level, evidence should be provided through methods such as a requirements-based test, fault injection test or back-to-back test (Table 1, Part 4). This evidence should be captured in an Integration Testing Report (Work Product 8.5.3, Part 4).

The implicit safety argument in ISO 26262 has two main categories of claim: product claims and process claims. Based on the hazard analysis and risk assessment, the product claims focus primarily on the safety goals and safety requirements (i.e. specifying and demonstrating behaviour which is free from unreasonable risk). The process claims focus on the adequacy of the organisations, people, lifecycles, methods and tools involved in the generation of the work products. The nature of these process claims and the rigour of the evidence needed to support them vary with the ASIL assigned to the safety goals and their corresponding safety requirements (i.e. high levels of risk require high levels of process rigour).

Compliance with ISO 26262 and the evaluation of the above implicit argument is demonstrated, in part, using two types of confirmation measures: functional safety audit and functional safety assessment. The requirements for both, and the necessary

independence, are specified in Part 2 of the standard. The functional safety *audit* is concerned with reviewing the implementation of the *processes* required for functional safety. Functional safety *assessment* is concerned with making a judgement on the functional safety achieved by the item and hence is concerned with the characteristics of the *product*. The assessment includes evaluating the work products specified in the item's safety plan, the required processes (i.e. the functional safety audit) and the appropriateness and effectiveness of the safety measures that are implemented.

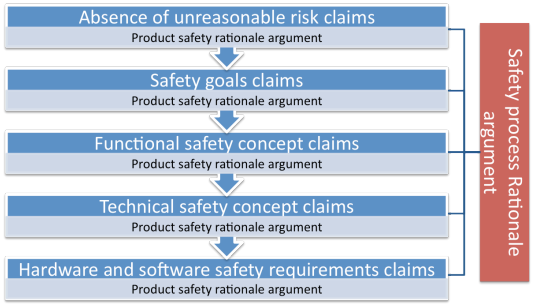

Fig. 1. Implicit ISO 26262 Safety Argument Structure

2.2 Product-Specific Safety Rationale

A legitimate question at this point should be: why is it necessary to document the above safety argument if it is common to all items compliant with ISO 26262? What is the added value of developing, reviewing and maintaining this common safety argument? In this paper, we contend that the challenge does not lie in merely capturing this common, implicit, argument. Instead, most of the effort should focus on justifying, through an explicit argument structure, how one hierarchy of claims, e.g. concerning absence of unreasonable risk, is supported by another hierarchy of claims, e.g. safety goals that address any unreasonable risk (Fig. 1). These sub-arguments should capture the product-specific safety rationale which typically varies from one item to another. That is, although the overall structure of the argument is stable (i.e. assessment of hazardous events and specification, development and assessment of safety goals and safety requirements), the assurance challenge lies in providing product-specific rationale, assumptions and justifications for why, given an operational environment, a vehicle configuration and the condition of other vehicle systems, the available evidence is sufficient to support the asserted claims. Typically, these claims, and their corresponding arguments and evidence, are the focus of the functional safety assessment process as they address the product-specific safety rationale that is often associated with unique characteristics of the system and its environment.

A claim is typically made that the absence of unreasonable risk of a hazardous event has been addressed by conforming to a safety goal. However, it would be naïve to define a safety goal as simply a negation of a hazardous event and simply to assign an ASIL to that safety goal. Although this approach is arguably valid from the perspective of literal ISO 26262 compliance, it is simplistic as it limits risk mitigation to

reducing the probability of the hazardous malfunction. Other risk reduction strategies related to reducing severity, improving controllability and/or reducing exposure (typically through a measure "external" to the item, which can be another E/E system) can be taken into account. For example, if a safety goal stipulates that the system shall transition to a safe state in the presence of faults that could otherwise cause the corresponding hazardous event, then an argument and evidence for why the specified safe state is considered to be adequately safe should be provided. This can be achieved by justifying that, were the vehicle behaviour in the safe state to be subject to ISO 26262 hazard classification criteria, then it would be classified 'QM' (Quality Management). QM in ISO 26262 denotes a risk that does not require the satisfaction of any specific safety requirements, thereby implying that the level of risk is reasonable and no further risk reduction is necessary. The main claim here would be that the *residual risk* associated with the hazardous event, after achieving the safety goal, has been reduced to a level that is *reasonable*. The subsequent argument used to support such a claim would then need to explicitly assert which risk parameters ('controllability', 'severity' or 'exposure') would be reduced if the residual risk were classified in this way.

A typical approach may be to provide an argument that some reconfiguration or degradation scheme is capable of placing a system into a safe state such that the controllability of any reaction, e.g. to an undemanded drive torque, is effectively C0 (controllable in general) whereas the hazardous event itself will have been classified with the controllability parameter taking a value of C1, C2 or C3. Another approach may be to place a system in a safe state by preventing a vehicle exceeding a speed threshold upon detection of a fault that can cause the hazardous event such that the exposure parameter that could be associated with the safe state is effectively E0 (incredible). Such reasoning is product-specific and the implicit safety argument in ISO 26262 does not prescribe any product-specific safety rationale.

The safety argument structure in Fig. 1 includes references to five product-specific safety rationale sub-arguments. These sub-arguments should provide justification for the inferential transition from one hierarchy of safety claims to another. For instance, the functional safety concept rationale argument should include a justification for why the deployment of safety measures such as fault detection, failure mitigation and/or driver warnings should lead to the satisfaction of the corresponding safety goals.

3 Industrial Case Study

This case study is based on a typical electric vehicle architecture (technology-specific details have been abstracted for reasons of commercial sensitivity), in which a basic Item Definition and hazardous event are considered. The purpose of the case study is to examine the product-based safety rationale arguments, discussed in Section 2, for the corresponding Safety Goal and Functional Safety Concept.

3.1 Item Definition

The Item Definition is shown in Fig. 2. The pertinent nominal operation is as follows:

- Driver requests positive longitudinal vehicle acceleration by depressing accelerator pedal
- Accelerator pedal provides a low voltage electrical signal to indicate pedal position to the Controller
- Controller reads this pedal signal and places a corresponding torque demand on the High Voltage Power Inverter (HVPI) via the Controller Area Network (CAN)
- HVPI converts a certain quantity of electrical energy from the High Voltage Battery to high voltage electrical power supplied to the Electric Machine, according to the torque demand from the Controller
- High voltage electrical power supplied to the Electric Machine induces a mechanical torque in the Electric Machine, which is transferred through the transmission to the vehicle's rear wheels.

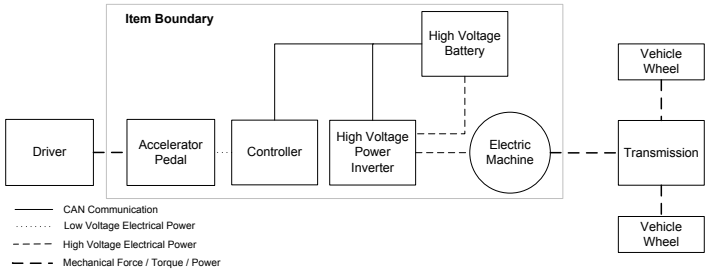

Fig. 2. Electric Vehicle Propulsion System

3.2 Hazard Analysis and Risk Assessment

This case study focuses on the Hazardous Event '*Unintended vehicle acceleration during a low speed manoeuvre amongst pedestrians*', which is classified as ASIL B based on values of E3 (medium probability), S2 (severe and life-threatening injuries, survival probable), C3 (difficult to control or uncontrollable) for the Exposure, Severity and Controllability parameters respectively. The rationale for this classification requires a detailed description of the vehicle, the operational and environmental constraints and peer systems and as such it has not been included for brevity.

3.3 Safety Goal and its Rationale Argument

The safety goal that has been defined to address the risk associated with the Hazardous Event is '*Vehicle positive longitudinal acceleration shall not exceed driver demand by > 1.5 m* s^2 *for longer than 1 s*'. However, the question that a safety assessor may rightly ask is why by meeting this safety goal is unreasonable risk avoided? It is not typical within industry for the answer to questions of this type to be documented, but doing so should help the engineer to be clear about why the safety goal achieves freedom from unreasonable risk, and to communicate that to the safety assessor.

The argument for this particular case study, presented in Goal Structuring Notation (GSN) [12] in Fig. 3, is based on improving *controllability*; specifically if the unin-

tended acceleration is kept below the stated threshold, the driver is able to slow and stop the vehicle before a collision with the pedestrian occurs. Within this argument, the '*Absence of Unreasonable Residual Risk'* strategy is generic, and could be applied to any safety goal, whereas the '*Residual Risk Controllability Classification*' strategy is specific to this particular safety goal.

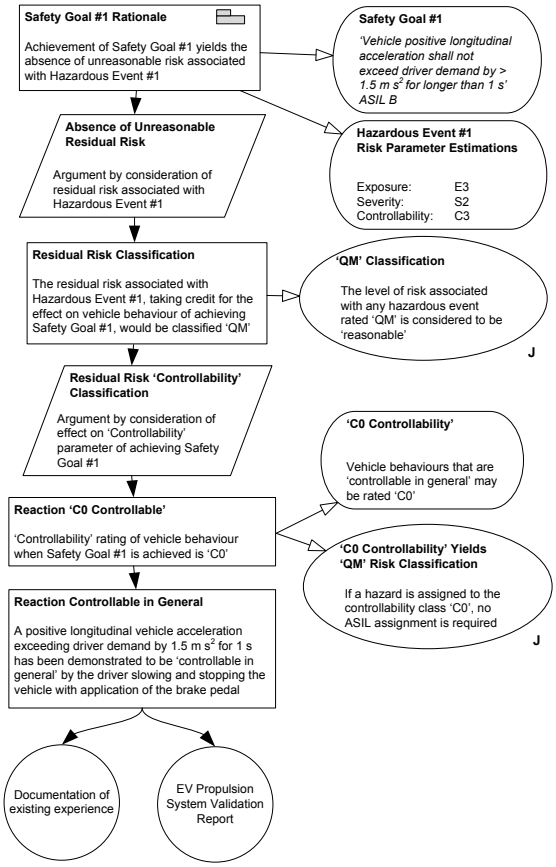

Fig. 3. Safety Goal Rationale Argument

3.4 Functional Safety Concept and its Rationale Argument

The functional safety concept that has been chosen to achieve the safety goal, named '*Distributed detection and mitigation of torque errors'*, is based on degradation; whereby all faults that can lead to excessive acceleration are detected within an acceptable time interval. On detection of a fault, the vehicle acceleration is limited to a value below that specified in the safety goal. The concept is based on the assertion that only malfunctioning behaviour of the Item that can violate the safety goal (which is specified in terms of *vehicle*-level behaviour; *acceleration*) is the delivery of exces-

sive *torque* to the Transmission; behaviour which is specified at the *Item-level*. The concept features are as follows (Fig. 4):

1. Detection of all faults that would otherwise lead to excessive torque delivery:
   (a) Controller detects accelerator pedal faults by comparing and arbitrating between the outputs from two independent pedal position measurement sensors
   (b) Controller self-detects torque-request errors by comparing its final torque request to the HVPI (output) with the accelerator pedal position (input)
   (c) HVPI self-detects torque-demand errors by comparing the quantity of high voltage electrical power supplied to the Electric Machine (output) with the torque request from the Controller (input)
2. Upon detection of errors, outputs are electronically 'limited' to a fixed value that ensures that the magnitude of excessive torque delivered to the Transmission is below that required to violate the safety goal's acceleration criteria.

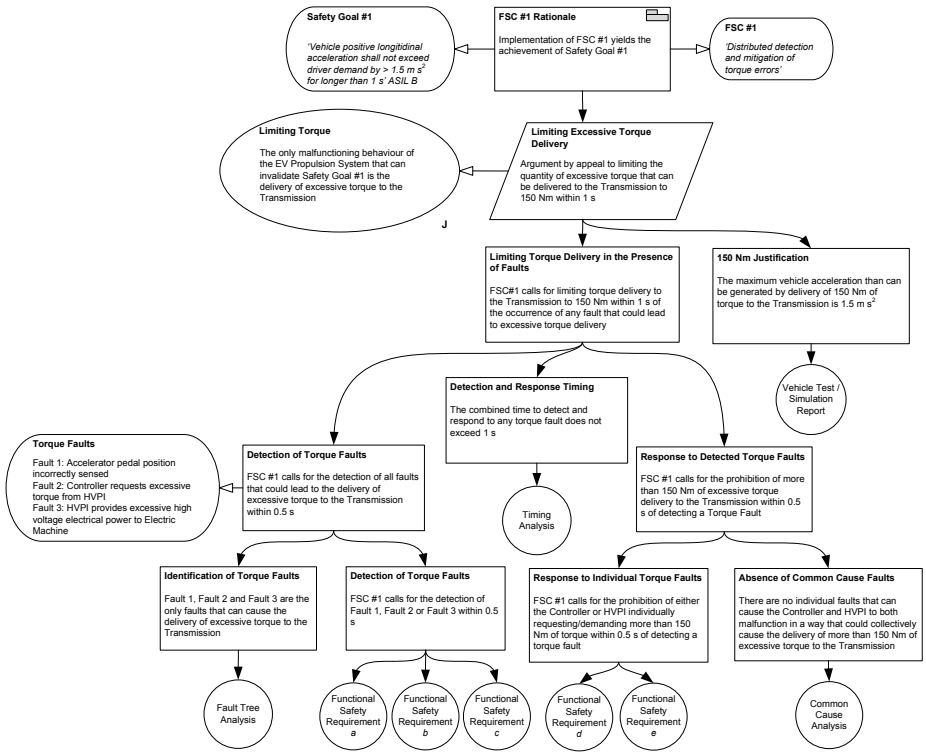

Fig. 4. Functional Safety Concept Rationale Argument

Typically, a company would document the failure modes of the concept in an analysis report, e.g. using Failure Mode and Effects Analysis (FMEA), and manage safety goal and functional safety requirements in a requirements database. It would also have vehicle test reports or simulations demonstrating that the safety goals had been met. However, the rationale explaining how this evidence fits together is not often docu-

mented. This means that whoever performs the Functional Safety Assessment has to deduce this for themselves by '*reading between the lines*', and for complex and highly interconnected systems, tenuous leaps may need to be made. The added value of formally documenting the rationale, as in Fig. 4, is not only that it helps the engineers to identify potential deficiencies in the safety argument, but also that it eases the subsequent task of performing the Functional Safety Assessment, and may highlight the need for further work.

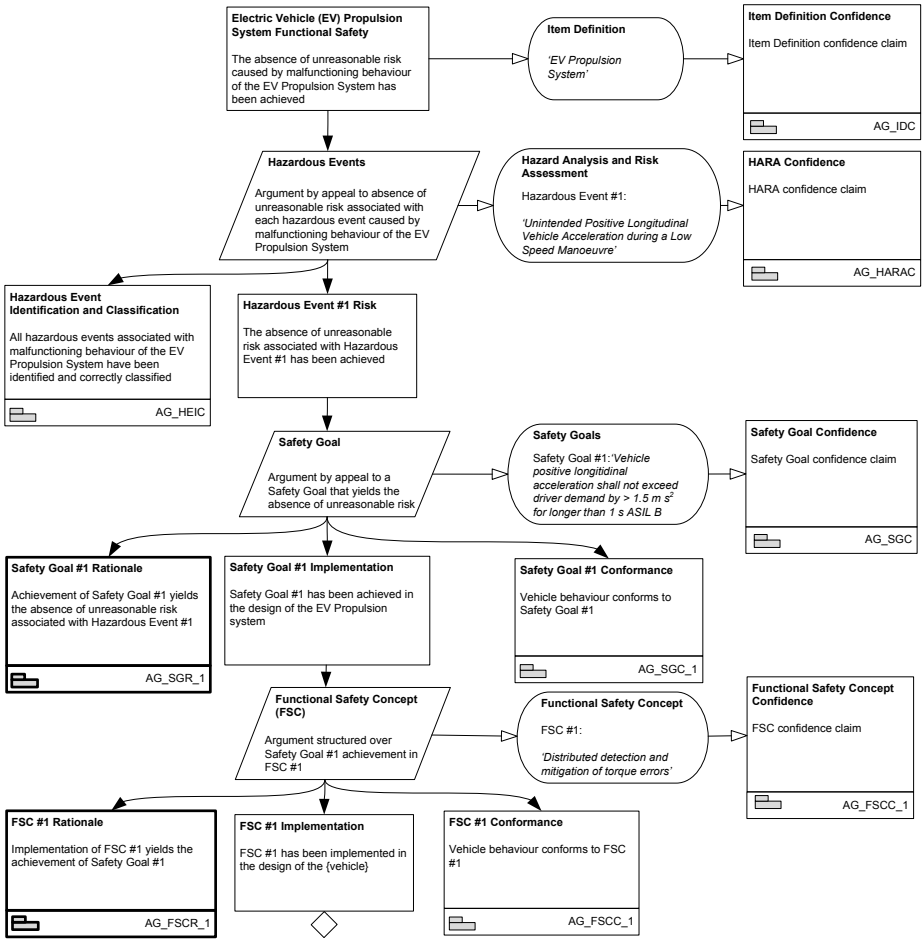

Fig. 5. Item Functional Safety Argument

3.5 Item Functional Safety Argument

The two arguments presented in Fig. 3 and Fig. 4 can be referenced as '*Away Goals*' [12] within the complete safety argument for the Item (Fig. 5). Other structures within the complete safety argument should include confidence-based process claims that

refer to the ASIL-specific processes used to develop the work products. The complete argument would also require further development to justify how the functional safety concept has been implemented by the chosen technical safety concept and subsequently by the hardware and software safety requirements. Although this has not been the focus of this paper, it has been found that the argument structure at the level of safety goals and functional safety concept can be successfully repeated at lower levels, with the capability to partition the argument to represent the distributed development commonly seen between a vehicle manufacturer and its suppliers.

4 Analysis and Discussion

Without a clear safety argument structure, a checklist approach to safety assurance, based on the creation of work products and requirements traceability, tends to be used. Important as this is, the rationale behind the requirements is often not documented. An important aspect of capturing the product-based safety rationale is that it helps the engineers identify potential deficiencies in the argument in a timely manner and supports the subsequent task of performing the Functional Safety Assessment. In this section, we reflect on the insights gained from different engineering perspectives.

4.1 Original Equipment Manufacturer (OEM) Perspective

Typically, a large list of safety requirements and work products is presented to an OEM, i.e. the vehicle manufacturer, for which there may be traceability to the safety goals but no, or only tenuous, basis for understanding whether and how the safety goals have been fully satisfied. Consequently the adequacy of the deliverables can only be determined by extensive question and answer sessions. This often reveals that the important safety rationale is not documented and only exists in the heads of the engineers. It also often reveals that there are many undocumented assumptions which need to be validated and would be better treated as safety requirements. Where the approach presented in this paper is adopted, engineers gain a deeper understanding of the system they are designing. Further, documentation is generated at a more appropriate lifecycle stage to enable effective and timely assessment.

Because of the hierarchical nature of the explicit safety argument, and the observation that its structure repeats between levels, the argument lends itself to being ‘split’ between organisations. For example, an OEM may typically develop the safety argument down to, and including, the level of functional safety concept. The supplier responsible for developing the technical safety concept and hardware and software safety requirements can then develop the relevant downstream rationale in a similar manner to the OEM in order to justify the safety requirements they have developed.

4.2 Supplier Perspective

E/E system suppliers are heavily dependent on requirements received from the OEM, as the OEM has a complete view of the vehicle, its systems and their dependencies.

However, by developing an E/E system and hardware and/or software components, a supplier will generally own a high proportion of the faults that can contribute to hazardous events. As such, suppliers will be responsible for the design requirements, safety analysis and verification of the E/E system to support the claim that vehicle level safety will be assured and that the requirements of the functional safety concept have been achieved. With this partitioning of responsibilities comes the need to demonstrate accountability i.e. the need for suppliers to provide a safety argument to justify that their design/implementation supports certain safety goals at the vehicle level. A structured argument provides this much needed visibility between parties at the different assessment stages.

Further, suppliers traditionally develop common E/E platforms prior to OEM engagement. For example, a supplier developing an engine management system for future vehicle emission legislation will identify requirements years before involving OEMs. It is important that any safety-related component/element, which is developed without a specific application context, is assessed. The supplier will need to capture assumptions, most likely based on previous experience, and possibly in isolation. These assumptions and the safety rationale are very well suited to an argument structure that clearly identifies product safety claims in relation to assumed hazards, safety goals and concepts. A clearly defined argument structure improves the engagement of customers with new applications not only to provide the safety justification but also to identify assumptions that require confirmation, redress and also the allocation of risk mitigation responsibilities to customers when needed.

4.3 Safety Assessor Perspective

In the infancy of ISO 26262, early project assessments have been based solely on work products and processes. This has resulted in lengthy protracted assessments, trawling through documentation and relying heavily on interviews to discover undocumented rationale. This has highlighted the need to have a safety case with a clear structure and purpose. It has also been found that it is both possible and beneficial to assess the 'top down' safety argument iteratively, as the design of an item evolves. For example, the safety assessor can review the safety goal rationale argument in Fig. 1 before the functional safety or technical safety concepts have been developed, rather than waiting until the later lifecycle stages. This helps to identify weaknesses in the eventual safety argument earlier on in the project lifecycle, reducing the cost and effort resulting from any subsequent rework.

Finally, the automotive industry like many other domains is driven by tight margins and time constraints. Once a project is underway, momentum increases quickly. It is therefore essential that the visibility of the project's technical and assurance attributes and any infringements identified early so that undesired consequences are addressed. This leads to the conclusion that the review and assessment of the safety case at key product gateways will not only keep focus on the emergence of the project's and product's safety attributes, but is more likely to have a safety case at the final functional safety assessment that is legible and more readily analysable.

5 Concluding Remarks

Safety case development is a relatively new concept for many safety practitioners in the automotive industry. The timely generation of well-focussed safety cases is capable of bringing considerable benefit in the context of development and assessment, and thus of contributing to the safety assurance of automotive E/E systems. Our experience to date suggests that the primary focus of many documented safety cases for ISO 26262-compliant systems and components remains on processes. In extreme cases, this can result in bulky documentation that does little more than render explicit the standard's implicit arguments or, even, recapitulate its requirements in a different form. Broadly, we perceive an educational challenge to exist in this area even among automotive safety engineers with considerable experience in other areas.

Other characteristics have reduced the effectiveness of certain safety cases produced to meet the requirements of ISO 26262. These include: lack of focus and brevity; unnecessary repetition of information available elsewhere; and the use of inappropriate means of expression (e.g. use of GSN where a table might be more effective and vice versa). Similarly, safety cases in the automotive industry are as susceptible as those in other industries to deficiencies such as fallacies and failures to acknowledge limitations. These weaknesses are found in safety cases in other industries but, we believe, may best be countered by didactic material that is targeted specifically at the automotive industry in order to improve outreach.

References

1. ISO: ISO 26262 Road Vehicles -- Functional Safety. ISO Standard (2011)
2. Graydon, P., Habli, I., Hawkins, R., Kelly, T., Knight.: Arguing conformance. IEEE Software, vol. 29, issue 3 (2012)
3. Bishop, P., Bloomfield, R.: A methodology for safety case development. In: Proc. 6th Safety-critical Sys. Symp. (1998)
4. Kelly, T.: A systematic approach to safety case management. In: Proc. Society of Automotive Engineers (SAE) World Congress (2004)
5. The Health Foundation, Using Safety Cases in Industry and Healthcare. ISBN: 978-1-906461-43-0, (2012)
6. Dittel, T., Aryus, H.: How to "Survive" a safety case according to ISO 26262. SAFECOMP 2010, Vienna, Austria, (2010)
7. Palin, R., Habli, I.: Assurance of automotive safety: a safety case approach. SAFECOMP 2010, Vienna, Austria, (2010)
8. Habli, I., Kelly, I.: Process and product certification arguments: getting the balance right. SIGBED Review, vol. 3, issue 4 (2006)
9. Langari, Z., Maibaum, T.: Safety cases: a review of challenges. International Workshop on Assurance Cases for Software-intensive Systems (ASSURE 2013) San Francisco (2013)
10. Denney, E., Pai, G., Habli, I.: Towards measurement of confidence in safety cases. In: Proc. 5th Intl. Symp. on Empirical Soft. Eng. and Measurement. pp. 380–383 (Sep 2011)
11. Ayoub, A., Kim, B., Lee, I., Sokolsky, O.: A systematic approach to justifying sufficient confidence in software safety arguments. SAFECOMP 2012, Magdeburg, Germany (2012)
12. Goal Structuring Notation Working Group: GSN Community Standard Version 1 (2011)