Is current incremental safety assurance sound ?

Valentin Cassano, Silviya Grigorova, Neeraj Kumar Singh, Morayo Adedjouma, Mark Lawford, T. S. E. Maibaum, Alan Wassyng

Is Current Incremental Safety Assurance Sound?

V. Cassano(✉), S. Grigorova, N.K. Singh, M. Adedjouma, M. Lawford, T.S.E. Maibaum, and A. Wassyng

McMaster Centre for Software Certification, McMaster University, Hamilton, Canada
{cassanv,grigorsb,singhn10,morayoa,lawford,wassyng}@mcmaster.ca
tom@maibaum.org

Abstract. Incremental design is an essential part of engineering. Without it, engineering would not likely be an economic, nor an effective, aid to economic progress. Further, engineering relies on this view of incrementality to retain the reliability attributes of the engineering method. When considering the assurance of safety for such artifacts, it is not surprising that the same economic and reliability arguments are deployed to justify an incremental approach to safety assurance. In a sense, it is possible to argue that, with engineering artifacts becoming more and more complex, it would be economically disastrous to not "do" safety incrementally. Indeed, many enterprises use such an incremental approach, reusing safety artifacts when assuring incremental design changes. In this work, we make some observations about the inadequacy of this trend and suggest that safety practices must be rethought if incremental safety approaches are ever going to be fit for purpose. We present some examples to justify our position and comment on what a more adequate approach to incremental safety assurance may look like.

Keywords: Incremental design improvement · Incremental safety assurance

1 Introduction

Incremental design improvement, a.k.a. *normal* engineering design [23], has a long history and proven value as a way for constructing improved versions of artifacts. This engineering praxis caters for time and budget constraints, ensures an artifact's effectiveness, fitness for purpose, and the reliability of its production.

We consider that the same considerations guiding incremental design improvement have fostered a practice of incremental safety assurance which relies heavily on the reuse of existing safety artifacts, e.g., safety related evidence and arguments. However, in contrast to incremental design improvement, we argue that incremental safety assurance, as presently viewed and practiced, is not necessarily sound. An important reason for this is the global nature of safety as a property of a system. Focusing the safety assurance efforts in a localized fashion, e.g., on the slice of the system where the design change occurred, may ignore newly created global hazards or the re-emergence of those that are otherwise

mitigated. Complicating things further, safety artifacts cannot be straightforwardly composed, as the context and the assumptions in one safety artifact may undermine safety claims established in another. Though these considerations may seem well-known, that certain current safety practices fail to tackle them properly indicates that they are neither entirely understood nor easily dealt with.

In this paper, we put forward some observations on why incremental safety assurance, when understood from the perspective of incremental design improvement, is problematic, and in fact inherently deficient (that is the bad news!). Our discussion hinges on two main points: compositionality and the defeasibility of safety arguments, and locality and emergent properties. By elaborating on these points, we hope to bring to the foreground what we believe is an important issue of safety practice: the reuse of safety artifacts. While we believe that there is certainly great practical value in the reuse of safety artifacts, we offer a view of what a more sound approach to incremental safety assurance might look like (that is the good news!), this needs of a great deal of caution.

Structure of the paper: In Sect. 2, we explain what incremental safety assurance means from the perspective of incremental design improvement, commenting on its underlying philosophy and the necessity for its existence. In Sect. 3, we elaborate on our reservations about such an incremental approach to safety assurance. In Sect. 4, we substantiate our claims by providing examples from the automotive and medical domains. In Sect. 5, we discuss the challenges and opportunities presented by an incremental approach to safety assurance. In Sect. 6, we comment on some related work. In Sect. 8, we offer some conclusions and talk about our next steps.

2 Incremental Safety Assurance

When faced with a problem, engineers tend to build on experience, best practices, and already existing artifacts, analyzing their pros and cons in order to try to adapt them (incrementally) to the problem at hand. This approach is key for guaranteeing an artifact's reliability and the reliability of its production. This commonly accepted view of engineering praxis is, among other places, discussed by Vincenti in [23] under the name of *normal design*. In Vincenti's terms, a design is normal if both the *operational principle*, i.e., "how the device works", and the *normal configuration*, i.e., "the general shape and arrangement that are commonly agreed to best embody the operational principle", are known and used [23, pp. 208–209]. If either the operational principle or the normal configuration are largely unknown, or, if known, are left unused, then, the design is *radical* [23, p. 8]. Radical design is then to be thought of as based on engineering principles that are wholly different from those guiding normal design. This said, Vincenti remarks that "though less conspicuous than radical design, normal design makes up by far the bulk of day-to-day engineering enterprise" [23, p. 8].

The difference between normal design and radical design is easily illustrated in the automotive domain. A case can be made that majority of current vehicles are based on, reuse, or extend, design elements existing in other vehicles of the

same kind, i.e., normal design. This applies both to the software and hardware components of a vehicle and enables the automotive industry to rely on well-tested systems while being up-to-date with technological advances. On the other hand, the development of an autonomous car would exemplify a radical design.

The inherent practicality of normal design, i.e., of incremental design improvement, has lead to its enduring prevalence. We consider that it is this prevalence, as well as the striving for efficiency and resource preservation, that has fostered an incremental approach to safety assurance. It is a given that designs often become more complex and sophisticated as they evolve from one version to the next. We are then naturally loathe to discard the immense amount of safety knowledge collected during the production of a previous version of the system, and documented in safety artifacts such as safety arguments, hazard analyses, test data, etc. In analogy with incremental design, it appears both reasonable and practical to take advantage of these safety artifacts and, whenever possible, e.g., if design changes are deemed "small" or systems are "sufficiently" similar, to reuse them so that safety engineers may focus their attention specifically on the effects of what has changed. This attempt to localize and focus safety assurance efforts by reusing safety artifacts is what we call an *incremental approach to safety assurance*, something that we further make clear in Sect. 4, where we present some real-life examples from the automotive and medical domains.

3 The Pitfalls of Incremental Safety Assurance

In this section we discuss some pitfalls associated with what we call an incremental approach to safety assurance. Our conclusion is that this approach to safety assurance cannot simply rely on principles analogous to those of incremental design improvement. If it does, it is unsound. This conclusion hinges on two main points. First, in contrast to what happens in incremental design improvement, safety assurance artifacts are not compositional. Second, while incremental design improvement is conducive to localization in terms of design parts, safety assurance requires a holistic view of the system. We elaborate on these points in Sects. 3.1 and 3.2, respectively.

3.1 Compositionality of Safety Artifacts

Regarding compositionality, the general idea of a safety argument provides us with a necessary context for discussion.

It is well-known lore that an argument is a series of assertions, in which the last element, the *conclusion*, *follows from* some foregoing assertions in this series, the *premisses*. More precisely, from an inferential standpoint, to 'follow from' means that the conclusion is obtained from the premisses by virtue of some judiciously chosen rules of inference. The bar against which an argument is then judged as being well-formed or not, i.e., right or fallacious, rests on an analysis of the properties that are satisfied by these rules of inference. In that respect, classical logical studies restrict their attention to the rules of inference

of the propositional and the predicate calculus, or some of their variants, such as those dealing with modalities. Rules of inference of this sort, henceforth called classical, enjoy the desirable property of being definite, i.e., they are not subject to rebuttal. This entails that if a conclusion follows from some premisses and some other conclusion follows from some other premisses, then, both conclusions follow from the union of their sets of premisses. In other words, if arguments are formulated in terms of classical rules of inference, then, they are *compositional*.

A safety argument is an argument whose main concern is the safety of an engineered artifact. Now, by looking at a safety argument, we can readily conclude that the rules of inference used in its formulation are far from being adequately captured as classical rules of inference (after all, we have yet to see definite safety claims). On the contrary, our view is that, whenever made explicit, safety arguments are formulated using defeasible rules of inference, i.e., rules of inference that are open to revision or annulment, e.g., as made precise in Toulmin's notion of a *rebuttal* [22]. This view of safety arguments makes them radically different from classical arguments; it makes them *non*-compositional. More precisely, as is well-known in the field of defeasible reasoning, in the presence of defeasible rules of inference, while a conclusion follows from some premisses, and while some other conclusion follows from some other premisses, neither of these conclusions may follow from the union of their sets of premisses [11].

In short, the preceding discussion indicates that composing safety arguments incrementally suffers from the inherent problem that this composition step is clearly unsound. In consequence, if safety arguments are built resorting to defeasible rules of inference, then, their compositionality requires principles that are radically different from those underpinning what can be done incrementally.

3.2 Localization of Safety Assurance Efforts

Regarding localization, the general idea of a safety goal decomposition provides us with the necessary context for discussion.

In essence, safety goal decomposition involves the mapping of safety claims across different levels of the design hierarchy. At the highest levels of design, some general safety claims are made. At lower, more detailed, levels of design, these general safety claims are refined into more specific safety claims, e.g., as safety claims concerning design parts. Fundamental to the soundness of safety goal decomposition is the assumption that any refinement step encompasses a full knowledge of the design elements it involves, how these elements interact, how these interactions may fail, and what measures can be put in place so that safety claims are not violated. When looked at from this perspective, safety goal decomposition requires a holistic view of the design at hand.

This said, the design hierarchy reflected in safety goal decomposition has led some to believe that design parts may be straightforwardly replaced by others which are substantially equivalent in terms of the safety properties they satisfy. For us, this is a serious misconception. What the previous chain of reasoning fails to take into account is that safety claims are not obtained in a localized fashion, but instead are the result of a refinement mechanism which accounts for

a holistic view of the design hierarchy. If any design part were to be replaced in any refinement step, not only would it be required to reassess the safety of the design parts involved in this refinement step, but also to reassess the safety of the design as a whole. The latter is largely due to the *emergent properties*, i.e., those arising from unexpected interactions between the replaced part and the rest of the system [17]. Because of their implications for safety, and given that they are not easily identified in the functional decomposition of a design, emergent properties are to be dealt with explicitly and seriously; failing to consider them is a serious omission in incremental safety assurance.

In short, safety assurance efforts cannot easily be limited to the modified design parts without considering a holistic view of the system. This means that, whether design parts and their corresponding safety artifacts may be replaced, or "plugged-in", modularly, without completely undermining what has thus far been deemed safe, requires principles that are radically different from those underpinning incremental design improvement. The approach is otherwise unsound.

4 Substantiating Our Claims

Focused on what we view are some of paradigmatic examples of safety gone wrong, in Sect. 4.1 we discuss the case of GM's faulty ignition switch, and in Sect. 4.2 we discuss the case of J&J's DePuy Orthopedics all-metal hip implants. We argue that these two real-life examples illustrate how what we call an incremental approach to safety assurance presents itself in practice.

4.1 Automotive Domain: The Ignition Switch Case

Not long ago, GM was faced with the recall of 2.6 million cars because of a defective ignition switch. The problem? The defective ignition switch would unintendedly move out of the "run" and into the "accessory" position during driving, leading to a partial loss of electrical power and turning off the car's engine. Why is this a problem? Under certain conditions, this accidental turn off of the car's engine resulted in an unfortunate series of events, which caused serious harm or death for car occupants; e.g., in a number of cases, this failure disabled the power steering, the anti-lock brakes, and the airbags, causing some fatal car crashes.

For us, GM's defective ignition switch problem is a glaring example of what may go wrong with an incremental approach to safety assurance. Why? GM found out that the problem with the ignition switch was the result of a new switch indent plunger that did not supply enough torque to keep the ignition from accidentally changing position [20]. It seems that, GM first became aware of the problem in 2001 and started to make incremental changes to the plunger part to address the issue in 2006 [13]. What went wrong with these changes? At least two things. First, our view is that when making the design change, GM engineers focused on meeting the specifications of an ignition switch, deeming unlikely that

this would introduce any new system level hazards. In a sense, the emphasis was placed on the physical and structural aspects of the design of the ignition switch. Second, this seemingly physical modification had a bearing on the the overall safety of the car. Most likely, the software requirements at the conceptual level of the car assumed that the car is not in motion when the key is in "accessory" mode. If the car is assumed not to be moving when the key is in "accessory" mode, it is reasonable to deactivate the airbags in order to prevent unintended deployment (in a parked car the accidental deployment of airbags could seriously injure passengers as they enter or exit the car). With the defective ignition switch, the assumptions underlying these software requirements are undermined. It was indeed possible for the car to be in motion with the key in "accessory" mode, e.g., as a result of hitting a bump on the road. (To be noted, the latter did not occur in cars prior to the problematic ignition switch design, where more torque was required to change the key position, virtually eliminating the possibility of the key accidentally changing position).

Can safety be assured locally? Replacing an indent plunger, a seemingly local issue, has global safety implications, exactly because of the intervention of a software-based control system. Concentrating solely on the physical or the software based aspects of the ignition switch may miss the real safety consequences. What the defective ignition switch misses is a global impact analysis of the design changes. This may have allowed an assessment of which other elements might have been affected and what new hazards this design change could have introduced. But this is easier said than done. It was not trivial for GM engineers to link the infrequent cases of airbags not deploying in an accident after loss of power steering and power brakes to the defective ignition switch [4,24].

As a final remark, touching on the notion of what has been *proven in use* and its potential contribution to the safety of a newer car, the determination of what is safe is intrinsically an evolving notion. Namely, small design changes, such as changing an indent plunger, may have worked well in the past. Yet, in the past, losing power to the car may have not been considered to be a catastrophic failure. E.g., in the past, failure of power steering and brakes would still leave the driver with some measure of control via manual steering and the mechanical connection to the brakes. In this past, an engineer dealing with mechanical components may view the change of the key position as an undesirable event that could result in a hazardous situation, but the hazard 'loss of control leads to an accident' would have been seen as being mitigated by the manual system. Supporting these claims, at a lawsuit resulting from a fatal accident, an engineer testified that the car was "safe" because it "could still be maneuvered to the side of the road" [13]. People have different expectations nowadays.

4.2 Medical Domain: The All-Metal Hip Implant Case

The FDA 510(k) substantially equivalent (SE) criterion for clearance of a new medical device is another example of what may go wrong with an incremental approach to safety assurance. Why? By its definition, the SE criterion relies on a comparison of a to-be-marketed with an already marketed medical device. In

essence, if changes in design are deemed to be "minor" or "small", inferences about the safety of the newer device can be made based on the safety artifacts of the device already marketed. Framed somewhat differently, the 510(k) SE criterion assumes that a small change in design will not likely bring about a major safety concern. As shown below, this assumption is, at least, problematic.

As reported in [12], in 2005, Johnson & Johnson's DePuy Orthopedics introduced a new all-metal design for their hip implants. A predecessor version of these hip implants was made of metal and plastic. The newer hip implants were cleared for market with the older hip implants being used as a predicate device using the 510(k) SE criterion. The new all-metal hip implants were cleared for market based on the fact that their predecessor had been cleared for market. No clinical trials nor additional tests were performed on the all-metal hip implants. Thus far, nothing seems to be wrong from an incremental safety assurance perspective (more so, a case can be made that the operational principle and the normal configuration are likely to be sufficiently similar, if not the same, for both the all-metal hip implant and its predecessor). The problem? It turned out that for the case of the all-metal hip implant "[t]he metal was eroding, releasing metallic particles into the blood and surrounding tissue of the joint and causing tremendous pain" [12]. This did not occur with the predicate device. It seems that drawing analogies between designs being substantially equivalent bears no obvious relationship to their safety. How can such a threat to safety be discovered if not by re-examining and carrying a thorough re-conceptualization of previously produced safety artifacts? Moreover, the all-metal hip implant is interesting for its ancestry, which can be traced back "more than five decades through a total of 95 different devices, including 15 different femoral heads and sleeves and 52 different acetabular components" [21]. It seems reasonable to assume that, even in the presence of impeccable initial safety artifacts, the compounded effects of design changes led up to a point where a new hazard was indeed present. This raises the question: do the small tweaks eventually get you?

5 Discussion

In hindsight, the threats to safety mentioned in Sects. 3 and 4 could have been mitigated with a proper preparation, revision, and perhaps re-conceptualization of the previously produced safety artifacts. Special attention must be given to impact that design changes may have on safety (potentially having to conduct new hazard analyses, reevaluate safety assumptions and the contexts in which these assumptions were made, etc.). Being able to count on a framework enabling the tracing of design changes to safety artifacts is a MUST, since it is precisely this framework that may enable the assessment of the effects of localized design changes on safety related artifacts. It is at this point where the notion of a safety case comes into play (a notion popular in some domains, but not so much in others). We believe that there is a version of incremental safety assurance that can take the necessary holistic view of safety assurance and perhaps offer a sort of middle ground between the present practices in many industries and the uneconomic approach of building all safety artifacts from scratch. Our hypothesis

is that this middle ground would need an explicit safety case in terms of which to assess the impact that an incremental design change may have on safety. Such an explicit safety case may then lead to some ability to localize required changes to safety artifacts, yet not necessarily in the sense of localization to design parts. The moral of the story? Reuse of safety artifacts can only be sound if we are able to trace the global effects that design changes may have on the system.

This said, having a well-defined notion of a safety case is only a part of the big picture. As we have argued above, an incremental approach to safety assurance cannot be based on principles similar to those of incremental design improvement. We are of the view that reusing safety artifacts requires *rely/guarantee*-like engineering principles, as understood by the formal methods community [7]. Intuitively, these principles may be understood as: the guarantee properties of this safety artifact are met only if the *rely* properties of a safety artifact are met. How hazard analyses, safety related evidence and arguments, test libraries, etc., are to be dealt with in a rely/guarantee fashion is something largely to be explored.

In summary, while we acknowledge that there is great practical value in the reuse of safety artifacts, this has to be done with a great deal of caution. We take as foundational that any incremental approach to safety assurance cannot be based on those engineering principles underpinning incremental design. Insofar as its soundness is concerned, what is then needed are engineering principles allowing for an analysis of the effect that a design change may have on safety artifacts. Among many things, these principles must involve a careful and thorough review of the validity of safety arguments. This would enable us to identify whether a safety argument contains some fallacious inferential steps and to assess the degree of certainty of the safety claims it involves. As usually conveyed in safety discussions, we view a safety argument as a cornerstone in safety assurance. Without a safety argument that links safety evidence with safety claims, it is well-nigh impossible to establish either the relevance and the sufficiency of the provided evidence, or how this evidence contributes to the safety claims. For us, this needs, as a first step, a precise definition of a safety argument, i.e., there is a need of a logic for safety argumentation (this logic need not be a formal logic, but it must be a logic nonetheless). Moreover, it is our view that emergent properties require special attention in safety assurance, as these pose one of the greatest threats to safety being assured in an incremental fashion. All in all, what is needed is a framework allowing for safety artifacts to be traced back to the design parts under consideration, enabling an analysis of effect propagation of localized changes, such as those caused by the addition of a new functionality or the replacement of a design part. An explicit safety case is a first step in the right direction.

6 Related Work

The need for an explicit and properly defined safety case is well-recognized in the safety community. There is, however, some disagreement regarding what counts as a "properly" defined safety case. In this respect, we are pluralists: maybe there is no THE properly defined safety case, but properly defined domain specific

safety cases. This said, we take as basic that a properly defined safety case shall consist of explicit safety goals, evidence of their fulfillment, acceptance criteria for the evidence, and a structured argument linking evidence to safety goals.

Among other places, the need for having an explicit safety case is commented on by Holloway in [15]. Holloway stresses that this is indispensable for evaluating the reasons why safety assurance practices in the aeronautics domain have thus far been adequate. Holloway makes this claim in reference to compliance with DO-178C, a standard which regulates the use of software on commercial airplanes, in an industry considered to be mature when it comes to safety matters. Our standpoint here is somewhat similar: without an explicit and adequate representation of a safety case, its analysis is close to impossible, as is the impact that design changes may have on safety. Works such as [2,3,8,19] also stress the importance of having an explicit representation of a safety case. However, in comparison to ours, these works are focused on what a safety case should look like, not on the problems with an incremental approach to safety assurance.

Particularly interesting in the context of incremental safety assurance is [18]. The authors of this work comment on how refinement, as understood by the formal methods community, allows for a much needed feature in incremental safety assurance: the introduction of more detail into the decomposition of safety goals. As a challenge of adopting such a technique for decomposing safety goals they point out that refinement leaves little room for revision. This is a consequence of refinement being conceived in a (logically) monotonic setting. The situation is radically different once one assumes safety properties are defeasible, as we have discussed in Sect. 3. In such a setting, the traditional ideas of refinement do not apply straightforwardly (e.g., it may be the case that refining a safety goal into two safety subgoals results in one of the subgoals undermining the other). Considering this phenomenon is crucial if safety assurance is to be thought of incrementally.

Works such as [1,6,9,14] are also related to incremental safety assurance. All of these works have in common with ours a discussion of safety being assured in an incremental fashion. However, in comparison with ours, their approach is presented from the point of view of techniques rather than principles. In that respect, they do not seem to discuss the issues that we have commented on in Sects. 3 and 4. Though they address and suggest a component based approach to safety assurance, they do not discuss how such components may be put together in a property preserving manner.

7 Some Final Remarks

Given that incremental design improvement is prevalent as an engineering practice, it is no surprise that matters related to the associated idea of incremental safety assurance appear in various safety standards and guidelines via the reuse of design elements.

The automotive domain incorporated the notion of *proven in use* in the recently published ISO 26262 standard for the functional safety of vehicles.

In ISO 26262 defines proven in use as "an alternate means of compliance [...] that may be used in the case of reuse of existing items or elements when field data is available" [16, Part 8, Clause 14]. ISO 26262 also introduces the concept of *safety element out of context* (SEOoC). A SEOoC is "intended to be used in multiple different items when the validity of its assumptions can be established during integration of the SEOoC" [16, Part 10, Clause 9]. Both 'proven in use' and SEOoC fall within an incremental approach to safety assurance under the assumption that they involve the reuse of the safety artifacts attached to a design element, with the purpose of contributing to the safety of a newly developed car.

The medical domain has its well-known 510(k) process. The US FDA defines the so called '510(k) program' as "a premarketing submission made to FDA to demonstrate that the [medical] device to be marketed is as safe and effective, that is, substantially equivalent (SE), to a legally marketed device that is not subject to premarket approval (PMA)" [5]. The SE condition indicates that the changes incorporated into the new medical device are somewhat "small" in relation to the already marketed medical device, from which the new device's safety follows. If looked at from this perspective, the FDA's 510(k) program is another instance of an incremental approach to safety assurance: small design changes cause no effect on the artifact's safety.

In avionics, an incremental approach to safety assurance may be seen as being present in the FAA's AC 20-148: Reusable Software Components [10]. In this advisory circular, the FAA comments that "because of economic incentives and advances in software component technology, software developers want to develop a *reusable software component* (RSC) that can be integrated into many systems' target computers and environments with other system software applications", all while still showing compliance with avionics safety regulations. As with ISO 26262's notion of a SEOoC, if we agree that a RSC involves the reuse of safety artifacts, then, it is more or less clear that this falls within the scope of what we call an incremental approach to safety assurance.

Following from the observations just made, to be noted is that, while the definitions and practices may vary across domains, a great deal of care should be taken so that these safety standards and guidelines are not undermined by the pitfalls and deficiencies that we discussed in Sects. 3 and 4.

8 Conclusions and Next Steps

Incremental design improvement, a.k.a. normal design, is a reliable and standard foundation for engineering practice. It is well understood, generally economic, and it supports the need and desire to see improvements in the artifacts that we use. When these are safety critical, the question becomes: how are safety related issues, arising due to changes in design, to be incorporated into the safety assurance scheme? We have argued that the obvious analogy to incremental design improvement encounters serious difficulties related to identifying new or re-emerging safety issues.

We have also discussed some of the principles and examples of an incremental approach to safety assurance. Resorting to the latter, we have shown that

safety related issues were missed, leading to some catastrophic results. In our view, shared by some, the fundamental problem, the root cause of the mistakes, is that, even though design changes might be local, as in the ignition switch example, their effects on safety assurance are of a global nature. More generally, mistaking incremental design change for limited effects on safety has resulted in essential difficulties related to safety, and serious damage to people, completely undermining claims about safety. This has been worsened by the fact that the safety cases often remain implicit, making it very difficult to determine the global safety effects of the localized design change. Of course, we do recognize that when safety engineers have a great deal of experience, and they devote sufficient attention to the effects of design modifications on safety artifacts, things appear to run smoothly, even if approached incrementally. The problem is that this is difficult to evaluate externally, i.e., without the inside knowledge these safety engineers may have. If looked at from this perspective, rather than an engineering discipline, safety assurance becomes something that falls within the realm of obscurantism and practiced by safety gurus.

Conversely, our position is that, incremental safety assurance needs principles other than those underpinning incremental design improvement. These principles will define the basis for analyzing how incremental design changes impact existing safety artifacts. Thus, our recommendation goes beyond that of producing an explicit safety case. This said, safety cases are definitely necessary. It is with respect to them that the effects that design changes may have on safety may be tracked down more easily, establishing a foundation for eliciting sound engineering principles for incremental safety assurance. In any case, our proposal is not the one usually put forward in the context of safety assurance: start afresh from the ground up. We recognize that while perhaps viable in domains where changes in design seldom occur, this is economically and logistically infeasible when changes in design are frequent, as is the case in the automotive and medical domains. As future work, we need to rigorously develop and systematize our hypotheses, so that they can be evaluated in carefully conducted experiments.

Acknowledgments. The authors wish to acknowledge the support of the Automotive Partnership Canada, the Ontario Research Fund, and the Natural Sciences and Engineering Research Council of Canada.

References

1. Althammer, E., Schoitsch, E., Sonneck, G., Eriksson, H., Vinter, J.: Modular certification support - the DECOS concept of generic safety cases. INDIN **2008**, 258–263 (2008)
2. Birch, J., Rivett, R., Habli, I., Bradshaw, B., Botham, J., Higham, D., Jesty, P., Monkhouse, H., Palin, R.: Safety cases and their role in ISO 26262 functional safety assessment. In: Bitsch, F., Guiochet, J., Kaâniche, M. (eds.) SAFECOMP. LNCS, vol. 8153, pp. 154–165. Springer, Heidelberg (2013)
3. Birch, J., Rivett, R., Habli, I., Bradshaw, B., Botham, J., Higham, D., Monkhouse, H., Palin, R.: A layered model for structuring automotive safety arguments. In: European Dependable Computing Conference (2014)

4. Bunkley, N.: GM engineer says he didn't remember changing ignition switch part. Automotive News 28 May 2014. http://www.autonews.com/article/20140528/OEM11/140529859/gm-engineer-says-he-didnt-remember-changing-ignition-switch-part
5. Center for Devices and Radiological Health: Device approvals, denials and clearances 4 June 2014. http://www.fda.gov/medicaldevices/productsandmedicalprocedures/deviceapprovalsandclearances/default.htm
6. Conmy, P., Nicholson, M., McDermid, J.: Safety assurance contracts for integrated modular avionics. In: 8th Australian Workshop on Safety Critical Systems and Software (SCS 2003). vol. 33, pp. 69–78. Australian Computer Society (2003)
7. de Roever, W., et al. (eds.): Concurrency Verification: Introduction to Compositional and Non-compositional Methods. North-Holland, Amerstadam (2007)
8. Dittel, T., Aryus, H.-J.: How to "Survive" a safety case according to ISO 26262. In: Schoitsch, E. (ed.) SAFECOMP 2010. LNCS, vol. 6351, pp. 97–111. Springer, Heidelberg (2010)
9. Elmqvist, J., Nadjm-Tehrani, S.: Tool support for incremental failure mode and effects analysis of component-based systems. In: DATE 2008. pp. 921–927 (2008)
10. Federal Aviation Administration: Ac20-148: Software reusable components (2004)
11. Gabbay, D.M., Woods, J. (eds.): Handbook of the History of Logic: The Many Valued and Nonmonotonic Turn in Logic, vol. 8. North-Holland, Amsterdam (2007)
12. Groeger, L.: Four medical implants that escaped FDA scrutiny 30 April 2012. http://www.propublica.org/special/four-medical-implants-that-escaped-fda-scrutiny
13. Gutierrez, G., et al.: GM chose not to implement a fix for ignition problem. NBC News 13 March 2014. http://www.nbcnews.com/storyline/gm-recall/gm-chose-not-implement-fix-ignition-problem-n51731
14. Hatcliff, J.A.L.K., Lee, I., Macdonald, A., Anura, F., Robkin, M., Vasserman, E., Weininger, S., Goldman, J.: Rationale and architecture principles for medical application platforms. In: ICCPS 2012. pp. 3–12 (2012)
15. Holloway, M.: Making the implicit explicit. In: ISSC 2013. Boston (2013)
16. Internatiional Standard Organization: ISO 26262: Road vehicles - Functional safety (2011)
17. Johnson, C.W.: What are emergent properties and how do they affect the engineering of complex systems? Rel. Eng. & Sys. Safety **91**(12), 1475–1481 (2006)
18. Lisagor, O., Kelly, T.: Incremental safety assessment: Theory and practice. In: Proceedings of 26th International System Safety Conference. Minneapolis (2008)
19. Palin, R., Ward, D., Habli, I., Rivett, R.: ISO 26262 safety cases - Compliance and assurance. In: Procedings of 6th IET International Conference on System Safety, pp. 1–6 (2011)
20. Spangler, T.: Delphi told GM Ignition Switch Didn't Meet Specs. Detroit Free Press, Michigan (2014). http://www.usatoday.com/story/money/cars/2014/03/30/gm-ignition-switches-recall-congressional-report/7085919/
21. Thompson, H.: Researchers say DePuy hip ancestry shows 510(k) flaws 19 February 2013. http://www.mddionline.com/article/researchers-say-depuy-hip-ancestry-shows-510k-flaws
22. Toulmin, S.E.: The Uses of Argument. Cambridge University Press, Cambridge (2003)
23. Vincenti, W.: What Engineers Know and How They Know It: Analytical Studies from Aeronautical History. The Johns Hopkins University Press, Baltimore (1993)
24. Wald, M., Vlasic, W.: 'Upset' GM engineer spoke in house inquiry. The New York Times 28 May 2014. http://www.nytimes.com/2014/05/29/business/upset-gm-engineer-spoke-in-house-inquiry.html?_r=0