CHALMERS

Applicability analysis of intrusion detection and prevention in automotive systems

Master's Thesis in Computer Systems and Networks

DANIEL FALLSTRAND

VIKTOR LINDSTRÖM

Department of Computer Science and Engineering
Chalmers University of Technology
Göteborg, Sweden 2015

Automotive IDPS
Applicability analysis of intrusion detection and prevention in automotive systems

DANIEL FALLSTRAND,
VIKTOR LINDSTRÖM

Examiner: ERLAND JONSSON

Chalmers University of Technology
Department of Computer Science and Engineering
SE-412 96 Göteborg
Sweden
Telephone + 46 (0)31-772 1000

Department of Computer Science and Engineering
Göteborg, Sweden June 2015

Abstract

The complexity of embedded automotive systems is rapidly increasing, both in terms of the number of ECUs (electronic control units), networks and the composite systems derived thereof. More electronic control in tandem with multiple novel interfaces for remote network communications with these control systems yield new security challenges. New networking capabilities might make the car a more viable target for malicious adversaries while more electronic control systems heighten the risks associated with a breach of security.

Several methods and measures for improving automotive security have been proposed, some of them tested and others purely theoretical. Among these are several calls for the design and development of intrusion detection/prevention systems (IDPS) for automotive applications.

We survey the state of the art in automotive security and attempt to establish a general model of the modern car. We proceed to evaluate the applicability of different intrusion detection and prevention methodologies in the context of embedded automotive systems.

We conclude that the diversity of automotive architectures make it difficult to produce a model of a complete car that is detailed yet generalizable and sufficient. Hence, barring any major consolidation of high level architectural principles, any successful automotive IDPS will have to be based on common ground at a lower level, e.g. the CAN or FlexRay buses.

Acknowledgements

We want to extend our thanks to our supervisors; David Lenntoft at Cybercom and Tomas Olovsson at Chalmers University of Technology. We would also like to thank Henrik Hallgren and Tad Heppner at Cybercom for the opportunity to carry out this thesis and the guidance along the way. Last but not least we want to thank Erland Jonsson, at Chalmers University of Technology, for being our examiner.

The Authors, Göteborg 4/6/15

Contents

1

Introduction

Embedded systems are becoming increasingly prevalent and important in the modern world and networking of embedded systems is also growing. Embedding computers and networking capabilities to products or systems has the potential to lower costs, increase performance and provide new functionality. The Internet of Things is a term that is now widely used to refer to the many machines and devices with embedded computers that interconnect in some way with other devices, vendors, services etc. over the Internet. A substantial amount of the embedded devices are used to monitor and control electro-mechanical systems. Examples of such embedded systems can be seen for example in the "smart home" or in modern cars.

These expansions of the domain of computers give rise to a number of computer security and network security concerns. Researchers have identified numerous security problems and vulnerabilities in automotive systems that are related to or even caused by (network-connected) embedded computer systems [1, 2, 3, 4]. Similar problems can certainly be found in other areas where networked computing has been rapidly introduced and where security has not been considered a primary factor in the design.

Automotive systems are safety-critical and protecting the driver and passengers from harm is central to the design and production of vehicles. The safety criticality coupled with the vast complexity of embedded automotive systems, a modern car runs millions of lines of code on around 70 special-purpose hardware units that in some cases physically control the car, means that addressing the computer/network security issues is highly important.

Previous research efforts have mainly been focused on identifying vulnerabilities or developing proof-of-concept exploits to highlight security concerns in the automotive setting. Surveying the current state of automotive security is important in order to produce a basis for research into protection mechanisms and identification of high-level architectural problems that might cause security problems. As of yet, few protection mechanisms have been described and evaluated in an academic context. Suggested so-

lutions are generally presented briefly in future work-sections of current research.

One commonly suggested protection mechanism is intrusion detection and prevention systems (IDPS), and our contribution consists of an applicability analysis of the employment of this type of system in an embedded automotive setting. Other suggested protection mechanisms are described and discussed briefly in Sections 2.2.5 and 5.2, but the scope of our thesis does not allow for in-depth evaluation of them. Brief descriptions of the alternatives are presented mainly to facilitate a comparison with IDPS solutions. The main reason for our focus on IDPS is the fact that they could be designed and deployed without major changes to existing hardware and software. Additionally, previous research shows that IDPSs have the potential to increase security in automotive systems [4].

Automotive systems differ greatly from the context in which IDPSs are generally developed and deployed. Network protocols and hardware platforms are not the same in cars as in more conventional PC/server networks. In addition, some considerations regarding e.g. cost, hardware performance, maintainability and service life are dependent on the setting in which a system is to be deployed.

In the conventional setting, diverging approaches to detecting and preventing or mitigating security breaches have been presented. These different approaches or design decisions are evaluated and discussed in order to create a theoretical basis for future practical implementation, assessment, testing and comparison of different types of IDPS methodologies in the context of automotive systems. Our work is general and not based on or aimed at any particular make, model or type of car.

Intrusion detection and prevention systems are commonly grouped or categorized based on four different properties [5]:

- **Scope** - What entity or entities does the system protect?
- **Location and Distribution** - Where and how are the system components deployed?
- **Detection method** - How does the system identify intrusions?
- **Post-detection** - How does the system respond to detected intrusions?

Analyzing different approaches and answering these questions in the automotive setting provides a high-level description of an ideal automotive IDPS.

Once this high-level description of our suggested system has been presented we proceed to discuss and analyze it further. It is important to identify the limitations associated with this particular setup and to review the IDPS as a whole. Next, we theoretically discuss alternative security measures, such as message cryptography and improved network architecture, and assess their potential benefits and weaknesses in comparison with an IDPS solution. In addition to this comparative analysis, we also discuss some more general questions related to computer/network security in the context of automotive systems.

2

Background

Below we describe the history and development of the two major concepts that this thesis is concerned with. Initially we outline the history of intrusion detection and prevention systems (IDPS) thereafter we go through the general characteristics. Second, the evolution of the car from a purely mechanical machine to the highly computerized systems of today is detailed.

2.1 IDPS

Protecting confidential data from unauthorized access has become an important task for most computer systems. Different types of tools have emerged that assist security officers and the like in the task of preventing attacks. Anti-virus systems and firewalls are examples of two such protection mechanisms that are now commonplace in most computers/networks. Another type of tool is intrusion detection and prevention systems (IDS/IPS/IDPS), these are less widespread and functionality usually overlaps with the two previously mentioned types of tools. The scope and objectives of intrusion detection is wider and more varied, and most modern intrusion detection systems also provide some mechanism for preventing the detected intrusions.

Below is a brief description of the history of IDPS, next characterization of IDPSs is described and concepts such as scope, location and distribution, detection method and post-detection are introduced. Finally, we outline the expansion of IDPS methods into novel domains and the new challenges that arises with it.

2.1.1 History

The foundation of Intrusion Detection research was developed in the early 1980s in a paper titled *Computer Threat Monitoring and Surveillance* [6]. Pathan et al. states "Audit trails provide valuable information that can be used for tracking misuses of information

systems and understanding user's behavior. This research work had established the foundation for the later design and development of intrusion detection systems (IDS)" [5].

Following that report, the first implementation of an intrusion detection system was built by SRI in 1984, the Intrusion Detection Expert System (IDES) [7]. A further development from SRI was made by creating a well-defined model of an IDES [8]. The model has served as a good basis for subsequent research within the field.

The next big leap made was the Haystack [9] System. This intrusion detection system was aimed at a multi-user system and for use in the American air force. The authors of Haystack improved the IDS by leveraging a distributed approach. Shortly after this, Haystack Labs released the first host-based IDS, Stalker. Stalker was able to determine rules automatically or manually through setup in form of a set of questions.

The first network intrusion detection system (NIDS) was described and discussed by Heberlein et al. in the early 1990's [10]. Before then IDSs were usually built for use in big mainframes or specific hosts in a network. The ability to instead analyze network traffic meant that the intrusion detection could be done regardless of operating system and setup of the individual hosts.

2.1.2 Categorization

As mentioned in the introduction, IDPSs are commonly categorized based on scope, location and distribution, detection method and post-detection. The most common categories derived from these 4 properties are detailed below.

Scope

IDPS systems can be categorized by their scope. Pathan et al. states that "In terms of the analyzed and monitored source of information, an IDS can be classified based on scope into host-based intrusion detection systems, network-based intrusion detection systems, or hybrid intrusion detection systems" [5]. **Host-based** IDS (HIDS) monitors a machine, **network-based** IDS (NIDS) monitors a network and **hybrid** systems were conceived to maximize the strengths of both HIDS and NIDS [5].

Location and Distribution

Depending on what the purpose is, different strategies on how and where to deploy the IDPS can be employed. Three general ways of modelling the placement strategies are as centralized, distributed or hierarchical systems. The centralized model consists of a single manager that performs all of the analysis. The distributed model on the other hand is a model where the IDPS is divided in to several different parts which cooperate. The hierarchical model consists of several different units cooperating with some hierarchical division of detection and/or prevention tasks [5].

Detection method

IDPSs are often divided into two main categories based on their detection characteristics, **misuse detection** and **anomaly detection**. Misuse detection is an approach where each suspected attack is compared to a set of known attack signatures [5]. If the IDPS matches an attack to a signature, the IDPS will log the event and depending on how the IDPS is set up it may take appropriate measures to prevent or mitigate the attack. In order to take advantage of misuse detection it is essential that a database of previously known attacks exists. It is exclusively the attacks in that database that can be detected, this detection method does not allow for detection of unknown/new attacks. Anomaly detection can be preferable when strict protocols are in place, since unconformity with protocol specifications can be used as a basis for signatures.

In order to mitigate against previously unknown attacks another method has to be used, i.e. anomaly detection. This detection method focuses on identifying anomalous behaviour rather than known misuse. A somewhat contrived comparison is that while misuse detection is a blacklist approach, meaning some behaviour is defined as malicious or bad, anomaly detection is based on a whitelist approach, where some behaviour is defined as normal and events or behaviour that differ from this baseline of normal behaviour is identified. The anomaly detection can be performed in a variety of ways but is often achieved by some kind of statistical analysis. Common to all anomaly detection systems is the need to create a baseline of ”normal behaviour” with which to compare data and/or network traffic.

Due to the statistical nature of anomaly detection, cases such as false positive, false negatives (and their negations; true positive and true negative) occur. This means that even if a suspected breach is reported, there is a possibility that it is not malicious (false positive). It is also possible that an attack conforms to the baseline and is not detected (false negative). Depending on for example how the anomaly detection is performed and how well the baseline describes the normal behaviour of the system, the number of false alarms and false positives will vary [5].

In 2001, E. Biermann et al. made a comparison between the main types of intrusion detection systems [11]. It can be observed that there are clear advantages associated with both misuse detection and anomaly detection. Misuse detection is accurate and efficient, but cannot detect new attacks and the signature database has to be kept up to date. Anomaly detection can detect new attacks and does not require updating, but it can produce false positives. Because of their different characteristics it is often proposed to have a solution that combines both misuse detection and anomaly detection.

One solution to the high rate of false positives produced by anomaly detection can be to use a specification-based anomaly detection method, as described by Sekar et al. [12]. One could view specification-based approaches as a third type of detection method, but here we consider it a type of anomaly detection because ”Specification-based techniques are similar to anomaly detection in that they also detect attacks as deviations from a norm” [12]. Sekar et al. describes this method as ”instead of relying on machine learning techniques, specification-based approaches are based on manually developed specifications that capture legitimate (rather than previously seen) system

behaviors" [12]. This makes the detection static (whereas methods based on e.g. machine learning are dynamic) and more transparent.

Post-detection

Another property upon which classification of IDPSs can be based on is the action taken in response to suspected intrusions, this is called post-detection. Post-detection can be classified into two different classes, **active** actions and **passive** actions [5]. Systems that use active actions are often called Intrusion Prevention System (IPS) or Intrusion Detection and Prevention System (IDPS). How the prevention is done differs widely depending on what the purpose of the IDPS is. Passive action is where the post-detection response is to only notify about the suspected breach. The notification can be issued in logs, emails or something similar to call for attention to the breach.

Active post-detection means that the IDPS actively takes steps to protect the threatened entity or entities, one example mentioned by Pathan et al. is that "the IPS can interact with edge routers (routers located at the perimeter between the local networks of an organization and external networks to the Internet) or firewalls and establish rules to filter certain packet flows originating from some specific source addresses" [5].

For a more detailed history and characterization of intrusion detection and prevention systems, see The State of the Art in Intrusion Prevention and Detection by Pathan et al. [5].

2.1.3 Novel domains and contexts

In recent years, we have seen the emergence of new types and domains of networked computing. In general, this evolution has been related to the increased networking and connectivity of embedded devices. Some examples of such domains include the Internet of Things, wireless sensor networks, mobile devices and the connected car. All of these new domains present new security challenges and properties that are unique to their respective settings. Different types of hardware, software and limitations imposed by for example embedded devices affect the way effective protection mechanisms are implemented.

This development has prompted researchers to investigate the applicability of intrusion detection and prevention methodologies in these novel domains. For example, IDPS in wireless sensor networks was discussed as early as 2004-2005 [13, 14]. Another example of IDPS applied in a new context is the SVELTE system for the Internet of Things proposed by Swedish researchers in 2013 [15].

2.2 Automotive systems

The modern car contains a complex network of electronic components that perform various tasks and provide functionality ranging from safety systems to infotainment. However, this influx of electronic control units and networking is quite recent. Below we describe the history of car electronics and the shift from almost completely mechanical

vehicles to vastly complex systems of embedded computers.

We proceed to describe the constituting entities or components of an automotive embedded systems architecture. This description covers a range of components commonly found in modern cars, but it does not completely match any real-world system. A modern car might have additional ECUs, some ECUs might be combined into one, etc. Nonetheless, this description provides a basis for analysis and discussion that is free from the limitations imposed by deriving conclusions from a particular system (or set of systems).

Next, safety criticality in automotive applications is described to show the context in which car computer- and network security considerations are made. Also presented below is a number of known attacks and vulnerabilities. A large portion of these are academic rather than examples of attacks in the wild, since real-world attacks have been, and currently are, scarce. This scarcity and its implications on research and implementation of automotive security is discussed in Section 5.1.3 and 5.3.4.

2.2.1 History

The first steps towards the embedded systems of modern cars were taken in the late 1970s with the introduction of the Engine Control Unit. This unit was used to increase the efficiency of the engine, reduce pollutants and it could also provide diagnostic information [16]. Digital control systems quickly spread to other parts of the car and the meaning of the acronym ECU changed from *Engine Control Unit* to *Electronic Control Unit*. [1]

In addition to providing improved control and diagnostic abilities, embedded computer systems facilitated the development and widespread use of advanced safety systems such as anti-lock brake systems (ABS) and electronic stability control (ESC).

An important factor in the evolution of embedded car electronics was the introduction of networking and communication buses. The controller area network (CAN) was introduced in the early 1980s and since then the use of networks has both reduced the need for dedicated cabling and enabled applications that input and output data to and from several different ECUs and sensors. Several other networking/bus standards aimed at different uses have been published and used since then, including FlexRay, LIN and MOST.

In recent years, novel connections and networks have been making their way into the automotive ecosystem, e.g. GPS, Bluetooth, 3G and WiFi. Navigation systems need to communicate with satellites, pairing a smartphone with in-car entertainment systems allow for hands-free calling and music playback, telematics systems can use 3G to send information about the car to the supplier etc. Contrary to the previously mentioned network protocols (CAN, FlexRay etc.) these new networks are not (primarily) used for communication inside the car but rather for communication with external services, devices etc. GPS, Bluetooth, 3G and WiFi introduce remote communication endpoints to the system.

2.2.2 Components and entities

Embedded automotive systems consist of ECUs interconnected with a number of networks or buses. The ECUs often rely on sensors and actuators to perform their various functions.

At a higher level, one can also identify a number of larger composite systems that are comprised of a subset of the ECUs, sensors and actuators of the car. E.g. a modern parking assist system needs access to data from cameras and proximity sensors and will both actuate the steering wheel and give the driver instructions in the instrument cluster.

Electronic Control Units

Contemporary cars contain somewhere in the order of 100 ECUs [17, 18] providing various functionality and controlling everything from engine operation to infotainment systems. Estimates of the number of lines of code in a car today range from tens to hundreds of millions [16, 18].

ECUs are generally produced and sold by subcontractors to the actual vehicle manufacturers. The code in the ECUs is oftentimes regarded as important intellectual property of suppliers and not made available to the vehicle manufacturers. Security problems arising from the integration of modules from different suppliers are discussed thoroughly by Amin and Tariq in their article titled *Securing the Car: How Intrusive Manufacturer-Supplier Approaches Can Reduce Cybersecurity Vulnerabilities* [19].

We will not delve into specifics regarding certain ECUs or types of ECUs here as it is sufficient for our purposes to conclude that electronic control units in cars are diverse and complex special purpose components whose production is often outsourced. Some more information on types of ECUs can be found in the article titled *Experimental security analysis of a modern automobile* by Koscher et al. [1].

Networks/buses

As summarized in Table 2.1 by Wolf et al. [20], different bus systems have different characteristics and are designed for different applications in the car. The table does not cover all of the protocols listed in Section 2.2.1 above but aptly shows some of the major differences between the different types of networks.

Table 2.1: Properties of selected automotive bus systems [20]

Bus	LIN	CAN	FlexRay	MOST	Bluetooth
Adapted For	Low-level Subnets	Soft Real-Time	Hard Real-Time	Multimedia Telematics	External Communication
Target Application Examples	Door locking Climate regulation Power windows Light, rain sensor	Antilock brake system Driving assistants Engine control Electronic gear box	Brake-by-Wire Steer-by-Wire Shift-by-Wire Emergency systems	Entertainment Navigation Information services Mobile Office	Telematics Electronic toll Internet Telediagnosis
Architecture	Single-Master	Multi-Master	Multi-Master	Multi-Master	Multi-Master
Access Control	Polling	CSMA/CA	TDMA FTDMA	TDM CSMA/CA	TDMA TDD
Transfer Mode	Synchronous	Asynchronous	Synchronous Asynchronous	Synchronous Asynchronous	Synchronous Asynchronous
Data Rate	20 kBit/s	1 MBit/s	10 MBit/s	24 MBit/s	720 kBit/s
Redundancy	None	None	2 Channels	None	79 Frequencies
Error Protection	Checksum Parity bits	CRC Parity bits	CRC Bus Guardian	CRC System Service	CRC FEC
Physical Layer	Single-Wire	Dual-Wire	Optical Fiber Dual-Wire	Optical Fiber	Air

CAN is currently the most common bus standard employed and it is often the main focus of discussions on automotive networks. The CAN protocol does not use an addressing scheme, instead packets are broadcast to all nodes on the bus. The real-time guarantees of CAN make it useful in time-sensitive applications.

FlexRay provides even stronger real-time guarantees than CAN and can be used in the same types of applications as CAN with higher speeds and reliability. Most safety critical parts of an automotive system are also time-sensitive and therefore connected to CAN or FlexRay buses.

In addition to the networks used for internal communication, a growing number of networks allowing remote communication can be found in modern cars. Protocols such as Bluetooth, 3G/4G, and WiFi are used in for example telematics and infotainment applications. As previously stated, remote connectivity is increasing and Vehicle-to-X (V2X) technologies is one example of systems that drive this development.

To summarize, an automotive system can contain a wide variety of networks and networking protocols, some of which are very different from the Ethernet/TCP/IP networks that are commonly found in conventional computer networks.

Cyber-physical systems

Edward A. Lee presents the following definition in a 2008 article: ”Cyber-Physical Systems (CPS) are integrations of computation and physical processes. Embedded computers and networks monitor and control the physical processes [...]” [21]. This sweeping definition covers everything from industrial control systems to smart coffeemakers, but it does pinpoint one very important property of some modern computer systems. The fact that CPSs have direct physical effect(s) on the world around them has a radical

effect on threat models, security measures and verification and validation-procedures.

Lee also states that "Transportation systems could benefit considerably from better embedded intelligence in automobiles, which could improve safety and efficiency". While it is certainly true that electronic control has and will continue to improve safety and efficiency of automotive systems, it is important to note that electronics and computers introduce new failure modes and scenarios.

Further discussion on the implications of computer systems with cyber-physical effects and the connection to safety-criticality is described in Section 2.2.3 below.

2.2.3 Safety-criticality

The term safety-critical is usually defined as follows: "Safety-critical systems are those systems whose failure could result in loss of life, significant property damage, or damage to the environment" [22]. Designing, producing and maintaining systems of this kind presents a number of challenges and increases the need for specification, testing and verification. Many different types of systems could be termed safety-critical, e.g. modern aircraft, nuclear plants and computerized medical equipment.

Clearly, failure of a car or individual systems within it could result in both loss of life and property damage. Hence, the car is a safety-critical system. As the degree of computer control in automotive systems increase, so does the need to ensure that the control systems are safe and secure. The introduction of external and wireless connectivity to cars further complicates matters by introducing a new potential source of faults and failures: malicious adversaries attacking the car remotely.

Failures of different components/ECUs/electronic systems in a car can impact safety in a range of different ways. One example where a breach might cause direct harm is electronic control over steering and throttle. As of today, electronic control systems, so-called drive-by-wire systems [23], are consistently backed up by mechanical linkages for safety reasons. However, sometime in the future this might not be the case [23]. Several of the attacks against automotive systems presented by Miller and Valasek [3, 4] actually have direct cyber-physical effects on the car, e.g. alteration of steering angle or engaging/disengaging the brakes.

Many of the ECUs and systems in modern vehicles such as ESC, ABS, etc. were developed with the intent to improve driver safety. Decreased functionality or loss of these systems will decrease safety of the driver and passengers to varying degrees depending on what device(s) are affected. If a malicious user should gain access to these systems, not only is the integrity of the system in danger but also the safety of the driver and the passengers. A situation where a safety system is disabled might not seem as hazardous as the cyber-physical attacks described above but could also have serious consequences.

There are other devices such as multimedia and infotainment systems that can also affect the safety of vehicle occupants but in a more subtle manner. Auditory distractions from the stereo or horn is one potential hazard, one could even imagine safety implications due to an attacker disabling the electronic window controls while all windows are rolled down.

2.2.4 Attacks

In order to discuss car security, we must discuss the threats and risks associated with the domain. Formal definitions of terms including risk, threat and related concepts can be found in the *SANS Glossary of Security Terms* [24]. However, for our purposes it is sufficient to say that risk is the combination of an outcome and its likelihood of occurring, a threat is a potential violation of security and an attack is an attempt to violate the confidentiality, integrity or availability of a system.

The intersection of classic computer security theory and the risk management and fault-tolerance fields when dealing with security and safety in cars has the potential of creating confusion among engineers with different backgrounds. It is therefore important to establish a terminology that is sufficient to specify and analyze systems of this kind in several different contexts. However, this problem is beyond the scope of this thesis (see 5.4 for more information on this).

The bulk of early work showing real world attacks against car computer systems was conducted and presented by a group of scientists from the University of California San Diego and the University of Washington (Checkoway et al. and Koscher et al.) [1, 2]. More recently Charlie Miller and Chris Valasek [3, 4] have presented attacks against several different car models, with a special focus on attacks that can be performed remotely, i.e. without physical access to the car. In addition to validating the results of Checkoway et al. the research by Miller and Valasek garnered a lot of attention from mainstream media and increased public awareness of the issues and dangers presented by the growing complexity of computer systems embedded in cars.

Miller and Valasek describe ”the anatomy of a remote attack” at length in their second article, their general definition of the term remote attack, consisting of a three stage process summarized below, is the one used in this report unless otherwise stated. A remote attack with this definition consists of three stages or phases: ”The first stage consists of an attacker remotely gaining access to an internal automotive network [...] a cyber physical attack usually requires a second step which involves injecting messages onto the internal automotive network in an attempt to communicate with safety critical ECUs, such as those responsible for steering, braking, and acceleration [...] The final step is to make the target ECU behave in some way that compromises vehicle safety.” [4].

The remote attacks described by Miller and Valasek are certainly highly critical, as they are designed to yield physical effects that compromise vehicle safety. It is however worth noting that this class of attacks is not the only thing to consider when it comes to car security. A host of other classes of attacks should also be examined and their potential severity or risk level evaluated. Note that the classes or types mentioned below are not necessarily mutually exclusive.

Denial-of-Service (DoS) attacks is a category of attacks targeting the availability of a system. Some examples of DoS attacks are described by Koscher et al. [1] and even in the instances where the attacks do not threaten safety directly, they can be highly problematic. E.g. ”Thus, we were able to easily prevent a car from turning on. We were also able to prevent the car from being turned off.” [1].

Another conceivable class of attacks is information stealing attacks. Access to a bus

or the internal state of one or more ECUs could provide attackers with sources of interesting data that should be confidential. GPS data, basic diagnostic information and the status of the infotainment system are some examples of data that could be transferred to an attacker. This data could then be used to for example monitor driving habits or by a burglar to identify times when the owner is not at home. Furthermore, an attacker could potentially target car makers and subcontractors rather than the car owner, for example by attempting to extract and steal proprietary software from ECUs.

Similar to the more conventional computer security scenario, an attack can start in a non-critical part of the system and then spread or escalate. Thus, it is important to consider the architecture of the system and the separation of different ECUs or buses from one another. This is done to some degree in most modern cars, but as shown by Koscher et al., gateways between different buses that should provide this separation can sometimes be bypassed or reprogrammed to effectively remove this obstacle for the attacker. The process of bypassing gateways and moving from one bus to another is sometimes called **bus-hopping**.

Examples of practically demonstrated attacks of different types (either by the researchers at the University of California San Diego and the University of Washington [1, 2] or by Miller and Valasek [3, 4]) can be seen in table 2.2 below.

Type	**Example**
Remote attacks	Checkoway et al. were able to utilize Bluetooth and Tire Pressure Monitoring System (TPMS) in order to gain access to the CAN bus and send arbitrary messages [2]
Cyber physical	Kosher et al. were able to both prevent braking and engage brakes using message injection on the CAN bus [1]
Denial-of-Service	By flooding the CAN network Miller and Valasek were able to render the network inoperable [3]
Bus-hopping	By having access to one network Kosher et al. were able to access a neighbouring network by re-programming a gateway [1]

Table 2.2: Examples of some attack types

2.2.5 Defenses

It is clear that input from both security researchers and the automotive industry is crucial when proposing or developing protection mechanisms and security solutions for use in vehicles. Checkoway et al. draws the conclusion that ”Developing security solutions

compatible with the automotive ecosystem is challenging and we believe it will require more engagement between the computer security community and automotive manufacturers (in the same way that our community engages directly with the makers of PC software today)" [2].

As previously stated, several protection mechanisms and tools have been proposed and some have even been tested and/or deployed. Suggested solutions include:

- Message cryptography (encryption and/or signatures).
- Cryptographically signed software.
- Packet filters or firewalls in bus gateways.
- Architectural measures such as improved network segmentation.
- Securing remote endpoints.
- Intrusion detection and prevention systems.

More on these ideas can be found in the articles by Miller and Valasek [4], Amin and Tariq [19] and Studnia et al. [17]. We discuss the suitability and some advantages and disadvantages of these different solutions in Section 5.2.

3

Method

The overall goal of this project was to produce a suggestion of an IDPS solution that is suitable for protecting the embedded automotive systems. The IDPS solution should be general, i.e. not constructed for a particular car, and its design and integration should aim to minimize the need for alterations of the target vehicles. In order to do this, first a model of the systems to be protected had to be established. This model would serve as a basis for the design decisions and the analysis of different intrusion detection methodologies and prevention strategies.

A number of key design decisions were identified and researched (discussed in Section 3.3 below), these where then analyzed in the context of the model produced. Additionally, a number of factors not directly covered by the model were considered in order to improve the analysis of our proposed IDPS architecture, see Section 3.2. The analysis was theoretical, but should serve as a good basis for practical tests of the different parts of the IDPS design, e.g. comparative tests of the performance of different detection methods.

3.1 Modelling

In order to make any sort of assessment of protective measures, the context in which it operates needs to be specified and modelled in some way. The modelling efforts of this project were not aimed at producing a mathematically rigorous theoretical model, but rather at constructing a conceptual image of a general in-car network (composed of several subnets or buses). The key issues were deducing what system components or entities are present in the systems in question and how they are interconnected.

Some of the components have already been enumerated in Section 2.2.2, but early on the large variations from vehicle to vehicle were discovered and the need for some sort of consolidation of the different architectures identified. Also, components present in different makes and models of cars with similar but not identical responsibilities and

functions needed to be described in some unified way.

The final and most important part of the modelling was to produce a map or blueprint of a generic automotive system, based on the identified components and research into the networking and buses used. The different bus standards and networking capabilities described in Section 2.2.2 have different intended uses and properties. The separation of ECUs on different networks also had to be evaluated and described in order to yield an accurate picture of the system.

3.2 Additional considerations

In addition to the basic architecture of the car, as described by the model (see Section 3.1), a number of other factors were taken into account when analyzing the different aspects of the IDPS design. These secondary factors include considerations that are related to the real world context in which the proposed IDPS systems will be deployed and used.

One important point was **economy** and price, as being economically viable is essential to ensure that a system actually gets implemented. E.g. a system that requires massive redesigns of existing hardware or software incurs additional costs and will hence be less appealing to car manufacturers (or car owners in case the IDPS solution is implemented as an aftermarket product).

Complexity was another aspect that was reviewed. The reason for this was twofold, first, price tends to increase with complexity and second, simplicity makes a security system easier to audit and evaluate.

Finally we looked at **hardware-imposed constraints**, the ECUs and networks embedded in cars are not comparable to e.g. modern desktop computers interconnected with high-speed Ethernet connections. Hence, strategies might be successful in other contexts could prove to be too resource-intensive to be deployed on current hardware.

3.3 Analysis of IDPS properties

Based on the derived model of an embedded automotive network and the considerations listed above, a design was proposed. The suggested system design consists of answers to a number of questions, decisions and methods to be used (many of which are described in Section 2.1.2 above). The method used to analyze and answer these questions is described in the following four sections (3.3.1 through 3.3.4).

It is important to note that the design decisions discussed in this thesis do not exist in a vacuum, completely decoupled from one another. E.g. a certain detection method may be the best choice in conjunction with host-based system while being a poor choice in a network-based system. Consequently, after considering a design decision it is important to analyze its effect on the IDPS as a whole and other decisions.

3.3.1 Scope

As previously stated, categorization of IDPSs can be done based on the scope. It was important for this project to establish whether a **host-based** or **network-based** solution would be most suitable for automotive applications. The two scopes have various benefits and disadvantages, some of which have different consequences in the automotive setting than in the more conventional computer networks.

In addition, the effect of scope on economy and complexity as well as the significance of current hardware capabilities was evaluated.

3.3.2 Location and Distribution

The question of placement was examined and the advantages associated with distributing the components of the intrusion detection/prevention system were evaluated.

Location and distribution has a profound effect on both the complexity and the price of a system. Hence, these factors were considered very important in the process of determining what model (**distributed**, **centralized** or **hierarchical**) is preferable.

As mentioned above, the design decisions affect one another, and scope, location and distribution are closely related. This means that in addition to analyzing each of these properties by themselves, the different combinations of them have to be considered.

3.3.3 Detection method

The detection is the essential task performed by the proposed system. It was crucial to identify an effective methodology for detecting intrusions that was also viable in light of other factors such as price and hardware-constraints.

The intended outcome of this effort was to determine which method is best suited for embedded automotive systems, **anomaly detection** or **misuse detection**. A secondary aim was to examine whether certain types of the chosen general method (anomaly- or misuse detection) were superior or preferable.

3.3.4 Post-detection

Different responses to detected intrusion attempts or attacks were discussed and evaluated. This entailed the enumeration and choice between possible **active** and **passive** actions.

Safety is the number one priority, and the post-detection action has to minimize the safety implications of detected intrusions or attacks. This is not a trivial task, but rather a complex question that was further complicated by the intention to provide general answers for a generic system. Extra attention was therefore paid to avoiding assumptions that were not generalizable.

4

Proposed system

In this section we present and briefly describe the results of our efforts. Section 4.1 covers the model of automotive systems upon which our proposed IDPS system is based. In Section 4.2, answers to the four major IDPS design decisions are provided. The proposed system design is analyzed and discussed further in Chapter 5.

4.1 Modelling

Our attempts at creating a model of the modern car in the shape of a map of interconnected components failed. Cars are simply too architecturally diverse to be covered by one simple model of this kind. Someone with more background knowledge about embedded automotive systems might have reached this conclusion earlier than we did, but as we gained insight into the systems at hand we soon realized the futility of our efforts to produce a general high level model or map. Furthermore, whilst studying the results of Miller and Valasek we saw that their *Survey of Remote Automotive Attack Surfaces* [4] contained simple diagrams of parts of the networks and ECUs in the examined vehicles. Miller and Valasek present one diagram per car and by simply reviewing these examples the heterogeneity of the systems was made abundantly clear.

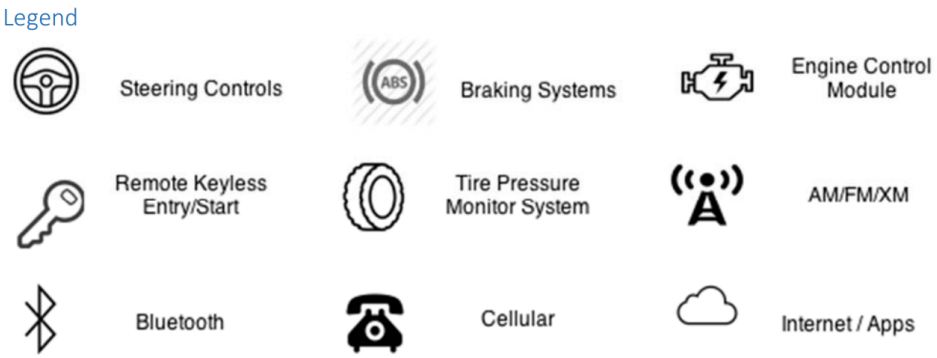

Figure 4.1: Legend [4], reprinted with permission from Miller and Valasek

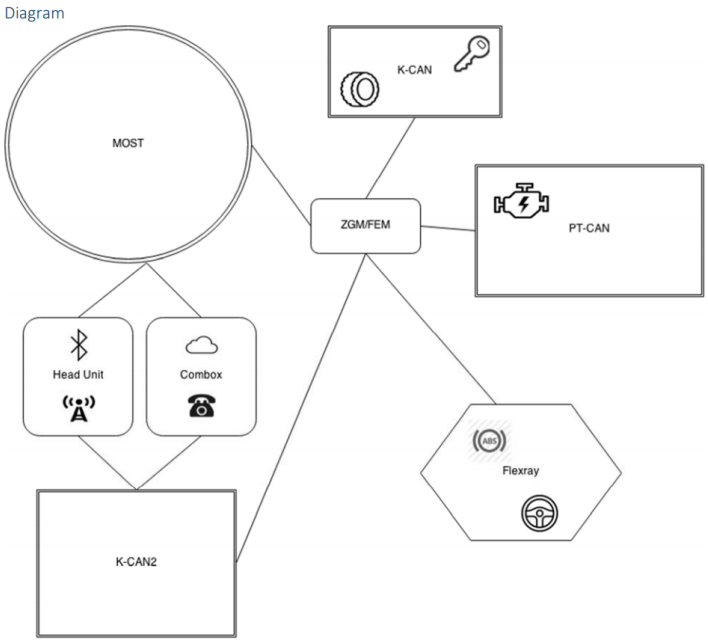

Figure 4.2: Diagram - 2014 BMW 3 Series (F30) [4], reprinted with permission from Miller and Valasek

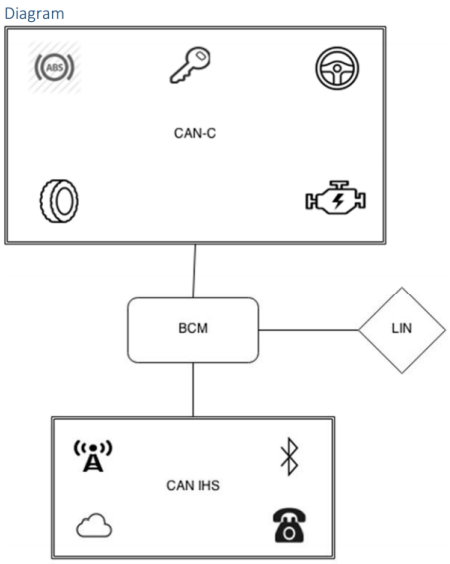

Figure 4.3: Diagram - 2014 Chrysler 300 [4], reprinted with permission from Miller and Valasek

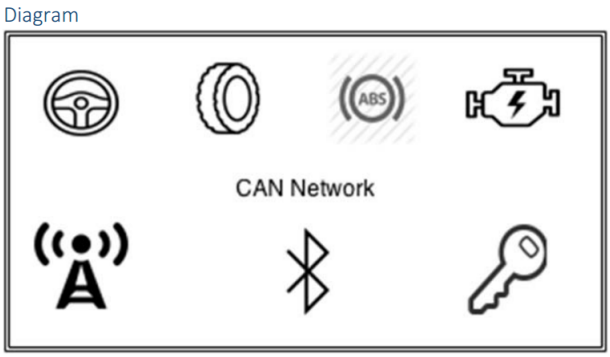

Figure 4.4: Diagram - 2010 Infiniti G37 (Sedan) [4], reprinted with permission from Miller and Valasek

Figures 4.2, 4.3 and 4.4 contain three of the diagrams mentioned above, several more are included in the original paper. The diagrams were created from wiring diagrams for the different cars and comparing them shows that the differences in network topology is huge. Figure 4.1 is a description of the various symbols used in the topology diagrams.

Another concern we identified was the dissimilarities between automotive systems and the conventional networks of PCs, servers, routers etc. where intrusion detection and prevention has been deployed previously. Almost all conventional networks use TCP/UDP over IP and the properties of these protocols are used in IDPSs. In compar-

ison, the networking in the car is very different, rather than addressed unicast internal automotive networks almost exclusively broadcast messages without specifying sender and receiver. As mentioned above, most modern cars also contain a number of different buses adhering to different standards potentially with some ECUs acting as gateways between two or more buses. Security features that are ubiquitous in the conventional setting are not present in automotive systems. These dissimilarities necessitate (re)evaluation of (sometimes implicit) assumptions and conclusions drawn in previous descriptions and discussions of intrusion detection and prevention.

In order to find a general IDPS solution we had to identify some lower level common denominator to base our reasoning on. Two potential candidates were identified, the bus or the ECU. Either a solution would be based around protecting an individual bus or an individual ECU, as most larger subsystems are vehicle specific and not general.

4.2 Analysis of IDPS properties

The following sections describe what we consider to be the best approach for designing an IDPS for embedded automotive applications based on the considerations described in Section 3.3.

4.2.1 Scope

The choice between a host-based or network-based system is closely related to the results of our modelling effort. Any host-based system would be deployed on individual ECUs whereas a network-based system would be deployed for a specific bus. Any interpretation of network as a larger subsystem consisting of several buses would yield results that are not general, but rather specific for a system with a particular set of buses that are interconnected in a particular way.

As previously mentioned, ECUs are rarely designed and developed by the car manufacturers themselves but more often by subcontractors. In some cases hardware and software are developed separately by different contractors. Some standardization is in place, specifically AUTOSAR [25], but the architecture, performance and properties of different ECUs vary greatly. This coupled with the increased complexity associated with deploying our system to all or a subset of ECUs (and possibly having to tailor it specifically for each one) makes a HIDS solution a poor choice.

A better choice would be a network-based IDPS. Standards are in place for the different buses (see 2.2.2 for more details) and the classification of malicious or normal activity can be based on messages conforming to these standards rather than the internal state of an ECU. Ultimately, a NIDS is superior in almost all respects, perhaps the need for dedicated hardware incurs an extra expense, but this should be negligible when compared to the cost of modifying existing ECUs.

The basic network IDPS presented by Miller and Valasek [4] was successful in detecting message injection attacks. As message injection is what enable bus-hopping, ECU reprogramming and remote attacks with cyber-physical effects (remote endpoints rarely

include functionality for actuating cyber-physical systems), this is a great leap forward in automotive security.

4.2.2 Location and Distribution

The placement of the IDPS is partially a consequence of architecture and the scope (HIDS or NIDS). The three different placement models have different advantages and disadvantages, distributed or hierarchical solutions enable fault tolerance by isolating faults and failures but also increase the complexity and the costs. If the number of networks and messages increase we might reach a point where distribution of work becomes cost effective or even necessary.

A centralized system with sensors on each of the different buses present in the car should lower costs by reducing the number of computing nodes needed. In addition, a centralized solution eliminates the need for messaging between computing nodes, thereby drastically reducing the complexity of the system. Analyzing the data from all covered networks together should also allow for more accurate detection methods. Thus, a centralized system is our recommendation.

4.2.3 Detection method

A combination of misuse detection and anomaly detection is usually preferred as this provides benefits of both approaches [11]. However, in order to detect misuse, or bad behaviour, a database with signatures of attacks has to be present. Due to the lack of known attacks against automotive systems in the wild (as of today) this approach does not hold much promise.

Another problem with misuse detection is the need to regularly update the database of signatures as new attacks are discovered. Updating over the internet etc. could be viable, but avoiding remote connectivity to the IDPS increases security by eliminating possibilities of remote reprogramming and tampering. Updates could also be installed by mechanics during services or downloaded and installed (offline) by the car owners themselves. Nonetheless, neither of the aforementioned strategies seem to solve the problem of secure and timely updates.

Automotive networks are a good setting for anomaly detection, as Miller and Valasek demonstrates in their 2014 survey [4]. The analyzed network traffic originates from hardware devices with predictable behaviour that should be relatively easy to distinguish from anomalous or bad traffic. The relative simplicity of automotive networks and protocols (when compared to conventional computer networks) is highly advantageous when implementing a statistical baseline, or specification-based model, of normal behaviour. It is important to ensure that this baseline covers all normal operation of the vehicle to minimize the number of false positives and false negatives.

The proof of concept IDPS described by Miller and Valasek effectively detect all of their own attacks, and in theory also the attacks presented by Koscher et al. [1, 2]. This is achieved by simple analysis of the frequencies of CAN IDs. Since all the attacks found

by these researchers rely on message flooding (which clearly increase the rates at which certain messages are sent) or abuse of diagnostic messages that should not occur during normal operation, the system easily detects them.

This rather simple method could be improved by also analyzing the relative frequencies of different messages/CAN IDs (perhaps even on different buses). Another option would be using a neural network that has been trained with normal traffic to identify anomalous traffic. However, using a specification-based strategy or simple statistics is preferable as these methods should produce less false positives and make detection more transparent.

With anomaly detection we get false positives and false negatives. In a safety critical automotive system it is important to thoroughly evaluate the effects of these undesired cases to ensure that for example false positives cannot be exploited to trigger unwarranted post-detection actions that could impede safety or reduce functionality.

4.2.4 Post-detection

As can be seen in the results of previous research into possible attacks on automotive systems [1, 2, 3, 4], the effects of malicious intrusions can be extremely dangerous and potentially lethal. Passive post-detection schemes that issues warnings, logs events and facilitates forensic examinations after the fact are of course useful, but the main concern must be actively maintaining vehicle safety. Thus, active post-detection, to the extent it is possible, is the superior alternative.

The range of active measures that can be employed to prevent intrusions depends on the architecture and functionality of the specific vehicle in question. A number of conceivable strategies are detailed in Section 5.1.4.

5

Discussion

This chapter begins with a discussion of the results presented in the preceding chapter. The design suggested in the results is critically reviewed in Section 5.1 and then compared to other alternative and/or complementary measures for increasing the security of embedded automotive systems in Section 5.2.

Furthermore, some general questions and topics related to automotive systems and research on safety and security are discussed. This discussion covers both academic and industrial challenges related to e.g. outsourcing, ethics and tools for practical research on vehicle bus systems.

5.1 Analysis of IDPS properties

This section elaborates on the results regarding the four central IDPS design considerations. The results are motivated further and their respective implications on the suggested IDPS design as a whole are examined.

5.1.1 Scope

The results of our analysis of scope (see Section 4.2.1) show that a network-based IDPS has several advantages over host-based counterparts in the automotive context. Differences between ECUs make software-only solutions difficult (if not infeasible) to implement and introducing new hardware for each ECU to protect drastically increases costs. The manufacturing process with outsourcing, subcontractors and intellectual property issues, discussed by Amin and Tariq [19] and Checkoway et al. [2], further complicates any attempts at modifying ECUs and adding new functionality to existing hardware.

A design that is based on addition of components rather than modification of the existing architecture also means that the IDPS could be sold as an aftermarket product or option. If the method for establishing a statistical baseline or training the system

(assuming the IDPS uses anomaly detection) is general and automated, it could conceivably be deployed in any car, regardless of make or model.

This approach could also enable third-party development of the IDPS as no knowledge of the inner workings of ECUs etc. is required. However, third-party development could be problematic for several reasons. As previously stated there are problems related to security reviews and verification and validation of "black boxes" containing hardware and software which is the intellectual property of a third-party. Also, baselines and coverage could probably be improved if the IDPS can leverage knowledge of the architecture of the specific car it is deployed in.

5.1.2 Location and Distribution

A single centralized IDPS node analyzing traffic from several (or even all) networks in a car seems to be the best choice. However, analyzing these large amounts of data might be to computationally intensive, depending on the algorithms used for detection. If this is the case, using one IDPS component or node per bus or network covered could be a viable option.

Depending on the variables and measurements used by the system, the advantages of centralization will probably vary. A system that only analyzes the frequencies of individual message types or CAN IDs, like the one shown by Miller and Valasek, could easily be distributed. However, if the system is able to identify anomalies based on correlations between traffic/messages, a centralized system might be more accurate. E.g. if a message occurs on the CAN bus that indicates that the car is moving forward at a 100 mph, and "simultaneously" a message transmitted over FlexRay contradicts this by indicating that the car is stationary. This behaviour could be detected and identified as anomalous by a computing node that is monitoring both of the networks, while independent analysis of each network could indicate normal operation.

The advantages and disadvantages of different approaches to location and distribution should be examined further and tested practically. Specific detection methods, such as a particular type of statistical analysis, will benefit (to varying degrees) from different IDPS architectures. Also, the architecture of the target vehicles might influence the decision on whether or not to employ a distributed system.

5.1.3 Detection method

Misuse/signature-based detection currently holds lite promise in and of itself, due to the lack of known patterns of misuse or attacks and the issues concerning up-to-date signatures. However, as attacks are discovered and documented it might become a valuable tool in securing automotive systems.

Perhaps facilities for signature-based detection could be included in automotive (anomaly detection-based) IDPSs to allow for activation and addition of signatures in case an attack is discovered. I.e. if a critical vulnerability is found the car manufacturer can protect customers with a cheap software update rather than a recall or expensive ECU replacement. However, issues surrounding updates remain and finding a good, se-

cure method for distributing and deploying a software update/activation like the aforementioned one is a non-trivial problem. We would argue that the risks associated with introducing a channel for updates outweigh any and all benefits.

Also worth noting is that the constant battle between attackers and authors of anti-virus signatures, IDPS signatures, security updates and patches in conventional computer networks is enabled by the infrastructure that allows for secure and timely distribution of patches and updates. This means that until that type of infrastructure is established in automotive systems, knowledge-based protection mechanisms will perform poorly.

Luckily, anomaly detection and statistical methods seem to provide adequate protection against the attacks demonstrated by researchers so far. A number of different types of anomaly detection, statistical and probabilistic methods as well as neural networks could be used. More research, practical tests and comparative studies is needed to determine if any specific method is better than the alternatives. However, we would hypothesize that advanced methods like neural networks increase complexity and cost without providing any significant benefits over simple statistical methods.

5.1.4 Post-detection

In the results section above (4.2.4), we conclude that active post-detection would be preferable. However, finding a good method or set of defensive actions that is applicable and safe in a generic car is not trivial. Possible options for containing or stopping a breach of security depend on the network architecture (segmentation etc.), what entity or system is being attacked and several other factors.

Typically, the options include shutting down individual ECUs, remote communications endpoints or buses. I.e. if an attacker is flooding a bus with anomalous messages we can either disable all communication on the bus, disable the targeted ECU (since it is hard to deduce the origin of messages in for example CAN networks) or disable the remote communication interfaces (WiFi, Bluetooth etc.). While all three of these methods might be successful in stopping the attack, they might also have a negative effect on safety. Depending on the responsibilities of the ECU or bus in question, shutting it down could disable safety features or disrupt the operation of the system in some other manner. Miller and Valasek [4] simply shut down the bus when their prototype IDPS detects an attack taking place. This might be a viable strategy in some cases and in certain architectures, but in other cases it might be problematic.

Hence, designers of automotive IDPSs must carefully consider the repercussions of any such action and pay extra attention to the effects of false detections. As stated in Section 4.2.3 it is also important to avoid introducing vulnerabilities or opportunities for attackers to trigger unwarranted post-detection actions. In conclusion, the inherent difficulties and risks associated with active post-detection methods might be an insurmountable obstacle to finding a general method.

This leaves us with two choices, either using a passive approach, issuing warnings and logging anomalous activity, or employing active actions tailored to specific cars or architectures. Further analysis of these issues, attempting to find a generalizable strategy for post-detection and real-world tests are some possible avenues for future research.

5.1.5 Limitations of the suggested system

The system suggested in our results (and discussed above) has a number of limitations. In this section we wish to elaborate on the most important of the limitations of a network-based, centralized anomaly detection system with active post-detection.

Focus on in-car networks

Our threat model is highly influenced by the definition of a remote attack as described by Miller and Valasek (see Section 2.2.4). This means that our recommendations regarding detection method and post-detection actions are primarily concerned with protecting safety-critical cyber-physical assets.

Thus, anomaly detection might not be the best fit for identifying intrusion attempts in WiFi/Bluetooth/3G traffic, which to a large degree is generated by user actions. But on the other hand it should be an excellent method of detecting problems in ECU to ECU communication, which is much more predictable and easier to produce a statistical model of normal behaviour for. Hence, our suggested system does not directly increase the security of the remote endpoints, such issues will have to be solved with complementary security mechanisms, but rather stops attacks from spreading to the in-car networks (CAN, FlexRay etc.).

Passive attacks

Information-stealing attacks based on passive monitoring of in-car networks cannot be prevented since they do not generate any traffic that the IDPS can detect. Some other security mechanism will be required to stop this class of attacks.

Anomaly detection

Anomaly detection is only as good as our model of normal behaviour. A poor representation of normal behaviour in the baseline will decrease the accuracy of the IDPS and increase the number of false positives. Miller and Valasek has shown that simple frequency analysis is a viable detection method in the CAN networks of their test vehicles. However, different architectures, network protocols and types of anomaly detection could conceivably cause problems.

Some architectural measures could be taken to improve the predictability of the analyzed traffic and thereby facilitate the creation of a good baseline. For instance, if we wish to make the frequencies of messages consistent we should avoid "triggered" messages in favor of periodic updates issued regardless of changes in state. Using this method we can ensure deterministic behaviour at the cost of an increased number of messages and possibly introducing some latency. Whether or not this kind of trade-off is acceptable will have to be decided by future research and the automotive industry.

One thing worth noting is that anomaly detection is sometimes associated with opaque machine learning and a constantly evolving systems. The approach that we would suggest is to make the system as static and transparent as possible. That is,

ensuring that the model of normal behaviour or baseline is transparent and that it does not evolve or change once the car has been delivered to the owner. By taking measures to increase predictability and using specification-based methods or perhaps simple statistical methods rather than machine learning, the model of normal behaviour becomes more transparent and predictable. By doing this, we should be able to avoid many of the pitfalls associated with anomaly detection, such as unexpected false positives for rarely occurring messages. A system using only specifications of normal behaviour (instead of a statistically derived or "learned" one) would probably be the optimal solution.

5.2 Other security measures

In addition to IDPS, a number of different tools and schemes could be used to protect embedded automotive systems and ensure confidentiality, integrity and availability. Using multiple tools or schemes in a layered fashion is often the most effective strategy to securing a system. Thus, we consider these other security measures to be complements rather than alternatives.

Below we describe three types or classes of protection mechanisms that are often mentioned in the context of automotive systems: **network segmentation**, **packet filters** and **cryptography**. We discuss the strengths and weaknesses of these three types of protection and also compare them to IDPS solutions. Lastly, we discuss the potential of avoiding vulnerabilities and problems by reviewing organizational questions regarding outsourcing and integration of third-party components.

Note that as these topics are not the main focus of this thesis, we do not present any in-depth technical analysis but rather a cursory overview of their respective functions and properties.

5.2.1 Network segmentation

Network architecture has a large influence of the security of a system. The placement of different ECUs on different subnetworks or buses helps to isolate the effects of attacks. A single compromised ECU should not have any effect on other subnets or buses. We have seen that this is not always the case, as reprogramming of gateways and ECUs that are connected to multiple networks enable bus-hopping. However, the security gains of good segmentation could be highly significant, assuming we can adequately protect the gateways.

Separating different types of ECUs and systems becomes even more important as we increase the number of remote endpoints. Effectively separating WiFi, Bluetooth, 3G etc. from the safety-critical control systems will make remote attacks much more difficult to carry out. New types of remote communications will continue to be introduced to cars. Vehicle-to-X communication is a hot topic in both research and the automotive industry. Vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) has the potential to increase safety and allow for new functionality in automotive systems. But, it could also introduce new threat surfaces and possibilities for remote attacks. Hence we would

argue in favor of separating these new communication interfaces from existing safety-critical control systems to the extent possible.

The diagrams by Miller and Valasek [4] show examples of the many different approaches to network architecture. Everything from flat architectures consisting of a single CAN-bus to advanced multi-bus systems exist in modern production cars. The fact that Miller and Valasek based their choice of cars to test based on research into different architectures is a good indication of the security implications of architectural choices.

Network segmentation is non-trivial, affects the cost and complexity of a system and is dependent on other defenses (gateway protection) to be an effective protection mechanism. It is not in and of itself an alternative to an IDPS solution like the one we propose, but is still an important factor in automotive security. The concept of network segmentation is closely related to packet filters which are discussed in the next section.

5.2.2 Packet filters

Using firewall-like functionality, or packet filtering, in bus gateways and filtering traffic that flows from one bus to another is a possible strategy to combat bus-hopping. In order for packet filters to be useful the in-car networks has to be segmented in some manner, i.e. we cannot filter traffic between networks if there is only one network. Specifying rules for what messages should be propagated across the border between two subnetworks gives us fine-grained control and allows us to enforce policies and specifications of in-car communication.

The functionality of a packet filter is partially similar to that of an IDPS solution, it identifies and filters out anomalous traffic (assuming a whitelist/default-deny policy is employed). Thus, packet filtering could be effective in isolating the effects of a security breach.

Also, packet filters could potentially be useful tools in IDPS post-detection. If an intrusion attempt is detected on a bus, the IDPS could alter the packet filtering rules to restrict or completely disable messages from the affected bus being propagated. This should reduce the risk of the attack spreading to other parts of the system, but is less drastic than shutting down an entire bus or ECU.

Wolf et al. also discuss packet filters or gateway firewalls [20] and in addition to basing rules on the authorizations of subnets/buses, they propose including authorization in cryptographic certificates in ECUs.

5.2.3 Cryptography

One approach to securing automotive systems is to use cryptography to provide confidentiality and integrity through encryption and cryptographic signatures, respectively. In this section we describe three different applications of cryptography in automotive systems: signed messages, signed software updates and encrypted messages. These three could be used separately or together and have different advantages and disadvantages. In general, secure cryptography is computationally expensive, which might be a problem

in automotive systems with limited computational power.

Miller and Valasek briefly describe cryptographically signing CAN messages as a way to ensure that "...only the ECUs (and mechanics tools) have the keys and so a random attacker wouldn't be able to send valid CAN messages on the compromised automotive network". By signing messages before sending them and having receivers discard messages if a valid signature from a trusted party is not present we can prevent message injection. However, Miller and Valasek go on to say: "This idea may present obstacles for attackers who add rogue devices to automotive networks, but in the context of a remote attack, the attacker is executing code on a compromised ECU. At this point, the keys are also compromised, or at least the ability to send valid CAN messages". This quote pinpoints one important limitation of cryptography in automotive systems. Another problem related to cryptographic signatures of CAN messages arises from the limited size of CAN packets. A signature is simply too long to fit in a single packet, which means signing a message requires several messages in and of itself. Cryptographically signing messages is also discussed, in less pessimistic terms, by Wolf et al. in their paper titled *Security in Automotive Bus Systems* [20], though their work is mostly theoretical.

Another use for cryptography could be cryptographic signatures on software and firmware updates. Signing the code running on the ECUs means that strict control of the code can be enforced which mitigates the problem of having an adversary running malicious code. Software updates will have their signatures checked before being installed and if the signature is not present and correct the update will not be performed. Signed software is already used in some automotive systems [26].

Signed software does not actively protect the car against malicious adversaries trying to gain access to the car, but instead provides passive protection against ECU tampering. It would protect against bus-hopping and prevent attackers from bypassing gateways which is very desirable. Without protection against bus-hopping, any attempts at increasing security through segmentation of networks and separation of ECUs on different buses (see Section 5.2.1) is futile.

The IDPS solution that we propose could be configured to disallow ECU software updates over the protected buses except under certain conditions. E.g. flashing new software into an ECU could be considered anomalous behaviour unless some physical switch is turned on. However, this does not render software signatures redundant. Thus these two methods could be used together to effectively protect ECU firmware integrity.

A third use for cryptography is message encryption, using this method we can ensure that messages cannot be read by an adversary that is monitoring a bus. This might become relevant if information-stealing attacks is considered a highly problematic type of attacks. Encrypting messages does not help us avoid message injection, which we consider the most immediate threat. It should also be noted that message encryption could complicate the use of packet filters or IDPSs. Both the aforementioned security measures needs to be able to read packet data, and if packets are encrypted they would need keys to decrypt the data in order to perform their respective tasks.

A number of issues surrounding key management and distribution have been identified and discussed, some of them have been solved while others remain. Problems could

also arise when combining encryption with the communication protocols used in automotive systems, e.g. can we fit a signature into the data field of a CAN frame? Further discussion on cryptography is beyond the scope of this work, but more information can be found in the texts by Wolf et al. [20] and Das et al. [26].

5.2.4 On outsourcing and integration

Outsourcing and using hardware and software from different suppliers is very common in the automotive industry. This has a profound effect on security and limits the possibilities for security reviews and audits. Checkoway et al. explains "... frequently manufacturers do not have access to the source code for the ECUs they contract for ... [T]hus, while each supplier does unit testing (according to the specification) it is difficult for the manufacturer to evaluate security vulnerabilities that emerge at the integration stage. ... [T]herefore, while this outsourcing process might have been appropriate for purely mechanical systems, it is no longer appropriate for digital systems that have the potential for remote compromise" [2]. A less radical proposition to solve the same problem is described by Amin and Tariq, whose work is focused at "identifying the manufacturer-supplier relationship that reduces the risk of vulnerabilities at the boundaries between electronic control units and thus protects the integrity of the car's critical software modules" [19].

Managerial and organizational issues of this type are very complex. The cost factor is important but we cannot overlook the effects on security and safety. In-depth audits enabled by transparency could potentially remove a large portion of the vulnerabilities whose exploitation we try to protect against by using IDPSs, packet filters etc.

5.3 General Discussion

5.3.1 Ethics and Safety

With all the modern vehicle safety systems a number of ethical dilemmas crop up. One example of this could be a collision avoidance system using proximity sensors and cameras, at some point a system like this could be required to choose between crashing the car and potentially harming the occupants, or running a person over. Dealing with this is not trivial, and we are moving closer to a world where these kinds of questions are no longer theoretical exercises.

Issues such as the one described above are not directly related to our work, but rather the work of engineers trying to implement active safety systems. However, some security-related mechanisms or systems can give rise to similar concerns. Miller and Valasek [4] for example suggest shutting down a bus once an intrusion attempt or attack is detected by their IDPS, this could lead to safety systems being disabled. The engineers tasked with protecting and securing safety-critical cyber-physical systems must take this into consideration and be aware of the possibly life-critical outcomes.

One extremely important aspect of safety and security in automotive systems is that independent researchers etc. get opportunities to audit and examine systems. Under

no circumstances can we allow the primary function of security professionals to become maintaining the reputation of manufacturers instead of ensuring the safety of vehicle occupants. Related issues are discussed further in Sections 5.3.2 and 5.3.4.

5.3.2 Ownership

One important issue when discussing cars is the ownership of the car, what rights do owners have to alter, modify, remove or add hardware- and software components to the car they bought? A number of laws and regulations have been put in place to guarantee some rights, while others are currently being discussed and contested [27, 28]. As technology evolves and the systems get more complex and gain more advanced functionality, the question of who has the right to do what with the car gets increasingly important. Software has generally been a problematic area when it comes to rights and ownership. Manufacturers and corporations tend to claim intellectual property while free software proponents argue for a higher degree of freedom and rights for users.

Further complications have been (and will continue to be) introduced with the advent of the Internet of Things (IoT), where more and more appliances, home electronics and various other things become computerized. Rules and regulations conceived before this explosion of embedded computing might not be useful or have the intended effects when applied in the modern world. Closed software can inhibit the possibilities for end-users to repair, modify and improve products that they buy.

Automotive systems is an interesting domain in which to discuss this, partly because there is an established culture of car tuning and car "modding", where car owners "hack" or modify their cars in different ways. Another factor that makes cars different to many other IoT-type devices or appliances is the safety-criticality.

There are essentially two different approaches to this, either the architecture and regulations need to ensure that systems are as safe as possible and that they restrict the opportunities for end-user/third party alterations, or the systems need to be open to allow for continuous improvement and audits. The first approach gives fewer rights to the owner in order to increase security and lock down the system. By disallowing (both technically and through regulation) modification we might gain some level of security. On the other hand, reverse engineering etc. will probably not be stopped altogether, perhaps making the second approach more tempting. By making the systems open, more people can review hardware and code and perhaps discover new vulnerabilities that can then be patched. However, openness and transparency could also help attackers to some extent and lead to an increase in attacks.

The introduction of security features like IDPSs and encryption can, as stated above, make modification of cars more difficult. This presents not just a technical challenge (if we wish to enable modification to some degree) but also questions of ethics, rules and regulations regarding ownership and the rights of car owners.

5.3.3 Privacy

Any system that analyzes messages in in-car networks has to be aware of the privacy of personal data. There is a lot of sensor and system data being communicated in a modern car. In addition to information-stealing attacks, it is also important to consider other, less acute, privacy problems.

Consider for example an IDPS solution where detection is moved from the limited hardware within the car to a remote cloud server for performance or cost reasons. We then have to transfer the traffic or data to be analyzed to a remote server. In addition to introducing a new possible attack surface (an adversary could attempt to gain access to this traffic), this traffic could be deemed to contain sensitive personal information. Perhaps it is possible to extrapolate information regarding the personal preferences and habits of the car owner. This gives rise to numerous questions regarding how this data is handled and what explicit permissions from car owners are needed.

The system that we propose should not present any privacy issues, as the only state information kept by the system is the statistical baseline and no data is sent from the IDPS to some other system or third-party. Any good statistical baseline should be independent of driver and only be based on properties of the car itself. If this is not the case, the system could produce false positives or false negatives based on the behaviour of the driver, which is highly undesirable. However, if we use a detection method with some sort of feedback to enhance the baseline, the internal state of the system could become sensitive information.

5.3.4 Real-world attacks

In what we have seen, practically all known attacks and vulnerabilities in embedded automotive systems have originated from security researchers. In addition to making signature-based IDPSs problematic (see Section 5.1.3) it poses a number of questions. Examining the reasons behind the low number of attacks in the wild is important, both in order to avoid inadvertently removing obstacles that have kept malicious actors out of cars and to design and implement practical and cost-effective protection mechanisms.

One possible reason is the obscurity and complexity of automotive systems. Protocols, architecture and hardware used in cars is dissimilar to traditional computer networks in a number of ways that might make it hard for attackers to apply their current skillset and technique in automotive systems. Another possible reason is the diversity of automotive systems, time and effort spent on devising a plan of attack might not be warranted if the resulting attack is only effective against a small subset of cars spread across the globe. Lack of motivation could also be a factor, modern cybercrime is to a large extent aimed at making money and a model for profiting from automotive attacks might not yet exist.

To the benefit of both security researchers and (potentially) attackers, hardware devices and software interfaces that simplify communication with in-car networks are being developed. One example of such a tool is CANtact [29], which is a cheap device that

connects to the OBD-II interface of the car and allows for Python scripting.

It is possible that there are attacks out there that we are not aware of, perhaps undetected or non-disclosed problems exist. Regardless, it is important that security professionals and researchers analyze these types of questions in parallel with technical and practical work.

5.4 Future work

In order to validate and support our arguments and suggestions, anomaly detection in LIN, MOST, FlexRay and other networks must be tested. The CAN IDPS presented by Miller and Valasek [4] could be used as a reference for implementation of similar tools for the other types of networks/buses in automotive systems. Possibly, specific details and properties of for example LIN makes it hard to effectively employ the same type of anomaly detection as in CAN. The cornerstone of anomaly detection is the ability to create a statistical model of normal system operation that can be used to distinguish anomalous traffic or behaviour. Verifying that this is viable and effective in networks other than CAN is important if one wants to cover these networks with an IDPS like the one we suggest.

Another avenue for future research is testing and comparison of different anomaly detection methods. Various statistical methods can be used and more advanced solutions based on e.g. neural networks could improve the detection. Also, taking factors besides message (CAN ID) frequency into account in the detection method could be beneficial.

Different types of post-detection methods should be evaluated and tested. In order to produce a viable IDPS solution, investigating different defensive actions, both in general and targeted at specific cars, is essential.

Further analysis of the problems related to outsourcing and the integration of code that the car manufacturers cannot review is also interesting. The different opinions of Checkoway et al. [2] and Amin and Tariq [19] lead to different ways forward for automotive architectures and should be discussed further in both academia and industry.

Elaborations on a unified terminology that enables good communication between researchers and engineers from different fields will be essential to the development of secure cyber-physical systems. Safety of physical or mechanical processes cannot be decoupled from the security of embedded computers and control systems, hence engineers and researchers will have to widen their scopes to include both of these aspects.

6

Conclusion

Our attempts to create a generic model of modern automotive systems failed, the architectural diversity made the construction of such a model impossible. Instead we base our analysis on the protection of one or several communication buses or networks.

Based on our "model" and a number important factors such as cost, complexity and hardware-imposed constraints we conclude that a **centralized network-based** system using **anomaly-detection** and **active post-detection** should theoretically be the optimal IDPS solution for modern automotive systems. This solution has some limitations, e.g. it cannot detect passive monitoring attacks and it is primarily focused on the protection of buses for internal communications. The focus on internal communication is motivated by the safety-criticality of cyber-physical systems that are connected to such buses.

It is important that our work is practically verified and that further research is conducted where different anomaly-detection methods and possible post-detection responses are evaluated and compared.

7

Glossary

ABS	Anti-lock Brake Systems
CAN	Controller Area Network
CPS	Cyber-physical systems
DoS	Denial-of-Service
ECU	Electronic Control Unit
ESC	Electronic Stability Control
HIDS	Host based Intrusion Detection System
IDPS	Intrusion Detection and Prevention System
LIN	Local Interconnect Network
MOST	Media Oriented Systems Transport
NIDS	Network based Intrusion Detection System
OBD-II	On Board Diagnostics Interface

Bibliography

[1] K. Koscher, A. Czeskis, F. Roesner, S. Patel, T. Kohno, S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, et al., Experimental security analysis of a modern automobile, in: Security and Privacy (SP), 2010 IEEE Symposium on, IEEE, 2010, pp. 447–462.

[2] S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, S. Savage, K. Koscher, A. Czeskis, F. Roesner, T. Kohno, Comprehensive experimental analyses of automotive attack surfaces, in: Proceedings of the 20th USENIX Conference on Security, SEC'11, USENIX Association, Berkeley, CA, USA, 2011, pp. 6–6.
URL `http://dl.acm.org/citation.cfm?id=2028067.2028073`

[3] C. Miller, C. Valasek, Adventures in automotive networks and control units, Last Accessed from h ttp://illmatics. com/car_ hacking. pdf on 13.

[4] C. Miller, C. Valasek, A survey of remote automotive attack surfaces (2014).

[5] A. K. Pathan, The State of the Art in Intrusion Prevention and Detection, CRC Press, 2014.

[6] J. P. A. Co, Computer Security Threat Monitoring and Surveillance, 1980.

[7] D. E. Denning, P. G. Neumann, Requirements and model for ides—a real-time intrusion detection expert system, Document A005, SRI International 333.

[8] D. Denning, An intrusion-detection model, Software Engineering, IEEE Transactions on SE-13 (2) (1987) 222–232.

[9] S. Smaha, Haystack: an intrusion detection system, in: Aerospace Computer Security Applications Conference, 1988., Fourth, 1988, pp. 37–44.

[10] L. Heberlein, G. Dias, K. Levitt, B. Mukherjee, J. Wood, D. Wolber, A network security monitor, in: Research in Security and Privacy, 1990. Proceedings., 1990 IEEE Computer Society Symposium on, 1990, pp. 296–304.

[11] E. Biermann, E. Cloete, L. Venter", A comparison of intrusion detection systems, Computers & Security 20 (8) (2001) 676 – 683.
URL http://www.sciencedirect.com/science/article/pii/S0167404801008069

[12] R. Sekar, A. Gupta, J. Frullo, T. Shanbhag, A. Tiwari, H. Yang, S. Zhou, Specification-based anomaly detection: a new approach for detecting network intrusions, in: Proceedings of the 9th ACM conference on Computer and communications security, ACM, 2002, pp. 265–274.

[13] A. Perrig, J. Stankovic, D. Wagner, Security in wireless sensor networks, Communications of the ACM 47 (6) (2004) 53–57.

[14] I. Onat, A. Miri, An intrusion detection system for wireless sensor networks, in: Wireless And Mobile Computing, Networking And Communications, 2005.(WiMob'2005), IEEE International Conference on, Vol. 3, IEEE, 2005, pp. 253–259.

[15] S. Raza, L. Wallgren, T. Voigt, Svelte: Real-time intrusion detection in the internet of things, Ad Hoc Networks 11 (8) (2013) 2661 – 2674.
URL http://www.sciencedirect.com/science/article/pii/S1570870513001005

[16] R. N. Charette, This car runs on code, IEEE Spectrum 46 (3) (2009) 3.

[17] I. Studnia, V. Nicomette, E. Alata, Y. Deswarte, M. Kaaniche, Y. Laarouchi, Security of embedded automotive networks: state of the art and a research proposal, in: M. ROY (Ed.), SAFECOMP 2013 - Workshop CARS (2nd Workshop on Critical Automotive applications : Robustness & Safety) of the 32nd International Conference on Computer Safety, Reliability and Security, Toulouse, France, 2013, p. NA.
URL https://hal.archives-ouvertes.fr/hal-00848234

[18] M. Broy, I. Kruger, A. Pretschner, C. Salzmann, Engineering automotive software, Proceedings of the IEEE 95 (2) (2007) 356–373.

[19] M. Amin, Z. Tariq, Securing the car: How intrusive manufacturer-supplier approaches can reduce cybersecurity vulnerabilities, Technology Innovation Management Review (2015) 21.

[20] M. Wolf, A. Weimerskirch, C. Paar, Security in automotive bus systems, in: Workshop on Embedded Security in Cars, 2004.

[21] E. Lee, Cyber physical systems: Design challenges, in: Object Oriented Real-Time Distributed Computing (ISORC), 2008 11th IEEE International Symposium on, 2008, pp. 363–369.

[22] J. Knight, Safety critical systems: challenges and directions, in: Software Engineering, 2002. ICSE 2002. Proceedings of the 24rd International Conference on, 2002, pp. 547–550.

[23] G. Leen, D. Heffernan, Expanding automotive electronic systems, Computer 35 (1) (2002) 88–93.

[24] Sans: Glossary of security terms, https://www.sans.org/security-resources/glossary-of-terms/, accessed: 2015-04-10.

[25] Autosar development partnership, https://www.autosar.org, accessed: 2015-04-21.

[26] S. K. Das, K. Kant, N. Zhang, Handbook on Securing Cyber-Physical Critical Infrastructure, Elsevier, 2012.

[27] Automakers say you don't really own your car, https://www.eff.org/deeplinks/2015/04/automakers-say-you-dont-really-own-your-car, accessed: 2015-04-13.

[28] K. Wiens, We Can't Let John Deere Destroy the Very Idea of Ownership, http://www.wired.com/2015/04/dmca-ownership-john-deere/, accessed: 2015-05-35.

[29] E. Evenchick, An introduction to the canard toolkit, in: Blackhat Conference 2015, Blackhat, Marnia bay sands, Singapore, 2015.