Cybersecurity Assurance Challenges for Future Connected and Automated Vehicles

Luis-Pedro Cobos

Vehicle Resilience Technologies, HORIBA MIRA Limited, United Kingdom.
E-mail: luis-pedro.cobos@horiba-mira.com

Alastair R. Ruddle

Vehicle Resilience Technologies, HORIBA MIRA Limited, United Kingdom.
E-mail: alastair.ruddle@horiba-mira.com

Giedre Sabaliauskaite

Institute for Future Transport and Cities, Coventry University, United Kingdom.
E-mail: ad5315@coventry.ac.uk

Increases in the connectivity of vehicles and automation of driving functions, with the goal of fully automated driving, are expected to bring many benefits to individuals and wider society. However, these technologies may also create new cybersecurity threats to vehicle user privacy, the finances of vehicle users and mobility service operators, and even the physical safety of vehicle occupants and other road users. Assuring the cybersecurity of future vehicles will therefore be key to achieving the acceptability of these new automotive technologies to society. However, traditional prescriptive assurance methods will not work for vehicle cybersecurity, due to the evolving threats, through-life software updates, and the deployment of artificial intelligence techniques. Cybersecurity regulations that are goal-oriented and risk-based, like those increasingly used in safety engineering for complex systems, are now mandated in recent vehicle type approval regulations. This results in many new assurance challenges, which will not be limited purely to cybersecurity. In particular, emerging standards have proposed that an assurance case approach should be adopted in relation to cybersecurity. This paper therefore proposes a novel cybersecurity case framework that adapts existing approaches from safety engineering, emphasizes the limitations of the analysis through eliminative argumentation, and merges in the attack-defence tree techniques used in cybersecurity engineering, with the aim of providing a better reflection of the some of the uncertainties in the cybersecurity risk analysis.

Keywords: Assurance, automated driving, connected vehicles, cybersecurity, risk, safety, software updates.

1. Introduction

The world is not shy of the emerging automated driving vehicles; they are becoming more of a reality with each passing moment. Although most of the established vehicle manufacturers have active development programs on this topic, it has also garnered considerable interest from several new market entrants. As the industry dives into the world of autonomous vehicles, where reliance on electronics is key to support vehicle autonomy, it opens the road to completely new issues to consider, like those of cybersecurity.

The emergence of cybersecurity threats in the automotive industry results not only from rising levels of electronic control, but also from the expansion of outward looking sensing and wireless connectivity that are being used to support and enable the automation of driving functions (El-Rewini et al., 2020) and the development of intelligent transport systems (Parkinson et al., 2017). It is expected that vehicles will become elements of a wider connected and automated mobility ecosystem. Software is a key enabler for advanced vehicle functions, and it is expected that in-vehicle software will in future be subject to through-life updates, using software changes, either to improve performance or to add new functionality. Despite their many benefits to individuals and society, these developments have also inadvertently created potential opportunities for malicious groups and individuals to encroach on privacy, attempt financial fraud and extortion, and even threaten the physical safety of vehicle occupants and other road users. Assuring the cybersecurity of future vehicles will therefore be key to achieving the acceptability of new automotive technologies to society.

These concerns have resulted in the United Nations Economic Commission for Europe (UN-ECE) formally adopting two new vehicle type

Proceedings of the 31st European Safety and Reliability Conference
Edited by Bruno Castanier, Marko Cepin, David Bigaud, and Christophe Berenguer
Published by Research Publishing, Singapore.
ISBN: 978-981-18-2016-8; doi:10.3850/978-981-18-2016-8_412-cd

approval regulations for new vehicles: Regulation 155, relating to automotive cybersecurity (UNECE, 2021a); and Regulation 156, regarding the safety and security of software updates (UNECE, 2021b). In order to be placed on the market in Europe, new vehicle models will need to comply with Regulation 155 by July 2022, while existing vehicle models will need to comply by July 2024. However, standards to support the implementation of Regulation 155 are currently still under development. These include ISO/SAE 21434 to support automotive cybersecurity engineering (ISO/SAE, 2021), and ISO/AWI PAS 5112 concerning the associated auditing requirements (ISO/AWI, 2021).

This paper considers the challenges and possible approaches for assuring the cybersecurity of vehicles, taking account of the legislative requirements and current automotive industry trends, as well as existing practice in the safety domain. In particular, a novel approach to the development of a cybersecurity case is proposed, which aims to provide a better reflection of the some of the uncertain aspects of cybersecurity risk analysis.

To start, Section 2 provides an outline of the difference between traditional prescriptive assurance methods and the newer goal-based approaches that are employed in vehicle functional safety, and now in cybersecurity, as well as an overview of the assurance case approach. A comparison between safety and cybersecuriy considerations is provided in Section 3, and Section 4 considers some of the assurance challenges that result from the emerging regulations, for both cybersecurity and safety. A brief overview of cybersecurity risk management is provided in Section 5 to introduce the notion of attack-defence trees, which are integrated into a graphical-style cybersecurity case along with eliminative argumentation approaches in Section 6. Finally, Section 7 outlines the conclusions and future outlook.

2. Automotive Industry Assurance Approaches

Ensuring the safety of vehicle occupants and other road users has long been established as an extremely important aspect of vehicle development. However, rising system complexity and the increasingly rapid pace of technological change have necessitated significant changes in the way that safety assurance is achieved.

2.1. *Traditional (Prescriptive) Methods*

The traditional approach to product assurance is highly prescriptive, detailing not only the required performance criteria, but also specifying exactly how compliance with these criteria is to be demonstrated. This type of approach is well-known and understood, with clear advantages in terms of simplicity and transparency. It is very well suited to relatively simple systems with few functions. When we consider the more traditional approaches of vehicle safety in passive and active safety it is highly reliant on a prescriptive approach, but as reliance on electronic control becomes more vital for the systems, delving into the safety of the electronics grows more complex. These are no longer independent mechanical systems; they communicate with each other and are highly automated.

However, the prescriptive approach becomes increasingly difficult as the complexity of the target system rises, resulting in a richer set of functions and such large numbers of states that comprehensive testing is no longer a practicable option (Kelly et al., 2005). In addition, the prescriptive approach is inevitably technology-centric, since it aims to specify the details of how and what are required to be done. This makes it difficult for prescriptive standards and regulations to adapt to new technology, since the acceptance criteria and validation methods are so closely related to the anticipated technology of the product. As different technological solutions emerge, the number of standards and regulations must multiply to accommodate the newer options. While the pace of technological change has been modest this burden has been manageable, but this is not expected to remain the case in future.

2.2. *Goal-Based Methods*

For more complex systems it becomes increasingly difficult to identify specific performance criteria and demonstration methods that can be used to assess satisfactory behaviour. The range of behaviours and operating states is too large to test comprehensively. A common response to these issues is to define goals, often described in terms of risk targets, that the system should meet. The manufacturer must then present an argument as to how the required goals have been achieved. A widely used approach to this is the construction of an assurance case, as used in automotive functional safety (ISO, 2018), and safety of the intended functionality (SOTIF), which is described in ISO/PAS 21448 (ISO, 2019).

2.2.1. *Assurance Cases*

There are different techniques and ways to demonstrate that the risks associated with using a system are acceptable. An assurance case approach is widely used to describe the links between claims, arguments and evidence. In safety assurance it has become common to require a safety case that explicitly states the safety claims, which are often risk-based, and documents the structured arguments that link those claims to the supporting evidence and associated assumptions.

The argument in assurance cases provides a logical link between the evidence and a claim. There

are two kinds of induction commonly used in argumentation: *enumerative induction* and *eliminative induction*. In enumerative induction, confidence increases as greater numbers of confirming examples are found. In eliminative induction, on the other hand, confidence in the validity of a claim increases as reasons for doubt are eliminated (through evidence or argument). The process is inevitably an idealization there is always some un-eliminated (residual) doubt in an argument, as can be seen in (Goodenough et al., 2015).

2.2.2. *Assurance Case Notations*

Assurance cases tend to follow a CAE (Claim-Argument-Evidence) notation. The Goal Structuring Notation (GSN) is a safety case form that visualizes an argument structure that supports a claim to be true. In industries in which safety assurance is critical, standards such as IEC 61508 for general process control (IEC, 2010), ISO26262 in automotive applications (ISO, 2018), DO-178C in the aircraft industry (RTCA/EUROCAE, 2011) etc., require suitable documentation and GSN is the standard format to document safety cases graphically as can be seen in (Kelly and Weaver, 2004).

The GSN safety case structure works towards a "goal", which is a claim, through a strategy and context, which are arguments, then leads to a solution, that is evidence. The "vanilla GSN" mentioned also has the possibility of extensions that make the safety case building more effective. These extensions include: maintenance of arguments, modular safety cases, assurance case patterns, eliminative argumentation and more. As a technique that is easily understood and accepts advanced concepts, it has been appearing with increasing frequency in the domains of safety and security, something mentioned in (Kelly and Weaver, 2004).

2.2.3. *Eliminative Argumentation and Confidence Maps*

In eliminative argumentation there are three potential types of doubt; doubts about the claim, doubts about the evidence, or doubts about the inference used to link the claim and the evidence. The objective is therefore to identify the relevant doubts about claims, evidence and inferences, and to provide counter arguments against the identified doubts where possible to increase confidence in the assurance case (Goodenough et al., 2015).

A graphical representation of an eliminative argument is described as a *confidence map*, as it details the identified doubts concerning an argument and also shows whether these doubts can be countered or if they remain, thus illustrating the confidence that can be attributed to the argument. It should be noted that not all doubts will have the same importance, and appropriate weightings should be. To maintain the clarity of the safety case it has been recommended (Hawkins et al., 2011) that the confidence map should be separate from, but linked to, the safety case.

3. Comparing Cybersecurity and Safety

There is a widespread view that cybersecurity and safety are essentially the same thing. This is perhaps partly due to the fact that many languages do not have separate words for safety and security. As (Kavallieratos et al., 2020) explain, there are over 86 methods of safety and cybersecurity co-engineering methods that mix and match various techniques. Out of these methods just 20 seem aware and able to follow safety-security regulations and standards. And not even half of the methods are capable of communicating the results clearly to stakeholders.

However, while it is undoubtedly true that some cybersecurity threats have potential safety consequences, they are not necessarily of the same nature as those that are considered in vehicle safety engineering. Furthermore, cybersecurity threats also have a much wider range of consequences beyond safety. For cybersecurity, safety hazards are in fact a subset of possible outcomes relating to interference with normal operation. Some attacks may interfere unacceptably with operational performance, in particular with functionality that may be regarded as mission-critical, without necessarily leading to any direct safety issues. Attacks that prevent authorised access to a vehicle or prevent a vehicle from starting, for example, do not create direct safety hazards for the authorised user. In addition, fraudulent financial transactions and compromised privacy are also possible and undesirable consequences of vehicle cybersecurity attacks that are not safety related.

With regard to product safety, manufacturers are not required to anticipate and take precautions against every conceivable use or abuse of their product. Nonetheless, manufacturers may have obligations with respect to uses of their product that, although not intended, are nonetheless "reasonably foreseeable" and some malicious attacks are reasonably foreseeable. However, the scope of functional safety considerations for vehicle programmable electronic and electrical systems, is limited solely to malfunctions (ISO, 2018). Although the scope of SOTIF does include reasonably foreseeable misuse (in the sense of unintended use), it specifically excludes intentional abuse (ISO, 2019). Thus, safety hazards that may result from cybersecurity threats are out of scope for mainstream vehicle safety engineering (see Fig. 1).

In other sectors, the boundaries between functional safety and cybersecurity are less rigid. For example, IEC 61508 (IEC, 2010) requires malevolent and unauthorised actions to be considered during functional safety hazard and risk analysis

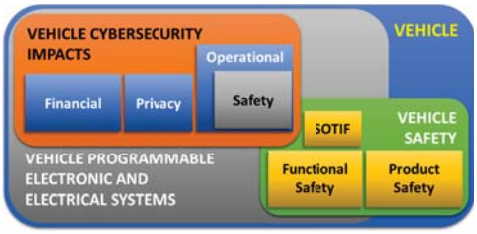

Fig. 1. Relationship between cybersecurity outcomes and safety considerations in relation to vehicle programmable electronic and electrical systems.

for industrial control systems. Although the safety hazards associated with vehicle cybersecurity are beyond the scope of automotive functional safety, there are similarities that can be usefully exploited in a unified approach that reflects the spirit of the analysis methods used in functional safety. However, adaptation is required, as impacts such as indirect safety threats (e.g. enabling occupant kidnapping or putting lives in danger during hijacking) or the potential to affect entire classes of vehicles are not considered in more traditional vehicle safety risk analyses.

Although safety and security considerations are not identical, there are some overlaps and also the potential for possible tensions between them. For example, safety requirements may include ready access to data that might help to avoid or reduce possible harm in certain situations, whereas security requirements may aim to severely limit access to this data. In addition to conflicts, there may also be conditional dependencies to be considered (Kavallieratos et al., 2020). Thus, there is certainly a need for cooperative *security-aware safety engineering* and *safety-aware security engineering* processes to be established. However, a fully integrated *safety-security co-design* approach may be difficult due to the significant differences that exist between the nature and scope of safety and security considerations.

Differences in the interpretation of the same or similar terminology between different disciplines can be a major cause of confusion and readily lead to misunderstandings and poor communication. Thus, the development of a common vocabulary and a unified approach, which are consistent across both disciplines whilst retaining their own unique features, would be significant enablers for such collaborative development activity.

4. Challenges for Cybersecurity and Safety Assurance

Conventional pre-launch assurance activities will remain essential to provide basic initial product assurance, including for cybersecurity. Cybersecurity threat and risk assessments are required to be carried out during the development phase, but this will not be sufficient in itself as existing cybersecurity threats evolve and new or unforeseen threats emerge. In addition, therefore, Regulation 155 (UNECE, 2021a) also requires vehicle manufacturers to establish an ongoing cybersecurity management process to complement the pre-launch activities and thereby ensure ongoing cybersecurity assurance throughout the operational life-cycle of the vehicle. Previously unforeseen attacks that are identified during the operational phase must be analysed, and where necessary, mitigated in order to maintain acceptable levels of risk.

Thus, a goal-based, risk driven approach, which is similar to that already used in automotive safety engineering (ISO, 2018), is now emerging in relation to vehicle cybersecurity. The draft standard (ISO/SAE, 2021) that is currently being developed to accompany Regulation 155 suggests the use of a *cybersecurity case*, which is analogous to the safety case approach already used to document automotive functional safety (MISRA, 2019), to provide assurance that the cybersecurity risks of using the vehicle are acceptable. Although the concept for a cybersecurity assurance case is not new (Armstrong et al., 2011), the results of a recent survey (Mohamad et al., 2021) over a wide range of application domains demonstrate that the practical implementation of cybersecurity cases remains relatively immature.

Various approaches have been developed for automotive cybersecurity risk analysis, as suggested in (SAE, 2016) and more recently reviewed in (Hao and Han, 2020). However, establishing reliable likelihood metrics for attacks remains a significant limitation in evaluating the associated risks. This is due to the wide range of possible attacker types and their specific motivations and access to resources (both technical and financial). Although technical difficulty and cost can be assessed in a relatively objective manner, the human motivation is a significant influence that is much more difficult to account for. This results in considerably more uncertainty than has historically been encountered in safety engineering. Thus, methods for making the limitations of the risk analysis more explicit are even more important in cybersecurity assurance.

Given that unforeseen cybersecurity threats are anticipated, and security patches may need to be deployed to mitigate the associated risks, the cybersecurity case would need to be a "live" document, to be updated throughout the operational life of the vehicle. Furthermore, this will also be the case for safety assurance in future, as Regulation 156 (UNECE, 2021b) requires both safety and cybersecurity risk assessments to be carried out for all in-service software updates, irrespective of their content and motivation.

The need to monitor and mitigate vehicle cyber-

security breaches through the operational life of vehicles also raises a need for the implementation of dedicated vehicle security operations centers (V-SOCs), which would provide monitoring of the operation of large numbers of vehicles as well as analysis capabilities for identifying, analysing and responding to emerging attacks. Although IT security operations centers (IT-SOCs) are already well established (Schinagl et al., 2015), the cybersecurity issues for cyber-physical systems such as road vehicles go beyond conventional IT security considerations. Many possible attacks on vehicles are more likely to be identified by their impact on physical traffic flow than by anomalous internal system behavior, since manipulation of the vehicle inputs (including environment sensors as well as GNSS and V2X signals) can modify vehicle behavior without the need to interfere with the internal vehicle systems. Nonetheless, in the connected and automated mobility ecosystem there will also be a need for interaction between the V-SOC and more conventional IT-SOCs.

In addition, both cybersecurity and safety assurance will require approaches to be developed that can accommodate non-deterministic AI and machine learning technologies, which are increasingly being used to support the higher levels of driving automation. However, the remainder of this paper concentrates on possible methods for addressing some of the issues that have been identified in relation to cybersecurity cases.

5. Risk Management for Automotive Cybersecurity

Although the elimination of all risks is impracticable (and would be unaffordable if practicable), risks can be managed to ensure that they are not unreasonable, by creating a risk management plan.

In the safety context, a hazard is the source of an accident or incident, and is something that may have repercussions. An incident is an event led by a hazard that does not cause losses, while an accident causes losses; these losses might be economic, health or life related. In cybersecurity, threats and attacks on the vehicle, perpetrated by malicious individuals, could lead to a variety of possible outcomes, including safety impacts as well as non-safety risks, as noted in Section 3.

5.1. *Risk Analysis*

Risk analysis is the process of identifying and analysing potential issues that could have a negative impact for the stakeholders. A risk analysis is a vital part of a risk assessment; it is a step in which all risks are identified and categorized, determining how significant each risk is, and hence the potential need for risk reduction measures.

Risk is a combination of the *likelihood* of an event and the *severity* of its impact on the stakeholders, such that low severity with a low likelihood represents a low risk and high severity with a high likelihood represents a high (and probably unacceptable) risk. Other combinations of severity and likelihood result in intermediate risk levels, and both the severity and likelihood are typically categorized in order to allow them to be mapped to risk categories.

A number of risk analysis approaches for automotive cybersecurity are outlined in (SAE, 2016) and (Hao and Han, 2020). It should be noted that the severity of impact can only be assessed at system level, where the impact on the stakeholders can be assessed, whereas the likelihood depends on the individual likelihoods of actions in the chain of events that lead to the specific outcome.

5.2. *Threat Modelling and Attack-Defence Trees*

Threat modelling must be appealing to the business objectives and security policies, but it also has to closely follow regulations and standards. When doing so it is important to consider the robustness of the vehicle, its surrounding environment and the motivations of attackers.

According to (Kordy et al., 2011) Attack-Defence Trees (ADT) provide a methodical, graphical way of modelling possible cybersecurity threats to systems. It is a popular technique for analysing threats in cybersecurity. The process to create ADTs is as follows:

- Develop a functional model of the system
- Propose possible attacker goals (i.e. illegal benefit to the attacker)
- Identify possible attack objectives that could allow the attacker to achieve these goals (i.e. possible harm or other loss to stakeholders)
- Identify attack methods that the attacker could use to implement the attack objectives
- Decompose the attack methods into lower level actions that would be required to achieve a successful attack.
- Identify opportunities to eliminate branches or mitigate the effects by reducing the likelihood of success.
- Let each goal form a separate tree (although they might share sub-trees and nodes)

ADTs are represented in a logical way and follow the node flow in one direction. As illustrated by the work of (Kordy et al., 2011), each node of an ADT that has more than one branch the relationships between the subsidiary branches may be either disjunctive (OR) or conjunctive, using either a simple AND or a sequential AND (SAND). The SAND approach provides a more compact representation a specific sequence of steps that may be needed to mount an attack.

ADT techniques have many advantages, for example they are easy to understand and can be

easily shared and explained to people with little experience in security, and can often be reused to address similar threats.

6. Proposed Approach Towards an Automotive Cybersecurity Case

The use of a cybersecurity case to document compliance with cybersecurity goals is suggested in ISO/SAE 21434 (ISO/SAE, 2021), but no further information is provided. As the safety case is already widely used for functional safety in the automotive industry, this provides a natural basis for developing an assurance case for cybersecurity.

A number of requirements for a cybersecurity case have already been identified in (Cobos et al., 2021). These were based on consideration of current automotive industry trends and the emerging type approval requirements relating to both cybersecurity and software updates. However, some of the requirements that have been proposed are also intended to respond to some of the criticisms of existing safety case techniques raised by (Leveson, 2011).

In particular, the perceived potential for undesirable confirmation bias and a lack of transparency concerning confidence and uncertainty seem even more relevant in the construction of a cybersecurity case than for a safety case, since there is even less certainty in relation to cybersecurity threats than there is in the safety domain. A possible solution to these aspects could be to take the attack-defence trees developed during cybersecurity risk analysis and integrate them into a GSN-type graphical assurance case argument using an eliminative style of argumentation. The aim of this is to produce a more explicitly "adversarial" case than has traditionally been used in functional safety, in a similar style to the way legal cases are presented and examined in a court of law. In fact, the recently launched version 3.0 of the GSN standard (SCSC, 2021) has also introduced new *dialectic extensions* in order to provide support this style of argument.

A generic Cybersecurity Assurance Case is proposed here that takes the general graphical approach of GSN, whilst applying the additional symbolism of confidence maps from eliminative argumentation as in (Goodenough et al., 2015), and also integrating the structure of ADTs to augment and amplify the arguments. The symbols used in the generic illustration presented here are summarized in Fig. 2, and their meanings are discussed further below.

Arguments taking account of the potential for unforeseeable cybersecurity risks and the requirements for a through-life cybersecurity management system (denoted by CSMS in the diagram) are presented in Fig. 3. The system description, which provides the context for the assurance case, is indicated by a white ellipse. The *claims* are represented by blue rectangles, which are justified by *inference rules* represented by green rectangles. The claims may be challenged by *rebutting defeaters*, the inference rules by *undercutting defeaters*, and the evidence by *undermining defeaters* (although the latter are not used in this illustration). These challenges may be responded to by using further claims, inferences and evidence. The need for a robust threat analysis approach is also included in Fig. 3, along with the treatment of a threat judged to be of inherently acceptably low risk, and considering a cybersecurity management system (CSMS). Lines of argument that are considered to have been acceptably resolved are indicated by the grey circles.

The diagram is continued in Fig. 4, which illustrates approaches for threats involving attack trees that could contain disjunctive and conjunctive relationships (the latter including both simple and sequential AND possibilities). These could be addressed either by outright elimination of possible attack steps (denoted by "ATK" and represented by circles with red boundaries), or at least mitigating the vulnerability to reduce the anticipated likelihood of success to a sufficiently low level to achieve an acceptable level of residual risk. These requirements for defense are indicated by dashed lines terminating in claims for possible successful elimination (denoted "ELIM" in a blue box) or mitigation countermeasures (denoted by "MIT" in a blue box), which therefore represent sub-claims that must be supported by appropriate evidence.

It should be noted that if an OR node occurs in an attack tree fragment then all of the options must be addressed with suitable countermeasures in order to achieve complete resolution. If there is an AND or SAND node, however, then mitigating any of the contributors could be sufficient to achieve the necessary risk reduction (in the ADT fragment denoted "ATK x" in Fig. 4, one of the required contributors is simply eliminated, thereby disabling that attack path). In practice, the requirements for mitigation might well be identified at lower levels of a specific attack tree network than is shown in the very abstract (non-specific) illustration presented here.

The idea is that this approach could be implemented into a future cybersecurity case frame-

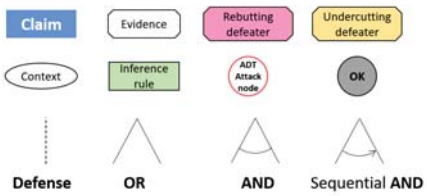

Fig. 2. Symbols used in proposed graphical cybersecurity assurance case (see Fig. 3–4).

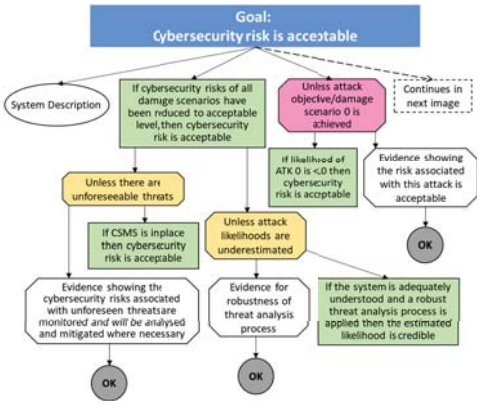

Fig. 3. Proposed approach (CSMS – Cybersecurity Management System; ATK – attack): part 1 (continued in Fig. 4).

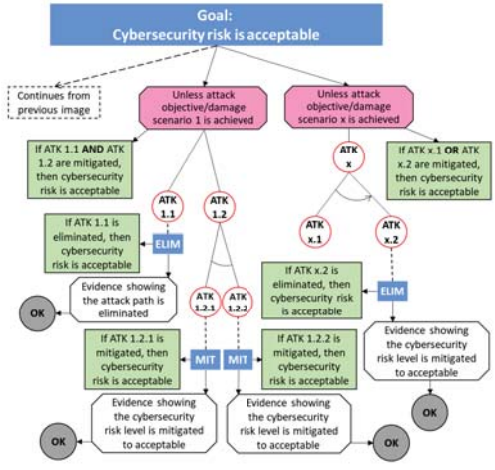

Fig. 4. Proposed approach (ELIM – elimination; MIT – mitigation): part 2 (continued from Fig. 3).

work, and more specific examples will be further developed in the line of the research.

7. Conclusions

Ensuring very high levels of dependability will be essential to achieve societal acceptability for automated driving. In order to address dependability, the assurance cases for automated driving will need to be extended considerably beyond those currently constructed for conventional automotive applications. This situation raises a wide range of challenges for validation, assurance and certification, which will need to become a much more wide-ranging and dynamic activity in future.

Traditional prescriptive assurance methods will not work for vehicle cybersecurity, which is not readily bounded and cannot be achieved with complete certainty. Cybersecurity therefore requires regulations that are goal-oriented and risk-based, like those increasingly used in safety engineering for complex systems. An assurance case approach is recommended, and constructing these assurance arguments is expected to help identify the requirements for the types of evidence needed to complete the assurance claims. However, recent surveys indicate that the practical implementation of cybersecurity cases remains relatively immature (Mohamad et al., 2021).

Criticisms of safety cases, including the perceived potential for undesirable confirmation bias and a lack of transparency concerning confidence and uncertainty, seem even more relevant in the construction of a cybersecurity case. Consequently, this paper has proposed a novel framework adapting existing approaches from safety engineering and mixing them with the specific analysis techniques used in cybersecurity engineering. The proposed approach integrates ADT from cybersecurity risk analysis and eliminative argumentation styles within a GSN-based diagram as a means to develop a more risk-focused assurance case for automotive cybersecurity. However, these are very preliminary ideas, which will be further refined and evaluated in subsequent work.

It is expected that such a cybersecurity case, prepared initially for product launch, would effectively be the first draft of a dynamic assurance case that would be updated through the operational life of the vehicle. Ongoing assurance activities will also be needed to complement the product launch assurance, in order to ensure that cybersecurity assurance is maintained over the operational lifetime of the vehicle as outlined in the UNECE regulations and ISO/SAE 21434. This will require the development of dedicated vehicle security operations centers to help ensure the through-life cybersecurity performance of vehicles, and methods that facilitate the construction and maintenance of dynamic assurance cases that can be readily modified as new threats are identified and the on-board vehicle software evolves. These requirements will also have an corresponding impact for vehicle safety assurance, in order to respond to a future in which through-life in-service software modifications become the norm, to implement new or improved features, correct faults and patch security. These software updates are expected to be delivered by over-the-air methods, which will themselves require safety and cybersecurity assurance.

Acknowledgement

The research leading to these results has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 812788 (MSCA-ETN SAS). This publication reflects only the authors' view, exempting the European Union from any liability. Project website: http://etn-sas.eu/.

References

Armstrong, R., R. Hawkins, and T. Kelly (2011). Security assurance cases: Motivation and the state of the art. University of York Report CESG/TR/2011/1.

Cobos, L.-P., A. R. Ruddle, and G. Sabaliauskaite (2021). Requirements for a cybersecurity case approach for the assurance of future connected and automated vehicles. In *Proceedings of 7th International Conference on Vehicle Technology and Intelligent Transport Systems (VEHITS 2021)*.

El-Rewini, Z., K. Sadatsharan, D. F. Selvaraj, S. J. Plathottam, and P. Ranganathan (2020). Cyber threats facing autonomous and connected vehicles: Future challenges. *Vehicular Communications 23*, 100214.

Goodenough, J. B., C. B. Weinstock, and A. Z. Klein (2015). Eliminative argumentation: A basis for arguing system confidence in system properties. Technical Report CMU/SEI-2015-TR-005, Software solutions Division, Software Engineering Institute, Carnegie Mellon University.

Hao, J. and G. Han (2020). On the modeling of automotive security: A survey of methods and perspectives. *Future Internet 12(11)*, 198.

Hawkins, R., T. Kelly, J. Knight, and P. Graydon (2011). A new approach to creating clear safety arguments. In *Advances in Systems Safety*, pp. 2253–2262. Springer.

IEC (2010). *IEC 61508 Functional safety of electrical/electronic/programmable electronic safety-related systems*.

ISO (2018). *ISO 26262 Road vehicles – Functional Safety*.

ISO (2019). *ISO/PAS 21448 Road vehicles – Safety of the Intended Functionality*.

ISO/AWI (2021). *ISO/AWI PAS 5112 Road vehicles – Guidelines for Auditing Cybersecurity Engineering*. In draft, expected Q3 2021.

ISO/SAE (2021). *ISO/SAE 21434 Road vehicles – Cybersecurity*. In draft, expected Q3 2021.

Kavallieratos, G., S. Katsikas, and V. Gkioulos (2020). Cybersecurity and safety co-engineering of cyber-physical systems—a comprehensive survey. In *Future Internet 12, No. 4*, pp. 65.

Kelly, T. and R. Weaver (2004). The goal structuring notation – a safety argument notation. In *Proceedings of the Dependable Systems and Networks 2004 Workshop on Assurance Cases*.

Kelly, T. P., J. McDermid, and R. Weaver (2005). Goal-based safety standards: opportunities and challenges. In *Proceedings of 23rd International System Safety Conference, San Diego, California*.

Kordy, B., S. Mauw, S. Radomirović, and P. Schweitzer (2011). Foundations of attack–defense trees. In P. Degano, S. Etalle, and J. Guttman (Eds.), *Formal Aspects of Security and Trust. FAST 2010. Lecture Notes in Computer Science, Vol 6561*, pp. 80–95. Springer, Berlin, Heidelberg.

Leveson, N. (2011). The use of safety cases in certification and regulation. Engineering Systems Division Working Paper ESD-WP-2011-13, Massachusetts Institute of Technology.

MISRA (2019). *Guidelines for Automotive Safety Arguments*. MISRA, ISBN 978-1-906400-23-1.

Mohamad, M., J.-P. Steghöfer, and R. Scandariato (2021). Security assurance cases—state of the art of an emerging approach. *Empirical Software Engineering 26*, 70.

Parkinson, S., P. Ward, K. Wilson, and J. Miller (2017). Cyber threats facing autonomous and connected vehicles: Future challenges. *IEEE Transactions on Intelligent Transportation Systems 18 (11)*, 2898–2915.

RTCA/EUROCAE (2011). *DO-178C/ED-12C – Software Considerations in Airborne Systems and Equipment Certification*. RTCA Inc. and European Organisation for Civil Aviation Equipment (EUROCAE).

SAE (2016). *J3061 – Cybersecurity Guidebook for Cyber-Physical Vehicle Systems*.

Schinagl, S., K. Schoon, and R. Paans (2015). A framework for designing a security operations centre (SOC). In *48th Hawaii International Conference on System Sciences*, pp. 2253–2262. IEEE.

SCSC (2021). *Goal Structuring Notation Community Standard (Version 3)*. The Assurance Case Working Group, Safety Critical System Club, SCSC 141-C.

UNECE (2021a). *Regulation No. 155, Cyber Security and Cyber Security Management*.

UNECE (2021b). *Regulation No. 156, Software Update Processes and Management Systems*.