Journal of Information Assurance and Security
ISSN 1554-1010 Volume 6 (2011) pp. 132-140

Anomaly Detection for In-Vehicle Networks using a Sensor-based Approach

Michael Müter[1], André Groll[2] and Felix C. Freiling[3]

[1]Daimler AG, Research and Development, GR/PTA
71034 Böblingen, Germany
michael.mueter@daimler.com

[2]University of Siegen, Institute for Data Communications Systems
57068 Siegen, Germany
andre.groll@uni-siegen.de

[3]University of Mannheim, Laboratory for Dependable Distributed Systems
68131 Mannheim, Germany
freiling@uni-mannheim.de

Abstract: **Modern vehicles made a meaningful development to more complexity as well as connectivity in the last decade. Therefore, they cannot be seen as a closed system anymore. As the openness of the vehicle is increasing, so does the security risk for the in-vehicle networks and its components. Apart from threats for comfort and confidentiality, these attacks can also affect safety critical systems of the vehicle and hence endanger the driver and other road users. This paper discusses anomaly detection capabilities for in-vehicle networks based on sensors monitoring the internal network traffic. With respect to characteristics of typical vehicular networks, like the Controller Area Network (CAN), a recognition model for potential attacks during the operation of the vehicle without causing false positives is presented. Moreover, important design and application criteria for such an extension of the vehicle's security architecture are explained and discussed.**

Keywords: Automotive Security, Vehicular Anomaly Detection, In-vehicle Networks, CAN.

I. Introduction and Background

The automotive industry has undergone a substantial development in the last decade: More and more electronics and software is integrated into the vehicle to provide more safety to the driver in case of new assistant systems or to add more functionality to the car in a cost-effective way. This means that the number of electronic control units (ECUs) has steadily increased. Modern upper class vehicles comprise up to 80 ECUs for different application areas and functions that pertain to different automotive networks and domains. Another important development is the interfacing with external networks for car-to-X communications (e.g., WLAN, DSRC, WAVE, IEEE 1609.2) and mobile communication networks (e.g., GSM, UMTS, Bluetooth) as well as storage media (e.g., USB, CD, DVD). Even the integration of nomadic third party devices like navigation systems, mobile phones, notebooks, etc. is possible, which in the future also may even gain access to the internal networks. The impact of the increasing complexity, the number of interfaces and communication possibilities is that modern and future vehicles are no longer closed systems like in the past; they have changed into open systems.

The downside of openness is exposure, and with exposure rises the risk of attacks for future vehicles. Consequently, the likeliness increases that vehicular IT systems will be subject to a similar amount of malice that original desktop systems are faced with [24]. Currently, it is a major trend that many recent research activities are launched, especially focusing on in-vehicle security. Since most in-vehicle networks are optimized with respect to safety and reliability, but do barely provide any explicit protection mechanisms against malicious attacks, it might be possible for attackers to inject or manipulate messages on certain automotive networks. These attacks could result in a negative impact for comfort and privacy, but also cause serious malfunctions of the vehicle and a threat for safety and human life – for instance if an attacker manages to inject packets into the powertrain network or manipulate messages for the Antilock Braking System (ABS) [21].

In vehicular security it is sometimes useful to look at the early developments in the security of desktop computers. There, several well explored measures for the mitigation of attacks have been analyzed. But the consideration of standard measures, e.g., firewalls and virus scanners, is not sufficient enough to provide useful protection for vehicular networks, because of their focus on a preventive approach and limited resources. Additionally, vehicles have a very long life span and are in use for decades in different conditions and locations. To provide an efficient protection, preventive measures only are not sufficient enough over this long period of time. One reason for this is the fact that regular updates of threat signatures – like they are known from virus scanners from the PC world – can not be guaranteed in this application domain. As a consequence, the vehicle's security system has to work autonomously without a necessity for user interaction.

In this paper, we focus on a *reactive* approach to in-vehicle network security. We present an approach to extend the se-

curity architecture of vehicles by implementing monitoring capabilities for traffic on vehicular networks in order to evaluate abnormal events and classify them as a threat or not.

A. Related Work

Previous research regarding in-vehicle networks has mainly focused on safety issues [2, 18], i.e., protection against unintentional, random events. More recent activities go beyond and consider security aspects as well [5, 23]. Different potential attack scenarios on future automotive systems have been presented [13] as well as the implementations of concrete attacks on the CAN bus [6]. In the world of desktop computers intrusion detection systems (IDS) are one well known countermeasure by now, and different concepts like misuse and anomaly-based detection have been developed. More details and comprehensive IDS surveys can be found in the literature [11, 22, 25]. In general, the challenge to anomaly detection systems is to achieve a low rate of false positives, since false alerts can be very costly. Additionly, it is well known that the fine-tuning of systems requires suitable test and training data which is often hard to obtain.

The first concept for in-vehicle intrusion detection was introduced by Hoppe et al. [7] with a presentation of three selected characteristics as intrusion detection patterns. This includes the recognition of an increased frequency of cyclic CAN messages, the observation of low-level communication characteristics based on typical properties of electric signals on the physical layer, and the identification of obvious misuse of message IDs. However, it remains unclear how the *obvious* misuse of a message ID is specified. The monitoring and analysis of electric signals on in-vehicle networks on the physical layer seems very difficult in the automotive domain. The signal characteristics on the physical layer can be subject to frequent change due to strong variations during the automotive life cycle with respect to the automotive environment, application fields, temperature ranges, and humidity. Moreover, the definition of these characteristics would involve extraordinarily high efforts. Besides, a first implementation of an attack on the electronic window lift has been published, including an approach for solving the given scenario [6]. The recognition is based on the same three properties and is directly adjusted to the selected attack scenario, a comprehensive or generic approach is not included.

Larson et al. [14] introduce an approach to specification-based attack detection for in-vehicle networks, which shows how to gain a description of the vehicle's normal behavior out of the network protocol and ECU specification based on the CANopen protocol [9]. Moreover, they discuss different aspects with respect to a meaningful IDS sensor placement. The paper reinforces the claim that the challenges intrusion detection systems (especially anomaly-based approaches) generally have to cope with become even more crucial in the automotive domain. While contributing the idea of specification-based attack detection, Larson et al. [14] do not attempt to classify types of sensors to give a broader picture of IDS in vehicular networks.

Hoppe et al. [8] take a more generic approach and ask whether notification concepts of intrusion detection from desktop computers can be applied to the automotive domain. They define an *adaptive dynamic reaction model* for the notification and reaction phase of an IDS, which describes different optical, acoustic, or haptic measures for the reaction to detected threats and the notification of the driver. Moreover, several other recent publications discuss anomaly detection as one potential security approach for future automotive systems, but leave the details to future work [12, 20, 23, 19].

B. Contributions

In this paper we attempt to go beyond present work by taking a first step towards an integrated and holistic approach to anomaly detection for in-vehicle networks. We present a threat detection scheme for in-vehicle networks that comprises nine fundamental types of attack detection sensors which serve as recognition criteria for automotive IT threats. We discuss several requirements that have to be fulfilled for an integration of the approach into the automotive security framework of future vehicles. Furthermore, we derive a classification of automotive attack detection sensors and present a first concept how to integrate our approach into a holistic intrusion reaction concept. Our approach can be regarded as a generalization of specification-based approaches that takes into account the typical characteristics of automotive networks like the Controller Area Network (CAN) [18] and their limitations.

C. Roadmap

Section II describes important design criteria of vehicular attack detection systems. In Section III we present the addressed set of automotive attack detection sensors as well as further aspects for a deployment. Afterwards in Section IV the integration into the vehicle is discussed, and we conclude the paper in Section V.

II. Automotive-specific challenges

For the development of an in-vehicle anomaly detection system several new issues arise, due to different constraints and the nature of automotive networks. In the following we discuss the major conceptual challenges that need to be considered for the design and the integration of an attack detection system into the vehicle.

A. Data Selection

A general issue for the development of an attack detection system is what kind of data the attack detection system needs to observe. In the area of desktop computing, intrusion detection systems are often separated into host-based and network-based approaches (or HIDS and NIDS [15]), depending on their data source. In the vehicle, data sources can be the different sensors and networks but also internal data of ECUs or gateways. Broadly speaking, the more data can be monitored and obtained for evaluation, the better the overall picture about the current situation of the system. However, the more information needs to be observed, gathered and evaluated, the more complex and costly the development and analysis process becomes. Although today's vehicles include several different networks, ECUs and communication sources, not all of these networks may be indispensable for the recognition of in-vehicle attacks.

B. Detection Methodology

One major question is how exactly the identification of in-vehicle attacks should be performed. This includes the vital question, which basic detection approach turns out to be most suitable for the automotive area. Misuse detection [17], sometimes also referred to as *signature detection*, promises a low false positive rate, which is important as numerous false alerts could question the usability of the entire concept in the vehicle and may negatively affect the driver's awareness. However, the focus on known attacks and the need for regular updates make the deployment in the automotive area difficult. At first, frequent updates require a communication channel. Mobile channels like GSM or UMTS cause extra costs and may not be available in every geographic region or country. Broadcast channels like RDS or future technologies like TPEG (Transport Protocol Experts Group) over DAB (Digital Audio Broadcasting) are a theoretic option but would have several technical challenges in this application domain. Updates could be included in the inspection service at the garage, but in this case the update frequency is fairly low and many car-owners worldwide do not rely on a garage service at all. Finally, the owner could install a special device at home which performs the update, resulting in high extra effort for the customer. Besides, this option may not be applicable for persons without technical skills. Second, signature-based detection approaches focus on known attacks and encounter problems as soon as attack patterns deviate from the original specification. In summary, all of the previously described solutions and aspects show serious drawbacks, which can make the signature-based approach fairly unattractive for automotive manufacturers.

Anomaly detection [4] promises to detect attacks, including novel attack patterns, that result in a system state which differs from the normal specification. However, in the past anomaly detection systems were typically prone to high false positive rates and the specification of the system's normal behavior has turned out to be a challenging and daunting task. Nevertheless, if the normal behavior of the vehicular networks can successfully be defined and adopted we consider anomaly detection to be the more promising approach to start with in the automotive domain as unknown attacks may be detected as well and no regular updates are necessary. In the future, hybrid approaches can be promising as well.

C. Sensor Placement

If the relevant data sources have been determined, the next question is where and how the acquired information is collected and evaluated. Two main concepts are possible: Simple sensors that just observe a special data source, e.g., by monitoring a certain bus system, and transfer the information to a central processing unit of the attack detection system, where the entire evaluation is performed. This keeps the sensors fairly cheap and simple but it either massively increases traffic on the automotive network or even requires a separate communication channel for each sensor to be built.

Alternatively, some intelligence of the attack detection system can be included into the sensors themselves. Each of these intelligent sensors can perform some pre-processing, data selection, or even parts of the threat detection [1]. Some data may be discarded because it is not considered relevant for attack detection, repeated data or signals could be summarized and compressed. This massively reduces the amount of traffic that needs to be transferred to a central attack detection unit but increases the costs per sensor.

D. Detection Performance

For a deployment in the automotive area, an attack detection system needs to fulfill real-time performance requirements [10]. Especially attacks which target the safety of the vehicle, e.g., by sending false messages to the brakes, engine, etc. can only be tackled if this requirement is fulfilled. However, the automotive environment is a network of embedded systems comprising highly specialized and cost-optimized components, which offer only limited computational power but are designed to work reliably under very different physical conditions, temperature ranges, etc. This means, for the implementation of attack detection methods, a reasonable balance between performance and costs has to be achieved while ensuring the physical hardware requirements are met.

E. Notification and Reaction

If an attack detection system continuously monitors the automotive network and starts to recognize an attack, immediately the next challenge turns up: What is an appropriate reaction for the system to carry out? In the world of desktop computers, a common response to a potential threat for an attack detection system is to pop up a message on the user's screen indicating the location, type and source of the attack and calling for user input what to do. In the automotive world, however, the situation is more difficult: Imagine a customer driving his car on a motorway with high speed when the vehicular attack detection system recognizes an attack. Displaying an alert message on the vehicle's instrument cluster and asking what action to perform, would cause high distraction for the driver and may also increase the chance for an accident. Moreover, the driver may not have sufficient technical knowledge or experience in order to know what reaction to decide for. Also, the time required for the user to decide and respond is too high to prevent the effect of an attack to prevail. Because of this, the design goals of an automotive IDS have to include an autonomous reaction concept in combination with a high detection reliability. Finally, only if no other option is left, the system should decide to interact with the driver. First approaches of such a user interaction have been discussed elsewhere [8].

F. Detection Reliability

The detection reliability of an automotive intrusion detection system can be described by the *detection rate*, which is the percentage of incidents, which have been successfully detected as an attack (*true positives*). The incidents spuriously marked as attacks are called *false positives*. The recognition rate an automotive IDS should strive to achieve needs to be much higher than the detection rate for common desktop computers, due to the immediate possible effects the attacks, but also the reactions of the automotive IDS, can have on the safety of the driver and other road users.

In general, the detection rate has to be correlated to the notification and reaction model of the IDS. Typically, most intrusion detection systems cause a notable number of false positives. Here a big challenge turns up: On the one hand, a vehicular IDS should act autonomously without asking the driver for feedback — especially not if the vehicle is in a driving situation, as described in Section II-E. On the other hand, an automatic intervention to an incident which turns out to be a *false positive* could lead to similar safety risks than a real attack does. Therefore, this work focuses on approaches which allow the identification of attacks without causing any false positives. In Section IV, we furthermore show how to evaluate the criticality of an incident detected by the approach introduced in this contribution. This focus on reliable detection measures is especially a requirement for the application of a notification and reaction model, which goes beyond a passive notification of the driver and also comprises active measures of intervention and response. Accordingly, if the IDS incorporates a detection approach which yields a higher number of *false positives*, this fact needs to be considered by the notification and reaction model. An exemplary consequence could be the avoidance of intervening measures and a restriction to passive notifications [8].

III. In-Vehicle Network Attack Detection

In this section we present a set of different network-based detection sensors, which allow the recognition of anomalies occurring inside the vehicular network. We point out the conditions that are required for each sensor type by introducing different applicability criteria. Afterwards we show how these criteria can be used to derive a structure for the sensor types.

A. Anomaly Detection Sensors

A major challenge in anomaly detection is to determine a reliable way how anomalies can be identified without generating too many false positives. Therefore, we present a set of different anomaly detection sensors for in-vehicle networks which comprise one major advantage: In contrast to other solutions in the area of anomaly detection [11] they do not produce any false positives. The reason for this is the fact that all sensors are based on unambiguous and reliable information only, namely, the network protocol specifications, the defined cooperative networking behavior of the devices (e.g., message duplication tables of ECUs), redundant data sources in the vehicle, or a combination of these. Therefore, if an incident is detected it is assured that the system is in an abnormal state, however, the sensors may not be able to detect all possible attacks (resulting in false negatives). Obviously, it cannot be determined if the anomaly is caused by a malicious attack or by other reasons, e.g., a hardware error. However, this is a general problem all anomaly detection systems of this type have to face in theory and it does not reduce the applicability of the approach. In fact, the detection of hardware errors results in a very worthwhile information for the driver as well. In this first approach, we assume that the IDS itself does not get compromised by an adversary. Future approaches may consider additional, technical measures, like trusted computing, to enforce this assumption [3]. All detection sensors we introduce are based on the typical behavior of automotive bus systems like CAN, but are described from an abstract point of view to allow an easy adaptation to other transfer media.

S-1: Formality Sensor Vehicular bus systems, like CAN, are very reliable and robust communication media. However, if we move forward from a strict reliability perspective and start to consider intelligent attackers, the standard measures of vehicular bus systems to ensure dependable communication are not sufficient any more. An intelligent attacker could add or manipulate devices in such a way that these components do not completely adhere to the protocol specifications any longer, e.g., in order to cause a *buffer overflow*. Therefore, a basic element for a vehicular anomaly detection system is a sensor which checks every message for formal correctness of the communication protocol, e.g., by verifying the packet header, delimiters, field sizes, checksums, etc.

S-2: Location Sensor For every message in an automotive network it is specified which sub-network this type of message is allowed in. Hence, even when a message is formally correct, it can still be part of an attack, e.g., if that type of message is not allowed within a given domain. For instance, a packet which adjusts engine settings in the powertrain domain is usually not allowed in the telematic domain.

S-3: Type Sensor The Type Sensor accesses the payload of the message and checks if the data type in the payload matches the expected type. A exemplary type mismatch would be a message comprising an *integer* value where a *boolean* is expected. The type sensor can only be sensibly implemented at an abstraction level, where type information is already available. In current automotive networks this is usually not the case. Nevertheless, if an abstraction level is chosen which provides type information, e.g., during data processing at ECU level, the sensor can be integrated.

S-4: Range Sensor The Range Sensor accesses the payload of the message and checks if the data range of the payload stays within the allowed boundaries. For instance, even if the data type *integer* is correct, in a message conveying the current vehicle speed, a value of > 300km/h usually indicates an anomaly (depending on the type of car).

S-5: Frequency Sensor Many messages in the automotive network are sent cyclically with fixed intervals, even when a function is not active or does not change its status. Other messages are only sent on demand cyclically or non-cyclically, e.g., when the driver presses a button to activate a function (like messages for the power windows). The frequency sensor checks if the interval between cyclic messages is within defined upper and lower bounds, but also verifies the interval between non-cyclic messages for realistic and feasible frequency. This type of sensor also ensures that a flooding attempt on the vehicular network in order to perform a *denial-of-service attack* can be detected.

S-6: Correlation Sensor Typically, the vehicular network is comprised of different domains and sub-networks, which are interconnected by dedicated automotive gateways. Often, several messages are not limited to a single bus system but are required by several devices in different sub-networks simultaneously. Therefore, for proper operation those messages are transcribed by the linking gateways. The correlation sensor is an independent entity, which verifies that mes-

Nr	Sensor	Description
S-1	Formality	Correct message size, header and field size, field delimiters, checksum, etc.
S-2	Location	Message is allowed with respect to dedicated bus system
S-3	Type	Compliance of payload in terms of data type
S-4	Range	Compliance of payload in terms of data range
S-5	Frequency	Timing behavior of messages is approved
S-6	Correlation	Correlation of messages on different bus systems adheres to specification
S-7	Protocol	Correct order, start-time, etc. of internal challenge-response protocols
S-8	Plausibility	Content of message payload is plausible, no infeasible correlation with previous values
S-9	Consistency	Data from redundant sources is consistent

Table 1: Automotive Anomaly Detection Sensors

sages which normally only occur in combination on specific sub-networks adhere to the defined specification. This allows recognizing attacks where the access of the attacker is limited to a particular bus system or domain.

S-7: Protocol Sensor Several devices in the vehicle implement small communication protocols on a challenge-response basis. Exemplary applications for such protocols are the diagnosis functions at system startup or the key exchange of the electronic immobilizer. Even without knowledge of the keys, the Protocol Sensor monitors the traffic with respect to the specification of these challenge-response protocols, e.g., by checking if somebody tried to tamper with the order of the messages in the protocol, if the timing (e.g., start- and end point-of-time) of the protocol is valid, etc.

S-8: Plausibility Sensor The Plausibility Sensor considers the semantics of the message payload and checks if the data content is realistic. Implausible data can be values which stay within their defined data range, but show infeasible correlation with previous values or other messages of that domain. An example would be a sequence of messages containing the vehicle speed which is shifting from 20 km/h to 200 km/h and backward immediately without sufficient intermediate values. A formal specification of such relations, which is applicable here, has been illustrated in the paper by Larson et al. [14]. In our case, a restriction to reliable and non-heuristic definitions ensures that only true positives are indicated by the system.

S-9: Consistency Sensor The Consistency Sensor examines the semantics of the message payload, but in contrast to the Plausibility Sensor it is not limited to a specific sub-network or domain. Instead, it can access various data sources in the car. The Consistency Sensor uses the fact that several events trigger consequences and effects which are noticed by different components, sensors or ECUs in the vehicle. In particular, the sensor operates in such a way that it verifies the correctness of the data by using redundant or duplicate information, which can be acquired from different sources in the vehicle. An exemplary event the Consistency Sensor would indicate, is the situation that the tire rotation sensors show the vehicle is standing, but the GPS sensor of the navigational system indicates a movement.

To summarize, the contribution of the anomaly detection sensors does not lie in the individual complexity of each detection criterion, but in the investigation and extraction of the critical factors a typical modern vehicular network is characterized by, and the combination of these factors into a holistic IDS scheme allowing the recognition of in-vehicle threats without generating false positives. An overview of the sensors is given in Table 1.

B. Applicability of Detection Sensors

A comparison of the different sensor types reveals that for each sensor different requirements, conditions and access options hold. For instance, whereas some sensors only require a single packet for a successful detection, others need a number of messages for being able to work.

This paper identifies six applicability criteria, which show the requirements and working conditions of the sensors. In the following we explain these criteria and discuss the consequences each criterion implies. An overview of the applicability criteria and the corresponding parameter values for each anomaly detection sensor is given in Table 2.

Criterion / Detection Sensor	Specification-Based	Number of Messages	Number of Bus Systems	Different Message Types	Payload-Inspection	Semantic-Based
Formality	$true$	1	1	$n.a.$	$false$	$false$
Location	$true$	1	1	$n.a.$	$false$	$false$
Type	$true$	1	1	$n.a.$	$true$	$false$
Range	$true$	1	1	$n.a.$	$true$	$false$
Frequency	$true$	n	1	$false$	$false$	$false$
Correlation	$true$	n	n	$true$	$false$	$false$
Protocol	$true$	n	n	$true$	$false$	$false$
Plausibility	$false$	n	1	$false$	$true$	$true$
Consistency	$false$	n	n	$true$	$true$	$true$

Table 2: Applicability of in-vehicle anomaly detection sensors

1) AC-1: Specification-Based

Vehicular networks have very strict specifications for the communication system including every message that is allowed on a bus system. For CAN, these specifications are covered in the CAN-Matrix of the specific network. Therefore, criterion *AC-1* describes if the result of the sensor can reliably be determined *only* with the help of the specification, like the CAN-Matrix. Otherwise, e.g., if further data sources are required or attack patterns have to be defined the value is *false*. For an integration into the vehicle this criterion means, that the specification needs to be included into the sensor but no further data, e.g., through the wiring to a redundant data source, is required.

2) AC-2: Number of Messages

This criterion refers to the minimum number of messages required for this sensor. We distinguish between one and many messages (n). A one here always implies a one for the criterion *number of bus systems* and makes criterion *AC-4* non-applicable *(n.a.)*. Sensors which require more than one message usually have higher hardware requirements with respect to performance, memory, etc.

3) AC-3: Number of Bus Systems

This criterion means the minimum number of bus systems the sensor needs access to in order to perform a detection. We distinguish between one and many bus systems (n). The integration of sensors into the vehicle which require access to multiple bus systems is more complex and requires higher efforts. The multiple access points can either be included into a central gateway or can be placed in a distributed manner (see Sect. II-C).

4) AC-4: Different Message Types

This criterion is *false* if one type of messages can be sufficient for a detection, and *true* if two or more message types are necessary. It is not applicable if criterion *AC-2* is one, indicated by *n.a.*. In the context of CAN two messages are of the same type if they have an identical identifier, meaning the ECUs addressed by this message are the same but the values transmitted can be different.

5) AC-5: Payload-Inspection

This criterion describes if at least one part of the payload of a message is taken into account. One major implication of this parameter value is, that if the value is *true* the sensor can only process unencrypted messages as in general no read access to an encrypted payload is possible. Although currently most in-vehicle networks do not use encryption, this might be a very important aspect in the future. Usually, a payload-based sensor implies higher performance requirements for the anomaly detection system since the entire payload needs to be read and processed.

6) AC-6: Semantic-Based

This criterion is *true* if semantic aspects of the payload are considered. Obviously, it can only be true, if the payload is taken into account. However, even when the payload of a packet is considered the semantic meaning of the data is not always relevant, e.g., when only a range check of the payload content is performed.

C. Towards a Classification

The applicability criteria can be used to organize and structure the different sensors we described for the detection of anomalies in vehicular networks. Therefore, we determine two key applicability criteria which are suitable to classify the set of anomaly detection sensors. Based on our first experiences with the sensors, we identify *AC-2* (Number of Messages) and *AC-5* (Payload-Inspection) as potential key criteria and receive the classification shown in Fig. 1.

Both applicability criteria are suitable for a classification because they do not influence each other and their values can be clearly and unambiguously determined. *AC-2* is a major criterion, because the minimum number of messages required for detection has proven to cause strong implications for the design and implementation complexity of a detection sensor. If the payload of a packet is inspected, the requirements for the performance of a sensor usually are much higher. This is a crucial fact which underlines the relevance of criterion *AC-5*, because performance, and especially its financial implications, are critical aspects in the highly cost-driven automotive industry [16].

Fig. 1 shows an arrangement into four classes: The two leftmost classes are packet-based, the two rightmost classes are stream-based as they consider multiple messages. If we assume an increasing complexity for payload-inspection, the classes *Packet-Inspection* and *Stream-Inspection* can be considered to have a higher complexity for implementation and realization in the automotive domain. Fig. 1 includes a mapping to the sensors introduced in Table 1, which serve as examples for each class. Consequently, the list of sensors in the classification might be supplemented at a later point of time, e.g., if new technical possibilities arise or the focus is driven towards another vehicular bus system.

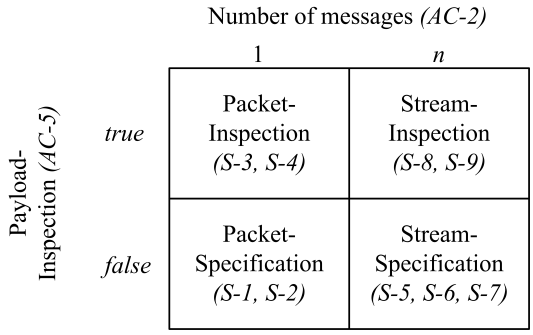

Figure. 1: Classification of Anomaly Detection Sensors

IV. Integration of Sensor Results

The anomaly detection sensors described in Sect. III-A bear the advantage that no false positives are produced. In this section, we show how the results of the sensor data can be evaluated furthermore, to facilitate a straightforward integration of the approach into a holistic IDS concept for the automotive domain.

A requirement for an integration of an IDS approach into the automotive domain is an evaluation of how severe a situation needs to be rated. Although a sensor recognizing an anomaly automatically implies that something is wrong within the vehicular system, a detection engine can support the notification and reaction phase of an IDS with additional information. This information can be used to support the decision of how to react when an anomaly has been detected. For instance, it may not always be inevitable to notify the current user of a successfully detected anomaly. In this approach, especially two criteria are considered for the evaluation:

- Number of recognized events

Under certain circumstances, a single anomaly might still be tolerable up to some degree, as it would usually just result in an error message and a retransmission by the sending ECU. An exemplary situation might be an incorrect CAN identifier or message checksum occurring in a specific CAN message, which is caused by disturbances or perturbation in regard to electromagnetic compatibility (EMC). A situation, however, where not just a single message is affected but suddenly the percentage of retransmissions in the network strongly increases, is considered much more critical as it can be an indicator of an attack. Consequently, the number of recognized events within a specific period of time is regarded in the evaluation.

- Type of sensors recognizing an event
  The impact an anomaly has for the system can be evaluated with the help of the type of sensor it was detected with. Some sensors can be more important than others, and therefore imply a higher criticality. A spurious message detected by a Location Sensor, for instance, may be more significant than a single incident recognized by a Frequency Sensor, due to the fact that an entirely dislocated message is much more unusual than a slight and singular drift in frequency.

In the following, the integration method is developed step by step and both aspects are incorporated into a holistic method for an estimation of the criticality of an incident.

An incident is a situation where at least one sensor S_i recognizes an anomaly at a certain point in time. Therefore we model the sensor S_i as a function of time to the range $\{0, 1\}$. Detection of an anomaly at time t is then denoted as $S_i(t) = 1$, where $i \in \{1, ..., n\}$ is the type of sensor (currently $n = 9$). We assume that the time domain is discrete and not continuous, i.e., there is a mapping of the time domain to the natural numbers.

We introduce a basic estimation of how critical an incident is by a separation into three classes with an increasing criticality:

$$C = \{important, critical, severe\} \tag{1}$$

These classes allow to differentiate between three basic criticality levels, which facilitate a reaction to anomalies and a suitable notification of the driver, e.g., by different optical, acoustic or haptic measures according to the *adaptive dynamic reaction model* proposed by Hoppe et al. [8].

If we define a weight w_i determining the impact for every sensor, the accumulated weights of all sensors at a time t can be acquired by

$$\mathit{Crit}(t) = \sum_{i=1}^{n} S_i(t) w_i. \tag{2}$$

This equation can be used to estimate the *criticality* of an incident at a given point in time t, based on the assumption that a larger number of sensors detecting an anomaly as well as higher weighted sensors result in a more critical classification.

Under certain circumstances, a single anomaly, like an incorrect checksum in a specific CAN message which is caused by disturbances in regard to electromagnetic compatibility (EMC), might still be tolerable up to some degree as it would usually just result in an error message and a retransmission by the sending ECU. A situation, however, where not just a single message is affected but suddenly the percentage of retransmissions in the network strongly increases, is considered much more critical. Therefore, we use a sliding window approach to include previous events into the evaluation and define $X_i(t)$ as the sum of all incidents for sensor S_i within the last window of size *SLW* up to time t:

$$X_i(t) = \sum_{j=t-\mathit{SLW}}^{t} S_i(j) \tag{3}$$

Note that we assume that time is not continuous, i.e., measurements are taken at discrete points in time and can therefore be summed up instead of using integration.

Equivalently to equation 2, we estimate the criticality of an incident with respect to previous events and define thresholds T for each criticality class C: $T_{\mathit{important}} < T_{\mathit{critical}} < T_{\mathit{severe}}$. This leads to the following equation for the detection of a critical event:

$$\sum_{i=1}^{n} X_i(t) w_i > T \tag{4}$$

We divide the equation through the arithmetic mean of all weights and adjust the thresholds appropriately (indicated by T'), in order to uncouple the threshold from the individual weights and number of sensors n. Hence, we receive

$$\frac{n}{\sum_{i=1}^{n} w_i} \cdot \sum_{i=1}^{n} X_i(t) w_i > T' \tag{5}$$

which gives an estimation of the criticality of an event at time t. Here, the weights allow a balancing of the different sensor types and the sliding window ensures that an accumulated situation is evaluated. A reasonable size for the sliding window still has to be identified. We expect a strict lower boundary size to be the highest loop period of all affected, cyclic CAN messages defined by the CAN-Matrix. However, the optimal size of the sliding window still needs to be determined by future investigations.

V. Conclusion

After describing challenges in the area of in-vehicle networks, this paper presented a threat detection scheme for such automotive networks – in this case examplified and oriented to the CAN bus – that is based on a set of sensors. These sensors can serve as real-time criteria for the recognition of IT-security related threats and do not cause false positives. We suggested to integrate this reactive approach into the security architecture of future vehicles in addition to preventive measures to build up a holistic attack protection. Moreover, different characteristics and typical limitations of automotive networks that affect the design and deployment of threat detection systems have been considered. The shown indicators and measures provide a good fundamental toolbox

for the detection of several threats in order to achieve a reasonable basic level of security for detecting attacks on in-vehicle networks. But of course, intelligent attacks are still possible without detection. This, for instance, is the case if an attacker is able to inject messages that are fully compliant to the network's normal behavior and plausible to previous values. So overall, there is still much future work to be performed.

References

[1] Ajith Abraham, Crina Grosan, and Yuehui Chen. Cyber Security and the Evolution of Intrusion Detection Systems. *Journal of Information Assurance and Security*, 1(1):74–81, 2005.

[2] Algirdas Avizienis, Jean-Claude Laprie, Brain Randell, and Carl Landwehr. Basic Concepts and Taxonomy of Dependable and Secure Computing. In *IEEE Transactions on Dependable and Secure Computing*, 2004.

[3] Andrey Bogdanov, Thomas Eisenbarth, Christof Paar, and Marco Wolf. Trusted Computing in Automotive Systems. In *Chapter in "Trusted Computing", N. Pohlmann and H. Reimer (Eds.)*. Vieweg, 2007.

[4] Hervé Debar, Marc Dacier, and Andreas Wespi. Towards a Taxonomy of Intrusion Detection Systems. *Computer Networks: The International Journal of Computer and Telecommunications Networking*, 31(9):805–822, Apr 1999.

[5] Mark Heitmann. Security Risks and Business Opportunities in In-Car Entertainment. In *Embedded Security in Cars*, pages 233–246. Springer, 2006.

[6] Tobias Hoppe and Jana Dittmann. Sniffing/Replay Attacks on CAN Buses: A Simulated Attack on the Electric Window Lift Classified using an Adapted CERT Taxonomy. In *CD-Proceedings of the 2nd Workshop on Embedded Systems Security (WESS 2007)*, 2007.

[7] Tobias Hoppe, Stefan Kiltz, and Jana Dittmann. Security Threats to Automotive CAN Networks - Practical Examples and Selected Short-Term Countermeasures. In *Proceedings of the 27th International Conference SAFECOMP 2008*, Newcastle, United Kingdom, September 2008.

[8] Tobias Hoppe, Stefan Kiltz, and Jana Dittmann. Applying Intrusion Detection to Automotive IT - Early Insights and Remaining Challenges. In *Journal of Information Assurance and Security*, volume 4, pages 226–235, 2009.

[9] CAN in Automation (CiA) e.V. CANopen application layer and communication profile, CiA draft standard 3.01, January 2005.

[10] Frank Kargl, Panos Papadimitratos, Levente Buttyan, Michael Müter, Björn Wiedersheim, Elmar Schoch, Ta-Vinh Tongh, Giorgio Calandriello, Albert Held, Antonio Kung, and Jean-Pierre Hubaux. Secure vehicular communications: Implementation, performance, and research challenges. *IEEE Communications Magazine*, 46(11):2–8, November 2008.

[11] Richard A. Kemmerer and Giovanni Vigna. Intrusion detection: A brief history and overview. *Computer*, 35(4):27–30, Apr 2002.

[12] Karl Koscher, Alexei Czeskis, Franziska Roesner, Shwetak Patel, Tadayoshi Kohno, Stephen Checkoway, Damon McCoy, Brian Kantor, Danny Anderson, Hovav Shacham, and Stefan Savage. Experimental Security Analysis of a Modern Automobile. In *IEEE Symposium on Security and Privacy*, 2010.

[13] Andreas Lang, Jana Dittmann, Stefan Kiltz, and Tobias Hoppe. Future Perspectives: The Car and its IP-Address - A Potential Safety and Security Risk Assessment. In *Computer Safety, Reliability, and Security, Proceedings of the 26th International Conference SAFECOMP*. Springer, 2007.

[14] Ulf E. Larson, Dennis K. Nilsson, and Erland Jonsson. An Approach to Specification-based Attack Detection for In-Vehicle Networks. In *Proceedings of the IEEE Intelligent Vehicles Symposium*, 2008.

[15] Jesus Molina and Michel Cukier. Evaluating Attack Resiliency for Host Intrusion Detection Systems. In *Journal of Information Assurance and Security*, volume 4, pages 11–19, 2009.

[16] Michael Müter and Felix C. Freiling. Model-based Security Evaluation of Vehicular Networking Architectures. In *Proceedings of the Ninth International Conference on Networks (ICN 2010), IARIA*, Les Menuires, France, April 2010.

[17] M. Murali, A. Rao. A Survey on Intrusion Detection Approaches. In *International Conference on Information and Communication Technologies, ICICT*, 2005.

[18] Nicolas Navet, Yeqiong Song, Francoise Simonot-Lion, and Céderic Wilwert. Trends in Automotive Communication Systems. *Proc. of the IEEE*, 93(6):1204–1223, 2005.

[19] Dennis K. Nilsson and Ulf E. Larson. Combining Physical and Digital Evidence in Vehicle Environments. In *SADFE '08: Proceedings of the 2008 Third International Workshop on Systematic Approaches to Digital Forensic Engineering*, pages 10–14, Washington, DC, USA, 2008. IEEE Computer Society.

[20] Dennis K. Nilsson and Ulf E. Larson. Conducting Forensic Investigations of Cyber Attacks on Automobile in-vehicle Networks. In *Proceedings of the 1st International Conference on Forensic Applications and Techniques in Telecommunications, Information, and Multimedia*, 2008.

[21] Chrisof Paar and André Weimerskirch. Embedded security in a pervasive world. In *Elsevier Science's Information Security Technical Report*, page 55–161, 2007.

[22] A. Qayyum, M.H. Islam, and M. Jamil. Taxonomy of statistical based anomaly detection techniques for intrusion detection. In *Proceedings of the IEEE Symposium on Emerging Technologies*, 2005.

[23] Marko Wolf. *Security Engineering for Vehicular IT Systems*. Vieweg + Teubner, 2009.

[24] Marko Wolf, André Weimerskirch, and Thomas Wollinger. Sate of the Art: Embedding Security in Vehicles. In *EURASIP Journal on Embedded Systems (EURASIP JES), Special Issue: Embedded Systems for Intelligent Vehicles*, 2007.

[25] Yan Zhai, Peng Ning, Purush Iyer, and Douglas S. Reeves. Reasoning about complementary intrusion evidence. In *ACSAC '04: Proceedings of the 20th Annual Computer Security Applications Conference*, pages 39–48, Washington, DC, USA, 2004. IEEE Computer Society.

Author Biographies

Michael Müter holds a degree in Computer Science from University of Technology, Aachen, Germany and joined Daimler AG as a security researcher after his studies. In recent years, he has been involved in different projects in the area of vehicular communication systems, like SeVeCom or SEIS, and several other security-related projects. His research interests include secure communication systems and vehicular security architectures, but also comprise more theoretical aspects of reliable and dependable systems.

André Groll studied information systems focussing on IT Security at the University of Siegen. In 2006 he received his diploma with a thesis on steganography and digital watermarking in the banking industry. Afterwards he started to work as a Scientific Assistant at the Institute for Data Communications Systems in order to continue his postgraduate studies to receive a PhD concentrating on automotive security. Mr. Groll is a member of the eSafety security WG of the European Commission and gained much experience in several automotive projects.

Felix Freiling is a full professor of computer science at the University of Erlangen, Germany. His research interests cover theoretical and practical aspects of computer security. Felix holds a Diploma degree and a PhD in computer science from TU Darmstadt, Germany. Before joining Erlangen he held positions at RWTH Aachen University and the University of Mannheim, Germany.