CHALMERS
UNIVERSITY OF TECHNOLOGY

A risk assessment framework for automotive embedded systems

A Risk Assessment Framework for Automotive Embedded Systems

Mafijul Md. Islam
Christian Sandberg
mafijul.islam@volvo.com
christian.sandberg@volvo.com
Advanced Technology and Research, Volvo AB
Gothenburg, Sweden

Aljoscha Lautenbach
Tomas Olovsson
aljoscha@chalmers.se
tomas.olovsson@chalmers.se
Chalmers University of Technology
Gothenburg, Sweden

ABSTRACT

The automotive industry is experiencing a paradigm shift towards autonomous and connected vehicles. Coupled with the increasing usage and complexity of electrical and/or electronic systems, this introduces new safety and security risks. Encouragingly, the automotive industry has relatively well-known and standardised safety risk management practices, but security risk management is still in its infancy.

In order to facilitate the derivation of security requirements and security measures for automotive embedded systems, we propose a specifically tailored risk assessment framework, and we demonstrate its viability with an industry use-case. Some of the key features are alignment with existing processes for functional safety, and usability for non-security specialists.

The framework begins with a threat analysis to identify the assets, and threats to those assets. The following risk assessment process consists of an estimation of the threat level and of the impact level. This step utilises several existing standards and methodologies, with changes where necessary. Finally, a security level is estimated which is used to formulate high-level security requirements.

The strong alignment with existing standards and processes should make this framework well-suited for the needs in the automotive industry.

ACM Reference Format:
Mafijul Md. Islam, Christian Sandberg, Aljoscha Lautenbach, and Tomas Olovsson. 2016. A Risk Assessment Framework for Automotive Embedded Systems. In *CPSS '16: Proceedings of the 2nd ACM International Workshop on Cyber-Physical System Security.* ACM, New York, NY, USA, 12 pages. https://doi.org/http://dx.doi.org/10.1145/2899015.2899018

1 INTRODUCTION

Conventionally vehicles are perceived as isolated, static and closed systems, but recently a paradigm shift towards autonomous and connected vehicles has begun, and vehicles are increasingly personalised and are becoming a part of the Internet of Things. Market research predicts that 85% of all cars will be connected to the Internet by 2020 [36]. As a result, the usage of electronics and the complexity of electrical and/or electronic (E/E) systems will continue to increase in the coming years, and future vehicles will introduce new safety and security risks.

Safety is regarded as a non-negotiable requirement in the automotive industry. Methodologies and processes to achieve a reasonable level of safety during the development of safety-critical systems have been developed and standardised, e.g., ISO 26262 [13], a functional safety standard for road vehicles, which is based on the more generic safety standard IEC 61508. On the other hand, security has only recently gained attention in the automotive industry and security risks have hardly been systematically addressed, even though security threats against a vehicle can potentially jeopardise the safety of drivers, passengers and road users. Researchers have already demonstrated that it is possible to mount attacks that can endanger safety significantly [6, 17, 21, 22]. Consequently, security threats and associated risks need to be addressed methodically to improve the quality and safety of vehicles.

Automotive security is rapidly converging with traditional information technology (IT) security due to advances in hi-tech electronic architectures and communication systems [36]. Alongside new challenges, this phenomenon opens up new opportunities to tackle security concerns with well-established IT security methodologies, such as Common Criteria [4] or the ISO/IEC 27000 series of standards [14].

Since safety engineering is already a well-established process, and since safety and security engineering are related, it makes sense to align new security processes with the existing safety processes. With this in mind, we propose a framework to perform risk assessment to derive security requirements specifically for automotive systems, which is well aligned with the functional safety standard ISO 26262. Another design goal was to make the framework easy to use for non-security specialists, and to make the results easy to understand, since the requirements usually have to be implemented by independent contractors.

We introduce the notion of "Security Level" which is conceptually similar to the notion of "Automotive Safety Integrity Level" (ASIL) as standardised in ISO 26262. The "*Security Level*" is an automotive-specific risk-based metric which is used to specify the level of risk reduction that must be employed during the development of automotive systems to manage security risks. To estimate the potential impact of a threat regarding particular security objectives, we adapt several industrial standards. However, in contrast to existing standards and frameworks, we do not use the elapsed time to an attack during the threat level estimation as a separate parameter, because it can be derived from the other parameters in the

CPSS'16, May 30-June 03 2016, Xi'an, China

framework, such as the available equipment and the attacker's skill level. Similarly, attacker motivation is not considered a separate parameter either, because it is implicit in the other parameters, and attacker motivation is notoriously hard to model. We demonstrate the applicability of our framework by using a real-life scenario from the automotive domain. This work is based on the "needs and requirements" [12] and the "security model" [11] developed in the HEAVENS project.

2 RELATED WORK

Several standards and frameworks for threat analysis and risk assessment are available for various industrial domains, but few are readily applicable to the automotive domain.

The newest standard SAE J3061 - Cybersecurity Guidebook for Cyber-Physical Vehicle Systems [29] was just released and is also aligned with the processes in ISO 26262. It is the first automotive standard solely concerned with security and it covers the entire process of security engineering for automotive systems, including recommendations for risk assessment. The "HEAVENS Security Model" [11] this paper is based on is mentioned in SAE J3061.

The standard for the evaluation of security products and systems, Common Criteria (CC, also known as ISO/IEC 15408), states that the security needs of the evaluation target are usually identified by means of risk assessment [4], but leaves the risk assessment methodology open, since its focus is the enforcement of security requirements.

The ISO/IEC 27005 standard [14] defines an iterative information security risk management process, and the process is very similar in nature to the process we propose. However, ISO/IEC 27005 does not integrate well with the automotive functional safety standard ISO 26262 which is already widely adopted. Moreover, ISO/IEC 27005 is also rather generic and leaves the details of the process to the implementer.

The European Telecommunications Standards Institute (ETSI) [7] provides a threat, vulnerability and risk analysis (TVRA) methodology to deal with security issues in the telecommunications industry. The approach taken in ETSI's TVRA is similar to our approach, but it requires a good understanding of Common Criteria, and is not aligned with any safety standards.

In addition to theses standards, some research has been done on mapping the relationship of safety and security in the dependability and critical infrastructure domains, e.g., by Pièrre-Cambacédès and Chaudet [26], Schoitsch et al. [34], Jonsson [16], Avizienis et al. [2], Line et al. [18] and Firesmith [9]. Although safety and security engineering have influenced each other [23, 25, 27], no single unifying concept has emerged, yet.

Burton et al. [3] proposed an integrated approach that extends the functional safety process by considering "security hazards", arising from intentional manipulation of the system, as a third major type of hazard during hazard analysis. In contrast, we propose an independent risk assessment for security purposes which run in parallel with the safety processes, because it requires a different set of expertise. Similar to Burton et al., Schmittner et al. extended Failure Mode and Effect Analysis (FMEA), which is used extensively in safety engineering, to include a security perspective [31]. Schmittner et al. also provided a comparison of Common Criteria's Evaluation Assurance Levels (EALs) and Automotive Safety Integrity Levels (ASILs) [32], and they also proposed an integrated safety and security life-cycle [33]. Another interesting approach to integrate security concepts into ISO 26262 was developed by Macher et al. [19, 20]. They also use STRIDE for threat analysis, but they reduce (to use our terminology) the threat level analysis to two parameters, expertise and equipment, and they consider impact a one-dimensional parameter, which may be too simplistic.

The pioneering risk rating methodology for automotive E/E systems stems from the EVITA project [28]. In the EVITA approach [10, 28], the estimation of threat level and attack potential is inspired by Common Criteria [5]. The EVITA approach considers four parameters for risk rating – safety, financial, operational and privacy. However, a detailed study of the different impacts in these areas is missing. Our main improvement with regard to their approach is that we define detailed impact level parameters based on security objectives and align all the parameters with existing standards and guidelines. Furthermore, they do not take legislation aspects into account, even though several laws regarding the environment and driver behaviour are already in effect, and there are threats that can potentially lead to the violation of those legislative requirements.

Finally, Wolf and Scheibel [37] refined the ideas by Henniger et al. [10], and also combine existing techniques into a risk rating framework for automotive systems. Our framework has many similarities with Wolf and Scheibel's work, but there are also significant differences. The approach and terminology in [37] is closely aligned with Common Criteria, whereas our approach aims to be compatible with ISO 26262 to ease industry adoption. In addition, we stress the modularity and adaptability of our framework. Another major difference is how threats are identified: Wolf and Scheibel use per-asset security objectives to define attack trees based on security questionnaires, whereas we propose to use STRIDE due to its easier use for non-security experts. Furthermore, in the attack potential calculation (Threat Level) we deviate more from Common Criteria [5] to adapt it for the automotive industry, and we settled on more intuitive names. They also do not consider the privacy and legislative impact.

3 WORKFLOW OF THE FRAMEWORK

In this section, we outline the main steps of our proposed framework for security risk management. One of the goals is that it should fit well into existing safety engineering processes, since they are fundamental to automotive engineering.

The framework uses the following four security objectives:

(1) Safety – to ensure the functional safety of the vehicle occupants and other road users
(2) Financial – to prevent fraudulent commercial transactions, theft of vehicles, damage to stakeholder reputation, and insurance and warranty fraud
(3) Operational – to maintain the intended operational performance of all vehicle and ITS functions
(4) Privacy and Legislation – to protect the privacy of vehicle drivers, and the intellectual property of vehicle manufacturers and their suppliers

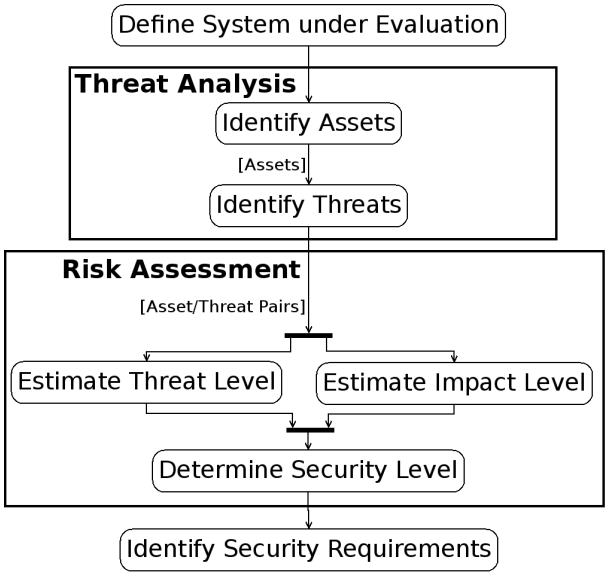

Figure 1: Workflow of the framework

The workflow of our framework is depicted in Figure 1, and starts with the ***definition of the system under evaluation***. Depending on the situation and what needs to be evaluated, a system can correspond to different levels of granularity: the entire E/E architecture, a vehicle function realised using several Electronic Control Units (ECUs), a single ECU, ECU hardware or ECU software. For simplicity, we use "system" to refer to "system under evaluation".

Once the system has been defined, the ***threat analysis*** starts, which is a two-step process. During the ***identification of assets***, the guiding question is which parts of the system have value and require security protection. Consider for instance privacy sensitive information such as vehicle location data or proprietary software on a specific ECU. In the next step, the ***identification of threats***, each asset is checked against a list of potential threats. Once all threats have been identified for all assets, the resulting asset/threat pairs are used as input to the risk assessment process. A detailed description follows in section 5.

During ***risk assessment***, the asset/threat pairs are analysed with respect to their likelihood of occurrence and severity of impact. These two activities consist of estimating the threat level, and estimating the impact level. The **threat level (TL)** is an estimate of the likelihood that a threat towards an asset is realised by an attacker, and the **impact level (IL)** is an estimate of the magnitude of harm to stakeholders resulting from threat realisation.

When the threat level and the impact level have been estimated for a particular asset/threat pair, the combination of the threat level and impact level can be used to ***determine a security level***. The **security level (SL)** is a measure of the level of protection particular asset needs and governs what security countermeasures should be taken to avoid unreasonable risk. It is similar to the way safety integrity levels (SILs) are used for functional safety. The entire risk assessment process, including the estimation of the threat and impact levels and the determination of the security level, is described in more detail in section 6.

Finally, high-level ***security requirements*** need to be identified that will guide the implementation of appropriate countermeasures in a later stage. However, the identification and implementation of appropriate countermeasures is out of the scope of the current framework. The identification of security requirements is explained in more detail in section 7.

4 THE SPEED LIMITER - A RUNNING EXAMPLE

In the remainder of this paper, we will use a running example to show-case the applicability of our framework to real-life scenarios. This section introduces the speed limiter use-case and the corresponding system model.

Commercial vehicles such as trucks are subject to legislative requirements in certain countries to limit the maximal vehicle speed, i.e., a road speed limit (RSL). Vehicle manufacturers aim to meet the legislative requirements by offering a vehicle functionality called the "speed limiter" to enforce the speed limit.

A system model for such a speed limiter is depicted in Figure 2, and it works as follows. A speed sensor (`SENSOR`) provides raw speed measurements to a Tachograph ECU (`TACHO`). In order to convert the data into speed, a `conversion factor` is needed. The Tachograph passes the raw data (Digital Input/Output pulses), together with the required conversion factor, to an ECU responsible for enforcing the set road speed limit (`RSLECU`). After calculating the current speed using the conversion factor, the `RSLECU` compares its current speed with the set speed limit parameters (e.g., factory settings) and chooses the lower speed. The `RSLECU` then sends an `RSL Request` with the chosen speed, along with the actual vehicle speed, to the `Engine ECU`. The `Engine ECU` compares the current vehicle speed with the requested speed, and calculates with the help of its own `RSL parameters` by how much the fuel supply needs to be cut, if at all, in order to achieve the requested speed.

5 THREAT ANALYSIS

Threat analysis is comprised of two steps: the first step is to identify all assets and the second step is to identify the corresponding threats to those assets. Several threat analysis frameworks for threat identification exist, and as long as the threat analysis results in a list of asset/threat pairs, any of them will work with our framework.

In recent years, Microsoft's STRIDE threat mnemonic [35] has been demonstrated to work well for automotive systems [19, 20, 30], so we will use it as an example. STRIDE is an acronym for six threat categories used: Table 1 briefly introduces the threats, and shows a static mapping of the threats to security attributes [35].

The goal of STRIDE is to identify all assets in the system which can be attacked. Tools to model the data flow in the system by creating data flow diagrams (DFDs) exist, which can be used for automatic threat identification with STRIDE [30]. From the data flow diagrams, a threat report is generated which lists asset/threat pairs. Threats exploit vulnerabilities, but there is no direct one-to-one mapping from vulnerabilities to threats. A particular vulnerability

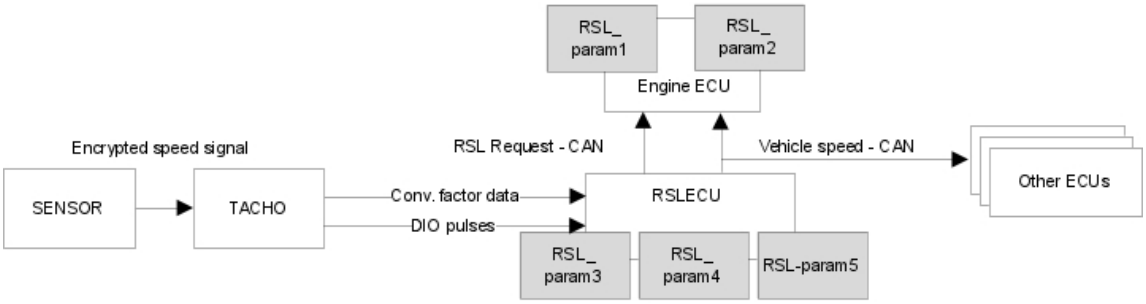

Figure 2: Model of a speed limiter

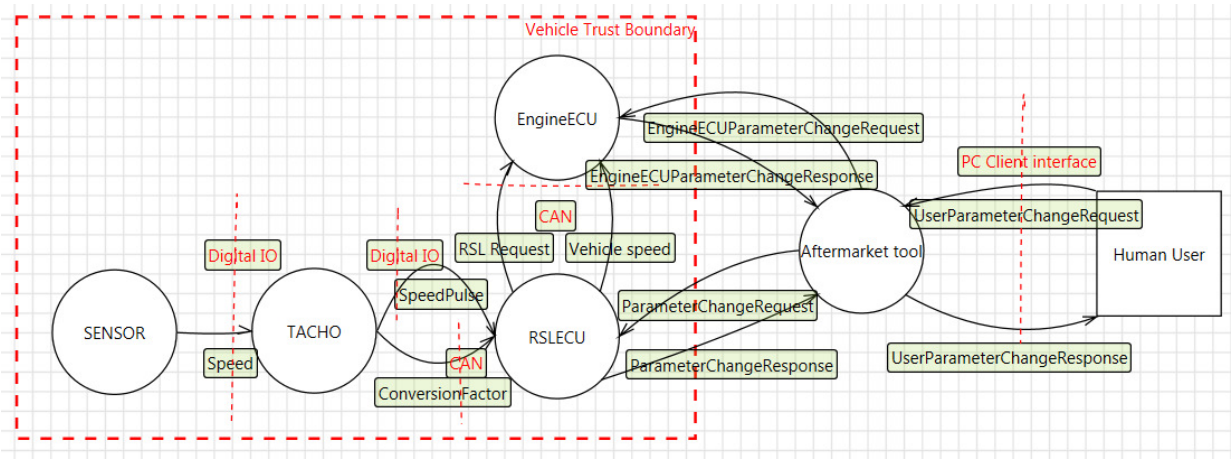

Figure 3: Data flow diagram of the speed limiter

Table 1: Microsoft's STRIDE methodology [35]

Threat	Violated Attribute	Explanation
Spoofing	Authenticity	Attackers pretend to be someone or something else
Tampering	Integrity	Attackers change data in transit or in a data store
Repudiation	Non-repudiation	Attackers perform actions that cannot be traced back to them
Information disclosure	Confidentiality/Privacy	Attackers get access to data (e.g. in transit or in a data store)
Denial of Service	Availability	Attackers interrupt a system's legitimate operation
Elevation of privilege	Authorisation	Attackers perform actions they are not authorised to perform

Table 2: Partial results from the speed limiter threat analysis

ID	Asset	Threat	Security Attribute
1	ConversionFactor	Tampering	Integrity
2	ParameterChange-Request	Spoofing	Authenticity
...	...	...	...

may lead to several threats and a particular threat may exploit several vulnerabilities [15]. The purpose of this step is to identify all such relations.

Applying this methodology to the speed limiter use-case, we construct a data flow diagram (DFD), depicted in Figure 3, using the Microsoft Threat Modelling Tool 2014. The DFD is slightly more detailed and concrete than the system model introduced earlier, since it also includes interactions for an `Aftermarket tool` which can be used by a `Human User` to change the RSL parameters.

The tool then generates a threat report, consisting of assets and associated threats. Finally, we extract the asset/threat pairs from the report to start the **risk assessment** process.

An example of such an asset/threat pairing is shown in Table 2. So the conversion factor being sent from the `TACHO` to the `RSLECU` is a data asset that can be changed in transit, i.e., tampered with, attacking its integrity. Similarly, the `ParameterChangeRequest`,

which can be made by an `Aftermarket tool` to the `RSLECU` to change the set speed limit, can be spoofed by an unauthorised user. Obviously there are many more such asset/threat pairs, but these two should suffice to illustrate the process.

6 RISK ASSESSMENT

Once all threats have been identified for every asset, the risk assessment process helps to prioritise the threats. Risk assessment consists of three steps: The determination of the threat level, the determination of the impact level, and finally by combining them, the determination of the security level. The security level governs the required level of protection. The threat level, impact level and security level will be discussed in detail in the following three subsections, respectively.

6.1 Threat Level

The threat level (TL) provides an estimate of the likelihood that this particular threat will occur. In order to estimate the threat level, we use four parameters similar to what is used the calculation of the attack potential in the vulnerability assessment of the Common Criteria [5].

The four parameters are:

(1) Expertise
(2) Knowledge about Target
(3) Window of Opportunity
(4) Equipment

Each of the parameters has four levels with an associated value, as shown in Table 3. The lower the value of the parameter, the more likely the occurrence of the threat. Unlike similar frameworks, we apply a linear scale for each parameter, which facilitates consistent reasoning about the different parameters while deriving the threat level for a particular asset/threat pair. However, the scales can easily be adjusted according to particular needs.

Before explaining the concrete meaning of the different levels, we will outline the process and outcome of estimating the threat level. After all parameter values have been estimated for a specific threat, a corresponding threat level can be computed, using the following simple linear equation:

$$T_{sum} = w_x t_x + w_k t_k + w_w t_w + w_e t_e \quad (1)$$

where w_i and t_i are the weight and the estimated threat level value of parameter i, and the indices x, k, w, e stand for the four parameters, respectively. We assume that the parameters are of equal importance, i.e., $w_i = 1\ \forall i$, so the equation is simplified to:

$$T_{sum} = t_x + t_k + t_w + t_e \quad (2)$$

However, the weights can easily be adjusted to the specific needs of an organisation.

Our framework provides five levels for a qualitative indication of the threat level: None, Low, Medium, High and Critical. After the value of T_{sum} has been calculated, the threat level is assigned according to a predefined scheme, as shown in Table 4.

The threat level parameters of the framework and their values are as follows. Since Wolf and Scheibel [37] base their attack potential calculation also on Common Criteria, we re-use several of their examples.

Expertise is the general level of knowledge required to carry out an attack:

- Layman. No particular expertise is required. Examples may include people who can follow simple instructions for existing attack tools, but who can not succeed if the instructions or the tools do not work as expected.
- Proficient. General security and domain knowledge is required. Professionals with knowledge about simple and popular attacks, are capable of mounting them with available tools, and if necessary, are able to improvise. For example workshop professionals who can install counterfeited parts without following step-by-step instructions developed by somebody else [37].
- Expert. Expert security and domain knowledge is required. Experts are familiar with underlying algorithms, protocols, hardware, software and concepts. They know techniques and tools of existing attacks and are able to create new attacks.
- Multiple Experts. Expert security and domain knowledge is required for several distinct domains. Allows for a situation in which different fields of expertise are required at an expert level to succeed with an attack.

Knowledge about target is the distribution of information about the target, i.e., the availability of information and the community size possessing that knowledge. This parameter points to the sources from where attackers can gain knowledge about the target and indicates how difficult it is for an attacker to acquire that knowledge:

- Public. The necessary information is public. Examples include information available on the Internet, in a bookstore or which is shared without non-disclosure agreements (NDAs), e.g. protocols like CAN or TCP/IP.
- Restricted. The information is shared with partners under non-disclosure agreements. For example requirements and design specifications or internal documentation which must be shared with suppliers or vehicle manufacturers.
- Sensitive. The information is shared between specific teams, but access is constrained to their members. Examples include restricted ECU configuration parameters, vehicle configuration databases or source code.
- Critical. The information is restricted to a few individuals. Access is tightly controlled on a strict need to know basis. Examples include root signature keys [37].

The first two levels, "Public" and "Restricted", specify knowledge distribution outside a single organisation, whereas "Sensitive" and "Critical" specify knowledge distribution within a single organisation. The attack potential decreases from "Public" to "Critical" due to the increasing difficulty for an attacker to obtain necessary information about the target.

Window of opportunity is the access type available to the attacker, and the time window the attacker has to mount a successful attack. The access type can be remote or physical:

Table 3: Threat level parameter values

Expertise	Value	Knowledge about target	Value	Window of opportunity	Value	Equipment	Value
Layman	0	Public	0	Unlimited	0	Standard	0
Proficient	1	Restricted	1	Large	1	Specialised	1
Expert	2	Sensitive	2	Medium	2	Bespoke	2
Multiple Experts	3	Critical	3	Small	3	Multiple bespoke	3

Table 4: Threat level calculation

Parameter Sum (T_{sum})	Threat Level (TL)	TL Value
> 9	None	0
$7 - 9$	Low	1
$4 - 6$	Medium	2
$2 - 3$	High	3
$0 - 1$	Critical	4

- Unlimited. Unlimited physical access, or network access for an unlimited time. Examples include always-on Internet access, or unlimited physical access to a vehicle for its owner.
- Large. High physical and/or remote availability with some time limitations.
- Medium. Low availability with severe time limitations. Limited physical and/or remote access to the target. Physical access to the vehicle interior or exterior without using any special tools (e.g., opening the hood to access wires) [37].
- Small. Very low availability. Physical access required to perform complex disassembly of vehicle parts to access internals to mount an attack on the asset.

Equipment is the equipment required to identify or exploit vulnerabilities. This can be hardware or software:

- Standard. The equipment is readily available to the attacker. The equipment may be part of the target itself (e.g. a debugger in an operating system), or is easily obtained. Examples include simple OBD diagnostic devices and common IT device such as notebooks [37].
- Specialised. The equipment is not readily available to the attacker, but could be acquired without undue effort. This could include the purchase of moderate amounts of equipment, or the development of more extensive attack scripts. Examples include in-vehicle communication devices such as network interface controllers, costly workshop diagnostic devices, power analysis tools or even resources on a computer cluster [37].
- Bespoke. The equipment is not readily available to the public as it may need to be specially produced, or because the equipment is so specialised that its distribution is controlled or restricted. Alternatively, the equipment may be very expensive. Multiple types of specialised equipment required for a successful attack also fall under this category.
- Multiple Bespoke. Multiple types of bespoke equipment are required for a successful attack.

In contrast to the attack potential calculation in Common Criteria, we do not consider the elapsed time required to mount a particular attack as a separate parameter, because it can be derived from other parameters. For example, depending on the attacker's skill level and the availability of the required equipment to mount the attack, the elapsed time may vary significantly – from less than an hour to several months. Similarly, we do not consider the motivation of the attacker as a separate parameter, since it is implicitly defined in the other parameters. For example, a highly motivated attacker may spend a lot of time to gain the necessary expertise to exploit a vulnerability, or spend a lot of money on the equipment needed.

It is noteworthy that the threat level parameters are highly dynamic in nature because they focus on the attacker, i.e., they can vary over time. For example, the development of better attack tools, disclosure of previously undiscovered vulnerabilities, etc. can change the estimate.

6.2 Impact Level

The impact level of a specific asset/threat pair is an estimate of the expected loss for different stakeholders when the threat is realised. In order to estimate the impact, we use four parameters which are directly related to the security objectives previously defined in section 3. The parameters are:

(1) Safety
(2) Financial
(3) Operational
(4) Privacy and Legislative

Each of the four parameters can be assigned one of four levels: None, Low, Medium, or High. All levels have a corresponding numerical value, as shown in Table 5.

Once again, we will first outline the process and outcome of the impact level estimation, before we discuss each of the four impact parameters in detail. After the impact parameters have been estimated, their overall sum can be computed to derive a corresponding impact level with the following simple linear equation:

$$I_{sum} = w_s i_s + w_f i_f + w_o i_o + w_p i_p \tag{3}$$

where w_j and i_j are the weight and the estimated impact value of parameter j, and the indices s, f, o, p stand for the parameters Safety, Financial, Operational and Privacy and Legislative, respectively.

Like with the threat level, the weights can be adjusted to accommodate specific needs. The default values have An unequal weight distribution of the impact parameters. Consider that safety and financial impacts can lead to the most severe consequences for

Table 5: Impact level parameter values

Safety	Operational	Financial	Privacy and Legislative	Value
None	None	None	None	0
Low	Low	Low	Low	1
Medium	Medium	Medium	Medium	10
High	High	High	High	100

Table 6: Impact level calculation

Parameter Sum (I_{sum})	Impact Level (IL)	IL Value
0	None	0
$1-19$	Low	1
$20-99$	Medium	2
$100-999$	High	3
≥ 1000	Critical	4

stakeholders, e.g., vehicle occupants may not survive or organisations may go bankrupt, whereas operational as well as privacy and legislative impacts are comparatively low in damages. We therefore propose weights of $w_s = w_f = 10$ and $w_o = w_p = 1$. So the equation can be simplified to:

$$I_{sum} = 10\ (i_s + i_f) + i_o + i_p \tag{4}$$

Finally, the resulting sum is used to derive the overall impact level, as shown in Table 6.

In the following, each impact level parameter is explained in detail.

(1) **Safety impact:** The safety impact refers to the safety of the vehicle occupants, road users and infrastructure. Safety is a first-order requirement in any automotive system. The four safety levels correspond to the ones in ISO 26262 [13]. A safety impact level of *None* means that there are no injuries, *Low* means light and moderate injuries, *Medium* means severe and life-threatening injuries with probable survival, and *High* means life-threatening injuries with uncertain survival and fatal injuries. More detailed explanations for each level can be found in ISO 26262 [13].

(2) **Financial impact:** The financial impact includes all direct and indirect financial damages of all stakeholders. *Direct financial damages* may include product liability issues such as penalties or recalls, legislation issues such as penalties due to non-conformance, or loss in revenue due to illicit activation of sellable features. *Indirect financial damages* on the other hand may include damages to reputation, loss of market share, IP infringement, and so on. Note that financial damage to customers, suppliers and other stakeholders must also be considered [8, 37].
Direct financial losses are comparatively easy to calculate; it is harder to estimate numerical values for indirect financial damages. For example, recent recalls of certain models of cars by several vehicle manufacturers due to various safety issues have had direct financial impact. At the same time, recalls also have an indirect financial impact because they hurt the vehicle manufacturers reputation. The total financial damage is the sum of direct and indirect costs.
The categorisation of financial damages depends on the financial strength of an individual stakeholder. For example, a loss of €100,000 may be relatively trivial to deal with for a large enterprise with billions of gross revenue, whereas even a loss of €10,000 may threaten the existence of a small enterprise. It may therefore be appropriate to express the limits as percentages of total sales, total profit, or on a similar base value as well as to classify the damages qualitatively into damage categories instead of calculating the damages quantitatively [8].
We propose to adopt the damage categories defined in the BSI standard 100-4 [8], with the names of the categories adapted to our framework. *No impact* means that there are no discernible effects or appreciable consequences for the stakeholders. *Low impact* means that the financial damage remains tolerable for the stakeholders. *Medium impact* means that there are substantial financial losses which do not threaten the existence of the stakeholders. Finally, *high impact* means that the financial damage threatens the existence of the stakeholders.

(3) **Operational impact:** Operational impact refers to operational damages which have little or no safety or financial impact, for instance the loss of secondary functionalities such as cruise control, or comfort and entertainment systems such as a CD-player or air-conditioning. In an operational context, a primary function is one which relates to driving, braking or steering, i.e., a function directly related to the vehicle's transport capabilities. Secondary functions are all other functions, e.g. comfort functions such as a media player or air-conditioning.
When operational damage causes safety issues or financial damages, this is also covered in those respective parameters. So the impact of a specific event can be cumulative by affecting several parameters.
We adapt the vehicular defect severity categorisation FMEA (Failure Mode and Effects Analysis) [1] to classify the operational damages. *No impact* means that there is no discernible effect. *Low impact* means that the appearance of an item or an audible noise annoys between 25% and 75% of customers. *Medium impact* corresponds to the degradation or loss of a secondary function, or the degradation of a primary function. Finally, *high impact* corresponds to the loss of a primary function which leaves the vehicle inoperable and potentially affects safety or legislative aspects.

Table 7: Calculation of security level from impact and threat level

Security Level (SL)	Impact Level (IL)					
Threat Level (TL)		0	1	2	3	4
	0	QM	QM	QM	QM	Low
	1	QM	Low	Low	Low	Medium
	2	QM	Low	Medium	Medium	High
	3	QM	Low	Medium	High	High
	4	Low	Medium	High	High	Critical

(4) **Privacy and Legislative impact:** Privacy and legislative impact deals with damages caused by privacy violations of stakeholders or violations of governmental regulations such as environmental or traffic laws.
The impact levels for estimating the privacy and legislative impact are similar to the ones defined in the German "Privacy Impact Assessment Guideline for RFID Applications" [24]. *No impact* means that there is no discernible effect. *Low impact* corresponds to privacy violations without direct potential for abuse, or legislative violations with no appreciable consequences, e.g., a warning without a fine. *Medium impact* corresponds to privacy violations which lead to abuse, or legislative violations with business and financial impact such as fines or reputation loss. Finally, *high impact* corresponds to privacy violations of multiple stakeholders which lead to abuse, or legislative violations with significant business and financial impact, such as significant loss of market share, trust or reputation.

Unlike the threat level parameters, the impact level parameters are primarily "stakeholder-oriented" and are therefore relatively "static", i.e., they are not expected to change significantly over time. It should also be noted that we do not consider the potential benefits for an attacker in the impact level, because the attacker's view is already considered heavily in the threat level.

6.3 Security Level

The security level guides the selection of the protection mechanisms and the required level of protection during system design and development. Table 7 outlines how the combination of threat level and impact level are used to derive the security level. In terms of goals and processes, this is similar to the assignment of automotive safety integrity levels (ASILs) [13]. However, it should be noted that, due to the dynamic nature of the threat level, the security level is also much more dynamic than an ASIL.

We propose the use of five increasing security levels: Quality Management (QM), Low, Medium, High and Critical. *Quality Management* is a term borrowed from ISO 26262 which means that no special controls for risk reduction are needed, the usual quality measures are sufficient. This also implies that for asset/threat pairs with a security level of QM, no security requirements will be formulated. For the remaining four security levels, *high-level security requirements must be formulated.* Note that the stringency of the security requirement is independent of the security level at this stage. This is similar to the way ASILs work in ISO 26262.

One asset may be associated with several threats, and as a result we may end up with multiple security levels for one asset. In that case, the security level for the asset is the highest security level of all the security levels associated with that asset.

6.4 Risk Assessment for the Speed Limiter

Performing the risk assessment for the speed limiter use-case, we choose relevant threat level and impact level parameters for each asset/threat pair, and establish a corresponding security level.

For the ConversionFactor/Tampering (CF/T) asset/threat pair, the *threat level parameters* are estimated as follows:

- The required expertise level is "Expert" ($t_x = 2$), because attackers need to know about the system architecture, the underlying algorithm, the communication protocol and the soft- and hardware.
- The knowledge about the target is "Restricted" ($t_k = 1$), since the required knowledge is primarily limited to the vehicle manufacturers and suppliers under non-disclosure agreements (NDAs).
- The window of opportunity is "Medium", since access to the asset is generally low, with time limitations ($t_w = 2$).
- Finally, the needed equipment is "Specialised" ($t_e = 1$), because special equipment such as a CAN network interface controller is needed, but it can be acquired without too much effort.

Using equation (2) to calculate the threat level, we get:

$$\begin{aligned} T_{sum}^{CF/T} &= t_x + t_k + t_w + t_e \\ &= 2 + 1 + 2 + 1 \\ &= 6 \end{aligned}$$

According to Table 4, this results in a "Medium" threat level.

Similarly, the *impact level parameters* for CF/T can be estimated as follows:

- The safety impact is "None" ($i_s = 0$), because the probability of injury due to tampering with the conversion factor is close to zero.
- The financial impact is "Low" ($i_f = 1$), because vehicle manufacturers and fleet owners may suffer some financial damage.
- The operational impact is "High" ($i_o = 100$), because the vehicle may not be in an operational state.
- Finally, the privacy and legislative impact is "Medium" ($i_p = 10$), because while there is no privacy impact, laws may be violated with potential financial penalties and loss of market share.

Table 8: Estimating threat level, impact level and security level for a subset of the asset/threat pairs

ID	Asset	Threat	Security Attribute	Threat Level (Value)	Impact Level (Value)	Security Level
1	ConversionFactor	Tampering	Integrity	Medium (2)	High (3)	Medium
2	ParameterChangeRequest	Spoofing	Authenticity	Low (1)	High (3)	Low

Using equation (4) to calculate the impact level yields:

$$\begin{aligned} I_{sum}^{CF/T} &= 10\ (i_s + i_f) + i_o + i_p \\ &= 10\ (0 + 1) + 100 + 10 \\ &= 120 \end{aligned}$$

According to Table 6, this results in an estimated impact level of "High".

The final step is to determine the *security level* by consulting table 7:

- With a "Medium" threat level, and a "High" impact level, the corresponding security level is "Medium".

The same principle is used to estimate the threat, impact and security levels of all remaining asset/threat pairs. For the two examples in the case-study, the results are summarised in Table 8.

7 SECURITY REQUIREMENTS

After threat analysis and risk assessment have been performed, what remains is to derive high-level security requirements for the identified asset/threat pairs. It should be noted that detailed and technical security requirements should be formulated at a later stage, but this is out of the scope of this paper.

The formulated high-level security requirements are independent of the security level at this point. The security level only indicates that a security requirement must be formulated: an interpretation or translation to countermeasures of the security level happens when the technical, hardware or software security requirements are formulated. Once again it should be stressed that no security requirements for asset/threat pairs with a security level of QM are formulated at all, while high-level requirements must be formulated for all other asset/threat pairs.

Conceptually, this step is closely related to processes in both ISO 26262 and in Common Criteria. In ISO 26262, the functional safety requirements are defined at the end of the concept phase when ASILs are determined based on the hazard analysis and risk assessment. The functional safety requirements have the same purpose as the security requirements, i.e., to formulate high-level requirements. In Common Criteria terminology, this step of identifying high-level security requirements corresponds to the formulation of the high-level outline of a solution for a specific security problem, and the translation of that solution into security functional requirements (SFRs). SFRs are implementation-independent security requirements formulated in a standardised language.

We already derived the security level for the speed limiter use-case in section 6, as shown in Table 8. For every asset/threat pair which does not have a security level of QM, a high-level security requirement is formulated which is independent of the concrete security level. Remember that the security level is only used at a later stage for the formulation of the technical security requirements. In our example, we derive the following high-level security requirements:

- Security requirement # 1: The integrity of the ConversionFactor signal shall be ensured.
- Security requirement # 2: The authenticity of the ParameterChangeRequest signal shall be ensured.

Threat analysis and risk assessment are generally performed during the concept phase of the development lifecycle and implementation details are often not available at this stage. Consequently, it may not always be possible to identify the concrete security mechanisms that need to be implemented to fulfil the derived high-level security requirement. However, the estimated security level for each asset/threat pair, along with its high-level security requirement, should be used during the product development stage of the lifecycle to choose an appropriate security mechanism to fulfil the requirements for a particular security level.

Our framework does not address security assurance, but Common Criteria can be used if needed. Common Criteria defines a well-established process within the security industry for quality assurance of IT security solutions. Security categories that are generic for a certain product type, called "Protection Profiles", enable comparison of security solutions from different vendors. We believe that protection profiles for common systems should be developed for and by the automotive industry.

8 PARALLELS TO ISO 26262

Many companies in the automotive industry have already developed process-, method- and tool-support to comply with the requirements of the automotive functional safety standard ISO 26262. Since functional safety is a core aspect of automotive engineering, and can be expected to be one for the foreseeable future, the alignment with safety processes is an important feature of our framework. In this section, we highlight the similarities and differences between our framework and ISO 26262.

The processes in ISO 26262 follow a V-development-model, progressing from *concept phase*, through three distinct *product development* phases (system level, hardware level, software level), to the *production and development* phase.

Table 9 provides a visual representation of the parallels between the different lifecycle phases in ISO 26262 and our framework. We will discuss each step in each lifecycle phase in turn.

8.1 Concept Phase

In this paper, in the context of ISO 26262, we primarily focus on the concept phase. The concept phase includes all activities which must happen before product development can start, such as system definition, hazard analysis and risk assessment and the definition

Table 9: Safety requirements in ISO 26262, and security requirements in our framework

Lifecycle phase	Functional Safety		Security	
Concept Phase	Item definition		Definition of system under evaluation	
	↓		↓	
	Hazard analysis and risk assessment		Threat analysis and risk assessment	
	↓		↓	
	Safety Goals & ASILs		Security Levels	
	↓		↓	
	Functional Safety Requirements		High-level Security Requirements	
Product Development Phase	↓		↓	
	Technical Safety Requirements (System Level)		Technical Security Requirements (System Level)	
	↓	↓	↓	↓
	Hardware Safety Requirements	Software Safety Requirements	Hardware Security Requirements	Software Security Requirements
Operational Phase	—		—	

of the functional safety requirements. It also includes the creation of all corresponding documentation.

8.1.1 Item definition vs System definition. According to the standard, the first step in safety requirements engineering is to define an item which provides a function at the vehicle level, for example, cruise control. Similarly, in our framework, the first step in security requirements engineering is to define a system. However, unlike an item in ISO 26262, our notion of system may represent any abstraction level: an entire E/E system, a vehicle function realised using multiple ECUs, a single ECU, software or hardware component, etc. As a result, our definition of system can correspond to any level of the hierarchy of the abstraction levels (item, system, hardware or software component, and hardware part or software unit) presented in ISO 26262. Furthermore, our definition of system corresponds to the concept of target of evaluation (TOE) consisting of a set of assets from the viewpoint of Common Criteria.

8.1.2 HARA vs TARA. After item definition, hazard analysis and risk assessment (HARA) is performed. During HARA, possible hazards are identified, and their corresponding risks are assessed, in order to avoid unreasonable risk. According to ISO 26262, a *hazard* is a "potential source of harm caused by malfunctioning behaviour of the item". When the hazards and all possible hazardous event have been identified, their impact must be estimated by considering the three parameters *severity*, *probability of exposure* and *controllability*.

In our proposed framework, once a system is defined, we perform threat analysis to identify the assets and potential threats that are associated with the assets. Next, we perform risk assessment to rank the threats. We apply a set of parameters to estimate threat levels and impact levels.

Risk assessment in security is more multi-dimensional than its counterpart in functional safety. We estimate threat levels by using the parameters expertise, knowledge about target, window of opportunity and equipment. We apply another set of parameters (safety, financial, operational, privacy and legislation) to estimate the impact level, which is conceptually similar to "severity" in ISO 26262. It is important to note that we use safety as one of the parameters to estimate severity during risk assessment.

Furthermore, threat level parameters (expertise, knowledge about target, window of opportunity, equipment) are relatively dynamic over time compared to the parameters (severity, probability of exposure, controllability) used for ASIL determination. For example, expertise of an attacker, equipment to mount an attack and knowledge about target can potentially change significantly over time – people may not be aware of a vulnerability for a long time.

8.1.3 ASILs vs SLs. In functional safety, based on the results of HARA, an ASIL is determined and safety goals are formulated for an item. In our case, we establish security levels based on threat levels and impact levels. It is notable that we do not formulate security goals as an end result of threat analysis and risk assessment. Instead,

we rely on the security objectives to estimate the impact level, and our notion of security objectives are similar to the safety goals of ISO 26262.

In our framework, a system can include several assets and each asset may be related to one or more threats. As a result, we may have several security levels for each asset and even for the system as a whole. When this happens, we choose the highest security level for the asset or system, similar to how this is handled for ASILs in ISO 26262.

8.1.4 Functional Safety Requirements vs High-level Security Requirements. Finally, in functional safety, functional safety requirements for an item are derived based on safety goals and ASILs. These requirements are in general high-level and can be implementation independent.

Similarly, in our framework high level security requirements are derived from the asset/threat pairs and their security level, which should be implementation independent.

8.2 Product Development Phase

In ISO 26262, the product development phase consists of three sub-phases, "Product development at the system level", "Product development at the hardware level" and "Product development at the software level". Each of them has its own V-model, including, among others, steps for requirements engineering, design, implementation and verification.

During the product development phase, taking the ASILs into account, the functional safety requirements are translated into technical safety requirements, which in turn lead to hardware functional safety requirements and software functional safety requirements.

Similarly, taking the security level into account, the high-level security requirements for each asset/threat pair are translated into concrete technical security requirements on system level, which in turn result in hardware and software security requirements.

8.3 Operational Phase

The operational phase includes steps for operational management and maintenance. Neither ISO 26262 nor our framework make any claims about new safety or security requirements in the operational phase.

9 CONCLUSIONS

The ongoing paradigm shift towards autonomous and connected vehicles, augmented with the growing usage and complexity of E/E systems, will undoubtedly introduce new safety and security risks. In the automotive industry, safety risk management is relatively well-known and standardised, whereas security risk management is largely missing.

In this paper, we present a framework to perform threat analysis and risk assessment to systematically address security risks. We propose a methodology to determine a "Security Level" and derive high-level security requirements.

Our framework has several advantages. It is easy to use, also for non-security specialists. Its strong alignment with ISO 26262 makes it easy to grasp for automotive engineers: The four security levels are similar to automotive safety integrity levels, and the resulting high-level security requirements correspond to functional safety requirements. That the requirements are clear is especially important, since their implementation is often done by sub-contractors.

Furthermore, our threat level estimation is a slightly simplified version of Common Criteria's vulnerability assessment [5], because elapsed time and attacker motivation are not considered separately.

Unlike any current approach, our estimation of impact levels is well-aligned with several industrial standards. Utilising inter-domain (e.g., IT security and Common Criteria) and intra-domain (e.g., functional safety and ISO 26262) standards and guidelines, our proposed framework facilitates the use of existing knowledge, competencies and processes to systematically assess security risks.

In future work, the formulation of technical security requirements on system, hardware and software level depending on the security level should be considered. Their translation to concrete security mechanisms, the implementation of those mechanisms, and finally, the security evaluation of the implementation, should also be considered in follow-up work.

REFERENCES

[1] Automotive Industry Action Group (AIAG). *Potential Failure Mode and Effects Analysis FMEA Reference Manual.* AIAG, 4th edition, 2008.

[2] A. Avizienis, J.-C. Laprie, B. Randell, and C. Landwehr. Basic concepts and taxonomy of dependable and secure computing. *IEEE Transactions on Dependable and Secure Computing*, 1(1):11–33, 2004.

[3] S. Burton, J. Likkei, P. Vembar, and M. Wolf. Automotive functional safety = safety + security. In *Proceedings of the First International Conference on Security of Internet of Things*, SecurIT '12, pages 150–159, New York, NY, USA, 2012. ACM.

[4] CCRA Members. Common Criteria for Information Technology Security Evaluation. CCMB-2012-09-00X, Version 3.1, Revision 4.

[5] CCRA Members. *Common Methodology for Information Technology Security Evaluation – Evaluation Methodology*, chapter Vulnerability Assessment (AVA), pages 404 – 433. September 2012. CCMB-2012-09-004, Version 3.1, Revision 4.

[6] S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, S. Savage, K. Koscher, A. Czeskis, F. Roesner, and T. Kohno. Comprehensive experimental analyses of automotive attack surfaces. In *Proceedings of the 20th USENIX Security Symposium*, pages 77–92, San Francisco, CA, USA, Aug. 2011.

[7] ETSI. Telecommunications and internet converged services and protocols for advanced networking (TISPAN); methods and protocols; part 1: Method and proforma for threat, risk, vulnerability analysis. Technical Specification TS 102 165-1, v4.2.3, ETSI, Mar. 2011.

[8] Federal Office for Information Security (BSI), Germany. *BSI-Standard 100-4 – Business Continuity Management.* 2009.

[9] D. G. Firesmith. Common concepts underlying safety security and survivability engineering. Technical report, Dec. 2003.

[10] O. Henniger, L. Apvrille, A. Fuchs, Y. Roudier, A. Ruddle, and B. Weyl. Security requirements for automotive on-board networks. In *Proceedings of the 9th International Conference on Intelligent Transport System Telecommunications (ITST 2009), Lille, France*, 2009.

[11] M. Islam, C. Sandberg, A. Bokesand, T. Olovsson, H. Broberg, P. Kleberger, A. Lautenbach, A. Hansson, A. Söderberg-Rivkin, and S. P. Kadhirvelan. Deliverable D2 - Security Models. HEAVENS Project, Version 1.0 (Release 1), September 2014.

[12] M. Islam, C. Sandberg, M. Olsson, A. Bokesand, T. Olovsson, H. Broberg, K. Calais, R. Svenningsson, A. Söderberg, M. Jidhage, M. Wallerström, J. Ekberg, T. Johansson, and A. Hansson. Deliverable D1.1 Needs and Requirements. HEAVENS Project, Version 1.0 (Release 1), February 2014.

[13] ISO. ISO 26262:2011: Road vehicles — Functional safety, 2011.

[14] ISO. ISO/IEC 27005: Information technology — Security techniques — Information security risk management, 2011. ISO/IEC 27005:2011.

[15] R. G. Johnston. Being vulnerable to the threat of confusing threats with vulnerabilities. *The Journal of Physical Security*, 4(2):30–34, 2010.

[16] E. Jonsson. Towards an integrated conceptual model of security and dependability. In *Availability, Reliability and Security, 2006. ARES 2006. The First International Conference on*, pages 646–653. IEEE, 2006.

[17] K. Koscher, A. Czeskis, F. Roesner, S. Patel, T. Kohno, S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, and S. Savage. Experimental security analysis of a modern automobile. In *2010 IEEE Symposium on Security and Privacy (SP)*, pages 447–462. IEEE, 2010.

[18] M. Line, O. Nordland, L. Røstad, and I. Tøndel. Safety vs. security. In *Probabilistic Safety Assessment and Management (PSAM), Proceedings of the 8th international*

Conference on, pages 685–699. IAPSAM, 2006.

[19] G. Macher, A. Höller, H. Sporer, E. Armengaud, and C. Kreiner. A combined safety-hazards and security-threat analysis method for automotive systems. In *Computer Safety, Reliability, and Security*, pages 237–250. Springer, 2015.

[20] G. Macher, H. Sporer, R. Berlach, E. Armengaud, and C. Kreiner. Sahara: a security-aware hazard and risk analysis method. In *Proceedings of the 2015 Design, Automation & Test in Europe Conference & Exhibition*, pages 621–624. EDA Consortium, 2015.

[21] C. Miller and C. Valasek. Adventures in automotive networks and control units. *Last accessed on 2015-12-18 from http://illmatics.com/car_hacking.pdf*, 2013.

[22] C. Miller and C. Valasek. Remote exploitation of an unaltered passenger vehicle. *Last accessed on 2015-12-18 from http://illmatics.com/Remote%20Car%20Hacking.pdf*, 2015.

[23] D. Nicol, W. Sanders, and K. Trivedi. Model-based evaluation: from dependability to security. *IEEE Transactions on Dependable and Secure Computing*, 1(1):48–65, Jan. 2004.

[24] M. C. Oetzel, S. Spiekermann, I. Grüning, H. Kelter, and S. Mull. *Privacy Impact Assessment Guideline for RFID Applications.* Federal Office for Information Security (BSI), Germany, 2011.

[25] L. Piètre-Cambacédès and M. Bouissou. Cross-fertilization between safety and security engineering. *Reliability Engineering & System Safety*, 110:110–126, 2013.

[26] L. Piètre-Cambacédès and C. Chaudet. The SEMA referential framework: avoiding ambiguities in the terms "security" and "safety". *International Journal of Critical Infrastructure Protection*, 3(2):55–66, 2010.

[27] C. Raspotnig and A. Opdahl. Comparing risk identification techniques for safety and security requirements. *Journal of Systems and Software*, 86(4):1124–1151, 2013.

[28] A. Ruddle, D. Ward, B. Weyl, S. Idrees, Y. Roudier, M. Friedewald, T. Leimbach, A. Fuchs, S. Gürgens, O. Henniger, R. Rieke, M. Ritscher, H. Broberg, L. Apvrille, R. Pacalet, and G. Pedroza. Security requirements for automotive on-board networks based on dark-side scenarios. EVITA Project, Deliverable D2.3, v1.1., Dec. 2009.

[29] SAE International. SAE J3061_201601 - Cybersecurity guidebook for cyber-physical vehicle systems, Jan. 2016.

[30] K. Schmidt, P. Tröger, H.-M. Kroll, T. Bünger, F. Krueger, and C. Neuhaus. Adapted development process for security in networked automotive systems. *SAE International Journal of Passenger Cars - Electronic and Electrical Systems*, 7(2):516–526, 2014.

[31] C. Schmittner, T. Gruber, P. Puschner, and E. Schoitsch. Security application of failure mode and effect analysis (FMEA). In *Computer Safety, Reliability, and Security*, pages 310–325. Springer, 2014.

[32] C. Schmittner and Z. Ma. Towards a Framework for Alignment Between Automotive Safety and Security Standards. In F. Koornneef and C. van Gulijk, editors, *Computer Safety, Reliability, and Security*, volume 9338 of *Lecture Notes in Computer Science*, pages 133–143. Springer International Publishing, Jan. 2015.

[33] C. Schmittner, Z. Ma, and E. Schoitsch. Combined safety and security development lifecycle. In *Industrial Informatics (INDIN), 2015 IEEE 13th International Conference on*, pages 1408–1415. IEEE, 2015.

[34] E. Schoitsch, C. Schmittner, Z. Ma, and T. Gruber. The need for safety and cybersecurity co-engineering and standardization for highly automated automotive vehicles. In T. Schulze, B. Müller, and G. Meyer, editors, *Advanced Microsystems for Automotive Applications 2015*, Lecture Notes in Mobility, pages 251–261. Springer International Publishing, 2016.

[35] A. Shostack. *Threat modeling – designing for security.* Wiley, 2014.

[36] N. Tare. *NE30-01: Cybersecurity in the automotive industry.* Frost & Sullivan, October 2014.

[37] M. Wolf and M. Scheibel. A systematic approach to a qualified security risk analysis for vehicular IT systems. In E. Plödereder, P. Dencker, H. Klenk, H. B. Keller, and S. Spitzer, editors, *Automotive - Safety & Security 2012*, Lecture Notes in Informatics, pages 195–210. Gesellschaft für Informatik, Bonn, 2012.