THREAT MODELLING FOR FUTURE VEHICLES

STIJN VAN WINSEN

S.VANWINSEN@STUDENT.UTWENTE.NL

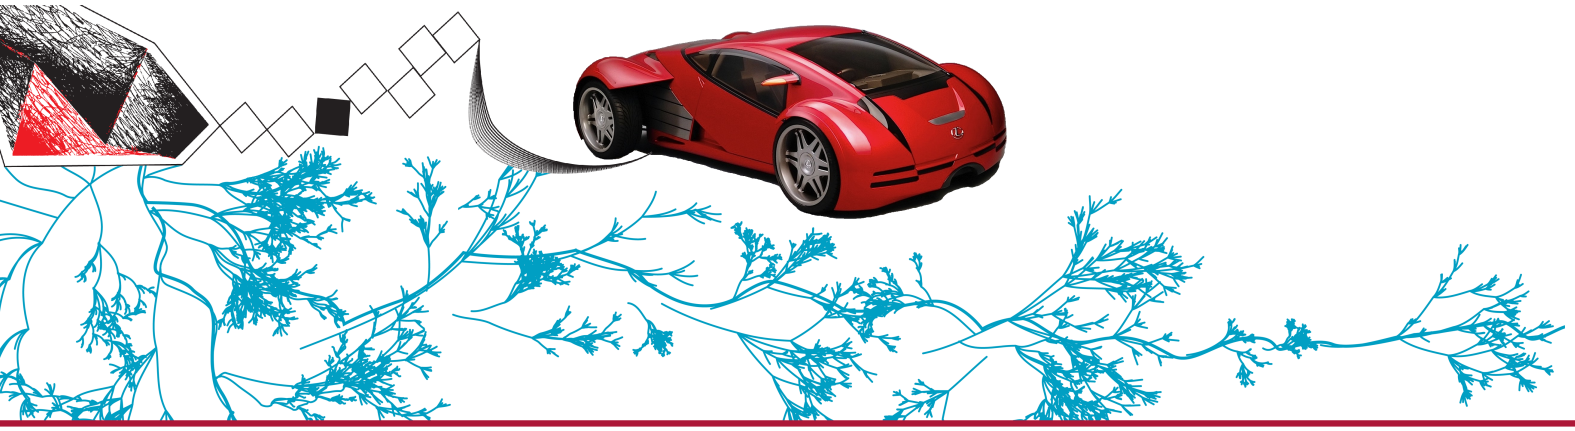

On Identifying and Analysing Threats for Future Autonomous and Connected Vehicles

Master of Science - Computer Science - Kerckhoffs Institute
Faculty of Electrical Engineering, Mathematics and Computer Science
University of Twente

January 2017 – Final Version

Stijn van Winsen: *Threat Modelling for Future Vehicles,* On Identifying and Analysing Threats for Future Autonomous and Connected Vehicles, © January 2017

SUPERVISORS:
Dr. N. (Klaas) Sikkel - University of Twente
Dr. Ir. G.J. (Geert) Heijenk - University of Twente
Ir. J.A.W. (Jeroen) de Wit - KPMG

LOCATION:
Enschede

SUMMARY

Modern day vehicles contain many IT components that have been designed for isolated vehicles. The shift towards vehicles that are connected to other devices creates an increased attack surface for attackers. Together with a shift towards more autonomous vehicles, which introduce more cyber-physical systems, it becomes evident that the IT in vehicles needs to be properly secured.

Vehicle manufacturers have only recently started to incorporate security in the design process, and lack techniques to do this. This thesis therefore proposes a composite threat model focused on identifying all threats under the assumption that the system is already breached.

The proposed composite threat model consists of two steps that a security expert within a vehicle manufacturer should follow for all relevant applications or systems.

First, a complete interconnections drawing should be created to get a complete overview of all relevant components in the system, including all entities and high level data flows.

Second, using the drawing from step 1, the STRIDE threat modelling technique is used to identify all possible threats. Then, for all threats on the list, the threats are analysed based on their consequences on two aspects: Severity and Controllability. Severity describes how severe a threat is if it occurs, this is analysed on four aspects: Safety, Operational, Privacy, and Financial. Controllability describes how controllable a threat is if it occurs.

Using these results, a security expert can reason about the different threats and prioritise them on how important they are, and use this to find mitigation techniques.

Since our model focuses on identifying all possible threats, the main recommendations from the validation include creating a tool that can help a security expert in identifying the possible threats, and in particular help in reducing irrelevant threats, as well as finding a way to make the results sellable to management.

With this composite threat model, security experts can identify and analyse their vehicles for possible threats and make their vehicles more secure, as is evidently needed for future vehicles.

ACKNOWLEDGEMENTS

This thesis marks the end of seven and a half years of studying Computer Science and Computer Security at the University of Twente. After almost a year of writing this thesis, which is a tad long, my life as a student is now over.

Although there have been some moments that things could have gone more smoothly, my supervisors were always there to get me back in track. I am therefore glad that Klaas en Geert agreed to supervise me during this research. You have always taken the time to discuss progress, read my thesis, provide me with valuable feedback and above all guide me in the process, even though you are not primarily engaged in the field of security. Thank you!

I am also grateful that Jeroen from KPMG has taken the time to discuss my progress from time to time and helping me out with finding experts in the field. Jeroen, thank you for sometimes giving me the room to find things out by myself, but also helping me whenever I got stuck or needed a discussion in a field that was also new to you.

I wrote my thesis as part of KPMG's Cyber team. I am grateful to them for providing me the opportunity to do so and for making me feel a part of the team so quickly. Both on a professional level, as on a social level, you have provided me a great time. In particular, I have enjoyed the discussions, drinks and time in the office with my partners in crime: the co-interns.

I would also like to thank my friends and family for supporting me, sometimes asking me how things were going, and sometimes keeping quite and being an outlet if things were going less smoothly.

Lastly, I would like to thank the guys from ShareLateX for providing me the tools to write this thesis and André Miede for providing this wonderful thesis template.

Stijn van Winsen

CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACRONYMS

ABS Antilock Braking System

ASIL Automotive Safety Integrity Levels

C2C-CC Car 2 Car Communication Consortium

CAN Controller Area Network

CHASSIS Combined Harm Assessment of Safety and Security for Information Systems

DFD Data Flow Diagram

EAL Evaluation Assurance Level

ECU Electronic Control Unit

ETSI European Telecommunications Standards Institute

HARA Hazard Analysis and Risk Assessment

HSM Hardware Security Module

ISO International Organisation for Standardisation

ITS Intelligent Transportation System

LIN Local Interconnect Network

MOST Media Oriented Systems Transport

NHTSA National Highway Traffic Safety Administration

NIST National Institute of Standards and Technology

OBD On-Board Diagnostics

OBU On Board Unit

OEM Original Equipment Manufacturer

SAHARA Security-Aware Hazard and Risk Analysis

SDL Security Development Lifecycle

SIL Safety Integrity Levels

SINA Security in Networked Automotive

TARA Threat Analysis and Risk Assessment

VANET Vehicular Ad-hoc Network

Part I

INTRODUCING THE RESEARCH

1

INTRODUCTION

The automotive market is undergoing rapid changes. Since the introduction of the first car in the late-1880s until now, a lot has changed inside the car, ranging from better engines to more comfort for the driver. One of the biggest changes however, is that of IT becoming increasingly integrated into the car to help it become faster and safer.

As early as the late-1970s, IT was added to the car to make it more fuel efficient in the form of Electronic Control Units. In the years to come, IT would take over a wide variety of functions, ranging from mirror-adjustments to Antilock Braking Systems.

The most recent development in the automotive industry is the introduction of self-driving technology. Major vehicle companies have announced to be working on autonomous vehicles, a technology to be ready in 2025. This trend is making IT take over even more functions of the driver. To achieve this, systems inside the vehicle are becoming more connected to each other, as well as to other vehicles and the internet.

That IT is becoming dominant inside the vehicle brings a lot of opportunities, but also a lot of threats. The fact that a vehicle is becoming connected to the internet means a hacker could possibly disable the brakes from anywhere in the world, with possibly severe consequences. That vehicles need to be properly secured seems evident.

This thesis looks at the security of vehicles. In particular, it focuses on how an organisation should act upon the threats possible future functionality might bring, by looking at how such threats should be identified and analysed.

2

BACKGROUND

This chapter provides some high-level background information on the topic of automotive cyber security to familiarise the reader with the subject and introduce the problem. This chapter will first provide a brief history of the application of IT in vehicles. Some aspects will be explained in more detail in upcoming chapters, however, some design decisions are better understood when knowing the history of IT in vehicles. A quick overview is therefore given to keep in mind when reading upcoming chapters.

2.1 HISTORY

From the creation of the first car in the late-1880s until now, many changes have been made to make the car faster, safer and more aesthetic. In the beginning, these changes were purely mechanical. However, as early as the late-1970s, intelligence was added to the car to make it more fuel efficient in the form of what was then still called an Engine Control Unit (ECU). By measuring the oxygen present in the exhaust fumes, the ECU could adjust the fuel/oxygen ratio before combustion, making it more efficient and reducing pollution. Since then, more intelligence has been added in a variety of other systems such as the door locking mechanism, light control, brake control, entertainment system and many more. With this change, the name of the ECU also changed to a more general Electronic Control Unit.

The digitalisation of the ECU meant that it basically became a small computer specialised in one task and operated on an individual basis. The Power Door Lock had no connection with the Anti Blocking System of the wheels. However, after a while, these systems became more and more intertwined. The modern day Electronic Stability Control system for example, combines steering angle, accelerometers, individual wheel speeds and throttle position for the stability of the vehicle [37].

Since more and more ECUs got connected, it was no longer cost efficient to connect every ECU with every other ECU. Therefore, Bosch developed the Controller Area Network bus (CAN bus), a vehicle bus standard designed for real-time communication to connect all ECUs in a vehicle via this single bus. Since then, other buses with other

properties have been designed and implemented, though, none of these have yet replaced the CAN bus as the default in-vehicle communication system.

Because the ECU became more and more dominant in vehicles, replacing other mechanical components, its functioning also became more important for the safety of the vehicle and its passengers. A faulty ECU could mean the life or death for a passenger. This started a movement by governments and vehicle manufacturers to integrate safety risk assessments into the development process of ECUs, and to formalise this process. The biggest standard being ISO 26262 [28] that defines the functional safety of electrical and/or electronic systems in production automobiles. Worth noting is that until this point, the in-vehicle communication system was still considered to be isolated from other systems. Therefore, security of these components has never been a part of the design process.

Although the vehicle became more and more digitalised in all these years, it was still an isolated systems of ECUs. However, starting in the mid-1990s, this changed. Cars began integrating GPS-modules, started providing On-Board Diagnostic (OBD) modules for diagnosing the internal network of the car via the so called OBD-II port, and in later years even let a passenger connect his devices via Bluetooth or even WiFi. This meant that the in-vehicle communication system in the car was no longer an isolated system, and with this, susceptible to outside attackers.

In the mean time, the ECUs themselves became more and more complex, making them more prone to errors. At first, ECUs could not be reprogrammed at all, meaning a defective unit had to be replaced on site with a newer version of the software, making it very expensive. In later years, ECUs became reprogrammable, but this often meant the manufacturer or a mechanic still had to physically connect to the ECU, either directly, or via the CAN bus. Replacing the unit was no longer necessary, but a recall because of a defective ECU was still very expensive.

In recent years, vehicles have become even more connected. Many new vehicles provide cellular communication systems meant for letting the vehicles communicate with its manufacturer. This new way of communicating provides the manufacturer with great capabilities. A vehicle can for example be monitored extensively to predict when a component is going to break down, or to provide over-the-air firmware updates of its ECUs, making recalls unnecessary. However, this also provides attackers access to the vehicle from an even greater distance.

In the near future, vehicles are very likely to communicate with each other and the surrounding infrastructure. To this end, vehicles will set up so called Vehicular Ad-hoc Networks (or VANETs) and exchange information such as upcoming collisions, road conditions

and current speed. Particularly challenging in this area is determining which information can and which information cannot be trusted.

2.2 AUTOMOTIVE SECURITY

The increased role of IT in the functioning of the vehicle brings a lot of possibilities; in the future, vehicles will probably drive autonomously without any intervention from a driver. However, this increased role of IT also has its downsides. Before the introduction of ECUs in the vehicle, malfunctions were often due to wear of certain parts, such as the brakes. Historical data of those parts could often be used to predict or calculate the chance of malfunction, and more importantly, be mitigated. However, with more and more ECUs in the vehicle, a malfunctioning part can no longer only be attributed to wear, but to programming errors as well.

To cope with these new kinds of risks, vehicle manufacturers began integrating functional safety processes in the design process to explicitly incorporate safety goals in the design of software and minimise the chance of malfunctions due to software errors. Many of these processes, however, have existed for quite some time and often do not incorporate any security goals. At first, these security goals were hardly needed; the software components in a vehicle were considered to be isolated; an attacker would need physical access to the vehicle, where cutting a brake line would have the same effect.

However, modern and future vehicles are more and more connected to each other and the internet. Due to this trend towards more connected vehicles, the IT within a vehicle can no longer considered to be an isolated system and properly securing these systems against adversaries has become critical.

2.3 PROBLEM STATEMENT

According to an initial literature search, no framework exists that describes a threat modelling technique that identifies risk for future vehicle functionality.

Older frameworks often only focus on the safety aspect of software in the vehicle, and do not consider security. Many newer frameworks that do consider security, either in an integrated approach with safety, or in a separate approach, have no way of incorporating the new attack vectors that have been introduced by all new communication channels such as the telecommunication module.

The problem considered in this thesis therefore focuses on: how to design a framework to help a security expert at a vehicle manufacturer to control threats and risks for future functionality in a vehicle?

3

DEFINITIONS

In order to better understand the research and subject, this chapter provides definitions for the key concepts of the research.

THREAT MODELLING

A procedure for optimising network security by identifying objectives and vulnerabilities, and then defining countermeasures to prevent, or mitigate the effects of, threats to the system.

THREAT MODEL

A model describing possible attack points/threats concerning a system or subsystem based on the resources within the (sub)system the designer cares about. This is often described in a high level model of the (sub)system.

ATTACKER MODEL

A model describing possible attackers of a system or subsystem, including their motives, capabilities, knowledge, window of opportunity, etc.

ATTACK SURFACE

The total sum of the different points (the "attack vectors") in a system where an unauthorised user (the "attacker") can try to access the system.

ATTACK VECTOR

A path or means by which a hacker (or cracker) can gain access to a computer or network in order to deliver a payload or malicious outcome.

AUTOMOTIVE SAFETY

Field of safety in the automotive industry that focuses on minimising the occurrence and consequences of traffic collisions and road safety.

AUTOMOTIVE SECURITY

Field in the automotive industry that focuses on securing vehicle components and communication systems.

DRIVETRAIN

Group of components that deliver power to the driving wheels, often consisting of the clutch, axles, gearbox, final drive, etc. The Drivetrain does not include the engine or motor that generates the power.

POWERTRAIN

The main components that generate power and deliver it to the road surface. The powertrain includes components such as the engine, transmission, differentials, final drive, etc.

ITS - INTELLIGENT TRANSPORTATION SYSTEM

Systems in which information and communication technologies are applied in the field of road transport, including infrastructure, vehicles and users, and in traffic management and mobility management, as well as for interfaces with other modes of transport.

4

RESEARCH DESIGN

The research problem stated in chapter 2 is a design research and can be divided into two problems: a design and knowledge problem, a distinction that comes from design science [70]. This chapter will first describe the research objective and translate these to research and knowledge questions. It will finish by describing the approach how these questions will be answered, and the contributions of this research.

4.1 RESEARCH OBJECTIVE

The objective of this research is to help a security expert within a vehicle, manufacturer to manage threats and possible risks for future functionality in the vehicle by designing a threat model.

This requires answering a few knowledge questions. The first two questions focus on getting a better overview of the IT landscape in which threats and risks are located in the future. The last questions focus on identifying and analysing threats and risks within this landscape.

4.2 RESEARCH QUESTIONS

The main research question is therefore formulated as:

> RQ: What could be a suitable framework that helps a security expert within a vehicle manufacturer to control threats and risks for the vehicle of the future?

In order to answer the research question, the following knowledge questions need to be answered:

1. What functionality will be present in vehicles in five to ten years' time?
2. What will a general IT architecture in a vehicle look like?
    a) What are the attack surfaces within this architecture?
3. What is a threat model for selected functions within vehicles in 5 to 10 years' time for a general IT architecture?

4.3 RESEARCH APPROACH

An initial framework will be created based on a literature research and an expert interview. The expert interviewed will be the following:

- The head of Information Security at a European truck manufacturer

The interview will be conducted in a semi-structured way by asking open ended questions to allow the interviewee to focus on areas where he/she wants to go in-depth. Partial transcriptions will be made during the interview and the interviewee will be given the option to review the transcript to ensure no confidential information is disclosed and the transcript resembles the interview.

The expert interviews and the literature research will be used to answer the knowledge questions and form the threat model that answers the main research question.

After creating the framework, the model will be applied to some use cases to illustrate the working of the model. An additional expert from another truck manufacturer will be interviewed to validate the designed framework on three areas: Completeness, Focus and Workability. The expert interviewed is:

- Cyber Security Systems Architect at a European truck manufacturer

Any feedback from the validation interviews will then be used to improve the framework.

This process is visually described in Figure 4.1, including references to the corresponding chapters.

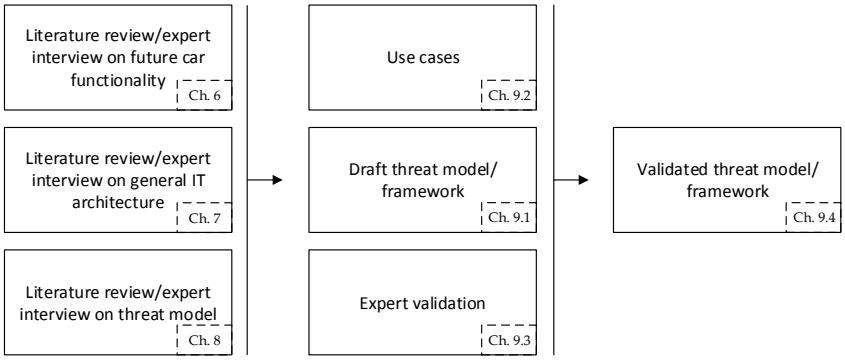

Figure 4.1: Phases, inputs and outputs of this research

4.4 CONTRIBUTIONS

This research contributes to the scientific and practical world by creating a composite threat model that focuses on identifying all possible threats in a vehicle under the assumption that the vehicle has already been breached and an attacker can perform any attacks he likes.

5

LITERATURE REVIEW

In order to get an overview of the state of the art of cyber security in the automotive industry, we performed a systematic literature review based on the approach of Kitchenham [35]. This approach consists of two phases, an initial search, and an iterative backward and forward citation search.

5.1 REVIEW QUESTIONS

Firstly, we define some review questions that define the search scope of the literature review.

1. How is IT integrated in a vehicle?
    a) What IT components does a vehicle consist of?
    b) How are these IT components connected?
    c) How are these IT components secured?
2. How are IT Security Risk Management and Threat Modelling techniques applied to the automotive domain?
3. What kind of attacks exist against vehicles?

5.2 REVIEW METHOD

Because of the fact that there are a lot of papers on the subject of IT in the automotive industry, we decided to focus the initial search on a combination of risk analysis and the automotive industry, and possibly include more detailed studies via the backward and forward citation search when needed.

Data sources and keywords

Based on these research questions, we define some keywords and in- and exclusion criteria. The keywords that were used are "risk analysis", "risk assessment", "secur*", "vehic*", "automotive" and "car". The exact search string that was used can be found in Appendix A. Since the field is relatively new, and to limit the useful results, we only focused on results that were published after the year 2011. Any relevant

papers published before will probably be found in the backward and forward citation search. Next to this, the paper would have to be in the 'Engineering' or 'Computer Science' area. For the initial search, only Scopus was used as database for its wide coverage.

Inclusion criteria

The following inclusion criteria are defined:

1. The study regards information on cyber security in vehicles
2. The study regards information on IT in vehicles
3. The study regards information on risk management in the automotive industry

Needless to say, if the study does not conform to any of these criteria, it is removed from the results.

Exclusion criteria

The following exclusion criteria are defined:

1. The study is not in English
2. The study is not accessible through the University of Twente library subscriptions
3. The study is reported several times

Needless to say, if the study conforms to any of these criteria, it is removed from the results.

5.3 FINDINGS

The initial search resulted in 114 papers. However, these papers still contained some duplicates, which, after being removed, resulted in 110 papers. After reading the abstract, and applying the in- and exclusion criteria, there were 27 papers left.

These 27 papers however, contained 6 proceedings which contain multiple papers. After manually looking through these 6 proceedings, we found another 7 papers, resulting in 28 papers.

These papers were read in full, after which 14 final results were left. This total process is visually described in Figure 5.1.

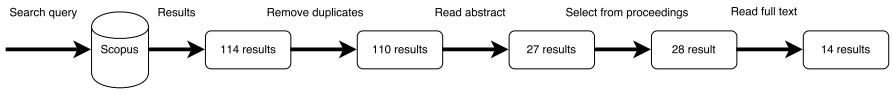

Figure 5.1: Visual representation of study selection

CATEGORY	NUMBER OF PAPERS
In-vehicle communication	15
Cyber risk management	27
Security architecture	8
Attacks	12
Other	4

Table 5.1: Overview of found articles per topic

5.4 BACKWARD AND FORWARD CITATION SEARCH

These 14 results were then subjected to a backward and forward citation search in an iterative process. In this search, mainly *Google Scholar* was used to also include 'grey' papers. This iterative process was repeated for some time, until either no new papers were found, or the found papers were not relevant or considered to be too old. Please note that during this search, papers older than 2011 are also considered.

This resulted in another 144 papers, of which, after reading the abstract, 82 results remained. These papers were read in full, after which 41 final results remained.

5.5 RESULTS

By performing this literature review, a total of 55 papers were found. To give an idea about the topics of these papers, we have categorised them according to the following topics:

IN-VEHICLE COMMUNICATION Papers concerning the internal IT components and internal communication networks of a vehicle

CYBER RISK MANAGEMENT Papers concerning risk management and threat assessment

SECURITY ARCHITECTURE Papers concerning security frameworks and architectures

ATTACKS Papers concerning attacks on vehicles

OTHER Papers that do not concern the other topics

Table 5.1 gives an overview of how many papers cover a certain topic and Table 5.2 shows in which years these papers have been published. Please note that the total amount of papers covering these topics exceeds the total amount of papers of 55, since some papers might cover multiple topics.

YEAR	NUMBER OF PAPERS
Before 2009	9
2009	3
2010	3
2011	4
2012	9
2013	5
2014	12
2015	10

Table 5.2: Overview of found articles per published year

5.6 DISCUSSION

There are some concerns for validity applicable in this literature review. First of all, due to there being only one author, this review has also been performed by only one person. This might introduce a bias in both the found results, as well as the classification of the papers on different topics. However, since the goal of the research is not to provide a good classification, the latter is considered to be not important.

The first bias however, is strengthened by the fact that there were many papers available on the subject, but might not be completely relevant for further research. That is why only newer papers were considered in the initial search and older papers were only added if deemed relevant.

However, since still many papers have been found concerning multiple different topics, and no topic is discussed in a lot of detail, it is considered the bias is minimal.

Part II

CREATING THE FRAMEWORK

6

FUTURE FUNCTIONALITY

Knowledge Question 1:
What functionality will be present in vehicles in five to ten years' time?

To answer this question, a literature review has been performed, an expert has been consulted and websites of vehicle manufacturers have been consulted. This chapter contains an integrated overview of the results. Most results follow from the literature review and looking at the manufacturer's sites.

As explained in chapter 2, vehicles are becoming more connected to one another, which opens a whole new set of opportunities that haven't been possible before. However, for cooperative vehicles to be able to communicate with each other, standardised protocols are needed. Organisations such as the Car 2 Car Communication Consortium (C2C-CC) (a consortium between many car manufacturers, Original Equipment Manufacturers (OEMs) and universities), and the Amsterdam Group (a strategic alliance for the development of ITS) have therefore worked on developing use cases and standardising communication.

Functionality that is not related to cooperative vehicles however is harder to predict, since no standardisation has been conducted and no general functionality that every manufacturer builds exists. Some manufacturers do reveal some futures they are working on or planning on building via their own websites.

The goal of this chapter is to provide a list of possible future functions that might be present in 5 to 10 years' time. It will therefore first provide an overview of future functionality that requires vehicles to more interactively communicate with one another in section 6.1. Section 6.2 provides an overview of functionality that does not require cooperation between vehicles. Finally, this chapter will be concluded in a brief summary in section 6.3.

6.1 COOPERATIVE FUNCTIONALITY

Many different applications have been proposed that vehicles could do in the future. However, to start small and make sure that there

is a basic set of applications that all vehicles support, the European Telecommunications Standards Institute (ETSI) has defined and standardised a basic set of use cases. An overview of these use cases can be found in Table 6.1. As can be seen, the use cases differ a lot in their function, ranging from traffic safety to being able to manage a fleet. To provide an insight in their functions, the classes the use cases fall in are briefly explained in the upcoming sections.

The complete implementation of these use cases takes multiple phases, therefore, together with the C2C-CC and the Amsterdam Group, the ETSI has identified day-one use cases that should be included in the first implementation of ITS [13]. These day-one use cases are highlighted in Table 6.1 and will be described in more detail in the following sections.

6.1.1 *Active Road Safety*

Applications within the Active Road Safety class have as primary objective to improve the road safety. However, it is noteworthy to mention that applications in this class may have secondary benefits that are not directly associated with road safety, such as improving the traffic flow by preventing collisions. This class includes functions such as warning for emergency vehicles, weather or road conditions, traffic condition warnings, and wrong way driving warnings. Within this class, a distinction is made between applications for Cooperative Awareness, which is about sharing information about the surrounding, and Road Hazard Warnings, which is about sharing information about possible hazards.

6.1.2 *Cooperative Traffic Efficiency*

Applications within the Cooperative Traffic Efficiency class have as primary objective to improve the traffic fluidity. This includes functionality such as traffic light optimal speed advisory (GLOSA), enhanced route guidance and navigation, and in-vehicle signage. Within this class, a distinction is made between applications for Speed Management, which is about communicating information about speed, and Cooperative Navigation, which is about communicating information for optimal navigation routes.

6.1.3 *Cooperative Local Services and Global Internet Services*

Applications within the Cooperative Local Services and Global Internet Services classes are meant to provide on-demand information to passing vehicles on either a commercial or non-commercial basis. Examples of these classes are providing notifications of Points of Interest, media downloading, and insurance and financial services. The

BASIC SET	APPLICATION	USE CASE
Active Road Safety	Driving Assistance - Cooperative Awareness	Emergency vehicle warning*
		Slow vehicle indication*
		Intersection collision warning*
		Motorcycle approaching indication*
	Driving Assistance - Road Hazard Warning	Emergency electronic brake lights*
		Wrong way driving warning
		Stationary vehicle - accident*
		Stationary vehicle - vehicle problem*
		Traffic condition warning*
		Signal violation warning*
		Roadwork warning*
		Collision risk warning*
		Decentralized floating car data - Hazardous location*
		Decentralized floating car data - Precipitations*
		Decentralized floating car data - Road adhesion*
		Decentralized floating car data - Visibility*
		Decentralized floating car data - Wind*
Cooperative Traffic Efficiency	Speed Management	Regulatory/contextual speed limits notification*
		Traffic light optimal speed advisory (GLOSA)*
	Cooperative Navigation	Traffic information and recommended itinerary
		Enhanced route guidance and navigation
		Limited access warning and detour notification
		In-vehicle signage*
Cooperative Local Services	Location Based Services	Point of Interest notification
		Automatic access control and parking management
		ITS local electronic commerce
		Media downloading
Global Internet Services	Communities Services	Insurance and financial services
		Fleet management
		Loading zone management
	Life Cycle Management	Vehicle software/data provisioning and update
		Vehicle and RSU data calibration

Table 6.1: ETSI basic set of applications, adopted from [13].
Starred use cases have been identified as day-one use cases. Note that some use cases have been merged in the list of day-one use cases

difference between Cooperative Local Services and Global Internet Services is that the former services are provided via or by the ITS infrastructure, whereas the latter is acquired from providers in the internet. The Global Internet Services class is furthermore divided into two, the Communities Services, which is about providing services for certain communities, and Life Cycle Management, which is about providing data and updates.

6.1.4 *Day-One Use Cases*

The Car-to-Car Communication Consortium and the Amsterdam Group have defined a set of day-one use cases that should be implemented in the first roll out. The following list provides these day-one use cases including a small description[1]:

HAZARDOUS LOCATION WARNING
: The hazardous location warning is designed to inform the driver of a vehicle about upcoming dangers on the road. This information can for example be about obstacles on the road or weather conditions such that the driver might slow down the vehicle in advance.

SLOW VEHICLE WARNING
: The slow vehicle warning is designed to warn the driver about slow vehicles in front of the driver to avoid or mitigate rear-end collisions. The system is not designed to act upon the warnings to avoid an impending collision, however, it will warn other vehicles on the potential danger.

TRAFFIC JAM AHEAD WARNING
: The traffic jam ahead warning is designed to warn a driver about an upcoming traffic jam to avoid rear-end collisions. By communicating this, a driver can be warned about this danger even before the traffic jam can be noticed by the driver himself.

ROAD WORKS WARNING
: The road works warning is designed to warn a driver about upcoming road works. Road side units, mounted on a road work warning trailer, send messages to approaching vehicles so that they are aware of potentially dangerous conditions at the road works.

STATIONARY VEHICLE WARNING
: The stationary vehicle warning is designed to warn the driver about possible disabled vehicles, or might serve as warning for vehicles that are about to break down. Vehicles that receive this information are to relay it to other vehicles.

1 More detailed descriptions can be found on http://www.drive-c2x.eu/use-cases. Please note that every instance provides other descriptions of the use cases, but they all boil down to the same information being processed

IN-VEHICLE SIGNAGE INCLUDING SPEED MANAGEMENT
: The in-vehicle signage including speed management is designed to make drivers aware of potentially dangerous conditions. Road side units mounted on traffic sign and other key points along the road send messages to approaching vehicles such as speed limit or other signs that the driver could have missed.

PROBE VEHICLE DATA
: The probe vehicle data is designed to inform other road users for use in traffic management. Road side units gather anonymised sensor data such as speed, braking force and weather conditions from passing vehicles to get knowledge about that part of the road.

SIGNAL PHASE AND TIME
: The signal phase and time is designed to inform the driver about the current status and next change of the traffic signal ahead. With this information, the vehicle can provide information about the best speed to approach the signal.

EMERGENCE VEHICLE WARNING
: The emergence vehicle warning is designed to inform the driver about an approach emergency vehicle that claims the right of way.

EMERGENCY BRAKE LIGHT
: The emergency braking light is designed to avoid rear-end collisions that occur after a vehicle driving ahead suddenly brakes. Especially in dense driving situations or situations with decreased visibility, the driver can be warned before he notices the sudden brake himself, especially if there are vehicles in between and the driver cannot see the braking vehicle directly.

MOTORCYCLE APPROACHING INDICATION
: The motorcycle approaching indication is designed to warn drivers about motorcycles. The motorcycle continuously provides movement and position information to nearby vehicles so that these vehicles can compare their movement with the motorcycle data and warn the driver for possible collisions.

6.1.5 *Truck Platooning*

A recent trend that has received a lot of attention in the media, especially in Europe, is truck platooning [34]. In truck platooning, a few trucks follow each other at a short distance in so called truck platoons, whilst communicating driving information with each other. Only the first truck is driven by a person, and communicates its speed, or braking information to the trucks following in the platoon, that alter their speed to match the first truck. This way, drivers from the other trucks can spend their time on other matters.

Currently, truck platooning pilots use an altered version of WiFi (802.11p or also called WAVE) [61] as means for communication. This protocol only describes the Physical and Medium Access Control (MAC) layers, and manufacturers have built their own communication protocols on top of it, meaning that trucks from different manufacturers cannot yet cooperative with each other [34].

However, in time, these truck platoons, as well as communication between other vehicles, will use the ITS G5 standard as described by the ETSI. This standard also uses the 802.11p standard for the Physical and MAC layers, but also standardises the communication on top of these layers, making sure vehicles from different manufacturers can communicate with each other [14].

Although part of the ITS G5 standard does include security, the expert interview revealed that during this first test with platooning, security has not yet been addressed (Appendix B). It is for example possible to jam WiFi signals, which could have potentially catastrophic consequences when driving on such short distances from one another.

According to Janssen et al., wide-scale usage of truck platooning will start from 2020 onwards [34]. The development path of platooning has three different paths, each with three different stages:

INFRASTRUCTURE At first, platooning will only happen on closed areas. In the second stage, it will be used on public main roads, but still only on a national level. Finally, trucks will also platoon on public main roads on an international level.

FORMATION At first, the forming of platoons will be self-organised, fleet owners might schedule two drivers with the same destination to depart at the same time and form a platoon. In the second stage, a Platooning Service Provider might couple trucks based on requests from fleet owners. Finally, truck platoons will form on-the-fly; whenever they are driving and see other trucks, they might form a platoon.

AUTOMATION At first, platooning will require every truck to have a driver that stays awake to take over control when necessary. In the second stage, trucks will still need a driver, but drivers in the back of the platoon might rest. Lastly, one driver will control the whole platoon, and no other drivers are needed.

6.1.6 *Virtual Traffic Lights*

Another development is the research of Virtual Traffic Lights (VTL) [38]. The idea behind Virtual Traffic Lights is that vehicles, possibly with a road side unit, communicate with each other to determine which vehicles are allowed to go first at an intersection. Instead of having physical traffic lights at an intersection, these traffic lights can for example be projected on the dashboard and are all virtual.

The advantage of Virtual Traffic Lights is two-fold. In rural areas, where, especially in the USA, there are intersections without traffic lights, the use of Virtual Traffic Lights can increase traffic safety where it would otherwise be to expensive to place a physical traffic light. Next to this, in urban areas, Ferreira et al. have shown that self-organising traffic that is facilitated by Virtual Traffic Lights can improve traffic flow by 60% during rush hours [16].

Although this optimisation of traffic flow is technically also possible by using adaptive traffic lights, a 2010 survey has shown that 70 to 90% of traffic lights in the USA is non-adaptive [63], and replacing these is costly.

6.2 INDIVIDUAL FUNCTIONALITY

The previous section has given an overview of the different future applications in the cooperative domain: where vehicles communicate with each other and road side units. Since these applications are shared between manufacturers, standardising how this communication has to take place is vital. However, manufacturers are also working on functionality that is not shared between vehicles, such as Parking assistance or Adaptive cruise control. Since this functionality is not shared between manufacturers, only little information can be found on it.

Below is a list of technologies that are currently finding their way onto the market. Since many truck manufacturers often lack a bit behind car manufacturers, the functionality that car manufacturers bring onto the market today might prove a good estimate for what truck manufacturers might bring in 5 to 10 years' time.

LANE KEEPING ASSISTANCE
: The lane keeping assistance is designed to continuously monitor whether the vehicle is beginning to move outside of its lane, and either warns the driver, or acts to ensure the vehicle stays in its lane. This technology is already on the market for both cars and trucks, however, for trucks, the system only warns the driver, and does not act itself. Found in [1, 3, 11, 19, 43, 52].

PARK ASSISTANCE
: The park assistance is designed to detect whether it is possible to park the vehicle in a parking spot, and to actively park the vehicle in that spot. This system is currently being brought on the market for cars, however, it is not available for trucks yet, although some experts have noted working on autonomous parking and (un)loading. Found in [4, 12, 18, 68].

PREDICTIVE CRUISE CONTROL
: The predictive cruise control is designed to use GPS data of the vehicle and map data to actively predict what will happen in two to three kilometres and adjust the speed of the vehicle. If

the vehicle for example has to climb a hill in two kilometres, it might already start building a momentum and use the mass of the vehicle to get over the hill. The system is mostly designed to improve traffic flow and save fuel. The system is being introduced in trucks, but has not found its way to cars yet. Found in [42, 53].

CONTINUOUS DAMPING CONTROL

The continuous damping control, or sometimes also called the vehicle stability control, is designed to monitor the stability of the vehicle in for example tight corners to make sure it doesn't role over. It actively monitors the role and pitch and adjusts the hardness of the dampers to make sure the truck doesn't flip over. The system is designed for use in trucks and has recently found its way onto the market. Found in [40, 54].

EMERGENCY BRAKE ASSISTANCE

The emergency brake assistance is designed to monitor other vehicles on the road and check if they are suddenly braking so that the vehicle can brake as well. It is already used in cars and is being introduced in trucks. Found in [41, 51].

REMOTE UPDATES

The remote updates are designed to be able to flash a new firmware on an ECU from a distance. Since recalling a vehicle is very expensive, both in time and money, being able to remotely update parts of the vehicle increases the robustness. Remote updates are already being used by some car manufacturers, but is not yet used in trucks. Found in the expert interview in Appendix B and [67].

DATA LOGGING AND REMOTE DIAGNOSTICS

The data logging and remote diagnostics are designed to be able see what is happening inside the vehicle. This can for example be used after an accident to check what happened, or to proactively monitor and predict the break down of parts. Right now, data logging is already done to some extent in both cars and trucks, but often only for a short amount of time. In some cars, the logging is done to more extent and even transferred to a central server owned by the manufacturer. In trucks, the last part is not happening yet. Found in the expert interview in Appendix B and [44].

FLEET MANAGEMENT

The fleet management is designed to be used by the manufacturer or the fleet manager to be able to monitor the location, speed, driving behaviour, fuel consumption, etc. of all trucks in the fleet and act upon that information. Some fleet management systems are for example used to monitor whether a truck is stolen and prevent the engine from starting again. Although there are quite some use cases for fleet management, the main

motivators remain increasing efficiency and preventing theft. Fleet management is still only used for trucks. Found in the expert interview in Appendix B and [44].

6.3 CHAPTER SUMMARY

This chapter set out to answer research question one:

> What functionality will be present in vehicles in five to ten years' time?

To this end, we looked at the possible future functionality that might be present in vehicles by looking at the relevant literature and functionality that different vehicle manufacturers are introducing. We have identified a trend towards more complex functions, ranging from autonomous vehicles to more cooperative functionality. For the coming five to ten years, we separate two types of functionality: cooperative functionality, and individual functionality.

Cooperative functionality are applications or systems that focus on cooperative driving between vehicles, for example sharing information about road conditions, traffic information, or emergency braking. Since cooperative functionality requires vehicles from different manufacturers to communicate with each other, the communication requires standardisation. The ETSI has therefore defined a clear road map of possible applications of cooperative driving and identified a list of day-one use cases that should be included in the first implementation of ITS. These day-one use cases are relatively simple use cases, and often only focus on alerting the driver of hazardous situations such as emergency braking or upcoming traffic jams. But applications in the far future might include vehicles to make decisions based on this information.

Besides these ETSI standardised applications, two other applications have received a lot of attention: Truck platooning, where trucks drive in platoons and only the first truck is controlled by a driver, and Virtual Traffic Lights, where traffic lights are displayed on the dashboard of the vehicle, and the vehicles decide who has green light, and who has to stop.

For individual functionality, vehicles do not need to exchange any information, and hence, not standardisation is required. Because of this, it is harder to predict what individual functionality might be present in future vehicles. We identified some possible future functionality by looking at what certain manufacturers are introducing in their vehicles now, which may be taken over by other manufacturers in the future. By for example looking at what car manufacturers are introducing now, we can make an educated guess on what some truck manufacturers might introduce in the future, and vice versa. These kind of functionality include autonomous driving capabilities

such as parking assistance and predictive cruise control, but remote software updates and fleet management as well.

7

VEHICULAR IT ARCHITECTURE

Knowledge Question 2:

What will a general IT architecture in a vehicle look like?

As described in chapter 2, the IT architecture within the vehicle has changed a lot in all those years. Where it was sufficient to connect one ECU with another in the early times, it now no longer suffices to connect 70 ECUs with one another. Even the ECU itself has changed in how it is built and in what way it communicates with other ECUs. This chapter will therefore give an overview of the vehicular IT architecture.

The goal of this chapter is to provide an overview of the different IT components in a modern vehicle, explain how these components communicate with each other, provide an overview of possible attack surfaces, provide a general IT architecture of a vehicle and show how this architecture is secured. To that end, this chapter will first provide an overview of the IT components in section 7.1. It will then provide an overview of the different communication networks inside the vehicle in section 7.2. Next, the different attack vectors present in a vehicle will be provided in section 7.3. After that, this chapter will construct a general IT architecture per region in section 7.4. Lastly, it will briefly explain how certain research project have proposed means to secure these architectures in the future in section 7.5.

7.1 VEHICULAR IT COMPONENTS

As explained in chapter 2, ECUs are specialised pieces of hardware and software designed to do specific tasks and therefore differ a lot from conventional PCs and networks as the internet that most people know and are used to. Where PCs and the internet have steadily grown in an open world where they have seen a lot of attacks and could be properly defended, vehicles have remained isolated during the biggest part of their lifetime. Only recently has the vehicle been opened up and connected to systems we use every day, bringing with it a lot of potential attacks, such as hacking the telematics unit [10]. To better understand what IT components are present in a modern

vehicle, a brief description of what an ECU does, and how it is secured now is provided, followed by an description of the new attack vectors that have emerged in recent years.

7.1.1 *Electronic Control Units (ECU)*

The modern day ECU has undergone a lot of changes since it first introduction as a controller to adjust the fuel/air mixture in the engine, to a component that has now become part of almost every aspect of the vehicle. Modern day vehicles can contain upto 100 different ECUs, ranging from performing simple tasks such as controlling the lights, to more complex driving functions such as parking assistance.

Many of the simpler ECUs have been around for some time and haven't changed much. These ECUs have often been standardised and can be found in vehicles from all kind of brands. They are therefore often not built by the manufacturers themselves anymore, but by suppliers such as Bosch. Since more advanced systems such as lane keeping assistance can still give a manufacturer an edge on the market, they are still built by the manufacturers themselves. Part of this standardisation of ECUs is the establishment of AUTOSAR, a standard for building ECUs.

AUTOSAR

Established in 2003 as a partnership by many companies in the automotive industry, AUTOSAR (AUTomotive Open Systems ARchitecture) aims at establishing an open and standardised software architecture for ECUs [7].

The architecture standard describes basic software modules, defines application interfaces and described a common development methodology for building ECUs. These different aspects are illustrated in Figure 7.1. As can be seen, there are no specific security modules in the AUTOSAR architecture. The security aspects that have been added mainly focus on reducing the amount of errors between modules and interfaces by standardising these and so reducing the amount of functional safety issues[1].

There are some security functions in place in the form of memory partitioning and protection against unauthorised flashing of ECUs [6]. Prior to flashing, ECUs must for example perform a challenge-response protocol, and should not perform a flash if it is deemed unsafe. However, Koscher et al. show that in practice, this is not always the case [37]. Some extensions to AUTOSAR to include more security features into the framework have been proposed, however, these have not been adopted yet [5].

1 Functional safety issues are safety issues that are caused by programming errors or bugs

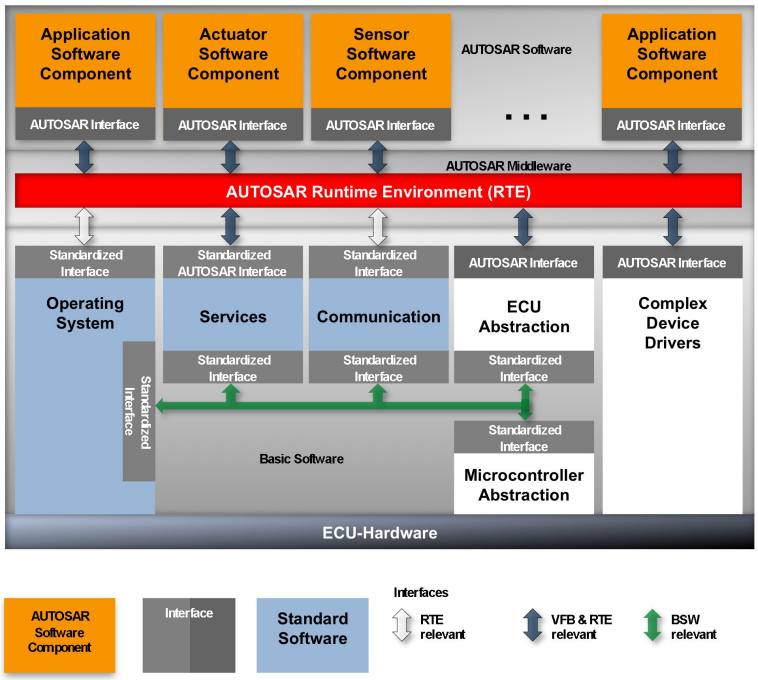

Figure 7.1: The AUTOSAR software architecture

7.1.2 *Entry Point Modules*

Within the vehicle, there are a lot of different modules or ECUs with a potential entry point from outside the vehicle network, ranging from the mandated OBD-II diagnostics port, to vehicles with an internet connection. Checkoway et al. distinguish three different categories of entry points [10]:

1. Physical access
2. Short range wireless access
3. Long range wireless access

Section 7.3 will give a more detailed overview of these modules that can serve as an entry point for an attacker.

7.2 IN-VEHICLE COMMUNICATION

ECUs first operated on a individual basis, however, after a while, they started communicating with one another. At first, it sufficed to connect all ECUs that needed to be connected via their own channel. However, when the functions became complex, and needed more connections, it became insufficient and cost-inefficient to connect them all together. Manufacturers therefore started standardising digital buses.

Because each ECU has a very specific function, ranging from safety critical driving functions to entertainment systems, they also have

different requirements for the communication system they are connected to. The braking system should work immediately when hitting the brake, whereas mirror-adjustment might take a second longer. Therefore, most modern day vehicles are equipped with multiple different buses, each with their own goals and costs. Koscher et al. distinguish five different categories for in-vehicle communication systems:

EVENT-TRIGGERED Used for real-time communication between controllers such as the Antilock Braking System (ABS)

TIME-TRIGGERED Used for time-critical and safety relevant communication such as drive-by-wire[2] systems

SUB BUS Used for small autonomous networks for less-critical systems such as door-locking and mirror-adjustment

MULTIMEDIA Used for high performance communication channels such as video data streaming and other media

WIRELESS Used to let the in-vehicle network communicate with external devices such as a mobile phone

GROUP	EVENT-TRIGGERED	TIME-TRIGGERED	SUB BUS	MULTI-MEDIA	WIRELESS
System	CAN	FlexRay	LIN	MOST	Bluetooth
	VAN	TTP	K-Line	D2B	GSM
	PLC	TTCAN	I²C	GisaStar	WLAN

Table 7.1: Grouping of selected automotive bus systems, adapted from [71]. Most used systems are on top of the list

Table 7.1 gives an overview of existing in-vehicle communication systems and in what category they fall. As stated above, wireless communication systems are often used for connecting external devices, such as a mobile phone, to the vehicle. There are some examples of wireless communication systems that are used within the vehicle's in-vehicle network, such as the tire pressure monitoring system. However, these protocols are well known, and won't be discussed any further. The next sections will give some more details on the most used communication systems of the other categories, as well as the recent development of using Ethernet as a replacement for these systems.

7.2.1 *Controller Area Network (CAN)*

CAN is the most widely used in-vehicle communication system available. It was originally designed by Robert Bosch GmbH in 1983 and later adopted into ISO standard 11898 [24]. In its basis, CAN is a serial, event-triggered messaging system that broadcasts messages to

2 Drive-by-wire is the name given to systems that use electronic pulses transmitted via a wire to an actuator rather than using for example brake fluid

all nodes. The newest version of CAN can achieve data rates of up to 1 Mbit/s. Every messages basically only contains a message identifier, the data and a checksum. The identifier is used by ECUs to see whether that message is important for it and whether it should be processed. CAN does offer a priority system for messages, in the form of the height of the message identifier. Due to it being an event-triggered messaging system, it can deliver real-time communication that is important for safety critical systems.

CAN however, does have its drawbacks. Since CAN was developed back in 1983, when vehicles were still considered to be isolated system, it does not feature any security features and hence, an adversary that gains access to the network, can inflict a lot of damage. For one, he could inject an infinite amount of messages on the network, performing a denial-of-service attack that could for example disable the brakes [37].

Since CAN is the most widely used bus available and is often used for safety critical systems, it has been researched extensively and quite some attacks against it have been found [9, 10, 23, 36, 37, 64]. These attacks are all practical examples of the same principle; get access to the CAN bus, and you can do almost anything.

7.2.2 *FlexRay*

FlexRay was designed by the FlexRay consortium, a group of car manufacturers and OEMs, including Robert Bosch GmbH. The consortium worked from 2000 until 2009, upon which version 3.0 of FlexRay was adopted as ISO standard 17458 [27]. It is designed for high data rates of up to 10 Mbit/s, making it particularly suited for x-by-wire applications.

FlexRay has not been as extensively researched as CAN. However, Nilsson et al. note security has not been a design aspect of FlexRay, and simulated attacks have shown that FlexRay indeed lacks security mechanisms. However, this attack was performed on the old 2.1 version of the protocol and no research has been performed on the 3.0 version of the protocol. Revision notes of the FlexRay protocol however, do not mention any overhaul for security [17], so it is plausible these simulated attacks still exist.

7.2.3 *Local Interconnect Network (LIN)*

The Local Interconnect Network (LIN) was designed by the LIN Consortium, a group of five car manufacturers, in the late-1990s. It was designed as an alternative for CAN where CAN would be too expensive to use and is therefore only a very simple communication system. It features small data rates of up to 20 kBit/s, making it extremely suitable for small tasks as mirror adjustment and lighting [71].

Not a lot of research has been done on attacks on the LIN protocol. This is probably due to the fact that LIN is not used for any critical systems. A hacked LIN might cause some annoyance to the driver, but it hardly causes any safety issues. However, since the design of LIN does not contain any security features, getting access to the bus gives access to any ECU on that bus.

7.2.4 *Media Oriented System Transport (MOST)*

The Media Oriented System Transport (MOST) was designed to be used for high-bandwidth media transport in the automotive industry. It features high data rates of up to 24 MBit/s, making it very suitable for transmitting audio, video, navigation, etc.

Again, not much research has been done on the security of MOST, probably because it does not access any safety critical systems. However, some attacks focus on using the media systems to deceive the driver by for example showing in the mp3 player that the engine is damaged and the driver needs to stop immediately [23]. This could create potentially dangerous situations, depending on what system is hacked, and what is done with it.

7.2.5 *Ethernet*

In recent years, many vehicle manufacturers have started looking into Ethernet as a replacement for in-vehicle communication [62]. Because of its wide use by consumers, Ethernet has become relatively cheap to implement. To what extent Ethernet will replace other standards still differs per manufacturer, some are researching whether Ethernet can replace all communication channels, whereas others only focus on using it for infotainment systems, or as a replacement for other non-critical systems such as a separate system for diagnostics and firmware updates [66].

7.3 ATTACK SURFACES

Knowledge Question 2a:

What are the attack surfaces within this architecture?

The previous section has provided an overview of the different components in a vehicle and how they communicate with each other. Before going to what a general IT architecture might look like, this section will first focus on identifying possible attack surfaces of a vehicle: what entry points an attacker might use to compromise a vehicle.

Not much extensive research has been done on automotive attack surfaces. Three papers provide some attack surfaces. Miller and Valasek

provides an extensive list of the possible attack surfaces on 21 different car models [48]. Checkoway et al. are the first to provide a classification of possible attack surfaces, along with some attack surfaces [10]. Lastly, Zhang, Antunes, and Aggarwal focus on attack surfaces particularly for malware [72].

The following paragraphs provide the found attack surfaces from the three papers combined in the classification given by Checkoway et al. [10].

7.3.1 *Physical Access*

Almost all vehicles provide several physical interfaces with either direct or indirect access to the vehicle's internal network.

OBD-II PORT Almost all modern day vehicles contain an OBD-II port, which is even mandated by the U.S. Government. It typically provides either direct access to the CAN-bus, or via a central gateway and is often used for diagnosing the car by for example garage employees. Originally, specific scanning tools for connecting to the OBD-II port were used, but nowadays, garages often use a PC to connect to the OBD-II port via a PassThru device. This means that a compromised PC could possibly compromise the vehicle.

ENTERTAINMENT/REMOVABLE MEDIA PORTS More and more vehicles contain entertainment systems such as CD-players, USB-ports or iPod connectors for playing music. Although compromising a CD-player is relatively harmless, these interfaces are becoming integrated with the in-vehicle networks to be able to deliver for example hands-free features. It is therefore possible that an iPod containing malware might compromise the vehicle.

7.3.2 *Short Range Wireless Access*

Newer vehicles often don't provide physical access for entertainment anymore, but use short range wireless access systems. It is for example used to connect external devices such as mobile phones, but also for remote key entry or tire pressure monitoring. In this category, 'short range' means around 5 to 300 meters.

PASSIVE ANTI-THEFT SYSTEM (PATS) Many newer vehicles contain a Passive Anti-Theft System, a sensor in the steering column that communicates with the ignition key. The on-board computer simply sends out an RF signal to which the key should respond. This way, the vehicle checks whether the key is in the proximity of the vehicle, and it is not started by an attacker. The attack surface is considered only small, an attacker might

use a Denial-of-Service to prevent the vehicle from receiving a response from the key.

TIRE PRESSURE MONITORING SYSTEM (TPMS) Some vehicles contain a Tire Pressure Monitoring System, a system that constantly measures the tire pressure and transmits this to an associated ECU. The attack vector is mainly focused on letting the vehicle believe it's having a tire problem or suppressing the warning message if it's having a real tire problem, the attack surface is therefore rather small. However, Ishtiaq Roufa et al. have shown it is possible to brick the associated ECU in some cases [32].

REMOTE KEYLESS ENTRY/START (RKES) Many vehicles contain a Remote Keyless Entry or even Remote Keyless Start system. Often, it is designed that the key will send some encrypted identifying information from which the ECU can determine if the key is valid. The attack surface is rather small, a DoS might again prevent the vehicle from (un)locking or starting. Although, since data processing takes place in the vehicle, code execution is technically possible if not programmed securely.

BLUETOOTH Many vehicles provide the ability to sync a device over Bluetooth with the vehicle. Often, this process is secured by demanding some user interaction to pair a device. However, the Bluetooth stack is quite large, and has contained vulnerabilities in the past. The fact that an external device is connected to the in-vehicle network creates a big attack surface via Bluetooth.

WIFI Many modern day cars and trucks contain WiFi that serves as a hotspot for mobile phones and other mobile equipment. These hotspots for example bridge the internet connection of the vehicle to the mobile phone, or serve as an entry point for multimedia.

EMERGING SHORT RANGE CHANNELS Especially now, a lot of research is begin done in Vehicular networks, where vehicles set up wireless networks to communicate with each other about upcoming traffic jams or road conditions

7.3.3 *Long Range Wireless Access*

Many vehicles are also integrating long range wireless access systems in the vehicle for communication over distances greater than 1 km. Two categories can be distinguished: broadcast channels, such as GPS and radio, and addressable channels, such as a 4G internet connection.

RADIO DATA SYSTEM Common radio systems no longer only receive audio signals, but data signals as well. Though no real parsing of data is present, it is possible such systems are susceptible to code executions. However, the attack surface is considered to be

small, mainly because the likelihood of such an attack is rather small.

GLOBAL POSITIONING SYSTEM A lot of vehicles contain one or more GPS receivers for use in navigation or internal automation. Connected vehicles for example, tell other vehicles about upcoming traffic jams or road conditions, which needs a location as well. GPS has been shown to be susceptible to spoofing attacks [57], making it possible for an attacker to divert a vehicle or render such systems inoperable.

TELEMATICS / CELLULAR Many new vehicles contain cellular radios or telematic units to gain access to for example traffic or weather information. This telematic unit gives a very broad attack surface, since it basically connects the vehicle to the internet. Although the telematics unit probably does not reside on the CAN-bus but for example on the media bus, it is often still connected to the CAN-bus via some other bridging ECU or gateway. Checkoway et al. have shown that some automotive telematic units are exploitable [10], and can for example be used to kill the engine or activate the windscreen wipers.

INTERNET / APPS Google and Apple have started building app stores and apps such as navigation apps specifically for use inside the car. This new trend brings a wide attack surface inside the car, ranging from malware in an app, to browser exploits. This code is often very complex, making it hard to secure properly.

An overview of the different attack vectors mapped on the wideness of the attack surface and the attack distance is given in Figure 7.2. In general, the further away an attacker, using a wider attack surface, the higher the potential risk.

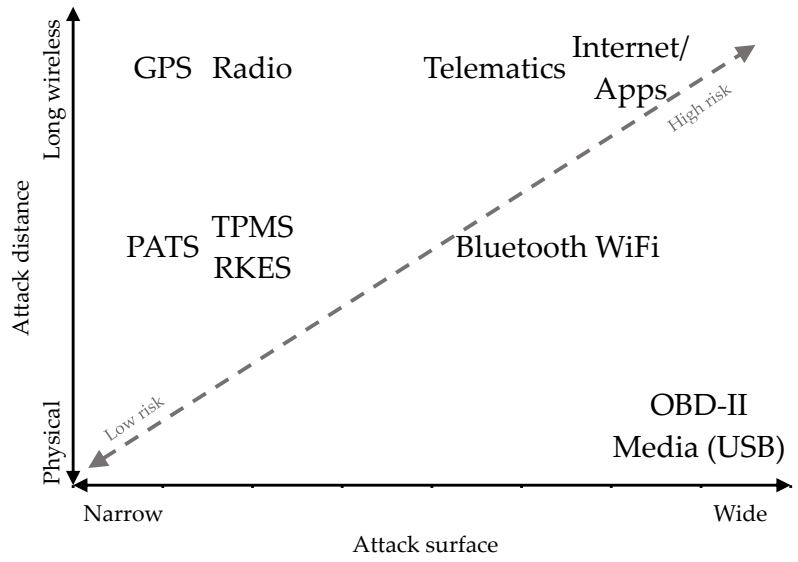

Figure 7.2: Mapping of attack surfaces

7.3.4 *Sensors*

A lot of newer functions, as already mentioned in chapter 6, require sensors to observe the physical world, such as cameras, radar or odometric sensors. Of course, these sensors can be tricked, providing false data [56]. However, since these sensors are outside of the digital world, they are considered outside of scope for this research.

7.4 IT ARCHITECTURE

Having looked at the different components a vehicle can consist of and via which buses these components can communicate with each other, there are numerous ways of how these are eventually connected in the vehicle. This section will look at the different IT architectures that each vehicle has.

In practice, almost each manufacturer, and each vehicle model has its own IT architecture. Some manufacturers, that are part of a larger manufacturing group, do have similar or the same architectures, though these examples are limited.

Even though almost every vehicle model has its own architecture, Miller and Valasek have examined the architecture of 24 different car models and concluded that manufacturers based in the same region have similar topologies [48].

These 24 car models have been analysed and used as a basis to create a general IT architecture of vehicles per region, which will be provided in the coming sections. Of each region, a schematic overview is provided of which the legend can be found in Figure 7.3. Special interest is also given to the location of the different attack vectors given in section 7.3.

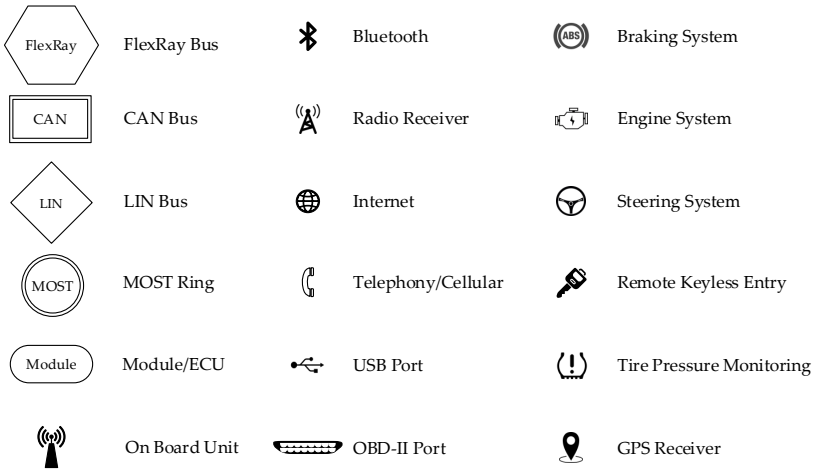

Figure 7.3: Vehicular IT Architecture Legend

In this overview, the safety critical functions will be highlighted, as well as the components that provide an attack surface as mentioned in section 7.3.

Please note that definitely not every vehicle has the same architecture. Even within the same manufacturer, vehicle's architecture can already differ a lot. However, some abstract (sub-)architectures do have similar characteristics. Also note that in these architectures, the location of the On Board Unit (OBU), the unit used for V2V Communication, is highlighted. This unit is not yet part of the vehicular IT architecture, but will be in the future. Based on assumptions of the EVITA project, the location of the OBU is added in this architecture.

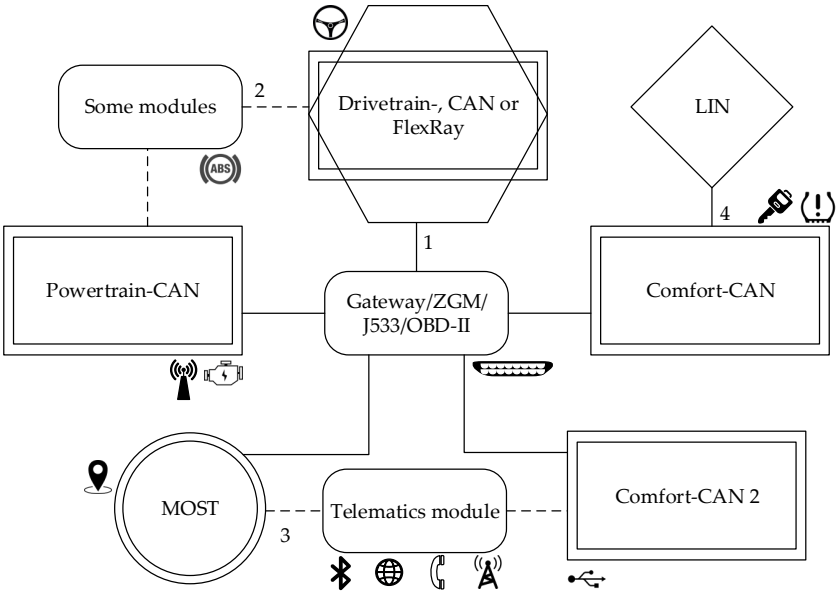

Figure 7.4: General IT architecture for a European vehicle. Please note:
1. The Drivetrain is sometimes a CAN bus, and sometimes a FlexRay bus.
2. Often, there are modules present that reside on both the Powertrain and the Drivetrain.
3. The telematics module sometimes resides on the MOST bus, sometimes on a Comfort-CAN, and sometimes on both.
4. Some modules on the Comfort-CAN have their own separate LIN buses for their functionality

7.4.1 *Europe*

The IT architectures in Europe are the most alike in the world. It is clearly visible that certain manufacturers have often discussed or researched IT architectures together in projects such as EVITA, SEVECOM, or sim^{TD}. Even though these collaborations have not been backed

by all manufacturers directly, they have always been backed by the C2C CC of which all European, and some non-European, manufacturers are part of. Figure 7.4 shows a general IT architecture for a European vehicle.

Characteristic for the European IT architecture is the division in multiple different buses, connected via a central module, that serves as a central module or gateway, and serves as a connector for diagnostics.

In the European IT architecture, there is a strict division between buses with safety-critical modules, and buses with comfort modules. Even within this division, the buses are divided again. The safety-critical bus often exists of two or three different buses, that divide it in a drivetrain bus, and a powertrain bus.

How many different buses there are in total differs per manufacturer and module, but in general, there is one drivetrain, one or two powertrains, and one or two comfort buses, that could possibly have multiple sub buses.

Attack Vectors

In Europe, the telematics module often contains the radio receiver, bluetooth connector, cellular unit and the internet connection. It is almost always separated from safety-critical functions and resides on either the MOST ring, the Comfort bus, or sometimes on both.

Next to this, the central gateway serves as the entry point for the OBD-II port, the Keyless Entry System resides on the Comfort bus, and the Tire Pressure Monitoring System often resides on the Comfort bus, but sometimes on the Powertrain bus.

7.4.2 *United States*

Particular for the IT architecture in the United States is that, more than in Europe, within the same manufacturer or manufacturing group, models tend to have similar IT architectures. The IT architecture of a Dodge Viper, Dodge Ram, Chrysler 300 and Jeep Cherokee (all part of the Chrysler Group) for example, all have (almost) identical architectures. The general IT architecture for an American vehicle can be found in Figure 7.5.

Characteristic for the American IT architecture is that they do not separate many buses. Often, a vehicle contains only two buses, a High speed bus for safety-critical functions, and a Medium or Low speed bus for other functions. Only some vehicles separate the High speed bus in a separate Powertrain and Drivetrain bus. However, even then, there are modules that reside on both these buses.

However, the division between a High and Low speed bus is not that strict as in the European architecture. There are often multiple modules that reside on both the High speed and Low speed bus.

Only some vehicles contain a MOST bus, which if present, is often connected to the gateway, but also connected via other modules to the Low speed bus.

Attack Vectors

The OBD-II port, that has been mandated by the US government is often a separate module that resides on both the High and Low speed bus, however, it is sometimes part of the gateway or BCM module.

The telematics module often resides on both the High and Low speed bus, although it is not necessarily connected to the High speed bus. In the case that the High speed bus is separated, the telematics unit is not connected to the drivetrain bus, but is to the powertrain bus.

The Keyless Entry System often resides on the High speed bus. The location of the Tire Pressure Monitoring System differs, sometimes it is on the High speed bus, sometimes on the Low speed bus.

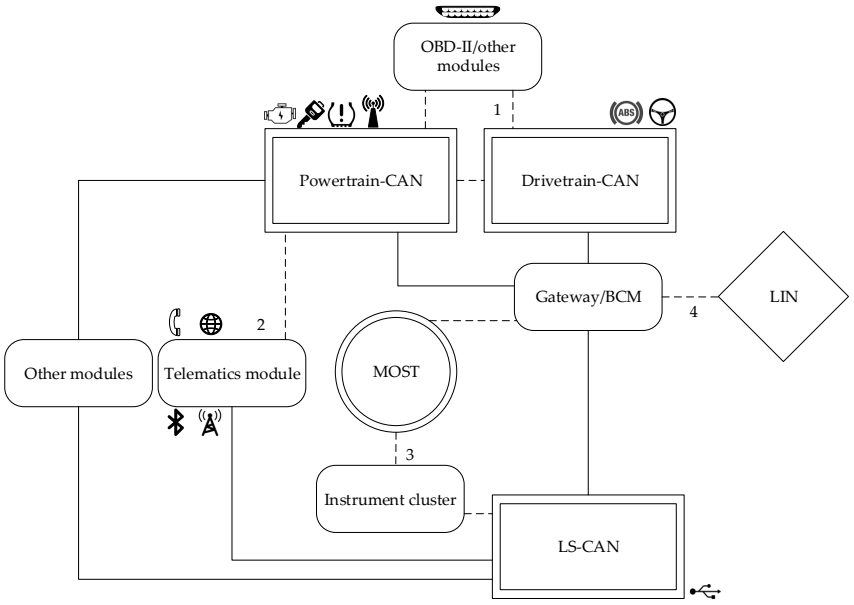

Figure 7.5: General IT architecture for an American vehicle. Please note:
1. The Powertrain and Drivetrain are often one bus, but are sometime separated. If they are separated, there are modules that reside on both buses.
2. The telematics module sometimes also resides on the Powertrain.
3. Not all vehicles contain a MOST bus. If they do, it is connected to the Gateway and other modules such as the instrument cluster that connects it to the LS-CAN.
4. Some less critical functions reside on separate LIN networks. Not all vehicles have this.

7.4.3 *Asia*

It is much harder to construct a general IT architecture for Asia. The architectures in the vehicles differ much more than in America or Europe. However, it is still possible to note a few general things in their architecture. Figure 7.6 shows a general IT architecture for Asia.

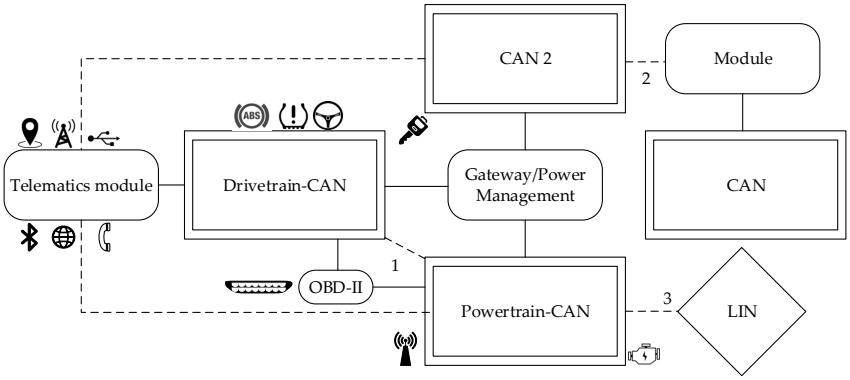

Figure 7.6: General IT architecture for an Asian vehicle. Please note:
1. The Powertrain and Drivetrain are often one bus, but are sometimes separated.
2. Some more advanced modules on the CAN 2 bus contain their own CAN bus for communication with sub-modules.
3. Some modules have their own separate LIN networks.

In Asia, there is a more layered approach in buses. There is often one main bus, where almost all critical functions reside on. In that case, some modules have a separate bus for specific functions, on which some modules can again have a separate bus.

In some cases, the main bus is divided in a Drivetrain and a Powertrain bus.

Attack Vectors

What is important to note is that the telematics unit is always connected to the Drivetrain bus, and often also on the Powertrain bus. Next to this, the Keyless Entry System is sometimes on the Drivetrain bus, and sometimes on the CAN 2 bus. The Tire Pressure Monitroing System is always on the Drivetrain bus. This means that in almost all cases, the attack vectors are directly linked to the safety critical functions.

7.4.4 *Trends in Time*

Looking at the the different IT architectures over time, there are a few trends visible of changes in the IT architectures. Whether it is for security reasons, or simply because there are too many ECUs to put

on a singe bus, but there is an ongoing trend to separate buses based on functionality. Safety-critical modules are becoming more separated from non-critical modules. Even within the safety-critical modules, there is a separation in for example the powertrain and drivetrain. It is important to note however that these buses are not strictly isolated, there are always some modules that reside on multiple buses.

Next to this, modules that can serve as an entry point for an attacker, such as the telematics module or the TPMS, are increasingly moved away from safety-critical functions. However, until now, they are sometimes still connected to safety-critical functions in some way.

Finally, where many manufacturers used to have one module responsible for all telematic communication, such as the radio, Bluetooth, calling and the internet connection, these are becoming more separated from each other, and added in modules where only that specific communication channel is necessary. This is however not the case for all manufacturers.

7.5 SECURITY

Until recently, security of the IT architecture has hardly been of any concern. As described in chapter 2, the vehicle has always been considered to be an isolated network, where an attacker would need physical access to the vehicle to for example alter the engine. However, in recent years, the European Union has subsidised projects to define security requirements, create a base of secure communication in the vehicle, and create secure communication between vehicles. Some projects really heavily on the result of the project EVITA. Therefore, the result of this project is briefly provided.

7.5.1 *E-safety Vehicle Intrusion Protected Applications (EVITA)*

EVITA is a European research project started to improve on-board network protection, consisting of a mix of universities, car manufacturers and equipment suppliers [2]. Amongst other things, they have created a basis of trustworthy communication between vehicles by designing a secure on-board network [22]. They argue that to be able to trust information, the communication within, but also between, vehicles should be secure. To this end, they have developed Hardware Security Modules (HSMs) that should be integrated in every ECU to secure communication. For economic reasons, they propose three different HSMs:

FULL HSM
: Designed for protecting the in-vehicle domain against vulnerabilities due to V2X communication. This offers the maximum

level of functionality, and is used on places where external information enters the in-vehicle network.

MEDIUM HSM
: Designed for securing on-board communication. It offers less performance than the Full HSM, contains no asymmetric cryptographic engine, and is used for communication between ECUs.

LIGHT HSM
: Designed for securing interactions between ECUs and sensors and actuators. Because it needs to be cheap, but fast, it only contains a symmetric cryptographic engine, and some optional other features.

This is demonstrated in Table 7.2. An example of how these different HSMs can be used is given in Figure 7.7.

	FULL HSM	MEDIUM HSM	LIGHT HSM
RAM (random-access memory)	x	x	optional
NVM (non-volatile memory)	x	x	optional
Symmetric cryptographic engine	x	x	x
Asymmetric cryptographic engine	x		
Hash engine	x		
Counters	x	x	optional
Random-number generator	x	x	optional
Secure CPU	x	x	
I/O component	x	x	x

Table 7.2: Components of automotive HSM classes, adapted from the EVITA project

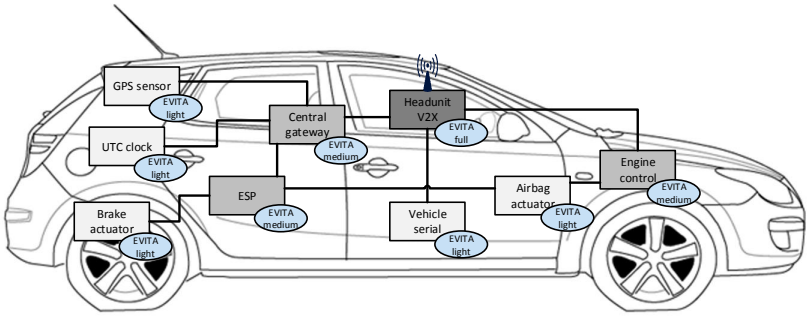

Figure 7.7: Instance of a secure on-board network with full, medium, and light HSMs attached to ECUs, sensors, and actuators, adapted from the EVITA project

Other researchers have devised similar approaches as EVITA. Moalla et al. also argue for the need of a trusted module for secure communication, however, they do not provide any particular specifications for the module itself [49]. Papadimitratos et al., as part of the SEVECOM project, also focus on more secure communication between vehicles by adding a trusted module to the On Board Unit (OBU, the module used for communication between vehicles) [55]. Although giving a different specification, the Trusted Computing Group [65] and the Herstellerinitiative Software [21] have devised their own versions of a HSM for use in automotive networks.

7.5.2 *Chapter Summary*

This chapter set out to answer research question two:

> What will a general IT architecture in a vehicle look like?

and as part of it, research question two a:

> What are the attack surfaces within this architecture?

To this end, we looked at what the IT architecture looks like in a vehicle by looking at the relevant literature and an expert interview. We have identified a trend towards more digitalised vehicles where more and more functions are taken over by IT. This, together with the fact that vehicles are becoming more and more connected to the internet makes vehicles very susceptible to attackers.

In this chapter, we first provided an overview of the different modules and IT components that are present in a vehicle. These so called ECUs are used for a wide variety of functions, ranging from small mirror-adjustments to more complicated systems such as ABS. We then provided an overview of the different communication buses that are used in vehicles, that showed that no currently existing bus provides any basic security mechanisms.

In more detail, we have looked at what a general IT architecture within a vehicle looks like. By analysing 24 different car models, we have created a generic IT architecture per region: Europe, the United States of America, and Asia. During this analysis, we in particular focused on identifying possible entry points for attackers and the locations of ECUs with a cyber-physical aspect.

We found that in Europe, IT architectures are more alike than in other parts of the world, possibly because they collaborate more in shared projects such as EVITA, SEVECOM, and sim^{TD}. However, IT architectures still differ a lot from each other, not only between different manufacturers, but within the same manufacturer as well. It also seems that no real thought has been given to how these networks are best secured. Although some models do separate critical functionality from non-critical functionality, it is more likely that this is due to costs, than because of security reasons.

Finally, we have looked at how security is integrated in the IT architectures, and in particular the EVITA project, that argues that for secure communication between vehicles, secure communication within vehicles is also required. To this end, they have designed HSMs that should be integrated in every ECU that makes communication between ECUs more secure.

Important to note however, is that it remains hard to predict what the architecture within a vehicle might look like in a few years' time. Projects such as EVITA (2012) and PRESERVE (2015) that address the security layout of vehicular architectures have only recently finished. These projects do however, focus on securing the existing architecture, and do not specify any specific architectures that should be used.

8

THREAT MODELLING

Knowledge Question 3:

What is a threat model for selected functions within vehicles in 5 to 10 years' time for a general IT architecture?

To answer this knowledge question, we first take a look at the relevant literature of threat modelling and its place within cyber risk management techniques for the automotive industry. The next chapter will thereafter describe a composite threat model based on the found literature.

The goal of this chapter is to provide an overview of existing automotive risk management techniques in the automotive industry, in particular cyber risk management and threat models. To this end, this chapter will first give a brief background of two commonly used techniques used in threat modelling in other areas than the automotive industry in section 8.1.1. It will then look at ISO 26262, a widely used framework for automotive safety risk management, in section 8.1.2. Next, two automotive cyber risk management techniques are discussed in section 8.1.3, followed by an overview of existing automotive threat models in section 8.1.4.

8.1 AUTOMOTIVE RISK MANAGEMENT TECHNIQUES

The literature research and the expert interview have revealed that until recently, not many automotive manufacturers use any cyber risk management techniques. Although many manufacturers have become aware that cyber risk management is becoming more important, they are only just starting to have a look at this.

Vehicle manufacturers do however, have a rich history in safety risk management. Before taking a further look into cyber risk management techniques for the automotive industry, this section will first provide some background information into Data Flow Diagrams and the STRIDE threat modelling technique. These techniques are not particular for the automotive industry, but will be referred to in upcoming sections. Next, section 8.1.2 will provide an overview of the already mentioned ISO 26262 standard. Section 8.1.3 will provide two

high level automotive cyber risk management frameworks. Finally, section 8.1.4 will provide some methods related to threat modelling, including some remarks on positive and negative aspects of these models.

8.1.1 *Background*

Techniques such as Data Flow Diagrams and STRIDE have been used as tools in cyber risk management techniques in other areas than the automotive industry for some time. We will briefly explain them to keep in mind when reading the upcoming sections.

Data Flow Diagrams

Data Flow Diagrams (DFDs) find their origin back in the 1970s as a way to visualise steps and data flows in software processes. It provides a tool to visually represent which entities are present in the system, and in which way they interact with other entities. An example of a DFD can be found in Figure 8.1. Such a diagram can consists of five different entity types:

PROCESSES (P) are tasks that operate on incoming data and potentially produce output

DATA FLOWS (F) are flows of information

DATA STORES (S) are physical or logical storage devices

EXTERNAL ENTITIES (E) are entities outside the target system upon which the system depends

COMMUNICATION ZONES (Z) are zones where entities can communicate with each other

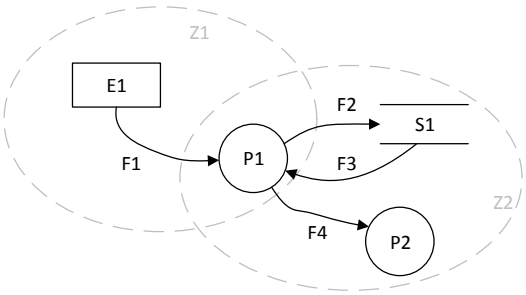

Figure 8.1: A simple data flow diagram, adopted from [58]

STRIDE

The STRIDE threat model was designed by Microsoft and used as part of their Security Development Lifecycle (SDL) to classify and identify

potential threats [60]. It is an acronym for the following six threat categories:

1. Spoofing identity
2. Tampering with data
3. Repudiation
4. Information disclosure
5. Denial of service
6. Elevation of privilege

The idea behind the STRIDE methodology is to provide a security expert or non-expert with the tools to think about security threats. Sometimes, the STRIDE methodology is referred to as the whole security development lifecycle, ranging from creating diagrams such as Data Flow Diagrams, to mitigating techniques. However, STRIDE is originally only part of the SDL process where threats have to be enumerated and helps at finding the correct threats for a particular element in a Data Flow Diagram. To save time, not all STRIDE classes need to be checked for all DFD elements. A Data flow for example, cannot be spoofed, only the originating process can. A mapping of which STRIDE classes should be checked again which DFD elements can be found in Table 8.1 [31].

ELEMENT	S	T	R	I	D	E
External Entity	x		x			
Process	x	x	x	x	x	x
Data Store		x	x[1]	x	x	
Data Flow		x		x	x	

Table 8.1: STRIDE categories mapped on Data Flow Diagram elements.
[1]For a Data Store, Repudiation only needs to be considered if it serves as a log

8.1.2 *Automotive Safety - ISO 26262*

In the early days of the vehicle, when all parts were purely mechanical, risk management in the automotive industry was focused on ensuring the safety of the vehicle if certain parts of the vehicle would malfunction. Such techniques would look at potential failures, such as a malfunctioning brake, and try to find possible causes of that malfunction: such as a broken brake line or a worn brake disk. Using historical data, experts could calculate the probability of certain malfunctions and use that information to mitigate high risks.

After a while however, with the introduction of the ECU and the integration of more IT into the vehicle, malfunctions were no longer

subjected to wear or a broken line, but also to programming errors and bugs. The fact that electronic components could influence the physical world started a movement by vehicle manufacturers and the government to define a standard on which vehicle manufacturers and OEMs should handle risks. To this end, in 2011 ISO 26262 [28] was adopted from the already existing IEC 61508 [25], which was designed for cyber physical systems, and altered to be more suitable for road vehicles.

ISO 26262 defines the functional safety of electrical and electronic systems in road vehicles and is considered to be the standard for automotive functional safety. Just as IEC 61508, it is a risk-based safety standard where the risk of hazardous operational situations is qualitatively assessed.

It provides an automotive safety lifecycle that ranges from development to decommissioning, where in particular, the functional safety aspects within that lifecycle are addressed. In short, the following steps are taken: during the concept phase, hazard analysis techniques are used to assess the safety risk of the system. Using these hazard scenarios, target Automotive Safety Integrity Levels (ASILs) are derived and used to drive the development of a system safety concept to include risk mitigation. The concept of ASILs will be explained below, suffice it to say that it determines the acceptable level of risk of a certain component of the vehicle. The ISO 26262 process is illustrated in Figure 8.2. As can be seen, the white process denotes the normal design process of a system, where the blue activities denote the associated safety activities.

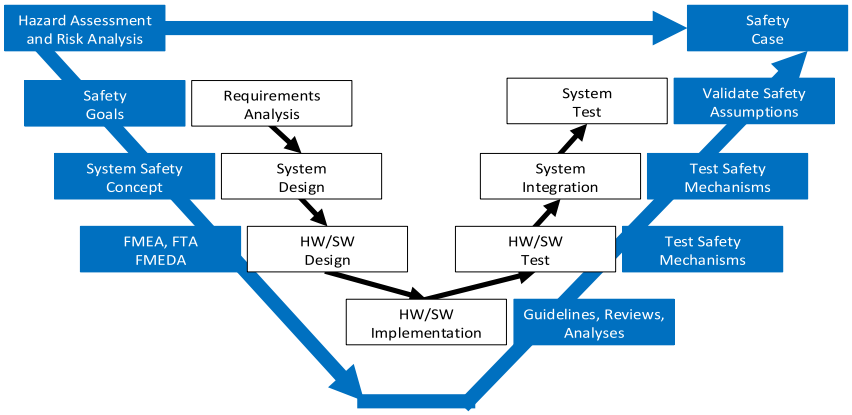

Figure 8.2: Overview of the functional safety development process in ISO 26262, adapted from [8]

Hazard Analysis and Risk Assessment

The first step in the ISO 26262 development process is the identification of hazards: The Hazard Analysis and Risk Assessment (HARA). HARA takes place in two step: analyse which hazards can occur, and to assess the risks of these hazards. Although ISO 26262 does not state any specific ways of performing HARA, in general, techniques that are used analyse threats by performing a brainstorm of possible hazards, and use techniques such as FMEA and FTA[1] to estimate the risk of a hazard by looking at three areas:

SEVERITY The severity of the hazard, for example whether it is life-threatening

EXPOSURE The exposure of the hazard or how likely it is that the hazard occurs

CONTROLLABILITY The controllability of the hazard, whether the driver is capable of preventing the hazard by for example braking

Together, these areas determine an ASIL for that specific hazard.

Before continuing with existing safety and security methods, I take a brief moment to highlight Automotive Safety Integrity Levels (ASILs) and Evaluation Assurance Levels (EALs) that are often used or referred to in other methods.

SIL/ASIL

The basis of ASILs dates back to IEC 61508, the previously mentioned ISO that concerns safety in all kinds of systems. In this IEC, Safety Integrity Levels (SILs) are defined, which provides a normative method for evaluating the safety of programmable electronic systems. It defines the chance that a certain component might fail per use (in the case it is used less than once a year) or per hour of use (if it is used continuously). Table 8.2 shows the different probabilities per SIL. This for example means that if a component receives a SIL4, on average, it brakes down less than once in every 10.000 times.

With the introduction of ISO 26262 for the automotive industry, these SILs were used for the definition of ASILs (Automotive Safety Integrity Levels). However, in this standard, no exact probabilities are given, but more general terms as *reasonable* or *life-threatening*. The ASILs range from ASIL A to D, where ASIL D means a reasonable possibility of causing a life-threatening or fatal injury, and ASIL A either means a low probability, or a not life-threatening event. Beneath ASIL A, there is also QM, which mean there is no safety relevance

1 Failure Mode and Effect Analysis (FMEA) and Fault Tree Analysis (FTA) are commonly used techniques in cyber physical systems to reason about the consequences and risks of failures or faults in a system

SIL	PROBABILITY OF FAILURE ON DEMAND	PROBABILITY OF FAILURE PER HOUR
1	$\geqslant 10^{-2}$ to 10^{-1}	$\geqslant 10^{-6}$ to 10^{-5}
2	$\geqslant 10^{-3}$ to 10^{-2}	$\geqslant 10^{-7}$ to 10^{-6}
3	$\geqslant 10^{-4}$ to 10^{-3}	$\geqslant 10^{-8}$ to 10^{-7}
4	$\leqslant 10^{-4}$	$\leqslant 10^{-8}$

Table 8.2: Failure rate for Safety Integrity Levels

to be considered and standard Quality Management processes are sufficient. For the exact calculation of the ASIL, I refer to ISO 26262 [28].

EAL

Evaluation Assurance Levels (EALs) find their basis in IEC/ISO 15408 [26], also called the Common Criteria, a widely used standard for computer security certification that is used in all kinds of computer areas. The Common Criteria is a framework where functional and assurance requirements can be specified for the desired security level. This means vendors can make claims on security attributes of their product, and testing laboratories can evaluate these claims and products. These claims are in the form of EALs, where each higher level provides additional assurance. Important to note is that where ASILs say something about the safety of a product, EALs say something about at what level a product has been tested. Each level in the EAL stands for the following:

EAL1 Functionally tested
EAL2 Structurally tested
EAL3 Methodically tested and checked
EAL4 Methodically designed, tested and reviewed
EAL5 Semiformally designed and tested
EAL6 Semiformally verified design and tested
EAL7 Formally verified design and tested

8.1.3 *Cyber Risk Management Frameworks*

Having looked at how automotive safety is ensured, we now take a look at risk management frameworks that look at the security of vehicles. Since this is a relatively new field, not many complete cyber risk management frameworks have been proposed. However, since quite some parallels are present with cyber risk management in other areas, some frameworks, or slight alterations of frameworks, have been proposed. This section will highlight two frameworks: an extension of ISO 26262, and a version of the NIST cyber risk management

framework from the National Highway Traffic Safety Administration (NHTSA).

8.1.3.1 *ISO 26262 Extended*

The first to extend ISO 26262 to include security aspects alongside the safety design process are Burton et al. [8]. As can be seen in Figure 8.3, the normal development process, along with the already defined safety activities, are extended with security activities. The idea behind this framework is that analog to the normal Hazard Analysis and Risk Assessment (HARA), a threat analysis is performed in the begin stage, which results in security goals. In the next steps, these security and safety goals should be used to make the system design.

Burton et al. however, do not provide any details on how such a Threat Analysis should be performed, but only mention the use of an attacker model in terms of motivation, skill level and available equipment, and the use of misuse case [8].

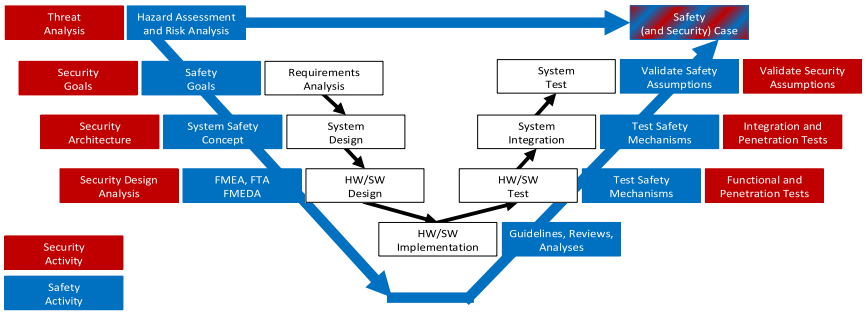

Figure 8.3: ISO 26262 safety process extended with security activities, adapted from [8]

8.1.3.2 *NIST/NHTSA Risk Management Framework*

Well known in the business world for cyber risk management is the National Institute of Standards and Technology's (NIST) Risk Management Framework [29]. The framework provides a disciplined and structured way to integrate information security and risk management activities into the development lifecycle. The framework includes the six following steps and is graphically represented in Figure 8.4:

CATEGORISE the information systems and define criticality/sensitivity of the information system based on an impact analysis

SELECT a set of baseline security controls and tailor or supplement controls as needed based on a risk assessment

IMPLEMENT and describe the security controls using sound engineering practises and apply security configuration settings

ASSESS the security controls to determine the control effectiveness. i.e. are the controls implemented correctly, are they operating as intended, and do they meet the security requirements of the information system

AUTHORISE information systems operations based on a determination of risks to organisational operations and assets, individuals, other organisations, and the Nation, if the risk is deemed acceptable

MONITOR the security controls continuously and reassess the control effectiveness, including reporting the security state to designated officials

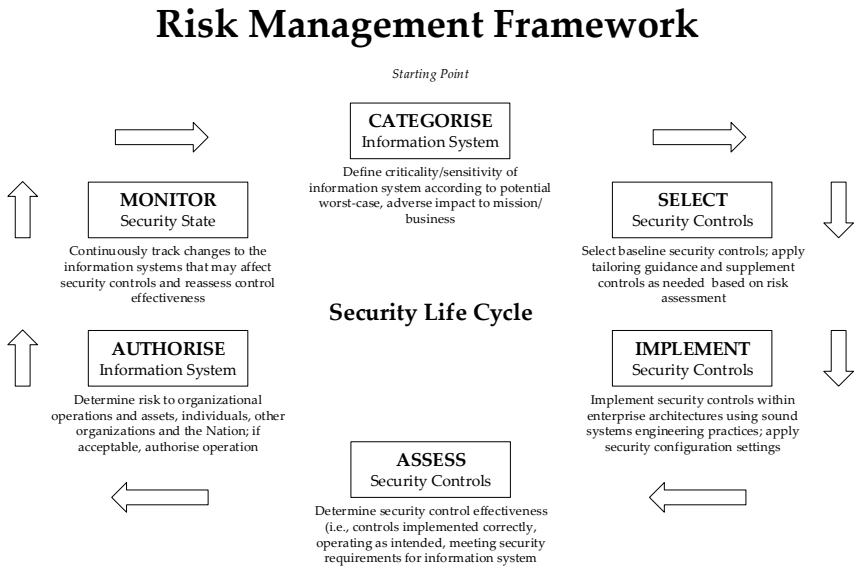

Figure 8.4: NIST Risk Management Framework, adapted from [29]

To be more suitable for the automotive industry, the NHTSA has released a modified version of this framework [46]. It is based on the existing frameworks NIST SP 800-37, NIST SP 80-39, and NIST SP 800-30 and can be found in Figure 8.5. Some important differences with the NIST Risk Management Framework is the removal of the step Authorise, since it only applies to Federal IT systems, and not vehicle control systems.

Next to this, they argue there is need for a step before categorisation that assesses the systems by means of a threat model and use cases. This first step should focus on hostile cyber or physical attacks, human errors, and natural and man-made disasters. For the hostile attacks, organisation should provide a detailed characterisation of tactics, techniques and procedures employed. With this in mind, experts can identify a set of representative threat (mis)use cases that can be graded on four possible areas: Privacy, Financial, Operational and

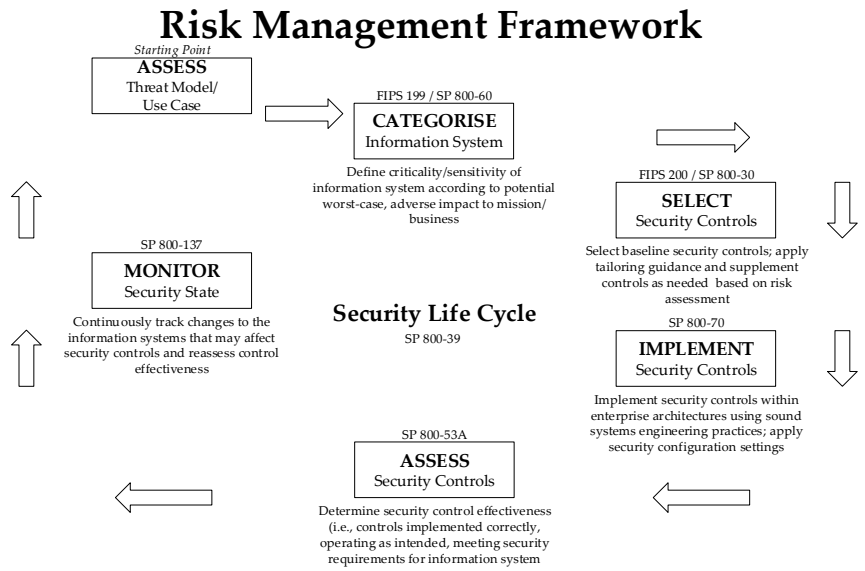

Figure 8.5: Modified NIST Risk Management Framework for the vehicle sector, adapted from [46]

Safety. The NHTSA however, does not provide any specific technologies to perform this step.

8.1.4 *Threat Modelling*

As explained in the previous sections, there exist some frameworks that focus on dealing with cyber security risks in the automotive sector. However, these frameworks are still quite high level and do not include any specifics on how certain steps in the framework should be performed. This section looks at the step of creating a threat model as described by both the NHTSA Risk Management Framework for the Vehicle sector and the Extended version of ISO 26262. It is important to note that the principle of threat modelling is not new, other areas such as cyber physical systems have a longer history of using threat models during the risk assessment. This section however, will focus on academic proposals for threat models in the automotive industry that often find their basis in existing threat models.

8.1.4.1 *TARA*

As described above, ISO 26262 contains a Hazard Analysis and Risk Assessment process that is well-established. This process however, is focused on hazards, and not completely suitable for threat analysis. Ward, Ibara, and Ruddle therefore argue that these methods need extension and adaptation to also be suitable for the cyber security do-

main, and propose the Threat Analysis and Risk Assessment (TARA) approach [69].

To this end, they propose extending the already existing severity class with two more levels, because they argue that a cyber security incident can influence multiple vehicles at the same time. Next to this, they argue that security covers more areas than safety alone, and therefore propose that the level of severity should also be measured on non-safety related areas. This results in four areas of severity: Privacy, Financial, Operational and Safety

They also propose to incorporate the concept of attack potential into the evaluation technique. The attack potential is a common concept in ISO 18045, a methodology for evaluating IT security that is an addition to ISO 15408 [26], or more commonly known as the Common Criteria. The attack potential is measured on five aspects:

ELAPSED TIME The time needed to perform the attack

SPECIALIST EXPERTISE The needed knowledge to perform the attack

SYSTEM KNOWLEDGE The needed knowledge of the system to perform the attack

WINDOW OF OPPORTUNITY The amount of access time to the system that is needed to perform the attack

EQUIPMENT The needed hardware or software that is required to perform the attack

Finally, the Controllability, Severity and Attack probability are combined as illustrated in Table 8.3 into a Risk Level that can be used to prioritise risks.

CONTROL-LABILITY	SEVERITY	ATTACK PROBABILITY				
		A1	A2	A3	A4	A5
C1	S_s1	R0	R1	R2	R3	R4
	S_s1	R1	R2	R3	R4	R5
	S_s2	R2	R3	R4	R5	R6
	S_s3	R3	R4	R5	R6	R7

Table 8.3: Risk graph fragment for safety-related security threats, adopted from [69]

8.1.4.2 *SAHARA*

Where TARA focuses on creating an analog process to the ISO 26262 HARA method, SAHARA focuses on expanding the HARA method [45]. The Security-Aware Hazard and Risk Analysis (SAHARA), uses

RESOURCES (R)	KNOW-HOW (K)	THREAT LEVEL (T) 0	1	2	3
0	0	0	3	4	4
	1	0	2	3	4
	2	0	1	2	3
1	0	0	2	3	4
	1	0	1	2	3
	2	0	0	1	2
2	0	0	1	2	3
	1	0	0	1	2
	2	0	0	0	1
3	0	0	0	1	2
	1	0	0	0	1
	2	0	0	0	1

Table 8.4: SecL Determination Matrix - ascertains the security level from R, K, and T values, adapted from [45]

the STRIDE threat modelling approach to quantify threats of the System under Development. The found threats are than given an ASIL on three different areas:

R - RESOURCES The required resources an attacker might need to pose the threat

K - KNOW-HOW The required knowledge an attacker might need to pose the threat

T - (THREAT) CRITICALITY The criticality of the threat[2]

Macher et al. do provide some guidance and examples on how these ASILs should be determined for each areas. For the exact determination of each level, I refer to their original paper [45].

The three ASILs combined determine the Security Level (SecL) of a a certain threat that can be used to prioritise the found threats and reason on where risks need to be mitigated. This determination matrix can be found in Figure 8.4.

8.1.4.3 *CHASSIS*

The Combined Harm Assessment of Safety for Information Systems (CHASSIS), is an approach for requirements engineering that has been developed for cyber physical systems. Schmittner et al. however, argue it can also be applied to the automotive industry [59].

2 The reason C is not used for Criticality is because in ISO 26262, C is used to denote Controllability

The approach is as follows: first, functional requirements are defined as a basis for elicitation of safety and security requirements. These requirements are supported by Use Case Diagrams and Sequence Diagrams that describe users and functions. Secondly, a brainstorming session is performed with safety and security experts combined to identify possible misuses and misusers of the system with help from the before created diagrams. These misuses are written down, as well as translated into Sequence Diagrams so that potential hazards, failures, threats, vulnerabilities and mitigation measures can be identified.

8.1.4.4 *SINA*

The Security in Networked Automotive (SINA) was developed by Schmidt et al. to make a security equivalent to the ISO 26262 HARA approach [58].

Their approach is as follows: at first, data flows of the system are modelled in Data Flow Diagrams (see section 8.1.1). Based on these DFDs, possible threat are identified and classified, similar to the STRIDE threat methodology. Schmidt et al. however, have created other threat classes. For Data Flows, one should use: Creation, Modification, Eavesdropping, and Blocking. For Processes, Data Stores and External Entities, one should use: Tampering, Denial of Service, and Information Disclosure. Although Schmidt et al. use different names, their classes map to the STRIDE classes as given in Table 8.5[3].

SINA CATEGORY	STRIDE CATEGORY
Creation	Spoofing Identity
Modification	Tampering with Data
Eavesdropping	Information Disclosure
Blocking	Denial of Service
Tampering	Tampering with Data
Denial of Service	Denial of Service
Information Disclosure	Information Disclosure

Table 8.5: SINA categories mapped to STRIDE categories

By checking every entity in the DFD against the relevant threat classes, every threat should be found. Next, these threats are evaluated and given a target ASIL, denoting how relevant they are. Finally, attack

3 Note that SINA does not offer a category that maps to repudiation. Although they do not give a specific reason, we assume this is because repudiation is not yet something that needs to be considered in an automotive setting. However, we believe that in the future, this will definitely become important in for example logging mechanisms for accidents and insurance.

trees are built for the threats with the highest ASILs. Using these attack trees, an expert can reason on where extra measures should be taken to mitigate risks.

8.1.4.5 *NHTSA Composite Modelling Approach*

To cope with the increasing threat of cyber security in the automotive industry, the National Highway Traffic Safety Administration started a cyber security research program. One of their goals is to help the automotive industry by creating a knowledge base and implementing an industry-based best practises for cyber security.

In one of their publications, they investigate STRIDE, Trike, and ASF to create their own composite threat model [47]. Their model consists of two phases and a few steps:

1. Identify critical applications/systems
    a) Application/system decomposition
        i. Create interconnection drawings
        ii. Create high level data flow diagrams
2. Determination and analysis of threats
    a) Threat identification
    b) Threat analysis
        i. Drawing review
        ii. Use case development
        iii. Vehicle threat matrix development and population

Identify Critical Applications/Systems

In the first phase, the model focuses on creating an understanding of the system. At first, critical applications or subsystems are identified (e.g. the brake or powertrain). These are components, that if compromised maliciously, could results in serious safety concerns. Next, for each of these systems, a complete interconnection drawing is created containing the following elements:

1. All relevant on-board components and systems
2. External interface connections
3. Data entry/exit points
4. Data types

Finally, a High level data flow drawing should be created that visually represent the data flows within this system. These two drawings will serve as a basis for phase two.

Determination and Analysis of Threats

In the seconds phase, the model focuses on identifying possible threats and analysing them. First, the threats need to be identified, however, the NHTSA does not propose any concrete way of doing this.

Secondly, the threats are analysed by reviewing the drawings. Typical questions in this analysis include: "What data paths are critical for the system?" and "What physical and wireless entry points can connect the vehicle to an external source?"

This analysis is used to create use cases, together with information about how an attacker can exploit a certain threat by looking at the following areas:

1. Entry Point
2. Access Method
3. Types of Attack
4. Outcome of the attack

Finally, these use cases are aggregated and used to populate a vehicle threat matrix which given time, will contain an extensive list possible (mis)use cases.

8.2 CHAPTER SUMMARY

This chapter set out to answer research question three:

> What is a threat model for selected functions within vehicles in 5 to 10 years' time for a general IT architecture?

To this end, we looked at the relevant literature and an expert interview. We have found that automotive manufacturers have had a rich history in making their products safe by using standardised techniques such as ISO 26262. However, these techniques have not been designed to incorporate security related safety in the design process. To be able to cope with these kind of threats, some frameworks such as an extension of ISO 26262 or the NHTSA's modified version of the NIST Risk Management Framework have been designed that do incorporate threat models and security objectives in the design process. These frameworks however, do not include any specifics on how such threat analyses should be performed.

Some literature can be found that does focus on building such threat models. Many of these models focus on getting a complete overview of the system under development. This is an important step, since it helps in creating an understanding of the possible consequences of certain threats, however small they may seem at first.

Next to this, the introduction of other classes of severity than safety illustrates what different kind of impact IT systems may have in a vehicle. Possible threats no longer only concern safety, but privacy,

operational, or financial consequences as well. As is the same in ISO 26262, it is also important to keep classifying threats on the ability to control them. Some threats, for example one that focuses on changing the set value of the cruise control, can be overridden by hitting the brakes. Although the consequences of setting the cruise control value to 300 km/h could be catastrophic, it is one that is easily controlled.

Microsoft's STRIDE threat modelling technique is already well known in the security industry and widely used to identify threats. The introduction of STRIDE in the automotive industry to secure vehicles therefore isn't strange. The threat modelling technique helps in reasoning about all categories of possible threats for all aspects of a system, making it a very exahustive way of identifying threats. However, as with all threat modelling techniques, the more time is used to identify possible threats, the more you will find.

It is however, also important to note that many of these threat modelling techniques for the automotive industry are very attacker centric. After focusing on potential threats, many models focus on how these threats can be exploited. The biggest concern with this approach is that the system is considered to be secure from the start and cannot be altered other than via entry points in the system. Modern day vehicles however, consist of many different ECUs from many different vendors, some of which can easily be flashed before being put in the vehicle. Already in the manufacturing phase, vehicles can be compromised without the manufacturer knowing.

Even after manufacturing, many vehicles visit garages and become connected to after market equipment that cannot easily be controlled. Making a claim that the vehicle is 100% secure to begin with, is therefore hard to make. Attacks such as Stuxnet have even shown that even if the system is air-gapped from any other system, it can still be compromised [39].

Lastly, some techniques require the making of Data Flow Diagrams, which can be a cumbersome process, and contain many elements that vehicles often do not have. Systems within vehicles are more about exchanging information, and less about storing confidential information, making them different from common IT systems.

9

PROPOSED THREAT MODEL

Knowledge Question 3:

What is a threat model for selected functions within vehicles in 5 to 10 years' time for a general IT architecture?

Having looked at the state of the art of threat modelling techniques, this chapter will propose a composite threat model that focuses on getting a complete image of the system under consideration, apply this model to selected future functionality use cases, and validate the this model by interviewing a subject matter expert.

To this end, section 9.1 will first provide the first version of our composite threat model as used in the validation phase. Then, section 9.2 will apply this model to some selected use cases to provide a better understanding of how the model works, and to identify possible threats in future functionality. Section 9.3 will then validate the composite threat model by means of an expert interview. Finally, section 9.4 will provide the improved version of step 2b of the composite threat model. The final version of the composite threat model can also be found in Appendix E.

9.1 COMPOSITE THREAT MODEL

As described in chapter 8.2, some threat modelling techniques already exist specifically for the automotive industry. However, these models are focused on what attack paths an attacker might use, and do not consider the system to have been breached already. We therefore propose a novel composite threat model that focuses on identifying all possible threats under the assumption that the system is already breached.

The composite threat model consists of three steps. In step 0, all critical applications and systems should be identified.

Then for each identified critical application and system, step 1; the applications and systems are decomposed to get a complete overview and understanding of the system, and step 2; threats are identified and analysed to determine their consequences. These steps are as follows:

0. Identification of critical applications/systems
   For all identified applications and systems:

1. Decomposition of the application/system
   a) Create interconnections drawing of the vehicle
   b) Create high level flows in interconnections drawing

2. Identification and analysis of threats
   a) Threat identification using STRIDE
   b) Determination of severity of threats

Please note that these steps are part of the modified NIST and NHTSA Risk Management Framework as given in chapter 8.1.3. These steps are visually described in Figure 9.1 and will be further explained in the coming sections.

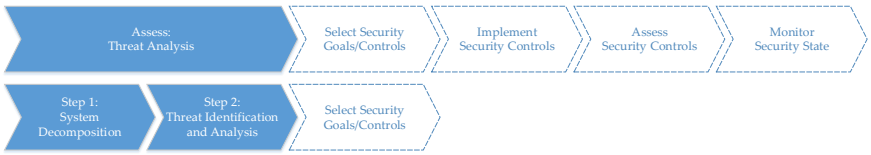

Figure 9.1: Composite Threat Model Steps as part of the NIST framework

Step 0: Identification of Critical Applications/Systems

For a complete overview of the vehicle, all critical applications or systems need to be identified and further investigated. However, applications that are deemed critical are more likely to result in serious threats, and could be investigated first when short on time. This step is therefore referred to as step 0, and could be skipped if all applications or systems are analysed anyway.

A critical application or system is in this sense, a function that if compromised maliciously, could result in serious consequences, either safety related, or in other ways.

Step 1: Decomposition of the Application/System

The goal of the first step is to get a complete overview of the application/system and its components. To this end, the decomposition of the application/system takes place in two steps: first, a complete interconnection drawing is created containing the following components:

- All (relevant) components, subsystems and buses: Identify all components that are connected to the system under consideration
- External connections: identify any external connections the system such as via WiFi, OBD-II, Bluetooth or 3G
- Data types: identify safety levels for data in components and on buses. Especially note systems and buses where a safety level tradeoff is present such as gateways

Note that for one vehicle model, such a decomposition can be made once and used for all relevant applications and sub systems. One may ignore irrelevant modules if not applicable for a certain application or system.

In the second step, high level data flows should be identified and added to the interconnection drawing, considering what data flows from which component to what other component.

Step 2: Identification and Analysis of Threats

Using the drawings from the previous phase, threats are identified and further analysed. This is again done in two steps.

a) Threat identification using STRIDE

During the first step, all threats are identified using the STRIDE threat modelling technique. This means that for each ECU, data flow and external entity from the interconnections drawing, the relevant STRIDE classes are used to identify possible threats, resulting in a list of threats for all components of the system where each STRIDE class is applied to each element as according to Table 9.1. This table is a modified version of Table 8.1, where it has been made suitable for our interconnections drawing. This means that the data stores have been removed, and processes have been replaced by ECUs, that have similar functions.

ELEMENT	S	T	R	I	D	E
ECU	x	x	x	x	x	x
Data Flow		x		x	x	
External Entity	x		x			

Table 9.1: STRIDE categories mapped on Interconnection drawing elements

b) Determination of severity of threats

In the second step[1], each entry in the list from the previous step is evaluated on two subjects: Severity and Controllability. For Severity, a distinction is made on four different areas: Safety, Operational, Privacy, and Financial.

To determine these classes, one must consider what would happen if the threat occurs. For example, in the case of a data flow and Denial of Service one needs to consider what would happen if the data flow is denied, and doesn't arrive at the next ECU. Would the vehicle for example crash? Will the driver notice? Does any sensitive data leak? Can he control the consequences somewhat? etc. Normally, such a determination should be done by a small group of experts, to get some discussions going on the consequences and improve the analysis, but could be done by a single person.

To determine the severity of these classes, one must use the classification from Ward, Ibara, and Ruddle [69] as a guideline as can be found in Table 9.3. As can be seen, Ward, Ibara, and Ruddle provide a clear distinction between consequences for one or multiple vehicles. For example, if certain messages that are shared between vehicles can be tampered with, the consequences could be greater than if only one message within a vehicle is tampered with.

Similar to the Severity classes, ISO 26262 provides a description for calculating the Controllability level [28], which is supplemented with another two classes by the MISRA standard [15]. This creates levels as described in Table 9.2.

CONTROLLABILITY	
LEVEL	DESCRIPTION
0	Controllable in general
1	Simply controllable
2	Normally controllable (most drivers could act to prevent injury)
3	Difficult to control or uncontrollable
4	Genuinely uncontrollable

Table 9.2: Controllability Level determination

1 Please note that in the validation, this step has changed somewhat, and an improved version can be found in section 9.4

SAFETY		OPERATIONAL	
LEVEL	DESCRIPTION	LEVEL	DESCRIPTION
0	No injuries	**0**	No impact on operational performance
1	Light or moderate injuries	**1**	Impact not discernible to driver
2	Severe and life-threatening injuries (survival probable) or light or moderate injuries for multiple vehicles	**2**	Driver aware of performance degradation or Indiscernible impacts for multiple vehicles
3	Life-threatening injuries (survival uncertain) or fatal injuries or Severe injuries for multiple vehicles	**3**	Significant impact on performance or Noticeable impact for multiple vehicles
4	Life-threatening or fatal injuries for multiple vehicles	**4**	Significant impact for multiple vehicles
PRIVACY		**FINANCIAL**	
LEVEL	DESCRIPTION	LEVEL	DESCRIPTION
0	No unauthorised access to data	**0**	No financial loss
1	Anonymous data only (neither specific driver not vehicle data)	**1**	Low-level loss ($10)
2	Identification of vehicle or driver or anonymous data for multiple vehicles	**2**	Moderate loss ($100) or low losses for multiple vehicles
3	Driver or vehicle tracking or identification of driver or vehicle for multiple vehicles	**3**	Heavy loss ($1000) or moderate losses for multiple vehicles
4	Driver or vehicle tracking for multiple vehicles	**4**	Heavy losses for multiple vehicles

Table 9.3: Severity Level determination for Safety, Operational, Privacy and Financial

Using the Severity and Controllability classes, one can determine or calculate a threat level such as via:

$$\mathrm{Threat} = w_s S + w_o O + w_p P + w_f F + w_c C$$

where classes are weighted according to how important a company might find the class. If it for example considers safety ten times more important than operational, it can give the safety class a weight of ten.

Important to note is that the exact calculation or determination of the threat level is highly dependent on wishes of the company. One can use the above calculation, or use a determination matrix as given in Table 8.3 or 8.4.

Example

Table 9.4 provides an example of the determination of Controllability and Severity levels for one data flow for a throttle message. Since it is a data flow, we only need to consider the Tampering with Data, Information Disclosure, and Denial of Service classes from STRIDE. For these relevant STRIDE classes for this data flow, the possible consequences are determined. For example, tampering with the throttle message could result in a serious safety concern where survival is questionable. However, most drivers could prevent a crash by hitting the brakes and killing the engine, resulting in the classes as given in Table 9.4.

ELEMENT	STRIDE CLASS	SEVERITY S	O	P	F	CONTROL-LABILITY
Throttle message						
	T	3	0	0	0	2
	I	0	0	3	0	0
	D	2	0	0	0	2

Table 9.4: Example of result of final step

9.2 USE CASES

To get a better feeling on how the composite threat model works, we provide two use cases of possible future systems, as described in chapter 6: Predictive Cruise Control and Emergency Brake Light. We chose these systems because together they cover many different aspects of vehicular systems. The Predictive Cruise Control is a complex system over multiple buses that combines a cyber physical aspect with information from the internet, and the Emergency Brake Light is a Vehicle-2-vehicle use cases where information is shared between vehicles and acted upon.

The first use case will be explain the steps in a bit more detail to better understand its working. The last use case will show the system applied to both a modern day vehicle, and a vehicle as if it were secured such as suggested by the EVITA project. Please note that since we have already identified these two use cases to use, step 0 of the model will be skipped.

9.2.1 *Predictive Cruise Control*

The first use case is Predictive Cruise Control as already explained in chapter 6.2. The idea behind this system is to combine map data and knowledge about what the road looks like ahead to adjust the speed of the vehicle to safe fuel and for example build up momentum for climbing up a hill.

In the design of this system, some assumptions have been made as to the exact working of the system, since no public data is available. However, a major part is taken from the result of a brainstorm session from the Software System Safety Working group organised by MIT that includes specialists from Ford. Although this does not give a detailed explanation of how Predictive Cruise Control works, it does provide a good basis [20].

The system works as follows: the navigation controller sends information about the upcoming few kilometres to the cruise control module, which uses this information, together with information from the Radar or LIDAR sensor, to control the engine and the brakes via the Engine Control Module and the Brake Control Module. The system can be turned on or off via buttons on the instrument cluster, and overridden by using the gas or brake pedal.

To extend this example somewhat, the over-the-air updating of the map information in the Navigation Controller has been added. Although these updates are also important for other applications, and should be handled separately, it has been added in this example to provide an insight in the connections of the vehicle with other entities.

Step 1: Decomposition of the Application/System

The first step of the composite threat model is to get a complete overview of the system by creating an interconnections drawing of the vehicle including all relevant components, external connections, and data types. This drawing is supplemented by high level data flows. The result of this step can be found in Figure 9.2. As can be seen, updates are forwarded to the Navigation Controller, that uses this map information to tell the Cruise Control module about heads up information. The Cruise Control Module uses this information to, depending on the current speed and the vehicle in front, speed up or slow down. For simplicity, all data flows and components have been numbered to refer to later.

Step 2: Identification and Analysis of Threats

In the second step, threats in the application/system are identified and analysed. This is done by applying the relevant STRIDE categories on each element in the interconnections drawing, and analysing

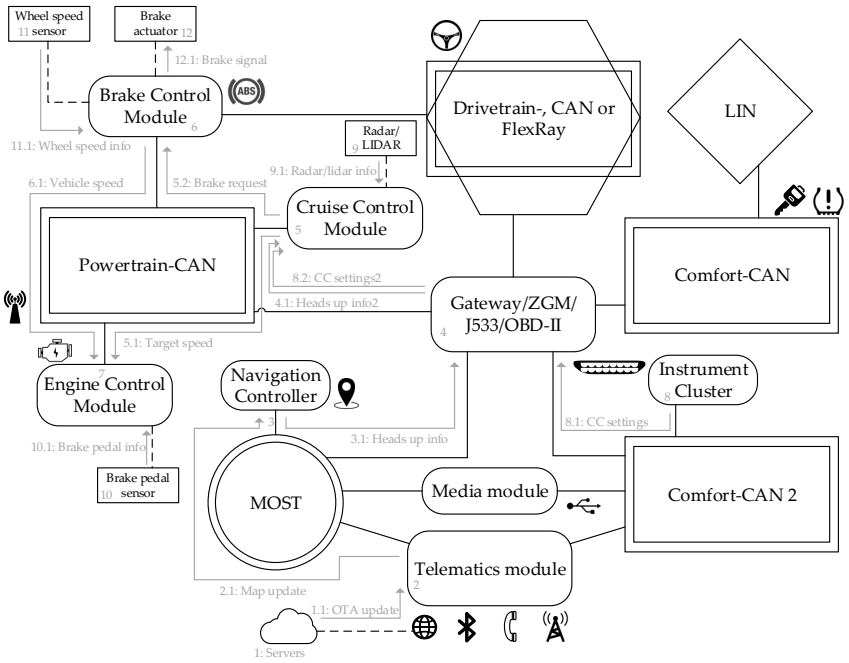

Figure 9.2: Interconnections drawing including high level data flows for Predictive Cruise Control

these threats for the Controllability and Severity classes. Ideally, a small group of experts should reason about this analysis to get a discussion going on the possible consequences. The result of this step can be found in Tables 9.5 and 9.6.

We will not discuss every item on the list, but give one example. Looking at the Brake Request message (Nr 5.2), we see that if the message is tampered with, the braking force that the Brake Control Module will use can be changed, to either braking too hard, or braking too soft. This could have possibly fatal consequences, and hence receives a level 3 Safety score.

Note that ideally, one would add a column in this table containing comments on why certain levels have been chosen. However, that does not fit onto this page, but a full threat list including comments can be found in Appendix D.

Also important to note is that specific modules and communication protocols always need a more detailed analysis if you want more specific threats. For example, to see if the OTA-update (nr 1.1) can be tampered with, you would need to know the details of the protocol. Still, the model provides a good idea about which modules are important to look at and secure, because certain consequences are higher.

NR	ELEMENT	STRIDE CLASS	SEVERITY S	O	P	F	CONTROL-LABILITY
1	Cloud	S	2	2	0	0	1
1.1	OTA-update	T	2	2	0	0	0
		I					
		D	0	2	0	0	0
2	Telematics Module	STRIDE	1	2	0	0	1
2.1	Map update	T	1	2	0	0	0
		I					
		D	0	2	0	0	0
3	Navigation Controller	STRIDE	1	3	3	0	1
3.1	Heads up information	T	1	3	0	0	1
		I	0	0	3	0	4
		D	0	1	0	0	0
4	Gateway	STRIDE	1	3	3	0	1
4.1	Heads up information 2	T	1	3	0	0	1
		I	0	0	3	0	4
		D	0	1	0	0	0
5	Cruise Control Module	STRIDE	3	2	3	0	3
5.1	Target speed	T	3	3	0	0	1
		I	0	0	1	0	4
		D	3	3	0	0	1
5.2	Brake Request	T	3	0	0	0	3
		I					
		D	3	0	0	0	1
6	BCM	STRIDE	3	0	1	0	4
6.1	Vehicle Speed	T	3	0	0	0	2
		I	0	0	1	0	4
		D	3	0	0	0	2
7	ECM	STRIDE	3	0	1	0	2
8	Instrument Cluster	STRIDE	2	2	0	0	2
8.1	CC settings	T	2	2	0	0	2
		I					
		D	2	2	0	0	2

Table 9.5: Threat list with determination of severity for Predictive Cruise Control

NR	ELEMENT	STRIDE CLASS	SEVERITY S	O	P	F	CONTROL-LABILITY
9	Radar/Lidar Sensor*	STRIDE	3	2	0	0	3
9.1	Radar/Lidar Info*	T	3	2	0	0	3
		I	0	0	1	0	4
		D	3	2	0	0	3
10	Brake pedal sensor*	STRIDE	3	0	0	0	4
10.1	Brake pedal info*	T	3	0	0	0	4
		I					
		D	3	0	0	0	4
11	Wheel speed sensors*	STRIDE	2	2	1	0	2
11.1	Wheel speed info*	T	2	2	0	0	2
		I	0	0	1	0	4
		D	2	2	0	0	2
12	Brake actuator*	STRIDE	3	0	0	0	4
12.1	Brake signal*	T	3	0	0	0	4
		I					
		D	3	0	0	0	4

Table 9.6: Threat list with determination of severity for Predictive Cruise Control, continued. *Threat requires physical access

9.2.2 *Emergency Brake Light*

The Emergency Brake Light, is, as described in chapter 6.1, a Day-One use case for Cooperative Functionality. The idea behind it is that vehicles share messages when a vehicle suddenly brakes. Especially in dense driving situations such as on a highway, where the driver has a decreased line of visibility, the driver can be warned by this system before he notices the braking himself. In future applications, the vehicle might decide to brake as well, without interference from the driver. Although this is not part of the system yet, this functionality has been added in this use case as means of illustration.

The system roughly works as follows: if a vehicle brakes hard, the vehicle decides to broadcast the emergency brake light message to other vehicles via the On Board Unit. Another vehicle notices this message, and forwards a emergency brake light message to the instrument cluster to display the message to the driver. In future applications, the On Board Unit may also send a brake request to the BCM to start braking itself.

Again, some assumptions have been made in the extend to how this systems might be implemented, since it is not used yet. The next

part will first apply the model to this Emergency Brake Light system as if it would be implemented in a modern day vehicle. Then, in the second part, the model will be applied to the system as if it would be implemented in a vehicle that has been secured as suggested by the EVITA project, as described in chapter 7.5.

9.2.2.1 *A Modern Day Vehicle*

Decomposition of the Application/System

The result of the decomposition step can be found in Figure 9.3. As can be seen, the system is pretty straightforward. The Emergency Brake Light message is forwarded via other modules to the instrument cluster that displays a message to the driver, who can decide to brake.

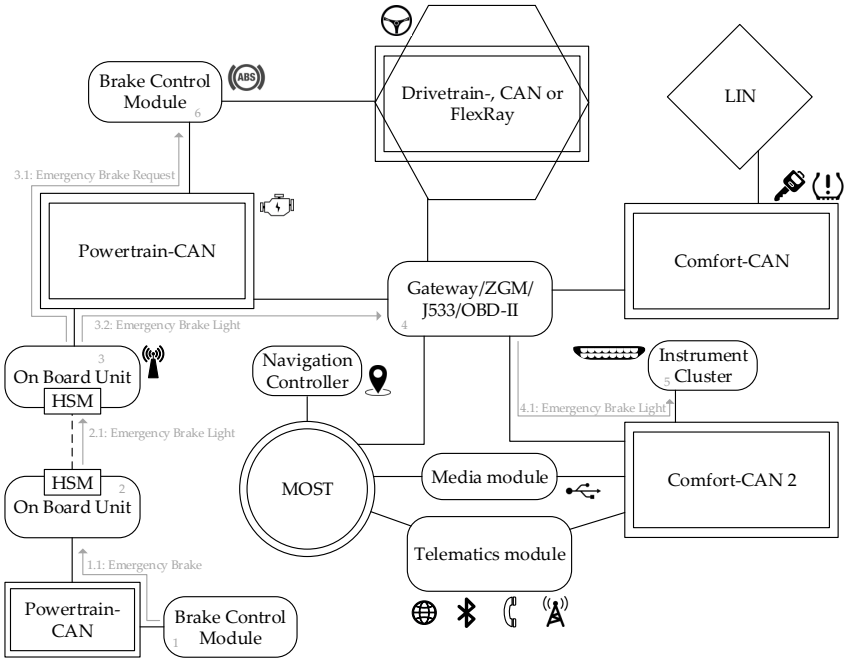

Figure 9.3: Interconnections drawing including high level data flows for Emergency Brake Light for a modern day vehicle

Identification and Analysis of Threats

Similar to the other use cases, STRIDE is applied to all elements of the interconnections drawing. The results can be found in Table 9.7.

As can be seen, since the CAN buses on both vehicles are not secured, messages can easily be tampered with, or modules spoofed. If this happens on the first vehicle, the Emergency Brake Light message will be broadcasted to multiple vehicles, causing potentially fatal con-

sequences for multiple vehicles, even though the OBUs communicate securely by use of an HSM.

Also note that since the message is forwarded by multiple modules, every module is a possible threat in the system. Message 4.1 for example, that is created by the Gateway, is not yet properly signed, meaning the gateway can be spoofed by any module on the Comfort-CAN 2.

This use case shows that current vehicles are not yet properly secured for cooperative functionality of vehicles.

NR	ELEMENT	STRIDE CLASS	SEVERITY S	O	P	F	CONTROL-LABILITY
1	Brake Control Module	STRIDE	4	3	0	0	4
1.1	Emergency Brake Light	T	4	3	0	0	4
		I	0	0	0	0	0
		D	4	3	0	0	3
2	On Board Unit	STRIDE	4	3	0	0	3
2.1	Emergency Brake Light	T	0	0	0	0	0
		I	0	0	1	0	0
		D	4	3	0	0	3
3	On Board Unit	STRIDE	3	3	0	0	3
3.1	Emergency Brake Request	T	3	3	0	0	3
		I	0	0	0	0	0
		D	3	0	0	0	3
3.2	Emergency Brake Light	T	3	3	0	0	3
		I	0	0	1	0	0
		D	3	0	0	0	3
4	Gateway	STRIDE	3	3	0	0	3
4.1	Emergency Brake Light	T	3	3	0	0	3
		I	0	0	0	0	0
		D	3	3	0	0	3
5	Instrument Cluster	STRIDE	3	3	0	0	3
6	Brake Control Module	STRIDE	3	3	0	0	4

Table 9.7: Threat list with determination of severity for Emergency Brake Light for a modern day vehicle

9.2.2.2 *An EVITA Secured Vehicle*

Having looked at how the Emergency Brake Light would function in a modern day vehicle, we now take a look at how it would work in a newer vehicle that has incorporated security measures as described in the EVITA project.

In the EVITA project, they argue that security between vehicles can not be secure, unless the communication within the vehicle is also secure, something that is supported by our findings in the previous part. They therefore designed the use of HSMs for secure communication between modules. For details, we refer back to chapter 7.5.

It is important to note that EVITA does not propose any specific architectures or communication channels one should use, but rather places the secure communications on top of it. In their examples, they however often use Ethernet as communication network. This example therefore uses Ethernet as well.

Decomposition of the Application/System

The result of the decomposition step can be found in Figure 9.4 and is quite similar to the previous part for the modern day vehicle. An important difference is that each module contains a HSM and communications between modules is signed.

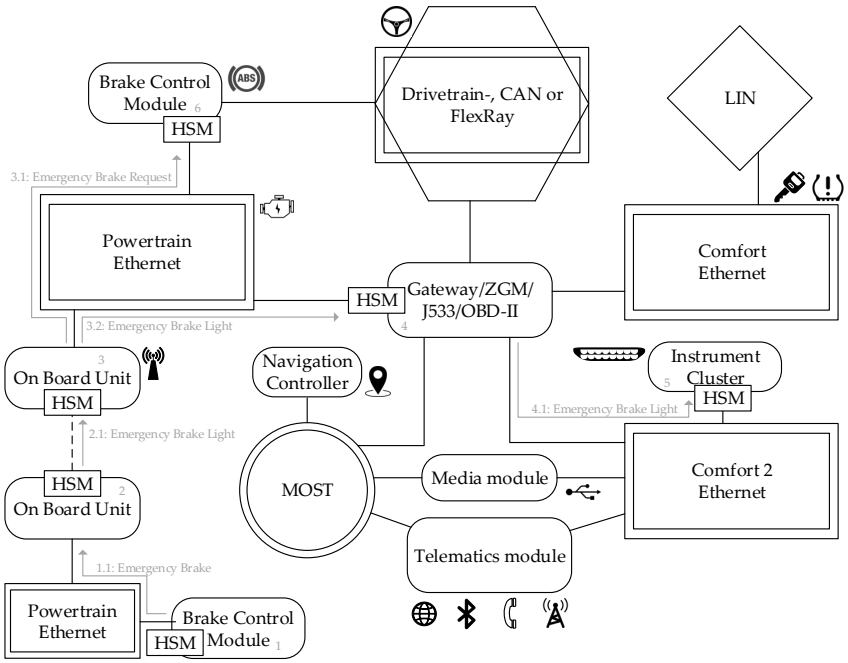

Figure 9.4: Interconnections drawing including high level data flows for Emergency Brake Light for an EVITA secured vehicle

Identification and Analysis of Threats

Similar to the other use cases, STRIDE is applied to all elements of the interconnections drawing. The results can be found in Table 9.8.

NR	ELEMENT	STRIDE CLASS	SEVERITY S	O	P	F	CONTROL-LABILITY
1	Brake Control Module	STRIDE	4	3	0	0	4
1.1	Emergency brake light	T	0	0	0	0	0
		I	0	0	0	0	0
		D	4	3	0	0	3
2	On Board Unit	STRIDE	4	3	0	0	3
2.1	Emergency Brake Light	T	0	0	0	0	0
		I	0	0	1	0	0
		D	4	3	0	0	3
3	On Board Unit	STRIDE	3	3	0	0	3
3.1	Emergency Brake Request	T	0	0	0	0	0
		I	0	0	0	0	0
		D	3	0	0	0	3
3.2	Emergency Brake Light	T	0	0	0	0	0
		I	0	0	0	0	0
		D	3	0	0	0	3
4	Gateway	STRIDE	3	3	0	0	3
4.1	Emergency Brake Light	T	0	0	0	0	0
		I	0	0	0	0	0
		D	3	3	0	0	3
5	Instrument Cluster	STRIDE	3	3	0	0	3
6	Brake Control Module	STRIDE	3	3	0	0	4

Table 9.8: Threat list with determination of severity for Emergency Brake Light for an EVITA secured vehicle

Notice the difference between the EVITA secure vehicle and the modern day vehicle. Since messages are now signed and encrypted, modules can no longer be spoofed by other modules on the bus, or tampered with. However, other threats do still exist. For example, controlling a module still gives the possibility to tamper with or spoof messages. Next to this, Denial of Service attacks are still possible, meaning that the arrival of fatal messages such as the Emergency Brake Request cannot be guaranteed.

It is also no longer possible to simple add or replace a module on a bus, since it would need proper keys to sign a message. However,

being able to hack or flash a module does again give access to the necessary keys.

9.3 VALIDATION

To validate the composite threat model as described in section 9.1, we conduct a validation interview. This interview focuses on two topics: getting an overview of the state of the art at the company of the interviewee, and validating our composite threat model.

The first part of the interview consists of an open interview to get a better understanding of what kind of techniques the interviewee has already worked with, and what kind of threats the company deems most important.

The second part of the interview consists of an introduction to the composite threat model and validating it on three different topics:

COMPLETENESS Does the model identify all important threats?
FOCUS Does the model focus on the threats that are most important?
WORKABILITY Is the model workable in practice?

The interview is conducted in a semi-structured way by asking open ended questions, allowing the interviewee to focus on areas where he/she wants to go in-depth. Partial transcriptions are made during the interview and the interviewee is given the option to review the summary to ensure no confidential information is disclosed and the summary resembles the interview.

The following part gives a summary of the finding of the interview on these three topics. Section 9.4 then provides an improved version of the composite threat model based on this feedback.

9.3.1 *Cyber Security Systems Architect of a European Truck Manufacturer*

The expert interviewed was a cyber security systems architect at a major European truck manufacturer that has worked at the company for quite some time. A summary of the whole interview can be found in Appendix C. This section briefly discusses the main findings.

State of the Art

The expert revealed that many aspects that we have seen at other vehicle manufacturers are also present in this company. There are hardly any standard risk management techniques used, but rather a mix of techniques that best suits that particular system, such as described in SAE J3061 [30]. Security it often not yet considered much in the design process, and the company still focuses a lot on safety. Some techniques that they do use, or have looked at, are the STRIDE threat modelling technique and the HEAVENS project [33].

The expert also revealed that their security team does not deem one area of threats the most important. Some examples of threats they are considering are theft, criminal hacking, or being liable if a client for example tunes his engine. However, in the rest of the company, safety is still considered to be an important focus area, because it has been around for some time and is more tangible.

Feedback on Composite Threat Model

Since the expert is already quite known with the STRIDE threat modelling technique, he agreed that our composite threat model will be very complete. He does however, also believe that the STRIDE threat modelling tool gives back too many possible threats, even the ones that are evidently mitigated. Since our composite threat model already discards non-relevant classes, this number of irrelevant threats is reduced somewhat. That is good, but it would also help a great deal if a good tool can automatically discard these kinds of irrelevant threats by for example supporting authenticated data flows.

Noteworthy however, is that by looking at some real-life examples, the interviewee has identified some hardware related threats that would not be identified by our composite threat model, since it is designed to look at the system from a high level.

The expert was also positive about the concept of having different severity classes: Safety, Operational, Privacy, and Financial. It provides an important sense that not only safety should be considered when designing a system. For example, under the new Data Privacy Regulation, companies can be fined up to 4% of their worldwide turnover if certain privacy sensitive data is leaked. Also important to note here is that it should be clear in what class a threat belongs. It can become unclear if for example a threat has operational and financial consequences and is mitigated, whether both consequences are mitigated, or only one of the two.

On the Controllability class however, the expert had a mixed opinion. Although it is interesting to add, since indeed some safety related consequences can be controlled, it can also give a false sense of security in its present form. Privacy related consequences for example is not something that can be controlled by the driver. The composite threat model has no good way of coping with this yet.

The expert also noted that weighting the the severity classes in calculating the threat level is a good idea, but he would propose already weighting the levels within the severity classes to have even more control. A company could in that way for example say that Safety 1 and Safety 2 are still not very important, but from Safety 3 onwards, it is.

Lastly, the expert noted that it would be a good idea to be able to present the threats in a simple and sellable way so that they can be better understood by management, for example by introducing a red/orange/green scale.

9.3.2 *Summary*

In short, we notice that the expert agrees with most steps from the model. Since step 1 and step 2a are already quite similar to what the expert uses in everyday life, he is quite used to those steps. The only addition in step 2a is to have a good workable tool that can be used to create the diagram, however this tool is out-of-scope for this research.

Aspects that do need to change are:

1. The Controllability class should not apply for Operational, Privacy, or Financial classes
2. It should be clearer what happens with a threat if it is mitigated when it has multiple consequences
3. The levels within the Severity classes should also have weights
4. The outcomes of the model should be more 'sellable' to management

9.4 IMPROVED COMPOSITE THREAT MODEL

As explained above, the improvements on the composite threat model all lie in step 2b) the determination of the severity of threats. This section describes a new improved step 2b as is must replace the step in the composite threat model, and is based on the feedback as discussed before.

Determination of Severity of Threats

In the second step, each entry in the list from step 2a is evaluated on the consequences in Severity. For Severity, a distinction is made between Safety, Operational, Privacy and Financial consequences, where in the analysis of Safety, on should also consider how Controllable the threat is.

To determine these classes, one must consider what would happen if the threat occurs. For example, in the case for a data flow and Denial of Service, one needs to consider what would happen if the data flow is denied, and doesn't arrive at the next ECU. Would the vehicle for example crash? Will the driver notice? Does any sensitive data leak? Can he control the consequences somewhat? etc. Normally, such a determination should be done by a small group of experts, to get some discussions going on the consequences and improve the analysis, but could be done by a single person.

To determine the severity of these classes, one must use the classification from Ward, Ibara, and Ruddle [69] as a guideline as can be found in Table 9.3. As can be seen, Ward, Ibara, and Ruddle provide a clear distinction between consequences for one or multiple vehicles. For example, if certain messages that are shared between vehicles can

be tampered with, the consequences could be greater than if only one message within a vehicle is tampered with.

Similar to the Severity classes, ISO 26262 provides a description for calculating the Controllability level [28], which is supplemented with another two classes by the MISRA standard [15]. This creates levels as described in Table 9.2.

The levels of these classes can be used to determine a single threat level. To make sure that a company can determine which kind of threats, and what kind of threat levels it deems most important, it is possible to weight the classes and levels. An example of how levels are weighted within a Severity class is is given in Table 9.9.

LEVEL	DESCRIPTION	WEIGHT
0	No injuries	0
1	Light or moderate injuries	2
2	Severe and life-threatening injuries (survival probable) or light or moderate injuries for multiple vehicles	10
3	Life-threatening injuries (survival uncertain) or fatal injuries or Severe injuries for multiple vehicles	15
4	Life-threatening or fatal injuries for multiple vehicles	20

Table 9.9: Example on how different Safety levels may be weighted within classes

Using these weighted levels, one can determine or calculate a threat level such as via:

$$\begin{aligned} \text{Threat} =& (S_{weighted} \times C_{weighted}) + O_{weighted} + P_{weighted} \\ &+ F_{weighted} \end{aligned}$$

Or for clarity, one could also weight the classes as well[2]:

$$\begin{aligned} \text{Threat} =& (w_s S_{weighted} \times w_c C_{weighted}) + w_o O_{weighted} \\ &+ w_p P_{weighted} + w_f F_{weighted} \end{aligned}$$

Important to note that is the exact calculation or determination of the threat level is highly dependent on wishes of the company. One can use the above calculation, or use a determination matrix as given in Table 8.3 or 8.4.

2 Please note that for clarity, the different classes can still have weights, but this is implicitly included when weighting individual levels: one could weight levels within the Safety class ten times higher than in the Operational class

9.5 CHAPTER SUMMARY

This chapter, together with chapter 8 set out to answer research question three:

> What is a threat model for selected functions within vehicles in 5 to 10 years' time for a general IT architecture?

Where chapter 8 has given an overview of already existing automotive risk management and threat modelling techniques, this chapter focused on designing a composite threat model that focuses on identifying all possible threats under the assumption that the system is already breached.

To this end, our composite threat model consists of two steps. The first focuses on getting an overview of the system by creating an interconnections drawing with high level data flows. By doing so, all relevant components and data flows are identified, creating a complete overview and understanding of the system. Next, this drawing is used to get an extensive list of possible threats by applying the STRIDE threat modelling technique and analysing these threats on Severity: how bad are the consequences of the threat, and Controllability: how controllable is the threat. This process is visually described in Figure 9.5.

Figure 9.5: Composite Threat Model Steps

We further provide three use cases based on the future functionality and the European general IT architecture from chapters 6 and 7, where we apply our composite threat model to. These use cases show that modern day vehicles are not yet secure for future functionality. The Predictive Cruise Control use case for example shows how threats on less critical parts of the system van still have Safety consequences. And the Emergency Brake Light use case shows that cooperative functionality, even though the communication between vehicles is secured, is still susceptible to threats in a modern day vehicle, supporting the claim from the EVITA project that secure communication between vehicles requires secure communication within vehicles. Finally, the use cases have also shown that even in an EVITA secured vehicle, threats still exist.

Lastly, we interviewed an expert from the field to validate our composite threat model. In short, the expert agreed with most of the steps from our model, finding it very complete and giving good focus on

threats that are most important with the introduction of different Severity classes. The expert did provide feedback on some aspects. The Controllability should only apply for the Safety class, and not for the other Severity classes, since Privacy for example, is not controllable by the driver. Next to this, the expert suggested weighting the levels within the Severity classes to give a company even more control over what kind of threats it deems most important or acceptable. A new version of this step is given in section 9.4.

Part III

CONCLUDING THE RESEARCH

10

CONCLUSION

This thesis set out to answer the problem statement: how to design a framework to help a security expert at a vehicle manufacturer to control threats and risks for future functionality in a vehicle.

This problem statement was phrased in the following main research question:

> What could be a suitable framework that helps a security expert within a vehicle manufacturer to control threats and risks for the vehicle of the future?

that has been divided into the following knowledge questions:

1. What functionality will be present in vehicles in five to ten years' time?
2. What will a general IT architecture in a vehicle look like?
   a) What are the attack surfaces within this architecture?
3. What is a threat model for selected functions within vehicles in 5 to 10 years' time for a general IT architecture?

This chapter will first briefly discuss the answers to the subquestions as addressed in previous chapters, and explain why we have created a new composite threat model. After which we will conclude with some remarks on our threat model.

Firstly, we identified a trend towards more and more complex functions where vehicles are increasingly taking over critical functions from the driver, such as parking assistance and lane keeping assistance. Often, these functions are accompanied by an increased information exchange both between vehicles, as well as with entities on the internet.

Secondly, we identified an increase of complexity inside the vehicle. More and more IT is being integrated, introducing many more Electronic Control Units (ECUs) as well as connections between ECUs. As a consequence, many ECUs are connected to one another via a web of buses, for example distinguishing in how critical they are.

Looking at how secure these vehicles are, we see that security is often still lacking. Communication protocols such as the CAN bus

lack basic security features and ECUs can often simply be flashed. Although some research projects have addressed this issue, these are still ongoing, and are not implemented yet. It is evident that this issue needs addressing.

We also see that although vehicle manufacturers have had a rich history in addressing safety issues, they have only recently started looking into making their products secure. Since IT systems are taking over more and more critical functions, we argue that your product cannot be safe, unless you also address security; a hacked ECU can have safety-related consequences.

To be able to cope with these kind of threats, some frameworks such as an extension of ISO 26262 or the NHTSA's modified version of the NIST Risk Management Framework have been designed that do incorporate threat models and security objectives in the design process. These frameworks however, do not include any specifics on how such threat analyses should be performed.

Some literature can be found that does focus on building such threat models. However, many of these models are very attacker centric. Having focused on identifying threats first, many models focus on finding ways how these threats can be exploited. A big concern with this approach is that the system is considered to be secure in the first place, and cannot be altered other than via an attack vector. Many modern day vehicles however, have so many different vendors of ECUs, some of which are not easily controlled, that ECUs can for example easily be flashed before being put in the vehicle. In short, it is hard to claim that the vehicle is 100% secure to begin with. Even if it is secure, attacks such as Stuxnet have shown that even air-gapped systems can be compromised. We therefore argue that threat models should not focus on how an attacker might be able to break into the vehicle, but rather assume the attacker is already inside the vehicle and design the system to cope with these threats. To our knowledge, such a threat model does not yet exist.

We have therefore designed a composite threat model, based on elements from other threat models, that focuses on identifying all possible threats as if the attacker has already breached the system. The model consists of two steps. First, a complete decomposition of the system including high level data flows is created in a so called interconnections drawing. Secondly, the STRIDE threat modelling technique (see chapter 8.1.1) is used to identify possible threats based on this interconnections drawing, and the threats are analysed for their Controllability, and Severity on four areas: Safety, Operational, Privacy, and Financial.

What makes our model different from other models is the focus on

identifying all threats as if the attacker is already inside, instead of looking at what actions an attacker needs to perform to exploit that threat. Next to this, our model combines the Controllability and Severity aspects for security. Some threats might have fatal consequences, but can be easily controlled by a normal driver, making the threat less probable.

Although not necessarily new, it is also worth noting that our model is capable of dealing with complex systems, by providing steps to decompose the system in a structured way. In our opinion, this is becoming more and more important since vehicles are also becoming more and more complex. By using the STRIDE threat modelling technique, our model could also be used by a non-security-expert to reason about possible threats. Nevertheless, in our opinion, it will be best to have a mixture of security and safety experts to perform this analysis to get input from both fields.

Finally, our model is designed to work with any architecture and system as input, which means that it can be used for both trucks and cars.

We validated this model by interviewing an expert from the field. We found that the expert agreed with most of the aspects in our model, and some feedback has been integrated in the final version of the model. By applying our model to two use cases of future functionality, we found that modern day vehicles are still very vulnerable to threats, and that vehicles are not yet ready for vehicle-to-vehicle communication if it is applied to modern vehicles. We conclude that for secure communication between vehicles, secure communication within the vehicle is required, supporting the claim from the EVITA project. Nonetheless, even with secure communication within a vehicle, other types of threats do still exist.

11

DISCUSSION

This chapter will discuss the contributions of this research, as well as some limitations and future work.

11.1 CONTRIBUTIONS

To the best of our knowledge, this thesis is the first to propose a composite threat model that focuses on identifying threats under the assumption that the attacker has already breached the system. Already existing models do often also focus on identifying threats first, but then rather look at finding a way an attacker might exploit that threat or enter the vehicle. By doing so, the results of the model often focus on making it hard for an attacker to enter the vehicle, instead of making the whole vehicle secure. Our model however, assumes that an attacker has already entered the vehicle and can perform any threat, thereby focusing more on the consequences of a threat.

Next to this, our model provides a step by step method to decompose the vehicle and its systems in a structured way and analyse these systems. In our opinion, this is becoming more and more important since vehicles are also becoming more and more complex. This step by step approach also provides non-security experts with the tools to reason about possible threats, something we believe is very useful in a world where security teams within vehicle manufacturers are still rather small. Nevertheless, we do believe that for a good threat analysis, both security and safety experts should be included.

To get to our model, this thesis has also provided a further analysis of 24 different car models to determine how certain buses are connected, as well as where systems with a cyber-physical aspects are located that could serve as an entry point for an attacker. This analysis resulted in a generalised IT architecture per region for Europe, the United States of America and Asia. To our knowledge, no such analyses and generalised IT architectures exist.

Lastly, by applying our model to vehicle-to-vehicle communication use cases on the generalised European IT architecture, we have concluded that many threats exist and implementing this functionality on a modern day vehicle may not be wise. Even if the communication between vehicles is secured, an attacker in one car can control

the consequences of many cars, supporting the claim from the EVITA project that secure communication between vehicles requires secure communication within vehicles.

11.2 LIMITATIONS AND FUTURE WORK

As already discussed in chapter 5, there are some concerns for validity applicable in the systematic literature review. Due to there being only one author, the literature has been performed by only one person. This might introduce a bias in both the found papers, as well as the classification of the papers on different topics. However, since the goal of the research is not to provide a good classification, the latter is considered to be not important.

The first bias however, is strengthened by the fact that there were many papers available on the subject, but might not be completely relevant for further research. That is why only newer papers were considered in the initial search and older papers were only added if deemed relevant.

Another limitation lies in the list of future functionality. Only a part of this list is standardised and from open sources, since it requires cooperation between manufacturers. The other part does not require any cooperation and is therefore often kept a secret, since it can give a manufacturer an edge on the market. This creates a bias in the list of future functionality, but still gives a good sense of the kinds of functions we might see in a few years' time.

The next limitation lies in the analysis of the generalised IT architecture. Although a part of it already comes from another study, this study proofed insufficient for our means and a further analysis has only been performed by one author. Next to this, the analysis has taken place on 24 different car models only. A further analysis should include truck models, as well as more car models for a better overview.

Having seen these limitations, we would like to add that in our model, the exact functionality and vehicular architecture are used as input. If in the future, either of the inputs changes, the model could still be used to identify and analyse threats. However, it is worth noting that the model might become unusable in the future if systems get very complex. Even though the model provides a step of decomposing the system, we can imagine that in the future, these systems become so complex, that decomposing them becomes very hard. Nevertheless, we believe that for a full threat analysis, every part of the vehicle should be analysed.

The biggest limitation lies in the validation of our composite threat model. During this validation, only one expert from a truck man-

ufacturer has been interviewed to validate our model. For a better validation, more interviews should be conducted, both with car and truck manufacturers.

Due to this limitation, saying how well validated the model is, is hard. For the expert we interviewed, the model proved workable and gave a complete list of threats. However, another manufacturer might not agree with this. Many manufacturers for example still have only relatively small security departments, making it very time consuming to fully analyse all systems within the vehicle.

Nevertheless, in our opinion, it is important get a complete overview of all possible threats, which is needed to fully secure a vehicle, and provide secure future vehicles.

11.2.1 *Future work*

Part of the feedback given by the expert during the validation interview has not been used to improve our model, but is left as future work. Firstly, the expert noticed that good tools are missing to support in identifying possible threats. A tool that could aid in the creation of the diagrams, and in particular, a tool that is capable of supporting elements such as 'authenticated data flow' to decrease the amount of irrelevant threats found, would make the whole step a lot easier. The creation of such a tool could be a whole new study.

Secondly, the expert noticed that it would be good to make the outcome of the model 'sellable' to management. There is already a lot of research outside the automotive industry going on to how you should report security risks to management and make these risks understandable. This is heavily dependent on the type of company, its wishes and how well known security is. Nevertheless, the automotive industry could benefit from such a research, since it in particular, is still trying to find a way of coping with security.

Part IV

APPENDIX

A

SYSTEMATIC LITERATURE REVIEW: SCOPUS SEARCH

The exact search query used in the systematic literature review.

```
(TITLE-ABS-KEY(
   ("risk analysis" OR "risk assessment")
   AND "secur*"
   AND ("vehic*" OR "automotive" OR "car")
))
AND (PUBYEAR > 2011)
AND (
   LIMIT-TO(SUBJAREA,"ENGI" )
   OR LIMIT-TO(SUBJAREA,"COMP" )
)
```

B

EXPERT INTERVIEW

This appendix is confidential in the public version of this thesis. A summary of the interview may be requested from the author.

C

EXPERT VALIDATION INTERVIEW

A summary of the expert interview as conducted in the validation step of the research in December of 2016.

C.1 INTERVIEW CYBER SECURITY SYSTEMS ARCHITECT OF A EUROPEAN TRUCK MANUFACTURER

The following is an anonymised report of a meeting between a Cyber Security Systems Architect of a European truck manufacturer and Stijn van Winsen regarding the validation of the threat model as described in this thesis. The meeting concerned the following two topics:

- The state of the art of risk management and threat modelling at the European truck manufacturer
- Feedback on the composite threat model as described in this thesis

Risk Management and Threat Modelling at the Truck Manufacturer

The first part of the interview concerned the state of the art of the manufacturer on risk management and threat modelling. Topics that were discussed include what kind of techniques are used, and what kind of threats and risks are deemed most important.

Currently, the manufacturer already uses some risk management techniques, however, there is no real standard that is used, but rather a mix of techniques. Some standards such as SAE J3061 give good guidelines on cyber security management, but still require an exact implementation. Whenever an assessment needs to be performed, they choose a technique that best suits that particular system.

Worth to note is that within the manufacturing company, safety related techniques have been used for a longer time and are better known by a wider range of people. However, security is sometimes still not considered much when designing a system. When the security team, that is relatively small, is included in the design process, common questions they consider include:

- What is the system meant to do? Does it for example regard any safety related issues?
- Where is the system located? Is it for example located on the safety critical bus?
- What is the business case? Can for example security measures be added without ruining the business case?

As explained before, the manufacturing company uses a mix of risk management techniques. An example is the use of attack-defence trees, however, these do not always work well since they get broad and big easily. Another technique that is used, is borrowed from the military: so called assurance trees, where the security of a system is based on proof that is built on sub-proofs.

Other techniques that are used include the techniques from the HEAVENS project, and the STRIDE threat modelling technique. However, worth noting is that the Microsoft Threat Modelling tool cannot handle big models and is very buggy and therefore doesn't always work well enough.

When looking at threats, there is not one threat that the manufacturer deems most important, but rather multiple areas of threats. Theft, future criminal hacking and being liable if a client tunes the engine of a vehicle are examples of threats that are considered important. However, it is still worth to note that safety is still considered by a lot of people to be very important since it has been around longer and is more tangible. It is therefore 'easy to sell' to management, especially when building ADAS functionality that has a cyber-physical aspect.

Feedback on the Composite Threat Model

The second part of the interview consisted of introducing the composite threat model, applying it to a case, and provide feedback on the model. The results are aggregated into the following four topics:

Completeness

Since the threat model has similarities to the STRIDE methodology that the company has already used, it will identify a near complete list of threats. However, since this composite threat model discards some of the STRIDE classes that are not important for all entities, it will return a less complete list than STRIDE would. STRIDE however, identifies too many unimportant threats that can be removed after a first analysis.

Next to this, this composite threat models looks at one function at a time. Although this is time consuming, it is good for completeness, and is what the manufacturers currently does as well.

Focus

The classification of severity on different classes (Safety, Operational, Privacy and Financial) is considered good, it provides a realisation that not only safety consequences are important. Having looked at two real-life examples, of which the exact details are not disclosed, the composite threat model will probably have identified the important threats of one example, but probably not the other one. Since that threat is more hardware related, it would need a far more detailed check, which the composite threat model does not cover.

Workability

As is also the case with the STRIDE threat modelling methodology, and even though the composite threat model identified less threats, the amount of identified threats is still quite big, making it less workable. One of the reasons STRIDE does not always work well is because the tool identifies many irrelevant threats. A good tool, in combination with the composite threat model could really help reduce these 'easy' threats that are irrelevant and make it more workable.

Other notes

Part of the composite threat model is to weight the severity classes, for example considering safety ten times more important than Privacy. However, the interviewee noted that weighting the levels within these classes gives even more flexibility of determining what kind of threats are more important than others.

Also important to note is that the definitions of the severity classes needs to be very exact. A threat such as handling private data that would fall under the General Data Privacy Regulation can have both Privacy consequences and Financial consequences (a data breach of this data can give a fine up to 4% of the worldwide turnover). When a threat falls under two of these categories, it becomes unclear when one, both or none of the threats are mitigated and what the exact severity of the threat is. You could say the severity is an addition of the categories and both threats are mitigated, or none, but the problem however remains that the composite threat model has no clear way of coping with this.

The interviewee also noted that adding the controllability class is interesting, because it gives a better notion about how much risk something actually is if it is still controllable. However, it can also be very misleading. Safety is something that can sometimes be controlled, but privacy isn't. The controllability class can therefore make a threat look less risky because the controllability is high, even though it could have a huge impact on privacy that isn't controllable.

Lastly, the interviewee noted that a good addition would also be to present the threats in a simple and 'sellable' way for management, by for example using a red/orange/green scale.

D

FULL ELABORATION USE CASE THREAT LISTS

D.1 PREDICTIVE CRUISE CONTROL

The full elaborate threat list from the Predictive Cruise Control use case.

Nr	Element	Stride Class	Notes	S	O	P	F	C
1	Cloud	S	If not properly authorized, anyone can give updates, leading to incorrect maps, which can cause annoyance that the driver will notice. However, driver is still in control of the vehicle.	2	2	0	0	1
1.1	OTA-update	T	If tampering possible, map updates can be altered. Replay attacks can also override newer updates	2	2	0	0	0
		I						
		D	Updates can be blocked, making the maps out-dated	0	2	0	0	0
2	Telematics Module	STRIDE	Module needs detailed check on all STRIDE classes	1	2	0	0	
			If controlled, map updates can be changed indefinitely, however, the driver will always remain in control					1
2.1	Map update	T	If tampering possible, map updates can be altered. Replay attacks can also override newer updates	1	2	0	0	0
		I						
		D	Updates can be blocked, making the maps out-dated	0	2	0	0	0
3	Navigation Controller	STRIDE	Module needs detailed check on all STRIDE classes	1	3	3	0	1
			If controlled, the map, and heads up information can be altered, and reveals continuous information about location (GPS). Driver still in control					
3.1	Heads up information	T	By tampering the heads up information, the cruise control can for example think it needs to brake really hard	1	3	0	0	1
		I	Leaks information about the location of the vehicle	0	0	3	0	4
		D	Denying this service to arrive, will cause some annoyance to the driver, but will not cause any safety incidents	0	1	0	0	0
4	Gateway	STRIDE	Module needs detailed check on all STRIDE classes	1	3	3	0	1
			By controlling this module, the heads up info can be altered, even if authentication is in place. Also reveals information on location of vehicle.					
4.1	Heads up information 2	T	By tampering the heads up information, the cruise control can for example think it needs to brake really hard	1	3	0	0	1
		I	Leaks information about the location of the vehicle	0	0	3	0	4
		D	Denying this service to arrive, will cause some annoyance to the driver, but will not cause any safety incidents	0	1	0	0	0
5	Cruise Control Module	STRIDE	Module needs detailed check on all STRIDE classes	3	2	3	0	3
			If controlled, can fully determine the set speed and brake requests, even if properly authorized. Braking request cannot be controlled by driver. (although exact implementation in the BCM might prevent hard brakes from the cruise control)					
5.1	Target speed	T	if not secured, can be set to dangerous speeds	3	3	0	0	1
		I	Reveals information about the set speed	0	0	1	0	4
		D	New speeds can be denied, preventing the setting of newer speeds that can be lower	3	3	0	0	1
5.2	Brake Request	T	Tampering with a brake request can alter the braking amount, when altering to brake really hard, hard to override	3	0	0	0	3
		I						
		D	Denying a brake request can result in going to fast and hitting another car	3	0	0	0	1
6	BCM	STRIDE	Module needs detailed check on all STRIDE classes	3	0	1	0	4
			Controlling the module gives complete control over the brakes, also revealing information about speed					

Figure D.1: Threat list with determination of severity for Predictive Cruise Control

Nr	Element	Stride Class	Notes	S	O	P	F	C	
6.1	Vehicle Speed	T	The messages can be changed, letting the ECM believe it is not going the required speed and increase the speed	3	0	0	0	2	
		I	Reveals information about the speed of the driver	0	0	1	0	4	
		D	If denied, will not change much but can let the ECM not now the correct speed	3	0	0	0	2	
7	ECM	STRIDE	Module needs detailed check on all STRIDE classes Controlling the module gives complete control over the engine, also revealing information about speed	3	0	1	0	2	
8	Instrument Cluster	STRIDE	Module needs detailed check on all STRIDE classes Controlling the module can change the cc settings, could in extreme cases lead to safety issues, but generally controllable	2	2	0	0	2	
8.1	CC settings	T	The settings can be changed, giving annoyance, and in extreme cases lead to safety issues	2	2	0	0	2	
		I							
		D	When denied, can prevent the setting of new settings	2	2	0	0	2	
9	Radar/Lidar Sensor	STRIDE	Module needs detailed check on all STRIDE classes If controlled, can fully determine the radar/LIDAR information, with the same consequences as Radar/LIDAR information	3	2	0	0	3	*
9.1	Radar/Lidar Info	T	Needs physical access, but can make the CC module believe it needs to brake really hard or can go faster if set is not yet reached	3	2	0	0	3	*
		I	Reveals information about tailgating	0	0	1	0	4	*
		D	Denying this service, can in extreme cases prevent the CC from braking if a car in front brakes as well	3	2	0	0	3	*
10	Brake pedal sensor	STRIDE	Module needs detailed check on all STRIDE classes If controlled, can fully determine the brake pedal message, even if secured communication	3	0	0	0	4	*
10.1	Brake pedal info	T	Can be tampered to brake harder or softer than needed	3	0	0	0	4	*
		I							*
		D	Brake messages can be denied, preventing braking	3	0	0	0	4	*
11	Wheel speed sensors	STRIDE	Module needs detailed check on all STRIDE classes If controlled, can be used to fake the wheel speed, and reveals information about driving behaviour.	2	2	1	0	2	*
11.1	Wheel speed information	T	Can be tampered to fake the wheel speeds	2	2	0	0	2	*
		I	Leaks information about driving behaviour (such as speeding)	0	0	1	0	4	*
		D	Denting this service can make other systems unaware of the actual speed and fall back to a slower mode	2	2	0	0	2	*
12	Brake actuator	STRIDE	Module needs detailed check on all STRIDE classes If controlled, can fully control the brakes	3	0	0	0	4	*
12.1	Brake signal	T	Message can be tampered with to not or brake harder than meant to	3	0	0	0	4	*
		I							*
		D	Brake message can be denied, needs physical access	3	0	0	0	4	*

Figure D.2: Threat list with determination of severity for Predictive Cruise Control, continued

D.2 EMERGENCY BRAKE LIGHT

The full elaborate threat list from the Emergency Brake Light use case.

Nr	Entity	Stride Class	Notes	S	O	P	F	C
1	Brake Control Module	STRIDE	Can trigger emergency brake lights impacting multiple vehicles with dangerous situations	4	3	0	0	4
1.1	Emergency brake light	T	Because it is not properly secured, can be tampered with. For example by changing the location of the vehicle	4	3	0	0	4
		I	No issues					
		D	If the message is denied, can cause serious injury because following cars wont be noticed of possible problems	4	3	0	0	3
2	On-Board Unit	STRIDE	If controlled, can trigger, or deny messages for multiple vehicles	4	3	0	0	3
2.1	Emergency Brake Light	T	Is properly secured, cannot be tampered with	0	0	0	0	0
		I	Messages are not encypted, leaks information about pseudonyms	0	0	1	0	0
		D	Messages can be jammed and denied	4	3	0	0	3
3	On-Board Unit	STRIDE	If controlled, can trigger or deny messages 3.1 and 3.2, causing serious injuries	3	3	0	0	3
3.1	Emergency Brake Request	T	Can be tampered with to alter the brake request message, braking less or harder then needed. Braking less can be overridden by the driver	3	3	0	0	3
		I	No issues	0	0	0	0	
		D	If the message is denied, car will not brake, returning to normal cicumstances where the driver needs to react	3	0	0	0	3
3.2	Emergency Brake Light	T	Because it is not properly secured, can be tampered with	3	3	0	0	3
		I	Leaks information about pseudonym's location	0	0	1	0	0
		D	Messages can be jammed and denied	3	0	0	0	3
4	Gateway		If controlled, Emergency Brake Light message can be spoofed, or denied	3	3	0	0	3
4.1	Emergency Brake Light	T	Because it is not properly secured, can be tampered with	3	3	0	0	3
		I	Leaks information about pseudonym's location	0	0	0	0	0
		D	Messages can be jammed and denied	3	3	0	0	3
5	Instrument Cluster	STRIDE	See Gateway	3	3	0	0	3
6	Brake Control Module	STRIDE	If controlled, emergency braking can be activated, or denied. In case of denying, driver can override	3	3	0	0	4

Figure D.3: Threat list with determination of severity for Emergency Brake Light for a modern day vehicle

Nr	Entity	Stride Class	Notes	S	O	P	F	C
1	Brake Control Module	STRIDE	Can trigger emergency brake lights impacting multiple vehicles with dangerous situations	4	3	0	0	4
1.1	Emergency brake light	T	Since messages are signed, tampering cannot happen unless keys are known	0	0	0	0	0
		I	No real issues	0	0	0	0	0
		D	Messages are authenticated, however, on ethernet, a DoS attack can still prevent messages from arriving	4	3	0	0	3
2	On-Board Unit	STRIDE	If controlled, can trigger, or deny messages for multiple vehicles	4	3	0	0	3
2.1	Emergency Brake Light	T	Signed and secured, cannot be tampered with unless keys are known	0	0	0	0	0
		I	Messages are not encypted, leaks information about pseudonyms	0	0	1	0	0
		D	Messages can be jammed and denied	4	3	0	0	3
3	On-Board Unit	STRIDE	If controlled, can trigger or deny messages 3.1 and 3.2, causing serious injuries	3	3	0	0	3
3.1	Emergency Brake Request	T	Messages are signed and encypted, tampering cannot happen unless kes are known	0	0	0	0	0
		I	No issues	0	0	0	0	0
		D	On ethernet, messages can still be DoS'ed	3	0	0	0	3
3.2	Emergency Brake Light	T	Messages are signed and encypted, tampering cannot happen unless kes are known	0	0	0	0	0
		I	No issues	0	0	0	0	0
		D	On ethernet, messages can still be DoS'ed	3	0	0	0	3
4	Gateway		If controlled, Emergency Brake Light message can be spoofed, or denied	3	3	0	0	3
4.1	Emergency Brake Light	T	Messages are signed and encypted, tampering cannot happen unless kes are known	0	0	0	0	0
		I	No issues	0	0	0	0	0
		D	On ethernet, messages can still be DoS'ed	3	3	0	0	3
5	Instrument Cluster	STRIDE	See Gateway	3	3	0	0	3
6	Brake Control Module	STRIDE	If controlled, emergency braking can be activated, or denied. In case of denying, driver can override	3	3	0	0	4

Figure D.4: Threat list with determination of severity for Emergency Brake Light for an EVITA secured vehicle

E

VALIDATED COMPOSITE THREAT MODEL

This appendix deescribes the final version of the composite threat model as given by chapters 9.1 and 9.4.

E.1 COMPOSITE THREAT MODEL

The composite threat model consists of three steps. In step 0, all critical applications and systems should be identified.

Then for each identified critical application and system, step 1; the applications and systems are decomposed to get a complete overview and understanding of the system, and step 2; threats are identified and analysed to determine their consequences. These steps are as follows:

0. Identification of critical applications/systems
   For all identified applications and systems:
1. Decomposition of the application/system
   a) Create interconnections drawing of the vehicle
   b) Create high level flows in interconnections drawing
2. Identification and analysis of threats
   a) Threat identification using STRIDE
   b) Determination of severity of threats

Please note that these steps are part of the modified NIST and NHTSA Risk Management Framework as given in chapter 8.1.3. These steps are visually described in Figure E.1 and will be further explained in the coming sections.

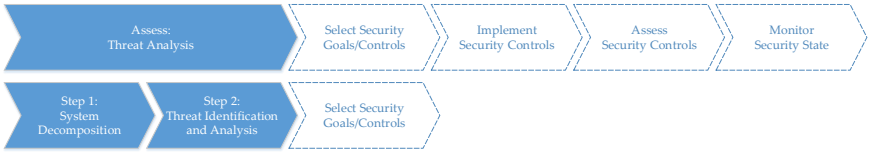

Figure E.1: Composite Threat Model Steps as part of the NIST framework

Step 0: Identification of Critical Applications/Systems

For a complete overview of the vehicle, all critical applications or systems need to be identified and further investigated. However, applications that are deemed critical are more likely to result in serious threats, and could be investigated first when short on time. This step is therefore referred to as step 0, and could be skipped if all applications or systems are analysed anyway.

A critical application or system is in this sense, a function that if compromised maliciously, could result in serious consequences, either safety related, or in other ways.

Step 1: Decomposition of the Application/System

The goal of the first step is to get a complete overview of the application/system and its components. To this end, the decomposition of the application/system takes place in two steps: first, a complete interconnection drawing is created containing the following components:

- All (relevant) components, subsystems and buses: Identify all components that are connected to the system under consideration
- External connections: identify any external connections the system such as via WiFi, OBD-II, Bluetooth or 3G
- Data types: identify safety levels for data in components and on buses. Especially note systems and buses where a safety level tradeoff is present such as gateways

Note that for one vehicle model, such a decomposition can be made once and used for all relevant applications and sub systems. One may ignore irrelevant modules if not applicable for a certain application or system.

In the second step, high level data flows should be identified and added to the interconnection drawing, considering what data flows from which component to what other component.

Step 2: Identification and Analysis of Threats

Using the drawings from the previous phase, threats are identified and further analysed. This is again done in two steps.

a) Threat identification using STRIDE

During the first step, all threats are identified using the STRIDE threat modelling technique. This means that for each ECU, data flow and external entity from the interconnections drawing, the relevant STRIDE classes are used to identify possible threats, resulting in a list of

threats for all components of the system where each STRIDE class is applied to each element as according to Table E.1. This table is a modified version of Table 8.1, where it has been made suitable for our interconnections drawing. This means that the data stores have been removed, and processes have been replaced by ECUs, that have similar functions.

ELEMENT	S	T	R	I	D	E
ECU	X	X	X	X	X	X
Data Flow		X		X	X	
External Entity	X		X			

Table E.1: STRIDE categories mapped on Interconnection drawing elements

b) Determination of Severity of Threats

In the second step, each entry in the list from step 2a is evaluated on the consequences in Severity. For Severity, a distinction is made between Safety, Operational, Privacy and Financial consequences, where in the analysis of Safety, on should also consider how Controllable the threat is.

To determine these classes, one must consider what would happen if the threat occurs. For example, in the case for a data flow and Denial of Service, one needs to consider what would happen if the data flow is denied, and doesn't arrive at the next ECU. Would the vehicle for example crash? Will the driver notice? Does any sensitive data leak? Can he control the consequences somewhat? etc. Normally, such a determination should be done by a small group of experts, to get some discussions going on the consequences and improve the analysis, but could be done by a single person.

CONTROLLABILITY	
LEVEL	DESCRIPTION
0	Controllable in general
1	Simply controllable
2	Normally controllable (most drivers could act to prevent injury)
3	Difficult to control or uncontrollable
4	Genuinely uncontrollable

Table E.2: Controllability Level determination

To determine the severity of these classes, one must use the classification from Ward, Ibara, and Ruddle [69] as a guideline as can be found in Table E.4. As can be seen, Ward, Ibara, and Ruddle provide a clear distinction between consequences for one or multiple vehicles. For example, if certain messages that are shared between vehicles can be tampered with, the consequences could be greater than if only one message within a vehicle is tampered with.

Similar to the Severity classes, ISO 26262 provides a description for calculating the Controllability level [28], which is supplemented with another two classes by the MISRA standard [15]. This creates levels as described in Table E.2.

The levels of these classes can be used to determine a single threat level. To make sure that a company can determine which kind of threats, and what kind of threat levels it deems most important, it is possible to weight the classes and levels. An example of how levels are weighted within a Severity class is is given in Table E.3.

LEVEL	DESCRIPTION	WEIGHT
0	No injuries	0
1	Light or moderate injuries	2
2	Severe and life-threatening injuries (survival probable) or light or moderate injuries for multiple vehicles	10
3	Life-threatening injuries (survival uncertain) or fatal injuries or Severe injuries for multiple vehicles	15
4	Life-threatening or fatal injuries for multiple vehicles	20

Table E.3: Example on how different Safety levels may be weighted within classes

Using these weighted levels, one can determine or calculate a threat level such as via:

$$\begin{aligned} \text{Threat} =& (S_{weighted} \times C_{weighted}) + O_{weighted} + P_{weighted} \\ &+ F_{weighted} \end{aligned}$$

Or for clarity, one could also weight the classes as well[1]:

$$\begin{aligned} \text{Threat} =& (w_s S_{weighted} \times w_c C_{weighted}) + w_o O_{weighted} \\ &+ w_p P_{weighted} + w_f F_{weighted} \end{aligned}$$

1 Please note that for clarity, the different classes can still have weights, but this is implicitly included when weighting individual levels: one could weight levels within the Safety class ten times higher than in the Operational class

Important to note that is the exact calculation or determination of the threat level is highly dependent on wishes of the company. One can use the above calculation, or use a determination matrix as given in Table 8.3 or 8.4.

SAFETY		OPERATIONAL	
LEVEL	DESCRIPTION	LEVEL	DESCRIPTION
0	No injuries	**0**	No impact on operational performance
1	Light or moderate injuries	**1**	Impact not discernible to driver
2	Severe and life-threatening injuries (survival probable) or light or moderate injuries for multiple vehicles	**2**	Driver aware of performance degradation or Indiscernible impacts for multiple vehicles
3	Life-threatening injuries (survival uncertain) or fatal injuries or Severe injuries for multiple vehicles	**3**	Significant impact on performance or Noticeable impact for multiple vehicles
4	Life-threatening or fatal injuries for multiple vehicles	**4**	Significant impact for multiple vehicles
PRIVACY		**FINANCIAL**	
LEVEL	DESCRIPTION	LEVEL	DESCRIPTION
0	No unauthorised access to data	**0**	No financial loss
1	Anonymous data only (neither specific driver not vehicle data)	**1**	Low-level loss ($10)
2	Identification of vehicle or driver or anonymous data for multiple vehicles	**2**	Moderate loss ($100) or low losses for multiple vehicles
3	Driver or vehicle tracking or identification of driver or vehicle for multiple vehicles	**3**	Heavy loss ($1000) or moderate losses for multiple vehicles
4	Driver or vehicle tracking for multiple vehicles	**4**	Heavy losses for multiple vehicles

Table E.4: Severity Level determination for Safety, Operational, Privacy and Financial

BIBLIOGRAPHY

[1] SKODA AUTO. *Assistants*. Accessed on: 2016-10-04. 2016. URL: http://www.skoda-auto.com/en/models/new-octavia/assistants/.

[2] Ludovic Apvrille, Rachid El Khayari, Olaf Henniger, Yves Rouder, Hendrik Schweppe, Hervé Seudié, Benjamin Weyl, and Marko Wolf. "Secure automotive on-board electronics network architecture." In: *FISITA 2010 world automotive congress, Budapest, Hungary*. 8. 2010.

[3] BMW. *Intelligent Driving*. Accessed on: 2016-10-04. 2016. URL: http://www.bmw.com/com/en/insights/technology/connecteddrive/2013/driver_assistance/intelligent_driving.html.

[4] BMW. *Intelligent Parking*. Accessed on: 2016-10-04. 2016. URL: http://www.bmw.com/com/en/insights/technology/connecteddrive/2013/driver_assistance/intelligent_parking.html.

[5] Martin Böhner, Alexander Mattausch, and Alexander Much. "Extending Software Architectures from Safety to Security." In: *Automotive-Safety & Security*. 2014, pp. 41–51.

[6] Catalin Virgil Briciu and Ioan Filip. "The challenge of safety and security in automotive systems." In: *Applied Computational Intelligence and Informatics (SACI), 2014 IEEE 9th International Symposium on*. 2014, pp. 177–181. ISBN: 9781479946945. DOI: 10.1109/SACI.2014.6840056.

[7] Catalin Virgil Briciu, Ioan Filip, and Franz Heininger. "A new trend in automotive software: AUTOSAR concept." In: *SACI 2013 - 8th IEEE International Symposium on Applied Computational Intelligence and Informatics, Proceedings*. 2013, pp. 251–256. ISBN: 9781467364003. DOI: 10.1109/SACI.2013.6608977.

[8] Simon Burton, Jürgen Likkei, Priyamvadha Vembar, and Marko Wolf. "Automotive functional safety = safety + security." In: *Proceedings of the First International Conference on Security of Internet of Things - SecurIT '12*. ACM, 2012, pp. 150–159. ISBN: 9781450318228. DOI: 10.1145/2490428.2490449.

[9] Paul Carsten, Todd R. Andel, Mark Yampolskiy, and Jeffrey T. Mcdonald. "In-vehicle networks: Attacks, vulnerabilities, and proposed solutions." In: *Proceedings of the 10th Annual Cyber and Information Security Research Conference*. 2015. ISBN: 9781450333450.

[10] Stephen Checkoway, Damon Mccoy, Brian Kantor, Danny Anderson, Hovav Shacham, Stefan Savage, Karl Koscher, Alexei Czeskis, Franziska Roesner, and Tadayoshi Kohno. "Comprehensive Experimental Analyses of Automotive Attack Surfaces." In: *USENIX Security Symposium*. 2011.

[11] Daimler. *Keeping in lane*. Accessed on: 2016-10-04. 2016. URL: https://www.daimler.com/innovation/safety/special/keeping-in-lane.html.

[12] Daimler. *Parking safely*. Accessed on: 2016-10-04. 2016. URL: https://www.daimler.com/innovation/safety/special/parking-safely.html.

[13] ETSI. *ETSI TR 102 638: Intelligent Transportation Systems (ITS); Vehicular Communications; Basic Set of Applications; Definitions*. TR 102 638. Sophia Antipolis, France: European Telecommunications Standards Institute, 2009-06.

[14] ETSI. *ETSI TS 102 636-3: Intelligent Transportation Systems (ITS); Vehicular Communications; GeoNetworking; Part 3: Network Architecture*. TS 102 636. Sophia Antipolis, France: European Telecommunications Standards Institute, 2010-03.

[15] Richard Evans, Paul Groves, Katrin Hartwig, Edith Holland, Peter Jesty, Keith Longmore, Frank O'Neill, Roger Rivett, and David Ward. "MISRA Guidelines for Safety Analysis of Vehicle Based Programmable Systems." In: (2007).

[16] Michel Ferreira, Ricardo Fernandes, Hugo Conceição, Wantanee Viriyasitavat, and Ozan K Tonguz. "Self-organized traffic control." In: *Proceedings of the seventh ACM international workshop on VehiculAr InterNETworking*. ACM. 2010, pp. 85–90.

[17] FlexRay Consortium. *FlexRay communications system protocol specification version 3.0.1*. Tech. rep. 2010, pp. 1–341.

[18] PSA Group. *City Park: PSA Group's intelligent and scalable park assist system*. Accessed on: 2016-10-04. 2016. URL: https://www.groupe-psa.com/en/newsroom/automotive-innovation/city-park/.

[19] PSA Group. *Highway Chauffeur: autonomous driving on motorways*. Accessed on: 2016-10-04. 2016. URL: https://www.groupe-psa.com/en/newsroom/automotive-innovation/highway-chauffeur/.

[20] U.S. Software System Safety Working Group. *5th Meeting of the U.S. Software Systems Safety Working Group: Adaptive Cruise Control System Overview*. Workshop. Anaheim, CA United Stated of America: Massachusetts Institute of Technology, 2005. URL: http://sunnyday.mit.edu/safety-club/workshop5/Adaptive_Cruise_Control_Sys_Overview.pdf.

[21] (HIS) Security Working Group. *Secure Hardware Extension*. Deliverable Version 1.1. Herstellerinitiative Software, 2009-10.

[22] Olaf Henniger, Alastair Ruddle, Hervé Seudié, Benjamin Weyl, Marko Wolf, and Thomas Wollinger. "Securing vehicular onboard IT systems: The EVITA project." In: *VDI/VW Automotive Security Conference*. 2009.

[23] Tobias Hoppe, Stefan Kiltz, and Jana Dittmann. "Security threats to automotive CAN networks-practical examples and selected short-term countermeasures." In: *Computer Safety, Reliability, and Security*. 2008, pp. 235–2478. ISBN: 3540876979. DOI: 10.1016/j.ress.2010.06.026.

[24] ISO. *ISO 11898. Road vehicles–interchange of digital information–Controller Area Network (CAN) for high-speed communication*. ISO 11898:1993. Geneva, Switzerland: Internation Organization for Standardization, 1993.

[25] ISO. *ISO/IEC 61508. Functional Safety of Electrical/Electronic/Programmable Electronic Systems*. ISO 61508:1998. Geneva, Switzerland: Internation Organization for Standardization, 1998.

[26] ISO. *ISO 15408. Common Criteria for Information Technology Security Evaluation*. ISO 14508:1999. Geneva, Switzerland: Internation Organization for Standardization, 1999.

[27] ISO. *ISO 17458. Road vehicles – FlexRay communications system*. ISO 17458:2009. Geneva, Switzerland: Internation Organization for Standardization, 2009.

[28] ISO. *ISO 26262. Road vehicles – Functional safety*. ISO 26262:2012. Geneva, Switzerland: Internation Organization for Standardization, 2012.

[29] Joint Task Force Transformation Initiative. *SP 800-37 Rev. 1. Guide for Applying the Risk Management Framework to Federal Information Systems: A Security Life Cycle Approach*. Tech. rep. Gaithersburg, MD, United States, 2010.

[30] SAE International. *SAE J3061 Cybersecurity Guidebook for Cyber-Physical Vehicle Systems*. WIP J3061. Warrendale, PA United States of America: Society of Automotive Engineers, 2013.

[31] *Introduction to Microsoft Security Development Lifecycle (SDL) Threat Modeling*. Presentation Slides. 2005. URL: {http://download.microsoft.com/download/9/3/5/935520EC-D9E2-413E-BEA7-0B865A79B18C/Introduction_to_Threat_Modeling.ppsx}.

[32] Rob Millerb Ishtiaq Roufa, Hossen Mustafaa, Sangho Ohb Travis Taylora, Wenyuan Xua, Marco Gruteserb, Wade Trappeb, and Ivan Seskarb. "Security and privacy vulnerabilities of in-car wireless networks: A tire pressure monitoring system case study." In: *19th USENIX Security Symposium, Washington DC*. 2010, pp. 11–13.

[33] Mafijul Md Islam, Aljoscha Lautenbach, Christian Sandberg, and Tomas Olovsson. "A Risk Assessment Framework for Automotive Embedded Systems." In: *Proceedings of the 2nd ACM International Workshop on Cyber-Physical System Security*. ACM. 2016, pp. 3–14.

[34] Robert Janssen, Han Zwijnenberg, Iris Blankers, and Janiek de Kruijff. *Truck platooning: Driving the future of transportation*. White Paper TNO 2014 R11893. Delft, The Netherlands: TNO Mobility and Logistics, 2014.

[35] Barbara Kitchenham. "Procedures for performing systematic reviews." In: *Keele, UK, Keele University* 33.TR/SE-0401 (2004), p. 28. ISSN: 13537776. DOI: 10.1.1.122.3308.

[36] Pierre Kleberger, Tomas Olovsson, and Erland Jonsson. "Security aspects of the in-vehicle network in the connected car." In: *IEEE Intelligent Vehicles Symposium, Proceedings*. 2011, pp. 528–533. ISBN: 9781457708909. DOI: 10.1109/IVS.2011.5940525.

[37] Karl Koscher et al. "Experimental security analysis of a modern automobile." In: *Proceedings - IEEE Symposium on Security and Privacy*. 2010, pp. 447–462. ISBN: 9780769540351. DOI: 10.1109/SP.2010.34.

[38] VTL LLC. *Virtual Traffic Lights*. Accessed on: 2016-10-04. 2011. URL: http://www.virtualtrafficlights.com/index.html.

[39] Ralph Langner. "Stuxnet: Dissecting a cyberwarfare weapon." In: *IEEE Security & Privacy* 9.3 (2011), pp. 49–51.

[40] MAN. *Continuous Damping Control*. Accessed on: 2016-10-04. 2016. URL: http://www.truck.man.eu/de/en/man-world/technology-and-competence/safety-and-assistance-systems/continuous-damping-control/Continuous-Damping-Control.html.

[41] MAN. *Emergency Brake Assist*. Accessed on: 2016-10-04. 2016. URL: http://www.truck.man.eu/de/en/man-world/technology-and-competence/safety-and-assistance-systems/emergency-brake-assist/Emergency-Brake-Assist.html.

[42] MAN. *GPS-assisted cruise control*. Accessed on: 2016-10-04. 2016. URL: http://www.truck.man.eu/de/en/man-world/technology-and-competence/efficiency-systems/gps-assisted-cruise-control/GPS-assisted-Cruise-Control.html.

[43] MAN. *Lane guard system*. Accessed on: 2016-10-04. 2016. URL: http://www.truck.man.eu/de/en/man-world/technology-and-competence/safety-and-assistance-systems/lane-guard-system/Lane-guard-system.html.

[44] MAN. *Solutions*. Accessed on: 2016-10-04. 2016. URL: http://www.truck.man.eu/de/en/solutions/overview/man-solutions.html.

[45] Georg Macher, Harald Sporer, Reinhard Berlach, Eric Armengaud, and Christian Kreiner. "SAHARA: A Security-Aware Hazard and Risk Analysis Method." In: *Design, Automation & Test in Europe Conference & Exhibition*. 2015, pp. 621–624. ISBN: 9783981537048.

[46] Charlie McCarthy and Kevin Harnett. *National institute of standards and technology (nist) cybersecurity risk management framework applied to modern vehicles*. Report No. DOT HS 812 073. Washington, DC, United States: National Highway Traffic Safety Administration, 2014.

[47] Charlie McCarthy, Kevin Harnett, and Art Carter. *Characterization of Potential Security Threats in Modern Automobiles*. Report No. DOT HS 812 074. Washington, DC, United States: National Highway Traffic Safety Administration, 2014.

[48] Charlie Miller and Chris Valasek. "A Survey of Remote Automotive Attack Surfaces." In: *BlackHat USA* (2014).

[49] R. Moalla, B. Lonc, H. Labiod, and N. Simoni. "Towards a Cooperative ITS Vehicle Application Oriented Security Framework." In: *IEEE Intelligent Vehicles Symposium, Proceedings* (2014), pp. 1043–1048. DOI: 10.1109/IVS.2014.6856548.

[50] Dennis K. Nilsson, Ulf E. Larson, Francesco Picasso, and Erland Jonsson. "A first simulation of attacks in the automotive network communications protocol flexRay." In: *Proceedings of the International Workshop on Computational Intelligence in Security for Information Systems CISIS'08*. Vol. 53. 2009, pp. 84–91. ISBN: 9783540881803. DOI: 10.1007/978-3-540-88181-0_11.

[51] DAF Paccar. *Emergency Braking System*. Accessed on: 2016-10-04. 2016. URL: https://www.daf.com/en/products/euro-6-range/comfort-and-safety-systems/advanced-emergency-braking-system.

[52] DAF Paccar. *Lane Departure Warning System*. Accessed on: 2016-10-04. 2016. URL: https://www.daf.com/en/products/euro-6-range/comfort-and-safety-systems/lane-departure-warning-system.

[53] DAF Paccar. *Predictive Cruise Control*. Accessed on: 2016-10-04. 2016. URL: https://www.daf.com/en/products/euro-6-range/comfort-and-safety-systems/predictive-cruise-control.

[54] DAF Paccar. *Vehicle Stability Control*. Accessed on: 2016-10-04. 2016. URL: https://www.daf.com/en/products/euro-6-range/comfort-and-safety-systems/vehicle-stability-control.

[55] Panagiotis Papadimitratos, Levente Buttyan, Jean-Pierre Hubaux, Frank Kargl, Antonio Kung, and Maxim Raya. "Architecture for secure and private vehicular communications." In: *Telecommunications, 2007. ITST'07. 7th International Conference on ITS*. 2007, pp. 1–6.

[56] Jonathan Petit, B Stottelaar, M Feiri, and F Kargl. "Remote Attacks on Automated Vehicles Sensors: Experiments on Camera and LiDAR." In: *Blackhat Europe, Amsterdam, Netherlands* (2015).

[57] Juha Saarinen. *Students hijack luxury yacht with GPS spoofing*. Accessed on: 2016-07-04. 2013. URL: http://www.itnews.com.au/news/students-hijack-luxury-yacht-with-gps-spoofing-351659.

[58] Karsten Schmidt, Peter Tröger, Hans-Martin Kroll, Thomas Bünger, Florian Krueger, and Christian Neuhaus. "Adapted Development Process for Security in Networked Automotive Systems." In: *SAE International Journal of Passenger Cars - Electronic and Electrical Systems* 7.2 (2014), pp. 516–526. ISSN: 19464622 19464614. DOI: 10.4271/2014-01-0334.

[59] Christoph Schmittner, Zhendong Ma, Erwin Schoitsch, and Thomas Gruber. "A Case Study of FMVEA and CHASSIS as Safety and Security Co-Analysis Method for Automotive Cyber-physical Systems." In: *Proceedings of the 1st ACM Workshop on Cyber-Physical System Security - CPSS '15*. ACM, 2015, pp. 69–80. ISBN: 9781450334488. DOI: 10.1145/2732198.2732204.

[60] Adam Shostack. "Experiences threat modeling at Microsoft." In: *Modeling Security Workshop. Dept. of Computing, Lancaster University, UK*. 2008.

[61] IEEE Computer Society. *IEEE Standard for Information Technology - Telecommunications and Information Exchange Between Systems - Local and Metropolitan Area Networks - Specific Requirements, Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, Amendment 6: Wireless Access in Vehicular Environments*. IEEE std 802.11p-2010. New York, NY, USA: Institute of Electrical and Electronics Engineers, 2010.

[62] Till Steinbach, Hyung-Taek Lim, Franz Korf, Thomas C Schmidt, Daniel Herrscher, and Adam Wolisz. "Tomorrow's in-car interconnect? A competitive evaluation of IEEE 802.1 AVB and Time-Triggered Ethernet (AS6802)." In: *Vehicular Technology Conference (VTC Fall), 2012 IEEE*. IEEE. 2012, pp. 1–5.

[63] Aleksandar Stevanovic. *Adaptive traffic control systems: domestic and foreign state of practice*. Project 20-5 (Topic 40-03). 2010.

[64] Ivan Studnia, Vincent Nicomette, Eric Alata, Yves Deswarte, Mohamed Kaaniche, and Youssef Laarouchi. "Survey on security threats and protection mechanisms in embedded automotive networks." In: *Dependable Systems and Networks Workshop (DSN-W), 2013 43rd Annual IEEE/IFIP Conference on*. 2013, pp. 1–12. ISBN: 9781479901814. DOI: 10.1109/DSNW.2013.6615528.

[65] TCG Working Group. *Trusted Platform Module - Main Specification*. Deliverable Version 1.2. Trusted Computing Group, 2007-07.

[66] Shane Tuohy, Martin Glavin, Ciarán Hughes, Edward Jones, Mohan Trivedi, and Liam Kilmartin. "Intra-vehicle networks: A review." In: *IEEE Transactions on Intelligent Transportation Systems* 16.2 (2015), pp. 534–545.

[67] Morning News USA. *Tesla Announces Software Update Following Remote Hack*. Accessed on: 2016-10-04. 2016. URL: http://www.morningnewsusa.com/tesla-announces-software-update-following-remote-hack-23107036.html.

[68] Volkswagen. *Parking Steering Assistance*. Accessed on: 2016-10-04. 2016. URL: http://www.volkswagenag.com/content/vwcorp/content/en/innovation/driver_assistance/parking_steering_assistance.html.

[69] David Ward, Ireri Ibara, and Alastair Ruddle. "Threat Analysis and Risk Assessment in Automotive Cyber Security." In: *SAE International Journal of Passenger Cars-Electronic and Electrical Systems* 6.2 (2013), pp. 507–513. ISSN: 1946-4622. DOI: doi:10.4271/2013-01-1415.

[70] Roel J. Wieringa. *Design science methodology for information systems and software engineering*. Springer, 2014.

[71] Marko Wolf, André Weimerskirch, and Christof Paar. "Security in Automotive Bus Systems." In: *Proceedings of the Workshop on Embedded Security in Cars*. 2004, pp. 1–13.

[72] Tao Zhang, Helder Antunes, and Siddhartha Aggarwal. "Defending connected vehicles against malware: Challenges and a solution framework." In: *IEEE Internet of Things Journal* (2014), pp. 10–21. ISSN: 23274662. DOI: 10.1109/JIOT.2014.2302386.