Date of publication xxxx 00, 0000, date of current version xxxx 00, 0000.

Digital Object Identifier 10.1109/ACCESS.2020.DOI

TIDAL-CAN: differential Timing based Intrusion Detection And Localization for Controller Area Network

PAL-STEFAN MURVAY[1] and BOGDAN GROZA[2], (Member, IEEE)

[1]Faculty of Automatics and Computers, Politehnica University of Timisoara, 300223 Timisoara, Romania (e-mail: pal-stefan.murvay@aut.upt.ro)
[2]Faculty of Automatics and Computers, Politehnica University of Timisoara, 300223 Timisoara, Romania (e-mail: bogdan.groza@aut.upt.ro)

Corresponding author: Pal-Stefan Murvay (e-mail: pal-stefan.murvay@aut.upt.ro).

This work was supported by a grant of the Romanian Ministry of Research and Innovation, CNCS - UEFISCDI, project number PN-III-P1-1.1-PD-2016-1198, within PNCDI III.

ABSTRACT Since the first reports on its lack of security, the Controller Area Network (CAN) was in focus for numerous research works. A specific area of research has employed physical layer characteristics that can be used to uniquely identify network nodes. But there are common downsides in existing approaches such as vulnerabilities in front of attacks involving node replacement or insertion or the inability to locate the intruder node within the network. In this work, we propose a new intrusion detection system for CAN which is based on monitoring the propagation time of the physical signals sent on the bus. Indeed, quite a number of recent works addressed the use of physical or timing characteristics to identify network nodes or to create covert channels. In our approach, by accounting for intrinsic delay characteristics of the bus and by monitoring the difference in signal arrival time at the two bus ends, we can identify nodes by location-related differential delays and provide relevant information for estimating the relative location of a transmitter node on the bus. The results of our experimental evaluation show that our approach provides very high identification rates and accurate localization in case of attacks from compromised nodes. The ability to detect attacks that replace an existing node or plug new adversarial nodes on the bus is also illustrated along with discussions on estimating sender location in these cases.

INDEX TERMS Automotive engineering, Communication system security, Intrusion detection, Physical layer

I. INTRODUCTION

Researchers were quick to react to reports such as [21], [4] or [24] that prove the existence of exploitable vulnerabilities in modern automobiles. As a result, many security mechanisms designed for protecting the in-vehicle network were proposed to date, e.g., a brief summary can be found in [12], and the industry also reacted by including security specifications in automotive standards, e.g., [2]. Some of these works tackle the issue by assuring authenticity and possibly confidentiality for in-vehicle communication while others focus on detecting intrusions alone. Since Controller Area Network (CAN) is the most commonly employed protocol for in-vehicle communication and given that it was targeted by the majority of the reported attacks, it is natural that most of the proposed security mechanisms for in-vehicle networks are designed for CAN, though other network types may exist inside cars.

Our work also focuses on the security of CAN and, in particular, on designing an efficient intrusion detection system (IDS). Based on the source of data employed for intrusion detection, existing proposals can be categorized either as application layer-based or physical layer-based IDSs. Application layer-based IDSs use traffic data made available at the application layer such as frame content or frame arrival timings. On the other side a physical layer IDS identifies intruders based on the physical characteristics of the signals captured from the communication line. Our line of work falls into this second category. Previous proposals have used unique characteristics of signals generated at transmissions by each individual node from the bus, which are caused both by node characteristics and the transmission medium depending on node's location, a comprehensive discussion on this follows in the related work section. Our current proposal is based only on characteristics of the transmission medium and network structure, by specifically using signal

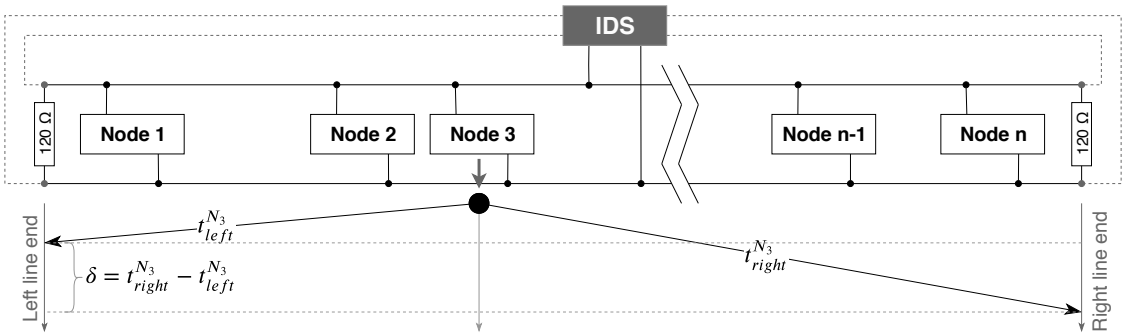

FIGURE 1. Example of a CAN bus with unevenly distributed nodes as considered by the addressed scenario.

propagation timings.

Targeted scenario and proposed solution. As recommended by its specification, CAN is most commonly implemented following a line/bus topology, although it allows implementations that use other topologies, e.g., passive star, which are less common. The scenario that we address in this work is illustrated in FIGURE 1 which depicts an example of a CAN bus, as it is commonly found inside vehicles, with nodes being usually unevenly distributed on the bus due to specific electronic control units (ECU) positioning restrictions and requirements. One of the bus nodes is responsible of identifying possible intruders (existing nodes that were compromised or additional nodes introduced on the bus, e.g., via the diagnostics port).

The mechanism that we envision for intrusion detection relies on monitoring signal arrival times at the bus ends. Since the propagation delay from the transmission point to each of the bus ends depends on the characteristics of bus nodes and the position of the transmitter, the IDS node uses the time difference between the two arrival times to identify the transmitter. By correlating differential delays associated to the transmitter with the type of the sent message the IDS node can identify illegal transmissions and estimate the location of the attack attempts source.

Paper contributions. The main contributions presented in this paper are listed as follows:

- We perform a comparative analysis of the existing physical layer-based CAN intrusion detection mechanisms.
- We propose a novel Intrusion Detection And Localization mechanism for CAN (TIDAL-CAN) which uses time differences in signal propagation from the transmission point towards the bus ends as a location-based characteristic of the sender node.
- The proposed mechanism is also able to distinguish between various types of attack approaches, i.e. compromised nodes, replaced nodes and node insertion.
- Experimental results are provided to prove the feasibility and efficiency of the proposed approach.

The rest of the paper is organized as follows. In Section II we provide a brief background on CAN and discuss related work. Section III discusses propagation delays and an electrical model of the bus. In Section IV we establish the system and threat model before proceeding to section V in which we give an overview of the proposed intrusion detection mechanism and account for various attack scenarios. In Section VI we present experimental results with the proposed methodology. Finally, Section VII holds the conclusion of our work.

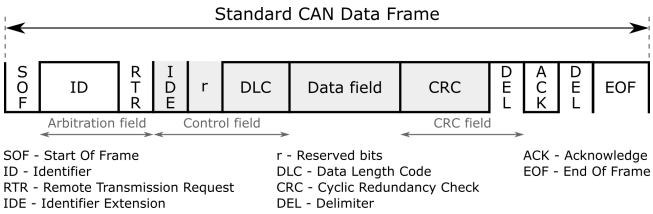

FIGURE 2. Structure of a standard CAN frame.

II. BACKGROUND AND RELATED WORK

This section provides a brief background on the CAN bus followed by related work on intrusion detection for CAN.

A. BRIEF BACKGROUND ON CAN

The CAN protocol was designed to provide reliable communication for noisy environments such as inside motor vehicles. It supports bit rates of up to 1 Mbit/s (although usually used at up to 500 kbit/s in real-life implementations) with a maximum per-frame payload of 8 bytes. FIGURE 2 illustrates the basic structure of a standard CAN frame. The identifier (ID) field is usually used to designate either the sender of the frame or the type of message being sent. The sender designated at network design time for a given frame can be determined at the application layer by ID and payload content. However, CAN offers no intrinsic mechanism to assure that the message originator is indeed as expected.

At the physical layer CAN uses a two-wire differential signaling model, each node being connected to the main bus line through a line stub. The CAN bus must be properly terminated at each end to suppress signal reflections with an 120 Ω resistor which can be either independently positioned at the line end as in FIGURE 1 or as part of the bus end nodes.

The physical layer signaling uses two states to encode logical information. The recessive state, present when none of the network nodes is actively driving the bus, represents a logical "1". The dominant state is the result of at least one node driving the bus and is interpreted as logical "0". As a consequence, dominant bits always overwrite recessive ones.

Physical layer signaling corresponding to some of the CAN fields can be the result of transmissions from multiple nodes due to two specific CAN mechanisms: arbitration and message acknowledgment. Access to the bus is gained based on the CAN arbitration mechanism. Therefore, when two or more nodes simultaneously start a frame transmission it is decided, based on the arbitration mechanism, which node is allowed to continue transmission past the arbitration field. Priorities are established mainly based on the frame identifier. Another CAN-specific mechanism is message acknowledgement. A sender of the frame transmits a recessive bit during the acknowledge field. All receivers on the bus must generate a dominant level during this bit to signal the correct frame receival.

B. RELATED WORK

As related work, we consider intrusion detection mechanisms based on both timing characteristics and physical layer characteristics. A good overview of recent approaches is given in [17]. In this section we separately address related work on the two categories of intrusion detection mechanisms.

1) Timing-based intrusion detection

Due to the periodic nature of the majority of CAN frames, it is no surprise that many IDS proposals for CAN have been focusing on frame arrival time, e.g., [32], [25]. In a somewhat distinct approach, the work in [22] uses timing differences between data frames and remote frames to detect intrusions. For this type of intrusion detection systems however, the measurement precision of the frame timing is not so critical in contrast to our work were the employed precision is in the order of several nano-seconds for correct identification. Better accuracy for frame timing is required by the proposal in [5] which uses clock skews to detect intrusions. This proposal is however vulnerable to cloaking attacks since an attacker can maliciously modify his clock skew to impersonate another node as proved by [31]. Our approach remains invulnerable to such attacks since we monitor delays at both bus ends. Frame arrival time has been also used to create time-covert channels that can carry additional authentication data in recent works such as [13] and [35].

2) Physical layer-based intrusion detection

TABLE 1 summarizes the comparison between basic features of existing intrusion detection approaches employing physical layer characteristics. In what follows, we present the basics of these mechanisms while leaving a comparison of their detection abilities to the experimental section of our paper.

TABLE 1. Comparative analysis of physical layer-based CAN intrusion detection mechanisms

Paper	Sampling rate	CAN bit rate (max)	Can detect attack by		
			Compromise	Replace	Insert
Murvay et al. '14 [26]	2 GS/s	125 kbps	✓	-	-
Cho et al. '17 [6]	50 kS/s	500 kbps	✓	-	-
Choi et al. '18 [7]	2.5 GS/s	500 kbps	✓	-	✓
Choi et al. '18 [8]	2.5 GS/s	500 kbps	✓	-	-
Kneib et al. '18 [18]	20 MS/s	500 kbps	✓	-	✓
Kneib et al. '20 [20]	2 MS/s	500 kbps	✓	-	✓
Foruhandeh et al. '19 [11]	50 MS/s	500 kbps[1]	✓	-	-
Rumez et al. '19 [27]	$\geq$ 2GS/s[2]	any	-	-	✓
TIDAL-CAN (our proposal)	250 MS/s	any	✓	✓	✓

[1] Extracted from the associated dataset [11].
[2] Estimated based on paper details.

The first study of the use of CAN bus voltage levels to uniquely identify the source of frame transmission is presented in [26]. This work briefly presents the idea of using voltage measurements of the signals transmitted by a node to extract unique transceiver characteristics. Simple signal processing mechanisms such as mean squared error, convolution and mean-value are applied over multiple samples from the arbitration field preamble of each frame.

Improvements of this previous approach are proposed by Choi et al. [7]. They use samples from extended identifiers from which they extract a set of 17 features. Classification algorithms are then used to fingerprint and identify nodes with a higher accuracy than the previous work.

Further improvements for the same general approach are provided by the Scission [18] intrusion detection system. Scission achieves a better detection accuracy than its predecessors by using only samples acquired after the arbitration field and focusing on signal characteristics arround rising and falling edges. The authors also use a lower sampling rate than previous work and simpler machine learning algorithm in comparison to [7]. To cope with the resource-constraint nature of automotive embedded platforms, Kneib et al. propose EASI, a newer embodiment of the approach presented in Scission that can be implemented using sampling circuitry commonly integrated in automotive-grade microcontrollers [20]. They achieve this by using random interleaved sampling with standard analog-to-digital converters reducing the actual sampling rate requirements to as low as 2 MS/s.

In another line of work, Choi et al. [8] also propose using only the rising and falling edges of the transmitted signal as a source of unique transmitter properties. Feature extraction is performed on the acquired samples and classifiers are used for attack detection. The authors also explore the possibility of using lower sampling rates but at the cost of considerable degradation of detection accuracy.

A common downside of most of the works presented so far is the lack of an actual implementation on an embedded platform and the high cost generated by the module required to provide the high sampling rate employed. In constrast, the Viden [6] attacker identification scheme uses a very low sampling rate (50 kS/s) to gather measurements of the dominant levels found throughout a frame transmission. Voltage profiles are then built based on the sampled data and used to identify the transmitter node. The authors of Viden provide

a proof of concept implementation and extensive analysis on experimental and real-world CAN networks demonstrating its high detection accuracy.

The accuracy of voltage based IDSs can suffer if signal variations caused by evironmental conditions are not considered. The effect of temperature variations is considered in a folllow-up paper of the authors of Scission [19] to improve detection accuracy of their initial proposal. SIMPLE [11] is another IDS that accounts for both temperature and voltage variations by implementing a secure update procedure for node fingerprints. While SIMPLE fits in the category of IDSs that report lower sampling rate requirements the actual sampling rate required increases along with the increase in the bit rate employed on the target CAN bus.

The work done by Rumez et al. [27] employs a completely different approach by using time domain reflectometry (TDR). They measure the network response to a pulse sent by the detection unit and determine if the network structure (i.e. connected nodes) is as expected at the moment of the measurement based on prerecorded reference responses. This approach allows the identification of disconnected nodes or introduction of new nodes on the bus and also permits the localization of the introduced nodes. However, since this approach does not make any correlation between a node and the messages that it transmits it would be unable to detect the case of existing network nodes that are compromised (e.g., by rewriting firmware) or to distinguish between a malign and a benign device connected on the diagnostics port.

III. CAN BUS TRANSMISSION DELAYS

The signals generated by a CAN node propagate from their point of origin toward the ends of the bus. The physical medium used to transmit information introduces a propagation delay, i.e. the amount of time needed for a signal to travel from its point of origin through the medium to its receiver(s). Physical characteristics of the transmission medium, such as its length and propagation speed, as well as the number and behavior of network nodes will influence the propagation delay. In this section we give an overview on the sources of propagation delays and discuss a theoretical model for bus electrical properties.

A. TRANSMISSION LINE PROPAGATION DELAYS

The time required for bus signals to reach a certain point on the bus mainly depends on the specific propagation speed of the line and the distance from the signal point of origin to the receiving node. This line dependent delay can be estimated based on the distributed model of the transmission line.

A lossy transmission line can be modeled as an infinite number of elementary components connected in series, as depicted in FIGURE 3 (a). Each elementary component represents an infinitely small segment of the line. Each such segment is characterized by a series resistance R, a series inductance L, a conductance G, caused by the imperfect insulation between two bus conductors, and a parallel capacitance C. The values of these parameters are given per

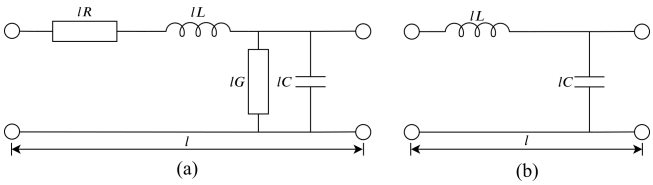

FIGURE 3. Infinitesimal elementary component of a lossy (a) and lossless (b) transmission line

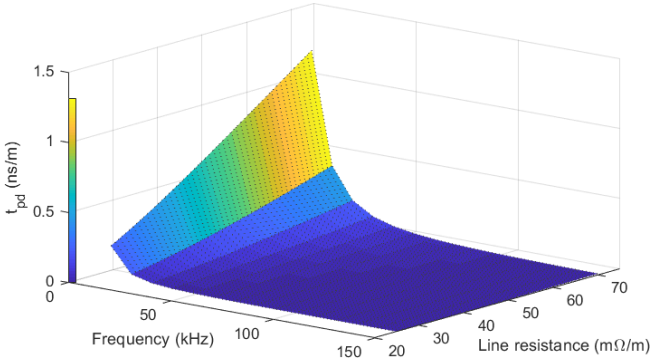

FIGURE 4. Line delay vs. resistance

line unit length. The complex characteristic line propagation constant can be calculated based on these parameters as $\gamma(\omega) = \sqrt{(R + j\omega L)(G + j\omega C)} = \alpha(\omega) + j\beta(\omega)$, where $\alpha(\omega)$ is the attenuation factor and $\beta(\omega)$ is the propagation coefficient of the transmission line ($\omega = 2\pi f$). The value of the conductance G can be safely assumed to be equal to 0 since it is very small in comparison to the ωC component. The nominal line resistance specified in ISO11898-2 is 70mΩ/m [16] while other standards such as the SAE J1939 require it to be even lower, i.e. 25mΩ/m [28], [29]. The additional delay introduced by a line resistance in the order of tens of milliohms, at frequencies above 50 kHz, is negligible as shown in FIGURE 4. Therefore, for most of the nominal 10k-1Mbps range of CAN operating bit-rate, the influence of line resistance on propagation delays can be considered negligible. Thus, we obtain the simplified lossless line model as depicted in FIGURE 3 (b). The propagation constant for the lossless line becomes $\gamma(\omega) = j\omega\sqrt{LC}$ which, being purely imaginary, characterizes only the line propagation coefficient. The resulting propagation delay characteristic of the line is given by $t_{pd} = \sqrt{LC}$ (s/m).

The nominal value for the specific line delay in a high speed CAN bus, considering a homogeneous medium, is 5 ns/m as specified by ISO11898-2 [16]. Standards regulating the use of CAN in the automotive industry, such as J2284 and J1939, set the upper bound for the specific line delay to 5.5 ns/m [29], [30].

Propagation delays, along with other signal distortion characteristics of the transmission medium, are commonly considered when designing a CAN bus [23]. Limitations are imposed by specification for the maximum bus length that should be used for a certain bit rate to assure propper operation. The typical in-vehicle bus length is around 5-6

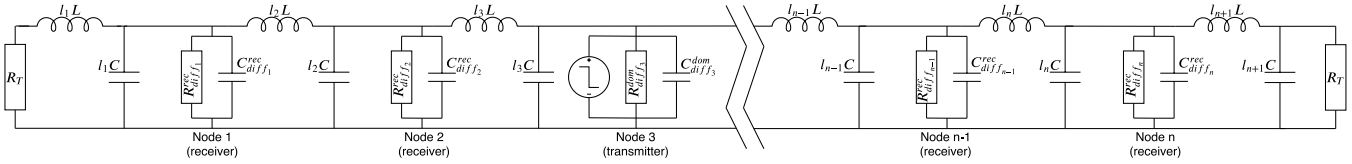

FIGURE 5. Equivalent model of a loaded CAN bus

meters with some works reporting the use of shorter networks down to 3 meters [10], or longer networks extending up to 15 meters [33]. Considering propagation speeds in the 5-5.5 ns/m interval, the transmission delay across the longest path (caused exclusively by the line delay) of a typical 5 m long CAN bus is between 25 and 27.5 ns.

B. DELAYS CAUSED BY LOADS ALONG THE BUS

Each CAN node connected to the bus acts as a load connected in parallel to the bus lines (i.e. CAN-Low and CAN-High) while sender nodes behave like voltage sources during message transmission. The load represented by each node is mainly caused by intrinsic characteristics of the CAN transceiver but can also include contributions from elements on the physical connection path (i.e. bus pins, connector contacts, connection stub, printed-circuit board traces, etc.). This load consists of a resistive component and a capacitive one. Therefore, each node can be modeled as a resistor-capacitor pair connected in parallel. The resistive load mainly consists of the transceiver differential input resistance R_{diff}, measuring in the order of tens of kilo-ohms (the CAN specification places R_{diff} in the 10-100kΩ interval). The capacitive load is mainly given by the transceiver internal differential capacitance C_{diff} which has a nominal value of 10pF during the recessive state [16]. Stub and connector capacitance also add to the load capacitance and can vary considerably [9]. The C_{diff} of each ECU, measured with the ECU disconnected from the bus, is limited, by standards, to a maximum of 50pF [28], [30].

FIGURE 5 depicts the equivalent model of a CAN bus. Each line segment connecting loads is represented by a lossless transmission line component, while nodes are defined as a parallel RC load. An additional voltage source depicts the transmitter node. The effects of receiver loads on propagation delay are illustrated in FIGURE 6 which shows the propagation delay over a 5 meter bus when no receiver loads are present and the effects of adding one receiver node with different load characteristics at various locations along the line. The propagation delay is measured at the 0.9V point (detection threshold for a dominant bit) on the rising edge of dominant bits. Signals were obtained from LTspice simulations of a 5 meter bus, built according to the model in FIGURE 5, with one transmitter node at one end of the line and a single receiver placed at various locations along the line. The plots show the influence of the resistive loads on the propagation delay in comparison to the effect of capacitive loads. Similar delays are obtained for distances up to around

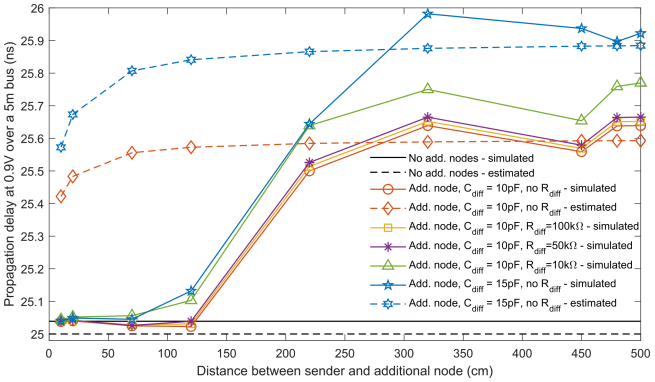

FIGURE 6. Influence of additional receiver node insertion on propagation delay (calculated as time difference at the 0.9V point of the rising edge between the signal at the transmission point and the signal at the end of a 5m bus)

2 meters when comparing the case of minimum specified R_{diff} to a 50% increase of the nominal C_{diff} (from 10 to 15pF). Still, it is visible that capacitive loads generally have higher impact on propagation delays than resistive loads. Increasing the distance between the added load and the sender node leads to increases in the propagation delay.

Since the R_{diff} component shows limited influence on the bus propagation delay, we can assume the load caused by network nodes as being capacitive. The delaying effect of the capacitive load comes in the form of an increase in the signal rise/fall time [1]. Therefore, the slope of the rising/falling edges is characterized by the loads along the signal propagation path.

C. ESTIMATING PROPAGATION DELAYS ON A LOADED LINE

In most transmission line modeling problems, loads are considered equal and uniformly distributed along the length of the line. In this case the capacitive loads are factored into the calculations of the specific line propagation delay as an additional distributed capacitance per unit length calculated as the total load capacitance divided by the line length [1].

However, for an in-vehicle CAN bus, nodes cannot be assumed to have equivalent loads, given the variation in functionality and manufacturer preferences, nor are they uniformly distributed, since node locations are restricted by physical locations of ECUs within the car. For estimating the propagation delay of a loaded CAN bus we consider that it consists of a number of segments. Each such segment spans the distance between a bus end and a node or between two

neighboring nodes. The delay from the transmission point to each of the bus ends is calculated as the sum of delays for each consecutive segment towards the bus end. The delay of each segment is estimated using the characteristic line delay expression and considering the characteristic line capacitance for the segment as $C_{seg_i} = C + C_{diff_i}/l_i$, where C_{diff_i} is the differential capacitance of the node that limits the segment toward the bus end, and l_i is the length of the i^{th} bus segment. If the bus terminator is not located on the last node, the delay for the last segment is computed as a pure line delay, i.e. $C_{diff} = 0$. Therefore, the estimated propagation delay from the transmission point to the left/right side end of the bus can be estimated as

$$t_{pd_{side}} = \sum_{i=1}^{n_{side}} l_i \sqrt{L(C + C_{diff_i}/l_i)}$$

FIGURE 6 shows the estimated delays corresponding to the same simulated 5 meter bus with one sender node and one receiver placed at various locations along the bus. The model provides estimates close to the simulated behavior for nodes that are placed farther from the transmission point. However, for receivers placed closer to the sender node the delay estimation is considerably higher than the simulation results. We will complement these theoretical estimations with concrete experimental measurements in the next section, our intention so far was to lay our theoretical background.

D. IDENTIFYING TRANSMITTERS BY PROPAGATION DELAYS

CAN signals transmitted by a node on the bus will propagate along the right and left-hand side of the bus with respect to the sender connection point. Based on the observations made in the previous sections we can state that the node's location on the bus directly influences the signal propagation delay to a fixed reception point. Therefore, by measuring the propagation delay of CAN signals, relative to a fixed receiver point, we could determine the transmitter location. Note that we consider the case of networks following a line/bus topology. However, using a single fixed receiver point to determine propagation delays of CAN messages would require the receiver knowing the actual message transmission start time and reliable time synchronization between the receiver and transmitter nodes. Even if these requirements are met, the receiving node would have to rely on the transmission timing information provided by the transmitter which could be falsified. Also, unless the receiver is always positioned at one of the bus ends, additional information is required to determine if the frame came from the left- or right-hand side of the bus relative to its location.

To alleviate these issues we envision a novel approach that does not require time synchronization nor knowledge about the message transmission start time. We propose the use of an intrusion detection node that monitors the bus traffic and the signals that arrive at each of the two ends of the CAN bus as suggested in FIGURE 1. An indication of the transmitter node location can be obtained by computing the

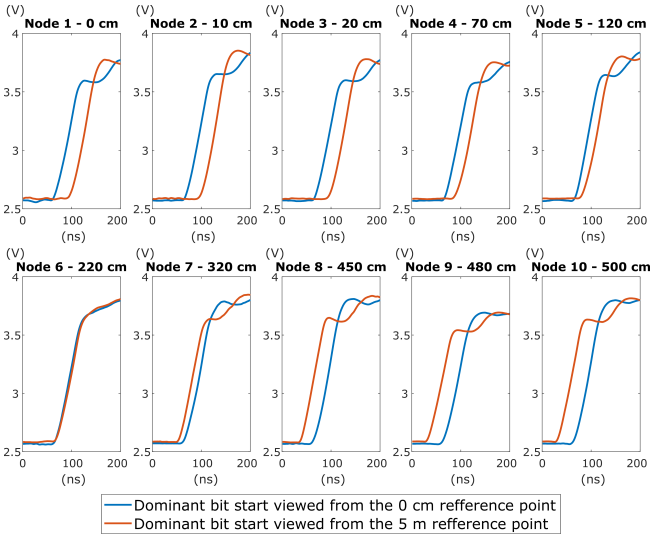

FIGURE 7. Variation of dominant bit arrival time at bus ends, on the CAN high line, depending on transmitter location

time difference between the arrival times of a CAN frame (or just a single specific signal edge within the frame) at the two bus ends. Therefore, if considering the right bus end as the reference point, the differential propagation time is $\delta = t^{N_i}_{right} - t^{N_i}_{left}, i = 1, n$, where $t^{N_i}_{right}$ and $t^{N_i}_{left}$ are the arrival times of the signal generated by node N_i at the right and left end of the bus respectively. This is equivalent with calculating $\delta = t_{pd_{right}} - t_{pd_{left}}$ without having to measure t_{pd} for the two bus sectors. The sign of δ hints on the location of the node against the point representing the bus center of mass with respect to propagation delay. For a given bus, the differential propagation time is only influenced by the transmitter node's connection point on the main bus lines. Therefore, since the location of a transmitter node on the bus uniquely characterizes this time difference it can be used to identify which was the transmitter node for a certain message.

FIGURE 7 illustrates the effect of bus propagation delays on arrival times at the ends of the bus. Plots show the start of a dominant bit captured from the CAN-High line, as it is received at the two ends of a 5 meter CAN bus populated by 10 nodes distributed unevenly along the line. Each plot represents a transmission made by a different node. Corresponding node locations are marked on the plots as relative distances to the reference end of the bus. As illustrated in the figure, the correlation between node positioning and bit arrival time differences is evident.

IV. SYSTEM AND THREAT MODEL

In this section we set the context that we target in our work by presenting the network and attacker model.

A. NETWORK MODEL

We consider networks or network segments following the CAN bus/line topology. Like all other physical layer-based

IDSs our proposal cannot be reliably applied over topologies that require signals passing through nodes since node delays can vary considerably according to loading. Therefore, when hybrid CAN topologies are employed to build the CAN network (e.g., star-bus topology commonly used to separate functional domains in vehicles), the IDS mechanism is applied independently on each bus sub-network. The operating principles that we employ are applicable to multidrop buses in general and are not limited to CAN or to a specific CAN physical layer. However, for simplicity of exposition, in this paper we illustrate the concepts considering their application for CAN with particular focus on high-speed CAN and CAN-FD. CAN-FD [14], is an extension of the standard CAN specification that provides larger payloads and higher bit rates to increase overall transmission bandwidth.

B. ATTACKER MODEL

We consider an attacker that has the ability to infiltrate the CAN network by one of three approaches: compromising a node, replacing an existing node or by introducing an additional node on the network. In-vehicle networks can no longer be assumed as being close to adversarial actions since, as proven by several lines of work, an existing network node could be compromised either through direct physical access to common in-vehicle interfaces (e.g., OBD, USB, Ethernet) [21], [34] or even remotely without having physical access to the vehicle [24], [3]. Node replacement or insertion can be done only by having physical access to the vehicle with the attacker knowingly placing the node programmed with malicious firmware. Using one of these approaches the attacker may perform spoofing, replay or DoS attacks on CAN communication. Our work mainly targets spoofing and replay attacks. It is well known that the arbitration mechanism makes CAN vulnerable to DoS attacks in which an adversary can prevent the transmission of other messages by continuously transmitting a higher priority message. Such attacks are out of scope for our current work.

An attacker is also aware of the IDS presence and its principle of operation. Hence, the adversary might attempt to mount attacks targeting the underlying IDS mechanism to evade detection. We assume that the attacker has the required knowledge on the network structure and characteristics which he gained through reverse-engineering or inside information.

V. INTRUSION DETECTION AND LOCALIZATION BASED ON PROPAGATION DELAYS

In this section we introduce TIDAL-CAN (differential Timing based Intrusion Detection and Localization for CAN), our proposal for intrusion detection on CAN, and its operating principle which uses characteristic delays of the CAN bus to identify network nodes and their location.

A. TIDAL-CAN OVERVIEW

The structure of the TIDAL-CAN intrusion detection and localization system is illustrated in FIGURE 8. TIDAL-CAN should be implemented on a network node that is continuously monitoring bus traffic. Each frame sent on the monitored bus is processed by TIDAL-CAN to determine if it was sent by a legit node.

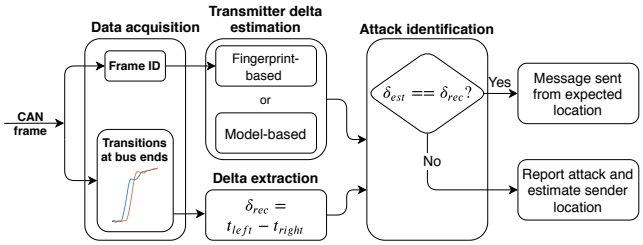

FIGURE 8. Schematic overview of the TIDAL-CAN intrusion detection and localization system

TIDAL-CAN first extracts the frame ID and samples the CAN signal, during a recessive to dominant or dominant to recessive transition, as it is seen at the two bus ends. The difference in the propagation times, i.e. δ_{rec}, is then calculated using the recorded signals from the bus. The expected transmitter of the frame corresponding to the received ID is identified and an estimate the propagation time difference δ_{est} is obtained based on its known location. Next TIDAL-CAN compares the estimated and the measured values and only declares the frame as transmitted by its legit node if $|\delta_{rec} - \delta_{est}| \leq \epsilon$, where ϵ is the error factor that accounts for small measurement and estimation errors. Upon detection of an illegitimate transmission, the frame can be directly classified either as coming from a known node that was compromised or as the result of bus modifications. We defer the discussion on identifying these cases to section V-D.

B. DATA ACQUISITION

During the data acquisition step TIDAL-CAN obtains the frame ID and signal samples required in the following steps. Obtaining the frame ID is straight-forward as it results from the standard frame decoding done by any CAN node. The signal sampling phase requires access to the signals as seen at the bus ends and careful selection of frame fields suitable for use.

1) IDS bus connections.

TIDAL-CAN must have access to the CAN physical layer signals as they arrive at the two ends of the bus. This is required in order to extract bus propagation delays. While using the differential CAN signal (i.e. CAN-High - CAN-Low) would be preferable, this would require sampling both lines at each end of the bus, thus increasing wiring requirements. TIDAL-CAN is able to efficiently extract propagation time differences using a single CAN wire (either CAN-High or CAN-Low). Therefore, the IDS node requires two signal sampling lines, each connected to one of the ends of the selected CAN signal line. The length and characteristics of the sampling lines will also have to be considered when processing sampled signals. The node also requires a classic

two-wire connection to the bus for normal message reception and transmission. Hence, an improvement, from the wiring point of view, would be to place the node at one end of the bus and run a single sampling line up to the other bus end. A two-node distributed signal acquisition mechanism could also be envisioned. However, such a solution would bring on new challenges such as secure transmission of sampled data to the processing node and time synchronization between nodes.

2) Sampling data.

For the purpose of measuring propagation delays we have to record rising or falling edges in the CAN bus signal. Therefore, the points of interest are recessive to dominant or dominant to recessive line state transitions. For a correct determination of transmitter location characteristics it is required that the sampled signals are the result of a single node actively driving the bus. Single node transmission can only be guaranteed during the *control*, *data* and *CRC* fields (gray fields in FIGURE 2) in a normal transmission of a CAN frame (i.e. in the absence of faulty or attacker nodes). Other nodes are allowed to transmit during the *arbitration* and *acknowledge* fields. Therefore, TIDAL-CAN will only perform signal acquisition of edges generated during the three aforementioned fields. In terms of number of edges required per frame, TIDAL-CAN provides good results with a single edge per frame. The use of multiple edges per frame could be used improve results in case of signals that exhibit higher noise levels.

C. EXTRACTION AND ESTIMATION OF TRANSMITTER δ

As a result of the second step in the mechanism, TIDAL-CAN uses the propagation time difference from the recorded signals δ_{rec} and an estimate on the expected δ_{est} for the transmission of the frame with the recorded ID.

1) Measuring the propagation delay δ.

The level on the rising/falling edge at which the δ_{rec} time difference is calculated must be selected to best reflect propagation delays so that results are consistent across devices. As presented in section III-B, the capacitive load consisting of nodes along the signal path affect the signal rise times. Therefore, the signals captured at the two bus ends might not have the same slope. The voltage level used by TIDAL-CAN to obtain δ_{rec} selected for each network node should provide maximum divergence between the two signals captured for each frame. Since the optimum level may not be the same for all transmission points, a common threshold value is established so that its usage assures best separation of differential delays for known network nodes. This threshold is established during production or as part of authorized network updates and does not change during normal system usage. The number of sample points also affects consistency of results. If required, TIDAL-CAN uses interpolation for a closer approximation of the sample time corresponding to the threshold value.

2) Estimating δ by frame type.

To determine if δ_{rec} corresponds to the expected value for the received frame, TIDAL-CAN needs to know the legit transmitter of frames of the received type. All frames that can be sent in a CAN network are defined during network design. Each frame is uniquely assigned an ID which can only be used by one specific sender. In some rare special cases part of the data field may be used for sender identification, still the frame is uniquely associated to a sender. Locations of existing nodes are also established at design time. Therefore, it can be considered that node locations and frame to node association are known and will not change under normal circumstances. We discuss circumstances for possible variations and the handling of these special cases in section V-E. TIDAL-CAN is able to detect the legit sender of a specific frame by using stored information about the network design. Also by having information about node location and network structure it can estimate propagation time differences by one of two methods. The first approach, which we cover in the experimental section, is fingerprint-based estimation. With this approach, δ_{est} values for each transmitter node can be recorded as fingerprints at vehicle production time. Post-production fingerprinting can be performed by authorized service personnel in case of any authorized changes are made to the network. An alternative approach would be model-based estimation in which δ_{est} is calculated using propagation delays along bus sectors that are estimated based on the network model. However, this requires building a model capable of accurately describing the bus delay behavior based on knowledge of bus structure along with line and node characteristics. Results obtained with a simplified model were previously illustrated in section III-C.

D. SENDER LOCATION ESTIMATION

For the case when frame transmissions fail to exhibit the propagation behavior expected from a known legit sender we envision a location estimation mechanism. The accuracy of the location estimation depends on the ability to distinguish between the three attack approaches considered in the attacker model: node compromising, node replacement, node insertion. For this purpose δ_{rec} values corresponding to several node locations are required since the measured differential delay that triggered the attack detection is not sufficient to characterize the type of attack.

1) Compromised node

Detecting a message transmission from an unexpected location, while δ_{rec} for all recent transmissions (preceding and following the illegal transmission) fall into known ranges, is an indicator that a compromised node was used for the transmission. Note that node replacing with nodes having very similar delay characteristics as the original node may induce the same behavior, as we show in the experimental section. Node localization in this case is a matter of comparing δ_{rec} to estimates for all known node locations on the bus obtained

either from fingerprints or bus model as described in the previous section. If a match is found the estimated location can be reported indicating that an illegal transmission was made from a node at one of the known network connection points.

2) Node replacement

Replacing a network node with one having different delay characteristics will induce noticeable impact on transmissions from other network nodes. Transmissions from the replacement node location tend to generate a differential delay very similar if not identical to the expected values from the original node. This is because signals originating from the replacement connection point will propagate through bus sections with the same delay characteristics as before replacement while any delay caused by the transmitter load is mainly symmetrical for the two propagation paths and will be canceled out when calculating the differential delay. This effect is illustrated in FIGURE 14 from our experimental section. This observation can be used to identify a node replacement. A good estimate on the location at which the replacement took place is the location for which differential delays are best matched after replacement was detected. It is possible (e.g., when replacement node characteristics resemble that of the original node) that, even though the replacement has been positively identified, more than one node is found to generate differential delays within expected ranges or very close. In this case the IDS can report several attack location estimates and assign a confidence factor to each.

3) Node insertion

Distinguishing a node insertion from a node replacement depends on the attacker actions. The effect of a node that is introduced simply for eavesdropping is hard to distinguish from that of a replaced node that is actively participating in the communication. However, when the inserted node is transmitting, the node insertion case can be detected since one extra differential delay range, in addition to the expected number of ranges, will be detected after its introduction. A location estimate for the inserted node could be given based on post-insertion delay ranges and know locations of initial bus nodes.

4) Other bus modifications

It is possible that other types of changes are made at the bus level, i.e. multiple node replacement, wire tampering or a combination of change types. In this case reliable estimates cannot be made about transmitter location without additional knowledge related to current bus structure.

E. VARIATION OF BUS CHARACTERISTICS

The structure of CAN networks inside cars is defined at design time. Bus lengths and location of node connectors are imposed by positioning requirements of the various ECUs. These characteristics are fixed and will not change during vehicle lifetime unless the network suffers modifications usually associated to tuning, fitting non-standard components or repairs. However, the number and characteristics of the nodes connected at available bus connection points may change, e.g., ECU replacement, fitting new ECUs as authorized car feature upgrades, connection of various diagnostic devices.

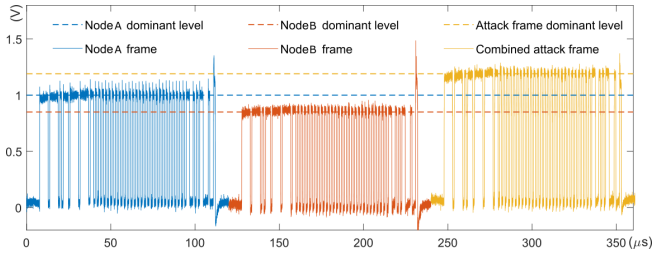

FIGURE 9. Individual attack frames vs. coordinated attack frame

With the exception of vehicle diagnostic devices usage, there are no other characteristics of the network that are expected to change intermittently. We can exclude any possible intermittent effects of diagnostic devices by isolating the diagnostic port from the in-vehicle network through a gateway as this is anyway a recommended security practice. The reliable operation of TIDAL-CAN relies on the knowledge of bus structure and characteristics. It would be possible to account for small propagation drifts caused by environmental conditions (e.g., by continuously updating stored fingerprints). However, more considerable deviations resulted from changes at the physical level of the bus cannot be automatically handled. Therefore, assuming that any such changes are done by an authorized service center, the manufacturer can mandate service centers as trusted parties to update required network information in TIDAL-CAN. This includes not only physical layer-related node and bus characteristics but also any changes regarding known frames since the set of frames and their legit senders are also fixed and and will not change under normal operating conditions (unless changes are made to bus nodes).

F. ATTACKING THE TIDAL-CAN MECHANISM

Besides the classic attacks on CAN communication, an adversary might target the TIDAL-CAN intrusion detection mechanism in an attempt to evade detection. In order to achieve successful attack frame transmission, without being detected, the attacker has to replicate propagation delay behavior of the legit frame transmitter. Accomplishing this using a single node be it compromised, replaced or inserted is not possible as previously discussed. However, the attacker might attempt a coordinated attack using two compromised nodes, A and B, each being preferably positioned closer to a different end of the bus so that corresponding signal propagation timings match the following requirement: $t^A_{left} < t^A_{right}$ and $t^B_{left} > t^B_{right}$. To replicate the delay behavior of a legit node the attacker must coordinate the two nodes to start the transmission of the same frame so that the resulting δ_{rec} matches the δ_{est} of the target node. Since each of the nodes

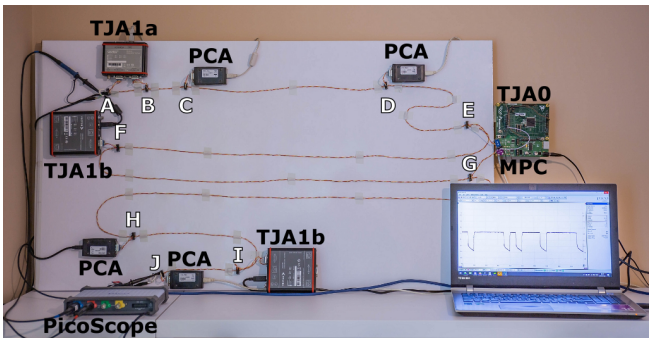

FIGURE 10. Experimental setup (depicting the CAN bus, connection points and nodes in one of the employed network configurations)

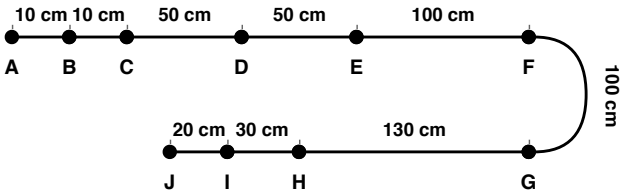

FIGURE 11. Location of node connections along the experimental bus

controlled by the attacker is closer to an opposite end of the bus, the resulting differential delay of the coordinated attack frame is $\delta_{rec} = t^A_{left} - t^B_{right}$. This approach could be used to replicate the differential delay behavior of virtually any network node by precise control of transmission start timings on the two compromised nodes.

Fortunately, this type of attack is easy to detect due to an intrinsic characteristic of the CAN physical layer. When two or more nodes superimpose dominant bit transmissions the resulting voltage level exhibits an increased voltage level for the dominant bits. This is illustrated in FIGURE 9 which shows the same frame transmitted independently by node A and B along with the superimposed transmission. Note, that the last dominant bit in each transmission characterized by a higher voltage level is the acknowledge bit which is the result of all receiver nodes in the network acknowledging the frame reception. TIDAL-CAN could easily detect this attack by using its sampling circuitry to get the dominant level of the analyzed frame and comparing it with pre-recorded expected levels for known transmitter nodes. Nonetheless, mounting the coordinated attack would be difficult since it requires knowledge of the network structure and node location, compromising two network nodes, synchronizing these nodes and precise timing control of transmissions.

VI. EXPERIMENTAL ANALYSIS

In this section we present results of the experimental evaluation of the proposed intrusion detection mechanism accounting for the effects of node placement, network layout and other factors that may influence the propagation delays.

A. EXPERIMENTAL SETUP

Our experimental setup, as depicted in FIGURE 10 consists of three main categories of elements: the bus, bus nodes and data acquisition system.

The bus consists of a 5 meter long twisted wire cable properly terminated with 120Ω resistors. Bus length was selected according to common lengths found in vehicles [10], [33]. A number of 10 connection points are positioned along the bus as illustrated in FIGURE 11 where each node connection point is labeled to simplify referencing in the following sections. Node inter-distances were chosen such that they include nodes located close to one-another as well as nodes that are farther apart. The smallest distance between nodes is set at 10cm according to standard requirements [16], [28]. Other bus characteristics are established in accordance to existing standard specifications [16], [29], [30].

TABLE 2. Devices and transceiver types employed in the experimental analysis

Abbrv.	Device	Transceiver	Channels	Amount
PCA_i	USB-CANmodul1	PCA82C251	1	10
$\mathrm{TJA1a}_i$	VN1630	TJA1051	2	1
$\mathrm{TJA1b}_i$ / $\mathrm{TJA1c}_i$	VN5610A	TJA1051	2	2
TJA0	EVBS12XF512	TJA1050	1	1
MCP	EVBS12XF512	MCP2551	1	1

Bus nodes are instantiated by PC-to-CAN devices or embedded development boards equipped with external CAN transceivers. The list of devices employed in our experiments as bus nodes is presented in TABLE 2 along with the transceiver type used in each device and the amount of available devices of each kind. The VN family of devices feature multiple CAN channels per device which is also indicated in the table. Each available connection through an individual CAN transceiver on the multi-channel devices is considered as a separate node. For easier referencing of individual nodes employed in our experiments TABLE 2 also lists the abbreviated notations assigned to each node type. Abbreviations are constructed according to the following format $\mathrm{T}[\mathrm{d}]_{[i]}$, where T represents the employed transceiver type, $d \in \{a, b, c\}$ is an optional parameter used to indicate the device if there are multiple device types using the same transceiver (a stands for VN1630 while b and c each indicate one of the two VN5610A devices) and i denotes the device number for single channel devices or channel number for multi-channel devices. A number of network configurations were built using the available nodes for providing a range of specific test beds

TABLE 3. Experimental network configurations

Nw. config.	Connection point									
	A	B	C	D	E	F	G	H	I	J
$\mathbf{Nw_1}$	$\mathrm{TJA1a}_1$	-	-	-	-	-	-	-	-	-
$\mathbf{Nw_2}$	PCA_1	PCA_2	PCA_3	PCA_4	PCA_5	PCA_6	PCA_7	PCA_8	PCA_9	PCA_{10}
$\mathbf{Nw_3}$	PCA_9	PCA_8	PCA_4	PCA_1	PCA_7	PCA_6	PCA_2	PCA_5	PCA_{10}	PCA_3
$\mathbf{Nw_4}$	$\mathrm{TJA1a}_1$	$\mathrm{TJA1a}_2$	PCA_3	PCA_4	TJA0	$\mathrm{TJA1b}_1$	MCP	PCA_8	$\mathrm{TJA1c}_1$	PCA_{10}
$\mathbf{Nw_5}$	PCA_1	-	PCA_3	-	PCA_5	-	PCA_7	-	PCA_9	-
$\mathbf{Nw_6}$	$\mathrm{TJA1a}_1$	$\mathrm{TJA1a}_2$	$\mathrm{TJA1b}_1$	$\mathrm{TJA1c}_1$	-	-	-	-	-	-

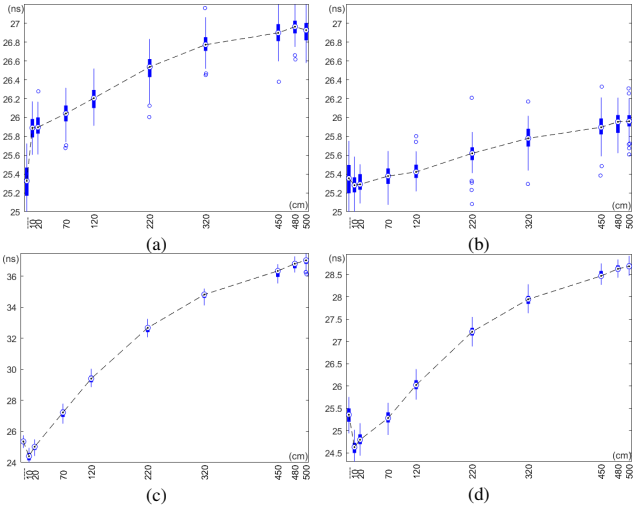

FIGURE 12. Variation of propagation delay over the test bus based on additional node location, measured for nodes of different types: (a) PCA_2 (b) $\mathrm{TJA1b}_1$ (c) $\mathrm{TJA0}$ and (d) MCP

for evaluating the TIDAL-CAN mechanism. The network configurations along with their designations are listed in TABLE 3 which indicates the node connected at each of the 10 available connection points in each configuration. In the next sections we use the following notation to indicate the point on the bus where a specific node is connected: $\mathrm{T[d]}^{\mathsf{L}}_{[i]}$, where $\mathsf{L} \in \{\mathsf{A}...\mathsf{J}\}$.

As the *data acquisition* system for the CAN signals we use a 5000 series PicoScope which samples data on 2 channels connected to the two ends (i.e. at points A and J) of the CAN-high line. The sampled data is then processed by a classification algorithm implemented MATLAB. Here the signals are passed through a de-noising step and then the differential propagation delay is calculated as $\delta = t_{\mathsf{J}} - t_{\mathsf{A}}$, where t_{J} and t_{A} are the sample times at which the signal recorded at point J and A respectively reach the selected voltage threshold level.

The data sets corresponding to each individual plot in the following sections are the result of processing 100 transmissions from each of the sender nodes pictured in the plot. Each data point corresponding to one of the nodes represents the differential propagation delay for one recorded CAN message. For testing the classification accuracy, 100 frames sent by the original network nodes, in each setup, are recorded to obtain δ_{est} and ϵ for each node location. A larger set, of 1000 frames for each node, recorded during the attack, are then used to test the intrusion detection mechanism.

B. EFFECT OF NODE LOCATION ON PROPAGATION DELAY

We analyze the effect of introducing network nodes at various distances from the transmitter. This is important for understanding how different types of nodes can influence propagation delays. Also, in case of using a model-based detection approach, this step is essential for building the network model or verifying its accuracy in estimating propagation delays. We employed the $\mathbf{Nw_1}$ configuration as a base for this analysis. In this setup, one channel of the VN1630 device is fixed at connection point A and set to transmit frames. A single additional node is placed at one of the remaining locations (B through J) for each signal acquisition round (100 signal pairs recorded per round) until the node has passed through all the 9 available connection points.

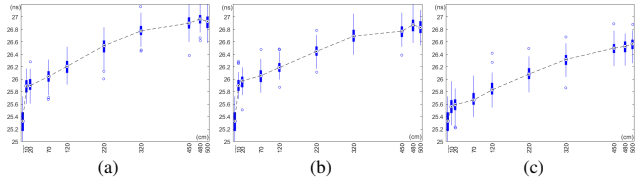

FIGURE 13. Variation of propagation delay over the test bus depending on additional node location, measured for 3 different USB-CANmodul1 nodes: (a) PCA_2, (b) PCA_9 and (c) PCA_{10}

FIGURE 12 shows the effect of introducing one receiver node at various locations along the bus, for each of the four transceiver types employed in our experiments. Box plots are used to represent the set of measurements done for each location. The first position in each plot representing the delay for the bus when the VN1630 transmitter device is the only connected node. The characteristic delay increase effect related to the increase in distance from sender is visible for all four transceivers and consistent with simulated results in FIGURE 6 in Section III-B. Actual delay ranges show variation between devices given differences in the differential load characteristics of each node type.

Nodes of the same type (i.e. identically manufactured devices) display very similar delaying characteristics as shown in FIGURE 13 that contains results for three USB-CANmodul1 devices. However, small variations are also visible here caused by various factors such as: uncontrollable variations in the production process or variations in connector and stub characteristics.

C. EFFECT OF NETWORK LAYOUT ON DIFFERENTIAL PROPAGATION DELAYS

TIDAL-CAN's detection mechanism relies on the ability to identify effects of various factors on differential propagation delays. Therefore, we analyze the differential propagation delay variation depending on changes in network layout. The actions that can be performed over nodes on a bus are: node compromising, replacement, insertion and removal. We look at the effects of these operations and account for the case of multiple network modifications. We also look at estimating sender location based on experimental observations.

1) Compromised node

When the structure of the network remains unchanged during the attack, the TIDAL-CAN mechanism provides 100% accuracy in identifying the source of a transmission. Differential delays recorded for each individual node form a cluster of values which can be clearly separated from clusters corresponding to other nodes. This is visible in FIGURE

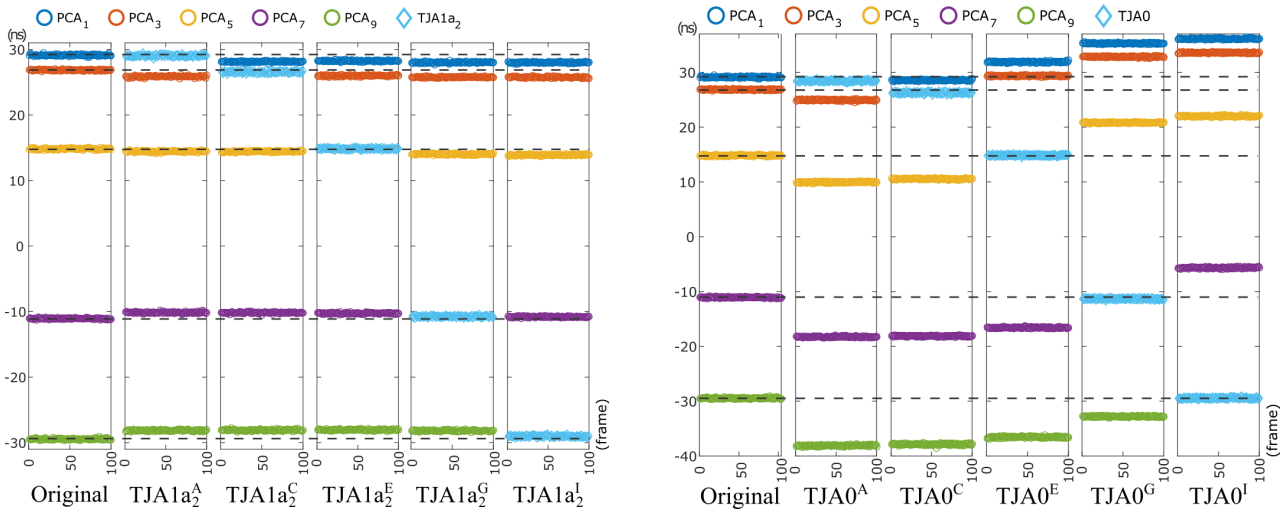

FIGURE 14. Differential delay variation when replacing one of the nodes in **Nw5** with a $\mathrm{TJA1a_2}$ (left) or a $\mathrm{TJA0}$ (right) device

TABLE 4. Confusion matrices for replacing one node in configuration 5 with a $\mathrm{TJA1a_2}$ or $\mathrm{TJA0}$ device (expressed as % of analyzed frames)

$\mathrm{TJA1a_2^A}$	A	C	E	G	I
A	97.8	0	0	0	0
C	0	0	0	0	0
E	0	0	50.7	0	0
G	0	0	0	0	0
I	0	0	0	0	0

$\mathrm{TJA1a_2^C}$	A	C	E	G	I
A	0	0	0	0	0
C	0	75.5	0	0	0
E	0	0	43.6	0	0
G	0	0	0	0	0
I	0	0	0	0	0

$\mathrm{TJA1a_2^E}$	A	C	E	G	I
A	0	0	0	0	0
C	0	0	0	0	0
E	0	0	99.1	0	0
G	0	0	0	0	0
I	0	0	0	0	0

$\mathrm{TJA1a_2^G}$	A	C	E	G	I
A	0	0	0	0	0
C	0	0	0	0	0
E	0	0	0	0	0
G	0	0	0	69.8	0
I	0	0	0	0	0

$\mathrm{TJA1a_2^I}$	A	C	E	G	I
A	0	0	0	0	0
C	0	0	0	0	0
E	0	0	0	0	0
G	0	0	0	93.1	0
I	0	0	0	0	39.2

$\mathrm{TJA0^A}$	A	C	E	G	I
A	2.1	0	0	0	0
C	0	0	0	0	0
E	0	0	0	0	0
G	0	0	0	0	0
I	0	0	0	0	0

$\mathrm{TJA0^C}$	A	C	E	G	I
A	23.6	0	0	0	0
C	0	32.7	0	0	0
E	0	0	0	0	0
G	0	0	0	0	0
I	0	0	0	0	0

$\mathrm{TJA0^E}$	A	C	E	G	I
A	0	95.7	0	0	0
C	0	0	0	0	0
E	0	0	99.7	0	0
G	0	0	0	0	0
I	0	0	0	0	0

$\mathrm{TJA0^G}$	A	C	E	G	I
A	0	0	0	0	0
C	0	0	0	0	0
E	0	0	0	0	0
G	0	0	0	90.5	0
I	0	0	0	0	0

$\mathrm{TJA0^I}$	A	C	E	G	I
A	0	0	0	0	0
C	0	0	0	0	0
E	0	0	0	0	0
G	0	0	0	0	0
I	0	0	0	0	99.9

14, 15, 16, 18 and 19 corresponding to the various tested network layouts. This detection accuracy for all illegal frame transmissions made by the attacker node is consistent across all the evaluated network configurations based on the test bus structure. Here we consider as being illegal all frame transmissions that should not have originated from the attacker location. Any frame that is assigned, by network design, to the node that was compromised will be regarded as a legitimate transmission even if its content is controlled by the attacker. This is a common problem of all IDSs that do not use the frame data field as a basis for intrusion detection. For this reason, an efficient IDS should also consider payload content, transmitted with the frame, as a basis for intruder detection.

2) Node replacement

The effect of replacing a network node on differential propagation delays greatly depends on differences in the delay characteristics of the replaced versus the replacement node. A commonly found attack case is that of a single node being replaced. We illustrate the effects of such an attack by using **Nw5** as a base on which a single node is replaced in each round of measurements with either $\mathrm{TJA1a_2}$ or $\mathrm{TJA0}$. FIGURE 14 presents the differential delay changes obtained by node replacements. To illustrate the classification results in each of the tested cases we employ the use of confusion matrices. We use the same matrix layout through the experimental section. Each matrix row represents one of the known/expected network nodes at specified locations as a predicted class instance, while each column represents actual node instances. TABLE 4 shows the results of attempting to classify frames received on the modified network as one of the known nodes. The replacement node and its location are indicated under each corresponding plot and confusion matrix with the classification success rates for attacker transmissions being marked on lighter red background for low rates or darker red background for higher rates.

While it is clear that the effect of using the $\mathrm{TJA0}$ as a replacement node has a more significant impact on the differential delays of all nodes, there are two common characteristics related to node replacements that can be observed in both cases. The first characteristic is that, when compared to delays from other nodes, the differential delays caused by transmissions from the replacement node are very similar (or even identical) to expected delays in the original bus layout. In FIGURE 14 averages of expected delays are represented as dashed lines extending from delays in the original data set. Attack frames associated to the original replaced nodes

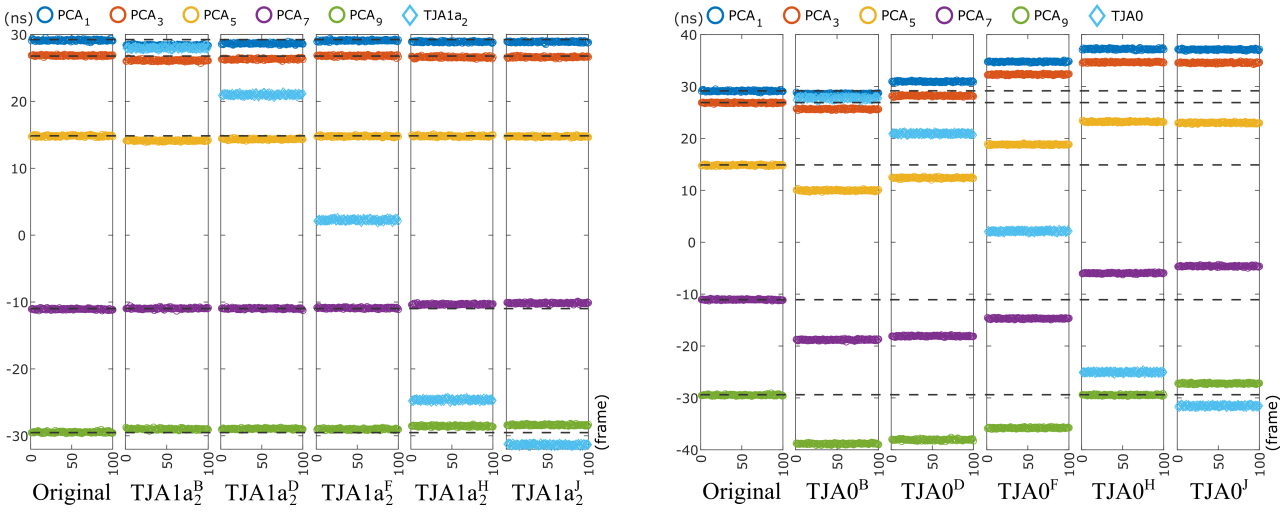

FIGURE 15. Differential delay variation when inserting a $\mathrm{TJA1a_2}$ (left) or a $\mathrm{TJA0}$ device (right) in one of the free connection points of $\mathbf{Nw_5}$

TABLE 5. Confusion matrices for inserting one $\mathrm{TJA1a_2}$ or $\mathrm{TJA0}$ device in configuration 5 (expressed as % of analyzed frames)

$\mathrm{TJA1a_2^B}$

	A	B	C	E	G	I
A	0.1	0	0	0	0	0
C	0	0	0	0	0	0
E	0	0	0	0.5	0	0
G	0	0	0	0	99.8	0
I	0	0	0	0	0	25.7

$\mathrm{TJA1a_2^D}$

	A	C	D	E	G	I
A	15.5	0	0	0	0	0
C	0	7.3	0	0	0	0
E	0	0	0	16.6	0	0
G	0	0	0	0	100	0
I	0	0	0	0	0	6.6

$\mathrm{TJA1a_2^F}$

	A	C	E	F	G	I
A	99.9	0	0	0	0	0
C	0	99.8	0	0	0	0
E	0	0	99.9	0	0	0
G	0	0	0	0	98.7	0
I	0	0	0	0	0	27.3

$\mathrm{TJA1a_2^H}$

	A	C	E	G	H	I
A	97.2	0	0	0	0	0
C	0	96.1	0	0	0	0
E	0	0	100	0	0	0
G	0	0	0	0.4	0	0
I	0	0	0	0	0	0

$\mathrm{TJA1a_2^J}$

	A	C	E	G	I	J
A	97.4	0	0	0	0	0
C	0	96.3	0	0	0	0
E	0	0	100	0	0	0
G	0	0	0	0	0	0
I	0	0	0	0	0	0

$\mathrm{TJA0^B}$

	A	B	C	E	G	I
A	5.2	0	0	0	0	0
C	0	1.1	0	0	0	0
E	0	0	0	0	0	0
G	0	0	0	0	0	0
I	0	0	0	0	0	0

$\mathrm{TJA0^D}$

	A	C	D	E	G	I
A	0	0	0	0	0	0
C	0	0	0	0	0	0
E	0	0	0	0	0	0
G	0	0	0	0	0	0
I	0	0	0	0	0	0

$\mathrm{TJA0_0^F}$

	A	C	E	F	G	I
A	0	0	0	0	0	0
C	0	0	0	0	0	0
E	0	0	0	0	0	0
G	0	0	0	0	0	0
I	0	0	0	0	0	0

$\mathrm{TJA0^H}$

	A	C	E	G	H	I
A	0	0	0	0	0	0
C	0	0	0	0	0	0
E	0	0	0	0	0	0
G	0	0	0	0	0	0
I	0	0	0	0	0	99.7

$\mathrm{TJA0^J}$

	A	C	E	G	I	J
A	0	0	0	0	0	0
C	0	0	0	0	0	0
E	0	0	0	0	0	0
G	0	0	0	0	0	0
I	0	0	0	0	0	0

are generally being classified as being legit with a higher probability when compared to legit transmissions which fail to be classified as any of the known nodes in most cases. The second common characteristic is that changes in differential propagation delay are larger as nodes are farther away from the replaced node which behaves as a change in delay characteristics at its connection point. This is caused by differences in delay characteristics of various nodes as well as the influence of transmitter-receiver distance on changes in propagation delay.

A side-effect of network modifications is that changes in delay characteristics of legit nodes could lead to their classification as other legit nodes. One such case was determined by replacing the node at location E from $\mathbf{Nw_5}$ with $\mathrm{TJA0}$. In this case, as shown in TABLE 4, transmissions from the legit node at location C are classified as originating from point A with a 95.7% rate. However, this will not affect intrusion detection as the mechanism will signal the presence of an intrusion as a result of failing to classify transmissions from most legit nodes. The number of differential delay clusters is the same as in the expected setup (corresponding to 5 nodes). Hence, since transmissions from more than one of these clusters will be considered illegal, the attack type can be categorized as either a replacement attack or an eavesdropper insertion.

3) Node insertion/removal

We exemplify the case of node insertion on the same setup using $\mathbf{Nw_5}$ as a base and devices $\mathrm{TJA1a_2}$ and $\mathrm{TJA0}$ as intruders that are inserted at available connection points. The results are presented in FIGURE 15 with $\mathrm{TJA1a_2}$ based tests on the left and $\mathrm{TJA0}$ based tests on the right. TABLE 5 presents the confusion matrices for the insertion attack with the column representing transmissions from the inserted node distinctly marked. The general effect of node insertion on the differential delays corresponding to legit nodes is similar to the effect observed in the replacement attack. There are only slight differences in the magnitude of the delay shifts between the two cases. A hint on the difference between node replacement and eavesdropper insertion could be that, in the insertion case, there seems to be no one single node that is generally classified as being legit with a higher confidence margin. However, without a clear indication on the presence of an additional node it is difficult to distinguish between the two cases based on delays alone.

Nevertheless, once the inserted node starts to transmit, the

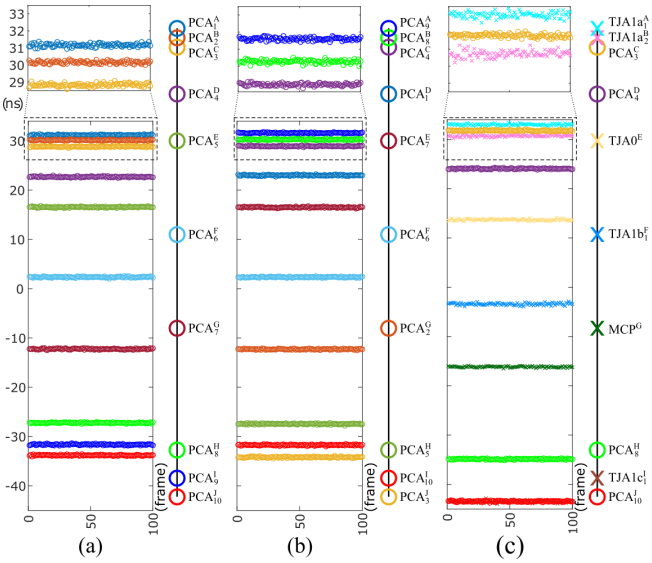

FIGURE 16. Differential propagation delays for each transmitting node in: (a) $\mathbf{Nw_2}$, (b) $\mathbf{Nw_3}$ and (c) $\mathbf{Nw_4}$

insertion attack case can be easily asserted as it will form a supplementary differential delay cluster. Since differential delays from each node form distinct clusters it is expected to identify as many delay clusters as the number of known nodes. This additional delay cluster is distinguishable as can be seen in FIGURE 15. Overlaps visible in the upper parts of the plots corresponding to insertion at position B are only caused by plot scale and marker size. The two neighboring bands are actually located at approximately 1ns from each other. This is also visible in corresponding confusion matrices as the column associated to location B shows 0% classification rate against predicted node at point A. However, to alleviate the possibility of not detecting nodes due to overlapping differential delay clusters (e.g., FIGURE 16) differential delays should be calculated at more than one level along the rising/falling edge as previously proposed. Transmissions from the inserted node generally fail to be classified as any of the prior known nodes while legit nodes tend to have higher successful classification rates. The IDS triggers the intrusion alarm due to the classification fails, the detection of more differential delay clusters than expected being a confirmation of the attack type.

The effect of node removal is also twofold: change in propagation delays as well as a reduction in the number of differential delay clusters. Hence, the TIDAL-CAN mechanism will also be able to detect this case as a result of failed node classifications.

4) Multiple network modifications

FIGURE 16 illustrates effects of two extreme cases of node replacement based on configurations $\mathbf{Nw_{2-4}}$ (node locations in each configuration are indicated alongside the corresponding plot). We consider $\mathbf{Nw_2}$ as the original bus layout, while $\mathbf{Nw_3}$ and $\mathbf{Nw_4}$ are extreme layout changes. $\mathbf{Nw_3}$ consists in

TABLE 6. Confusion matrices for the cases of modifying $\mathbf{Nw_2}$ into $\mathbf{Nw_3}$ (a) and $\mathbf{Nw_4}$ (b)

	A	B	C	D	E	F	G	H	I	J
A	21.9	0	0	0	0	0	0	0	0	0
B	0	99.9	0	0	0	0	0	0	0	0
C	0	0	99.9	0	0	0	0	0	0	0
D	0	0	0	75.7	0	0	0	0	0	0
E	0	0	0	0	100	0	0	0	0	0
F	0	0	0	0	0	99.9	0	0	0	0
G	0	0	0	0	0	0	99.9	0	0	0
H	0	0	0	0	0	0	0	95.9	0	0
I	0	0	0	0	0	0	0	0	99.9	0
J	0	0	0	0	0	0	0	0	0	43.6

(a)

	A	B	C	D	E	F	G	H	I	J
A	0	35.9	0.5	0	0	0	0	0	0	0
B	0	15.8	0	0	0	0	0	0	0	0
C	0	0	0	0	0	0	0	0	0	0
D	0	0	0	0	0	0	0	0	0	0
E	0	0	0	0	0	0	0	0	0	0
F	0	0	0	0	0	0	0	0	0	0
G	0	0	0	0	0	0	0	0	0	0
H	0	0	0	0	0	0	0	0	0	0
I	0	0	0	0	0	0	0	0	0	0
J	0	0	0	0	0	0	0	0	0	0

(b)

reordering the original set of nodes in $\mathbf{Nw_2}$ while in $\mathbf{Nw_4}$ 6 out of the 10 nodes are replaced with different types of nodes. The voltage level used as a threshold for calculating the differential delays is the one selected for best node separation in $\mathbf{Nw_2}$. As a result, differential delays for all nodes in the three configurations are clearly distinguishable with the exception of nodes located at points I and J in $\mathbf{Nw_4}$. We later show that changing the voltage level used for differential delay measurement allows for clear identification of all nodes in $\mathbf{Nw_4}$. Changing the bus layout to $\mathbf{Nw_3}$ brings virtually no visible change in differential delays except for the nodes located towards the ends of the bus which show skews of up to 0.5ns. This close resemblance in behavior is caused by the practically identical load behavior of same type nodes from the same manufacturer. On the other side, considerable differences can be observed when comparing the results for $\mathbf{Nw_2}$ and $\mathbf{Nw_4}$. Delay characteristics of nodes in $\mathbf{Nw_4}$ are not identical as in the homogeneous $\mathbf{Nw_2}$ and $\mathbf{Nw_3}$.

TABLE 6 (a) and (b) represents the confusion matrices resulted from applying the classification mechanism to the two extreme modification cases. The results are in accordance with the visual observations that can be made in FIGURE 16. Randomly rearranging the same set of nodes leads to a large percentage of false positives in node classification. Still, the inability to classify more than 50% of transmissions from nodes placed at the ends of the bus as one of the known nodes is an indicator of changes made at the bus level and will trigger an intruder alarm. In the second extreme replacement case, 8 out of the ten nodes always fail to be classified as a known node. Among transmissions from the node at location B 35.9 % are classified as originating from location A and 15.8% from location B while for the node at point C only 0.5% fall into predicted classes. The inability to classify most of transmissions as any of the predicted classes is a clear indicator of intrusion.

D. OTHER FACTORS INFLUENCING DIFFERENTIAL PROPAGATION DELAYS

We also analyze the effect of various factors related to signal characteristics and sampling on the ability to distinguish node transmissions based on the corresponding differential delays. The factors that we consider are: signal rise/fall time, sampling rate, CAN bit rate and delay measurement threshold.

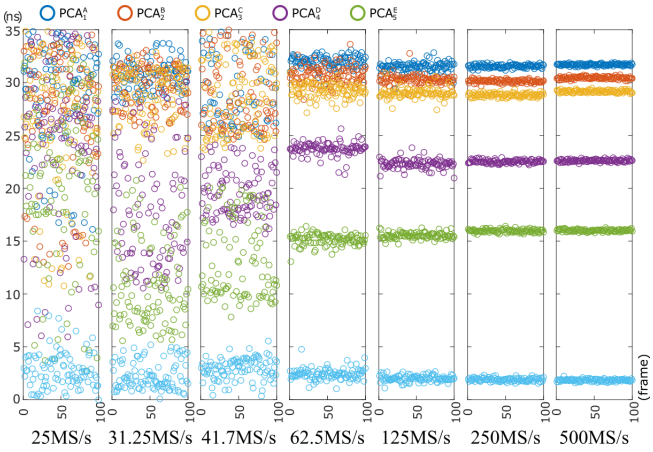

FIGURE 17. Effect of sampling rate on differential delays

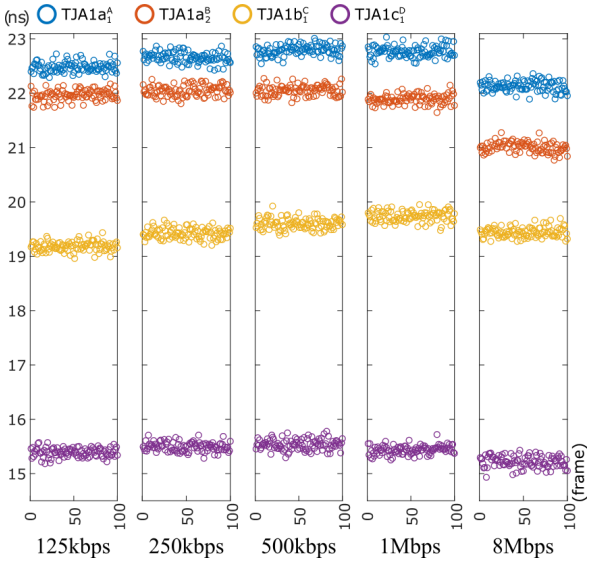

FIGURE 18. Effect of CAN bit rate on differential delays tested on network configuration 6

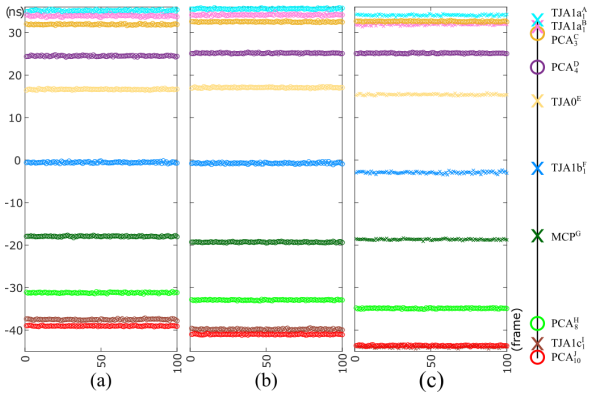

FIGURE 19. Differential delay cluster changes tested on $\mathbf{Nw_4}$ when setting measurement threshold for rising edge on CAN-high to (a) 0.25V, (b) 0.45V and (c) 0.65V.

1) Signal rise/fall time and sampling rate

Signal rise/fall time and the ability to accurately extract transition slope is one of the main factors influencing the accuracy of intrusion detection. Rise and fall times will vary depending on transceiver model, specific usage and employed bit rate. Various approaches for slope control can be found in available transceivers including internal and external slope control or even uniform slope prescription regardless of bit rate. The steeper the edge the higher the sampling rate required to accurately describe the slope. We analyzed the rise/fall times (measured from 10% to 90% of the signal) for the devices used in our experiment and found them to vary in the 30-45ns interval depending on model and bit rate. We evaluated the effects of sample rate on the ability to distinguish between differential delays of distinct nodes. FIGURE 17 illustrates these effects on nodes in $\mathbf{Nw_2}$ (only connection points A-F represented in the plots for increased visibility) at various sample rates between 25 and 500MSps. At around 125MSps some transmissions from nodes placed 10cm from one another are miss-classified and by 62.5MSps it becomes impossible to distinguish transmissions from nodes placed closer than 20cm to each other. The detection capability degrades even more towards 41.7MSps making it impossible to differentiate nodes that are closer than several meters from each other. For our intrusion detection tests sample rates no lower than 250MSps were used.

2) CAN bitrate

Commonly employed bit rates for standard high-speed CAN communication range between 10kbps to 1Mbps. The more recently introduced CAN-FD [15] extension to the standard CAN protocol allows the use of even higher bit rates during the data phase. Currently available transceivers support bit rates of up to 8Mbps. Since we only use signal edges as a base for intrusion detection it is important to identify relevant impact of bit rate on edge characteristics. As already pointed in the previous section, one characteristic that generally varies with signaling rate is duration of the rising and falling edges. Faster rates will require shorter rise/fall times while longer transition times would be preferred at lower bit rates. We evaluated bit rate effect on distinguishing differential delays of different nodes using configuration $\mathbf{Nw_6}$ which consists only of nodes based on the TJA1051 transceiver, the only one from our lineup capable of bit rates higher than 1Mbps. The results, illustrated in FIGURE 18, demonstrate that bit rate has little influence on the ability to separate transmissions by source node. Differential delay ranges for each device tends to vary slightly when using different bit rates. However, this would not effect the detection capability since, for a specific CAN bus, bit rates are fixed at design time. CAN-FD uses two bit rates, one for the arbitration phase and one for the data phase but this is not a problem since the IDS is not required to sample during the arbitration phase.

3) Measurement threshold

The selection of the threshold used for measuring the differential delay also depends on the slope of the rising/falling

edges. FIGURE 19 shows how different threshold levels influence the positioning of differential delay clusters based on **Nw$_4$**. When using the threshold established as optimum for **Nw$_2$** (FIGURE 19 (c)), nodes placed at connection points I and J cannot be distinguished. However, selecting a more suitable threshold enables the clear separation of differential delays from each node. In the case of **Nw$_4$** setting the threshold for the CAN-high rising edge to 0.25V (FIGURE 19 (a)) provides the best results. Using a threshold that provides best separation of differential delays from individual nodes for the designed network configuration is required for best results in attack identification. Classifying the attack type and estimating attack location require extracting differential delays based on at least two thresholds (positioned towards the start and ends of the rising/falling edge) for obtaining the characteristic slope.

E. DETECTION ACCURACY

Most of the existing IDS proposals based on physical layer characteristics only cover the case of compromised nodes. Mechanisms covering this type of attack provide best node identification rates (true positives) that are close to 100%, e.g., 99.8% for Viden [6] and 99.85% for Scission [18]. Notably, the authors of SIMPLE [11] report 100% true positives for the laboratory setup. All proposals present a certain level of false positives with the exception of Scission which was optimized in this regard. By using the proposed TIDAL-CAN mechanism we were able to obtain 100% correct node classification, without false positives, during attacks from compromised nodes, in all of the tested network configurations. As shown, these results are reproducible for networks using virtually any bit rate and having neighboring nodes placed as close as 10cm from one another.

When looking at other attack approaches, only three of the related works account for the possibility of inserting of an attacker node. A threshold based attack detection approach requiring multiple frame transmissions is employed in [7] and [18] for transmissions from inserted nodes while the work in [27] is only able to detect the existence of an additional node without identifying its transmissions. With TIDAL-CAN, the presence of an attack can be reported from the first frame that cannot be classified as originating from one of the known nodes. However, for positively indicating the node insertion case, multiple transmissions are required. In addition to the compromised and inserted node attack scenarios, by using the TIDAL-CAN mechanism it is also possible to detect node replacement.

F. ESTIMATING SENDER LOCATION

The differential delays associated to a node can be used to estimate the node location. A good estimate can be given by using simple linear interpolation based on differential delays of two nodes for which the location is known if transmissions from these two known nodes can be deemed as being authentic. TABLE 7 illustrates location estimates obtained using this approach for our basic network configurations considering as known nodes the first and last node in each configuration. In most cases estimation errors are below 10cm while in several instances that exceed this level the error increases to around 40cm. While using this approach may seem straightforward it may not be so when such changes at the network level impede identification of trusted reference nodes that are used for interpolation. Therefore, we discuss next particularities of each attack approach that could be used for node localization.

TABLE 7. Node location estimations represented as distance from the reference bus end based on mean differential delay values for network configurations **Nw$_2$**-**Nw$_6$**.

	Actual Distance (cm)									
	0	10	20	70	120	220	320	450	480	500
Config.	**Estimated Distance (cm)**									
Nw$_2$	0	7.7	18.5	65.5	112.5	221.9	335.1	450.7	483.8	500
Nw$_3$	0	9.9	20.5	65.3	114.6	221.5	333.1	447.6	480.3	500
Nw$_4$	0	8.8	22.2	72.1	125.3	240.6	357.8	447.7	489.9	500
Nw$_5$	0	n/a	18.8	n/a	117.9	n/a	329.8	n/a	480	n/a
Nw$_6$	0	8.5	29.3	70	n/a	n/a	n/a	n/a	n/a	n/a

Compromised node. When attacks are mounted using compromised nodes without any network changes, identifying transmitter location is straightforward. Accurate localization is assured by always identifying the transmitter as one of the known nodes for which bus locations are known.

Node replacement. For any node, differential delays resulted after node replacement tend to be closer to the ones expected for the actual node location than those of other locations. Therefore, in most cases the known node location for which $|\delta_{rec} - \delta_{est}|$ is minimum provides a good estimate based on a single frame transmission. When, based on this rule, the same location can be attributed to more than one node, transmissions from all nodes are required to establish correlations between new differential delays and node locations. Attackers can be more easily pinpointed as $|\delta_{rec} - \delta_{est}|$ for the δ_{est} associated to the actual node location is generally very small (smaller than 0.7ns in our experiments) when compared to the case of other nodes.

Node insertion. Estimating the location of the inserted node requires correlating post-insertion delay values to known locations on the original bus layout. This requires the analysis of multiple frame transmissions from all nodes.

Multiple modifications. Even if comparing expected differential delays with actual ones makes it clear that the layout of the bus has changed, pinpointing the attacker location in this case is clearly impossible.

While our current approach does not pinpoint the precise location of the attacker node for all types of attack, the classification that we provide is sensitive to the location of the attacker as proved by the experiments. For example, in FIGURE 15 the differential time shifts with the intruder localization on the bus between point B, D, F ,H, J.

VII. CONCLUSION

In this work we introduced TIDAL-CAN, a novel mechanism for intrusion detection on CAN. Differential delays of bus signals, which are influenced by bus characteristics and sender location, are used by TIDAL-CAN to identify transmitter nodes. We show that the proposed mechanism can successfully identify attacks based on compromised nodes. In addition, the presence of attacks involving node replacement and insertion can also be detected. The efficiency of the proposed mechanism is experimentally validated using different node types and multiple network configurations. As a result, we found that TIDAL-CAN provides reliable results for virtually all high-speed CAN and CAN-FD bit rates while using sample rates as low as 250MS/s.

We also identify features useful to estimate sender location depending on the attack method. TIDAL-CAN is able to accurately locate transmitter nodes when there are no changes on the network structure performed by the attacker. When changes are made at the physical network level it is still possible to extract useful data from transmissions to allow for transmitter location estimation. In this work we focus on identifying features that can be used for node localization and interpreting them in a numerical computing environment, i.e., MATLAB. This environment is commonly used to output industry standard source code that is further integrated on automotive-grade embedded platforms. However, due to specific computational and signal acquisition demands, achieving a real-life implementation may require more specialized embedded devices or even a dedicated hardware implementation. Since this is quite a demanding task, we leave it as future work and limit our current work to testing the feasibility of such an approach.

REFERENCES

[1] Line Driving and System Design. Application note, Texas Instruments, Apr. 1995.

[2] AUTOSAR. Specification of Secure Onboard Communication, 4.3.1 edition, 2017.

[3] Z. Cai, A. Wang, W. Zhang, M. Gruffke, and H. Schweppe. 0-days & mitigations: Roadways to exploit and secure connected bmw cars. Black Hat USA, 2019:39, 2019.

[4] S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, S. Savage, K. Koscher, A. Czeskis, F. Roesner, T. Kohno, et al. Comprehensive experimental analyses of automotive attack surfaces. In USENIX Security Symposium. San Francisco, 2011.

[5] K.-T. Cho and K. G. Shin. Fingerprinting electronic control units for vehicle intrusion detection. In 25th USENIX Security Symposium, 2016.

[6] K.-T. Cho and K. G. Shin. Viden: Attacker identification on in-vehicle networks. In Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS '17, pages 1109–1123, New York, NY, USA, 2017. ACM.

[7] W. Choi, H. J. Jo, S. Woo, J. Y. Chun, J. Park, and D. H. Lee. Identifying ecus using inimitable characteristics of signals in controller area networks. IEEE Transactions on Vehicular Technology, 67(6):4757–4770, June 2018.

[8] W. Choi, K. Joo, H. J. Jo, M. C. Park, and D. H. Lee. Voltageids: Low-level communication characteristics for automotive intrusion detection system. IEEE Transactions on Information Forensics and Security, 13(8):2114–2129, Aug 2018.

[9] S. Corrigan. Controller area network physical layer requirements. Application Report - slla270, 2008.

[10] P. Degauque, I. Stievano, S. Pignari, V. Degardin, F. Canavero, F. Grassi, and F. J. Canete. Power-Line Communication: Channel Characterization and Modeling for Transportation Systems. IEEE Vehicular Technology Magazine, 10(2):28–37, June 2015.

[11] M. Foruhandeh, Y. Man, R. Gerdes, M. Li, and T. Chantem. Simple: Single-frame based physical layer identification for intrusion detection and prevention on in-vehicle networks. In Proceedings of the 35th Annual Computer Security Applications Conference, ACSAC '19, page 229–244, New York, NY, USA, 2019. Association for Computing Machinery.

[12] B. Groza and P. Murvay. Security solutions for the controller area network: Bringing authentication to in-vehicle networks. IEEE Vehicular Technology Magazine, 13(1):40–47, March 2018.

[13] B. Groza, L. Popa, and S. Murvay. INCANTA - intrusion detection in controller area networks with time-covert cryptographic authentication. In International Workshop on Cyber Security for Intelligent Transportation Systems (ESORICS'18 Workshops), 2018.

[14] F. Hartwich et al. Can with flexible data-rate. In Proc. iCC, pages 1–9, 2012.

[15] ISO. 11898-1–road vehicles–controller area network (can)–part 1: Data link layer and physical signalling. Technical report, International Organization for Standardization, 2015.

[16] ISO. 11898-2, road vehicles controller area network (can) part 2: High-speed medium access unit. Technical report, International Organization for Standardization, 2016.

[17] M. Kneib. A survey on sender identification methodologies for the controller area network. In Sicherheit 2020, Lecture Notes in Informatics (LNI). Gesellschaft für Informatik eV, 2020.

[18] M. Kneib and C. Huth. Scission: Signal Characteristic-Based Sender Identification and Intrusion Detection in Automotive Networks. In Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, CCS '18, pages 787–800, New York, NY, USA, 2018. ACM.

[19] M. Kneib, O. Schell, and C. Huth. On the robustness of signal characteristic-based sender identification. arXiv preprint arXiv:1911.09881, 2019.

[20] M. Kneib, O. Schell, and C. Huth. EASI: Edge-Based Sender Identification on Resource-Constrained Platforms for AutomotiveNetworks. In Proceedings of the 2020 Network and Distributed System Security Symposium, San Diego, CA, 2020.

[21] 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 2010 IEEE Symposium on Security and Privacy, pages 447–462. IEEE, 2010.

[22] H. Lee, S. H. Jeong, and H. K. Kim. Otids: A novel intrusion detection system for in-vehicle network by using remote frame. In Procedings of PST (Privacy, Security and Trust), 2017.

[23] H. Lim, G. Kim, S. Kim, and D. Kim. Quantitative analysis of ringing in a controller area network with flexible data rate for reliable physical layer designs. IEEE Transactions on Vehicular Technology, 68(9):8906–8915, Sep. 2019.

[24] C. Miller and C. Valasek. Remote exploitation of an unaltered passenger vehicle. Black Hat USA, 2015:91, 2015.

[25] M. R. Moore, R. A. Bridges, F. L. Combs, M. S. Starr, and S. J. Prowell. Modeling inter-signal arrival times for accurate detection of can bus signal injection attacks: a data-driven approach to in-vehicle intrusion detection. In Proceedings of the 12th Annual Conference on Cyber and Information Security Research, page 11. ACM, 2017.

[26] P. Murvay and B. Groza. Source identification using signal characteristics in controller area networks. IEEE Signal Processing Letters, 21(4):395–399, April 2014.

[27] M. Rumez, J. Dürrwang, T. Brecht, T. Steinshorn, P. Neugebauer, R. Kriesten, and E. Sax. Can radar: Sensing physical devices in can networks based on time domain reflectometry, 2019.

[28] J1939-11 – Physical Layer, 250K bits/s, Twisted Shielded Pair. Standard, SAE International, Sept. 2006.

[29] J1939-14 – Physical Layer, 500 kbps. Standard, SAE International, Oct. 2011.

[30] High-Speed CAN (HSC) for Vehicle Applications at 500 kbps. Standard, SAE International, Mar. 2002.

[31] S. U. Sagong, X. Ying, A. Clark, L. Bushnell, and R. Poovendran. Cloaking the clock: emulating clock skew in controller area networks. In Proceedings of the 9th ACM/IEEE International Conference on Cyber-Physical Systems, pages 32–42. IEEE Press, 2018.

[32] H. M. Song, H. R. Kim, and H. K. Kim. Intrusion detection system based on the analysis of time intervals of can messages for in-vehicle network. In Information Networking (ICOIN), 2016 International Conference on, pages 63–68. IEEE, 2016.

[33] L. Thomas, C. Óscar, S. Naoshi, and D. Sam. Networked Automotive Subsystems Becoming the Rule of the Road. https://www.powerelectronics.com/power-electronics-systems/networked-automotive-subsystems-becoming-rule-road, March 2014. [Online: accessed 28-Jun-2019].
[34] J. Van den Herrewegen and F. D. Garcia. Beneath the bonnet: A breakdown of diagnostic security. In Computer Security, pages 305–324, Cham, 2018. Springer International Publishing.
[35] X. Ying, G. Bernieri, M. Conti, and R. Poovendran. Tacan: Transmitter authentication through covert channels in controller area networks. In Proceedings of the 10th ACM/IEEE International Conference on Cyber-Physical Systems, pages 23–34. ACM, 2019.

PAL-STEFAN MURVAY is a Lecturer at Politehnica University of Timisoara (UPT). He graduated his B.Sc and M.Sc studies in 2008 and 2010 respectively and received his Ph.D. degree in 2014, all from UPT. He has a 10-year background as an embedded software developer in the automotive industry. He worked as a postdoctoral researcher in the CSEAMAN project and is currently a senior researcher in the PRESENCE project. He also leads the SEVEN project (2018-2020) related to automotive and industrial systems security. His current research interests are in the area of embedded systems security in general with particular focus on automotive security.

BOGDAN GROZA is a Professor at Politehnica University of Timisoara (UPT). He received his Dipl.Ing. and Ph.D. degree from UPT in 2004 and 2008 respectively. In 2016 he successfully defended his habilitation thesis having as core subject the design of cryptographic security for automotive embedded devices and networks. He has been actively involved inside UPT with the development of laboratories by Continental Automotive and Vector Informatik. Besides regular participation in national and international research projects in information security, he lead the CSEAMAN project (2015-2017) and currently leads the PRESENCE project (2018-2020), two national research projects dedicated to automotive security.