Source Identification Using Signal Characteristics in Controller Area Networks

Pal-Stefan Murvay and Bogdan Groza

***Abstract*—The CAN (Controller Area Network) bus, i.e., the de facto standard for connecting ECUs inside cars, is increasingly becoming exposed to some of the most sophisticated security threats. Due to its broadcast nature and ID oriented communication, each node is sightless in regards to the source of the received messages and assuring source identification is an uneasy challenge. While recent research has focused on devising security in CAN networks by the use of cryptography at the protocol layer, such solutions are not always an alternative due to increased communication and computational overheads, not to mention backward compatibility issues. In this work we set steps for a distinct approach, namely, we try to take authentication up to unique physical characteristics of the frames that are placed by each node on the bus. For this we analyze the frames by taking measurements of the voltage, filtering the signal and examining mean square errors and convolutions in order to uniquely identify each potential sender. Our experimental results show that distinguishing between certain nodes is clearly possible and by clever choices of transceivers and frame IDs each message can be precisely linked to its sender.**

***Index Terms*—source identification, physical fingerprinting, CAN bus.**

I. Motivation and related work

The Controller Area Network (CAN) is a broadcast serial bus initially designed for in-vehicle communication. The ever growing design complexity of automotive embedded systems makes it difficult for the manufacturers to anticipate all possible threat scenarios. As a result, vulnerabilities in automotive systems are highlighted by an increasing number of recent papers [6], [1], [4]. All these prove that in-vehicle communication, in the absence of source authentication, is an easy target even in front of some of the most basic attacks, e.g., replays, packet injections, etc.

The CAN physical layer is typically implemented as a two wire differential bus as presented in Figure 1. Each frame begins with an identifier (ID) which determines the priority of the frame and carries up to 64 bits of data followed by a standard 15-bit CRC. The 64 bit data field of the CAN frame gives the first hints on why assuring cryptographic security is uneasy. Clearly, it is not feasible to fit an authentication tag plus the message in the 8 byte block that is carried by each frame. Adding separate authentication frames increases the bus load and CAN is limited to 1 Mbps, an upper-bound that is already reached in many practical scenarios. Message Authentication Codes (MACs), the cryptographic tool for assuring message authentication, are usually in the order of 128 bits and while truncating them to a particular length is an option one still needs to fit them along with the message within the 64 bits (clearly, this is not possible). Van Herrewege et al. propose in [8] the use of CAN+ in order to hide the authentication bits within the bits of a regular CAN frame. But CAN+ capable controllers are not yet available on the market and it is not clear if they will be produced in the near future (currently, as a more expensive alternative, the industry is migrating to other layers such as CAN-FD or FlexRay). In [7] Szilagy and Koopman allow each node to vote on the authenticity of the message and each vote consists in several MACs that are truncated in order to fit them inside each frame. The suggested value from [7] is 8 bits for each MAC, this is clearly too low to assure real-world security.

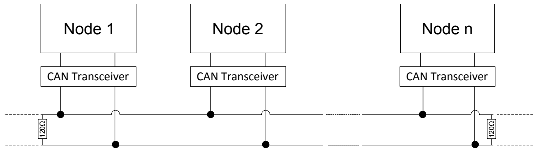

Fig. 1. Typical CAN bus topology

Obviously, the alternatives for assuring source identification on CAN are limited. To alleviate this, here we take an entirely distinct approach by trying to identify the nodes based on the signal patterns. The CAN specification allows great freedom in the implementation of the physical layer. As a consequence, signals produced by transceivers from different manufacturers are not identical. Moreover, as each electronic component gathers unique physical characteristics, even signals generated by transceivers from the same manufacturer show up unique peculiarities that can help to distinguish between senders.

Related work. Several lines of work were already focused on physical layer security. Hall et al. [3] used radio frequency fingerprinting for intrusion detection in wireless networks. Beamforming and artificial noise were used in wireless networks in a physical layer approach to provide secure communication in the presence of eavesdroppers [5]. Investigations were done in the case of wired buses as well. The work of Gerdes et al. [2] is focused on identifying Ethernet cards by studying the synchronization signal (found at the beginning of each Ethernet frame) with the help of a matched filter. Reported experimental results show that Ethernet cards of different models can be easily distinguished, while for cards of the same model an acceptable degree of accuracy can be

The authors are with the Faculty of Automatics and Computers, Politehnica University of Timisoara, Romania (Bd. Vasile Parvan nr. 2, 300223, Timisoara, Email: stefan.murvay@gmail.com, bogdan.groza@aut.upt.ro, phone: 0040721525705).

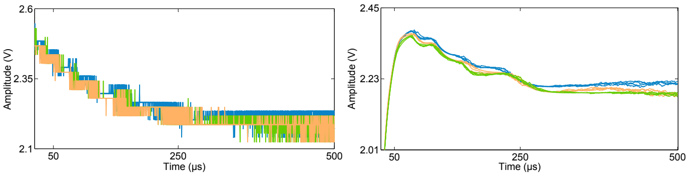

Fig. 2. Arbitration fields from three transceivers before low-pass filtering (left) and after low-pass filtering (right)

achieved.

II. Signal processing tools

We briefly describe the mathematical tools that we used in order to obtain a finer grain analysis of transceiver's characteristics. The following notations are used: the *fingerprint* $\mathcal{F}$ is the reference data stored for each device, the *signature* Sig represents a fresh data pending for verification and ℓ is the length of the previous two vectors. Subscripts α and β to Sig and $\mathcal{F}$ denote the index of a collected frame that is used either as signature Sig or fingerprint $\mathcal{F}$ (whenever $\alpha \neq \beta$ distinct frames are taken into account).

1) Low-pass filtering: Signatures extracted from CAN frames have to be compared with stored fingerprints. As shown in Figure 2 the signal is noisy and a simple bit-by-bit comparison between two signals may not be even possible as signals greatly overlap. As a first step, we filter the acquired signal in order to remove, as much as possible, from the unwanted noise. Figure 2 illustrates signals representing the start of the arbitration fields from 3 transceivers with very similar behavior both in an unfiltered and filtered form. The overlapping is abundant with the unfiltered signals. The filtering that we use is a low-pass filter as described by the following equation: $\widetilde{Sig}_\alpha[i] = \lambda \cdot Sig_\alpha[i] + (1-\lambda) \cdot \widetilde{Sig}_\alpha[i-1], i = 1..\ell$ (λ is the smoothing factor). This basic filtering technique produces a noticeable difference between the signals, but it is not yet enough to distinguish between very similar ones. The following techniques will be also applied on the unfiltered signal (in this case the tilde notation is omitted).

2) Mean squared error: For applying MSE in our case we consider the stored fingerprint as the reference set and the signal to be checked as the second set. The mean squared error is computed as: $MSE(Sig_\alpha, \mathcal{F}_\beta) = \frac{1}{\ell}\sum_{i=1}^{\ell}(Sig_\alpha[i] - \mathcal{F}_\beta[i])^2$. MSE can either be computed over the signal to be verified for all the stored fingerprints corresponding to the specified CAN ID, or else, the signature can be compared with the average of the stored fingerprints. If the signal being verified comes from one of the authorized nodes then the MSE computed with the node's fingerprint should have the lowest value while all other MSE computations should result in greater values.

3) Convolution: One common issue in signal processing is signal misalignment. Even when using a high-end oscilloscope (as we did for our tests) one cannot guarantee that the sampled signals are perfectly aligned. This problem can be alleviated by convolving the compared signals. The result of this operation is a vector of length $n+m-1$, where m and n are the lengths of the two input vectors. These vectors have the same length ℓ, thus, in our case the length of the result is $2\ell - 1$. Each element of the output is defined as: $CONV_k(Sig_\alpha, \mathcal{F}_\beta) = \sum_{i=1}^{\ell}(\mathcal{F}_\beta[i] \cdot Sig_\alpha[k-i+1])$. We consider the maximum value of the convolution vector as representative for signal similitude to the reference since it marks the point of best alignment. Therefore, the value used for comparison is: $MAXCONV = \max_{k=1}^{2\ell-1}\left(\sum_{i=1}^{\ell}(\mathcal{F}_\beta[i] \cdot Sig_\alpha[k-i+1])\right)$.

III. Experimental results

We performed our tests on CAN frames produced by two different types of devices: USB-to-CAN adapters and embedded development boards. The acquisition was done using an Agilent MSO6012A oscilloscope with a sample rate of 2 GSa/s and a resolution of up to 12 bits. The acquired CAN signals were saved on a PC using Matlab R2012 and Agilent device drivers to control the acquisition.

A. Source identification

In order to test the source identification techniques, we collected thousands of CAN frames from a series of 10 USB-to-CAN devices (sysWORXX series from Systec Electronics) and 5 development boards (ZK-S12-B from SoftecMicro) with S12 cores (each board has 2 transceivers and each of them was sampled). Differences between the signals produced by different types of transceivers are clearly visible even without applying additional processing techniques. For transceivers of the same type, signal similarities are considerable and require additional processing of the acquired data. The main set of signals was generated with CAN ID set to 0x000 at 10Kbaud.

1) MSE based separation: Figure 3 presents MSE values computed for a series of PCA82C251 transceivers (from the USB-to-CAN modules), numbered from 1 to 10, using as a reference the fingerprint for the fourth module $\mathbf{T}^4_{\mathrm{USB}}$. For each transceiver 20000 frames were captured. The band containing MSE values for signals produced by $\mathbf{T}^4_{\mathrm{USB}}$ (red) is situated in the lower part of the plot suggesting good similarities. The majority of transceivers, with two exceptions, produce less similar signals. One exception comes from $\mathbf{T}^1_{\mathrm{USB}}$ (purple) which generates signatures that overlap with values from the band for $\mathbf{T}^4_{\mathrm{USB}}$. The signature band for $\mathbf{T}^6_{\mathrm{USB}}$ (orange) stays lower on the graph and could be falsely deemed as the better fitting for the fingerprint of $\mathbf{T}^4_{\mathrm{USB}}$. In the case of TJA1054T transceivers (Figure 4) there are three transceivers that produce signals very similar to the reference signal (produced by $\mathbf{T}^{2''}_{\mathrm{S12}}$ – the second transceiver from the second development board) as displayed in the zoomed window. The difference between the other four transceivers and the reference is in this case greater than for the USB-to-CAN modules.

2) Convolution based separation: Figure 5 presents values obtained using convolutions over the same data set as in the previous section. The band containing values for signals produced by $\mathbf{T}^4_{\mathrm{USB}}$ is constantly overlapped by the band generated for $\mathbf{T}^6_{\mathrm{USB}}$ while the band for $\mathbf{T}^7_{\mathrm{USB}}$ is just occasionally overlapping it. The rest of the transceivers produce results that can be clearly distinguished from the target $\mathbf{T}^4_{\mathrm{USB}}$. When applying the convolution on the signals from the S12 board, results are similar as in the case of computing the MSE. Three transceivers are hard to distinguish from the true target based on the signals that they generate, while the other six transceivers clearly generate different signals as can be seen in Figure 6.

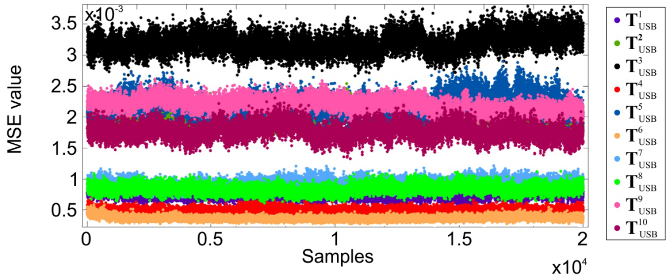

Fig. 3. 2×10^4 MSE values for PCA transceivers with $\mathbf{T}^4_{\mathrm{USB}}$ as fingerprint

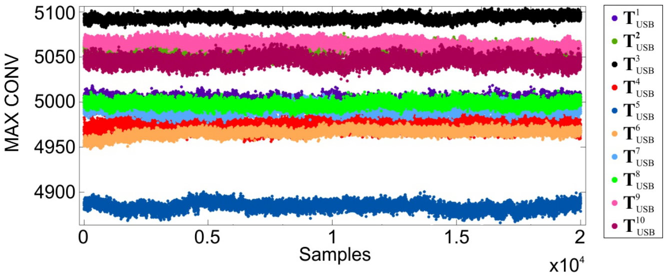

Fig. 5. 2×10^4 convolved values for PCA transceivers with $\mathbf{T}^4_{\mathrm{USB}}$ as fingerprint

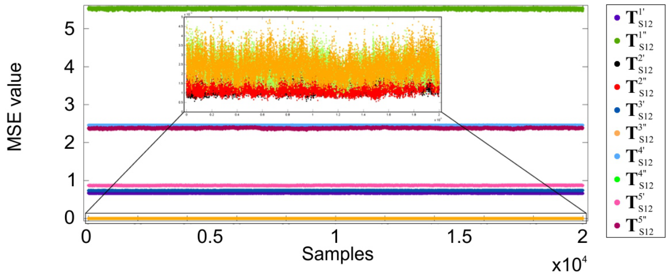

Fig. 4. 2×10^4 MSE values for TJA transceivers with $\mathbf{T}^{2''}_{\mathrm{S12}}$ as fingerprint

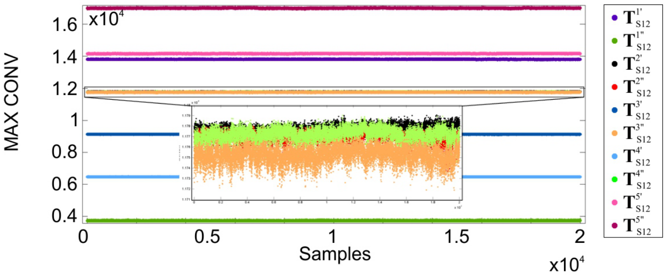

Fig. 6. 2×10^4 convolved values for TJA transceivers with $\mathbf{T}^{2''}_{\mathrm{S12}}$ as fingerprint

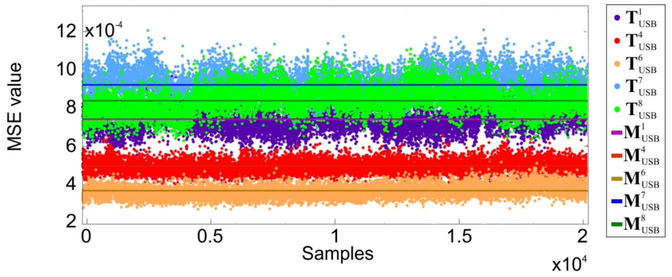

Fig. 7. Mean of 2×10^4 MSE values for PCA transceivers with signatures similar to the $\mathbf{T}^4_{\mathrm{USB}}$ fingerprint

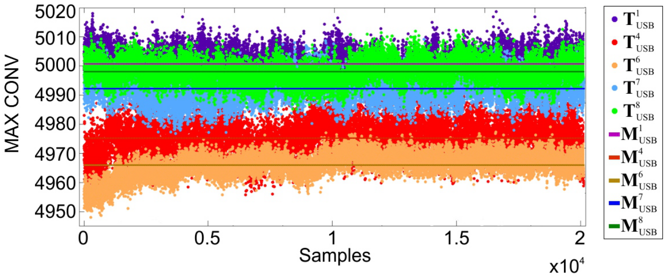

Fig. 9. Mean of 2×10^4 convolved values for PCA transceivers with signatures similar to the $\mathbf{T}^4_{\mathrm{USB}}$ fingerprint

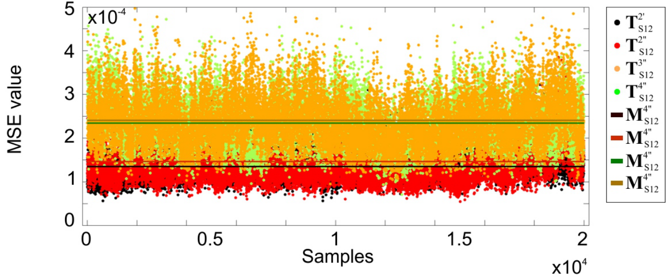

Fig. 8. Mean of 2×10^4 MSE values for TJA transceivers with signatures similar to the $\mathbf{T}^{2''}_{\mathrm{S12}}$ fingerprint

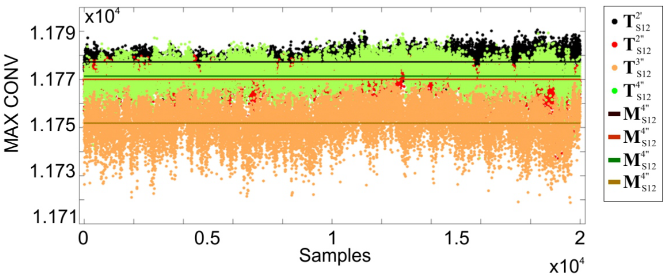

Fig. 10. Mean of 2×10^4 convolved values for TJA transceivers with signatures similar to the $\mathbf{T}^{2''}_{\mathrm{S12}}$ fingerprint

3) Mean-value based separation: As reflected by the plots, in some cases the signature of a single signal might not be enough to accurately determine its source. To increase detection accuracy we tested a multi-signature based identification by computing mean values for signatures and comparing it with a fingerprint obtained in the same manner. We superimposed a line representing the mean value of collected signatures over each set presented in the previous plots. Figures 7 through 10 contain signatures of the reference transceiver and signatures very similar to them along with the mean values for each data set.

B. Success rate of identification

As the identification success rates for using MSE are very similar to the ones obtained when using convolution based separation we include only results for the MSE case. Mean-value based separation increases the identification accuracy but only at the cost of more signal acquisitions which makes it less appealing for practical scenarios. For example, computing the mean on only 100 signals did not bring any noticeable improvements, these were visible only for means computed on 1000 frames or more.

Tables I and II present a measure of the detection accuracy when using MSE based detection. Each cell contains the identification result for a transceiver against a target transceiver designated as the first entry in each row. Cell values range from 0 (transceiver was not identified as being the target transceiver) up to 1 (tested transceiver is the target) and were computed from 20000 signals for each transceiver. For easy reading, we use the mark ✓ whenever there was no confusion between the two transceivers. Thresholds were empirically defined for

TABLE I
IDENTIFICATION RATES FOR PCA82C251 (ID SET TO 0X000)

Target	$\mathbf{T}^{1}_{\mathrm{USB}}$	$\mathbf{T}^{2}_{\mathrm{USB}}$	$\mathbf{T}^{3}_{\mathrm{USB}}$	$\mathbf{T}^{4}_{\mathrm{USB}}$	$\mathbf{T}^{5}_{\mathrm{USB}}$	$\mathbf{T}^{6}_{\mathrm{USB}}$	$\mathbf{T}^{7}_{\mathrm{USB}}$	$\mathbf{T}^{8}_{\mathrm{USB}}$	$\mathbf{T}^{9}_{\mathrm{USB}}$	$\mathbf{T}^{10}_{\mathrm{USB}}$
$\mathbf{T}^{1}_{\mathrm{USB}}$	0.995	✓	✓	0.001	✓	✓	✓	✓	✓	✓
$\mathbf{T}^{2}_{\mathrm{USB}}$	✓	0.920	✓	✓	✓	✓	✓	✓	0.695	0.003
$\mathbf{T}^{3}_{\mathrm{USB}}$	✓	✓	1	✓	✓	✓	✓	✓	✓	✓
$\mathbf{T}^{4}_{\mathrm{USB}}$	0.001	✓	✓	0.985	✓	0.151	✓	✓	✓	✓
$\mathbf{T}^{5}_{\mathrm{USB}}$	✓	✓	✓	✓	1	✓	✓	0.002	✓	✓
$\mathbf{T}^{6}_{\mathrm{USB}}$	✓	✓	✓	0.001	✓	0.999	✓	✓	✓	✓
$\mathbf{T}^{7}_{\mathrm{USB}}$	✓	✓	✓	✓	✓	✓	0.941	0.004	✓	✓
$\mathbf{T}^{8}_{\mathrm{USB}}$	✓	✓	✓	✓	✓	✓	0.002	0.967	✓	✓
$\mathbf{T}^{9}_{\mathrm{USB}}$	✓	0.437	0.001	✓	✓	✓	✓	✓	0.994	✓
$\mathbf{T}^{10}_{\mathrm{USB}}$	✓	0.047	✓	✓	✓	✓	✓	✓	✓	0.997

TABLE II
IDENTIFICATION RATES FOR TJA1054T (ID SET TO 0X000)

Target	$\mathbf{T}^{1'}_{\mathrm{S12}}$	$\mathbf{T}^{1''}_{\mathrm{S12}}$	$\mathbf{T}^{2'}_{\mathrm{S12}}$	$\mathbf{T}^{2''}_{\mathrm{S12}}$	$\mathbf{T}^{3'}_{\mathrm{S12}}$	$\mathbf{T}^{3''}_{\mathrm{S12}}$	$\mathbf{T}^{4'}_{\mathrm{S12}}$	$\mathbf{T}^{4''}_{\mathrm{S12}}$	$\mathbf{T}^{5'}_{\mathrm{S12}}$	$\mathbf{T}^{5''}_{\mathrm{S12}}$
$\mathbf{T}^{1'}_{\mathrm{S12}}$	1	✓	✓	✓	✓	✓	✓	✓	✓	✓
$\mathbf{T}^{1''}_{\mathrm{S12}}$	✓	1	✓	✓	✓	✓	✓	✓	✓	✓
$\mathbf{T}^{2'}_{\mathrm{S12}}$	✓	✓	0.908	0.876	✓	0.204	✓	0.029	✓	✓
$\mathbf{T}^{2''}_{\mathrm{S12}}$	✓	✓	0.831	0.901	✓	0.118	✓	0.121	✓	✓
$\mathbf{T}^{3'}_{\mathrm{S12}}$	✓	✓	✓	✓	1	✓	✓	✓	✓	✓
$\mathbf{T}^{3''}_{\mathrm{S12}}$	✓	✓	0.999	0.991	✓	0.901	✓	0.300	✓	✓
$\mathbf{T}^{4'}_{\mathrm{S12}}$	✓	✓	✓	✓	✓	✓	1	✓	✓	✓
$\mathbf{T}^{4''}_{\mathrm{S12}}$	✓	✓	0.527	0.934	✓	0.029	✓	0.900	✓	✓
$\mathbf{T}^{5'}_{\mathrm{S12}}$	✓	✓	✓	✓	✓	✓	✓	✓	1	✓
$\mathbf{T}^{5''}_{\mathrm{S12}}$	✓	✓	✓	✓	✓	✓	✓	✓	✓	1

each target transceiver based on the MSE value ranges so that the detection rate of each target transceiver is not smaller than 90%. We allowed rates smaller than 100% for transceivers with similar behavior when the bands of MSE values overlapped.

The results shown in these tables come as a support for the previously depicted graphical representations. The batch of PCA82C251 transceivers seem to be easier to distinguish with a maximum false detection rate of 69.5% ($\mathbf{T}^{2}_{\mathrm{USB}}$ vs. $\mathbf{T}^{9}_{\mathrm{USB}}$) while in the case of some of TJA1054T transceivers this range can go up to 99.9% ($\mathbf{T}^{2'}_{\mathrm{S12}}$ vs. $\mathbf{T}^{3''}_{\mathrm{S12}}$) which means that similarities in signalling behavior are too big to make a correct identification. However, this confusion can be apparently circumvented by clever allocation of the IDs as we discuss next.

Table III shows the identification rate while changing the ID to 0x555 and the baud-rate to 125k. At a glance, identification rate seems to worsen, however, on a closer inspection it turns out that for transceivers $\mathbf{T}^{2'}_{\mathrm{S12}}$ and $\mathbf{T}^{3''}_{\mathrm{S12}}$, that were hardly distinguishable, the confusion rate now drops to 0%. This strongly suggests that clever allocation of IDs for specific transceivers can yield extremely high identification rates. Changing the data-rate back to 10 Kbaud preserved the changes in the overlapping pattern which proves that this is determined by the ID alone.

TABLE III
IDENTIFICATION RATES FOR TJA1054T (ID SET TO 0X555)

Target	$\mathbf{T}^{1'}_{\mathrm{S12}}$	$\mathbf{T}^{1''}_{\mathrm{S12}}$	$\mathbf{T}^{2'}_{\mathrm{S12}}$	$\mathbf{T}^{2''}_{\mathrm{S12}}$	$\mathbf{T}^{3'}_{\mathrm{S12}}$	$\mathbf{T}^{3''}_{\mathrm{S12}}$	$\mathbf{T}^{4'}_{\mathrm{S12}}$	$\mathbf{T}^{4''}_{\mathrm{S12}}$	$\mathbf{T}^{5'}_{\mathrm{S12}}$	$\mathbf{T}^{5''}_{\mathrm{S12}}$
$\mathbf{T}^{1'}_{\mathrm{S12}}$	1	✓	✓	✓	✓	✓	✓	✓	✓	✓
$\mathbf{T}^{1''}_{\mathrm{S12}}$	✓	0.916	✓	0.004	✓	0.369	✓	✓	✓	✓
$\mathbf{T}^{2'}_{\mathrm{S12}}$	✓	✓	0.950	0.250	✓	✓	✓	0.452	✓	✓
$\mathbf{T}^{2''}_{\mathrm{S12}}$	✓	0.078	0.263	0.950	✓	0.777	✓	0.767	✓	0.72
$\mathbf{T}^{3'}_{\mathrm{S12}}$	✓	✓	✓	✓	1	✓	✓	✓	✓	✓
$\mathbf{T}^{3''}_{\mathrm{S12}}$	✓	0.426	✓	0.784	✓	0.950	✓	0.527	✓	0.817
$\mathbf{T}^{4'}_{\mathrm{S12}}$	✓	✓	✓	✓	✓	✓	1	✓	✓	✓
$\mathbf{T}^{4''}_{\mathrm{S12}}$	✓	✓	0.545	0.817	✓	0.580	✓	0.950	✓	0.510
$\mathbf{T}^{5'}_{\mathrm{S12}}$	✓	✓	✓	✓	✓	✓	✓	✓	1	✓
$\mathbf{T}^{5''}_{\mathrm{S12}}$	✓	0.002	0.003	0.750	✓	0.879	✓	0.484	✓	0.950

Since collisions between transceiver patterns appear to be random, they should be precisely linked to the resolution of the measurements (in our case, a 12-bit resolution). Consequently, this will bound the success probability of an adversary that inserts its own device on the network. If the available resolution is too low, one should consider increasing the size of the ID or even extending the identification to other portions of the frames for better identification rates.

To address potential signal drifts in time, we repeated our experiments over a period of several months and found the signatures to remain within their boundaries. One of our tests consisted in continuous signal acquisition from devices over a period of roughly one month of uninterrupted functioning. No noticeable drift was visible during these tests. Indeed, over longer periods of time, it is possible for drifts to appear but we consider that this effect can be compensated by continuously updating the fingerprint of each device with newly received signals.

IV. CONCLUSIONS

The methodology described here proved to be workable in distinguishing the source of messages without any modifications on the software that the node is running or of the network that it is part of. This may set a distinct perspective on assuring broadcast authentication in CAN networks, an environment where cryptographic techniques are currently absent and will be hard to implement due to the various constraints. Assuring security on this kind of bus will likely become mandatory in the near future due to many safety-critical applications that rely on it: both automotives and industrial networks recently proved to be constant targets of cyber-attacks. Besides detecting intrusions, this technique may also be useful for forensic purposes as event data recorders are closer to be mandatory for newly produced automobiles.

Acknowledgment. This work was partially supported by the strategic grant POSDRU 107/1.5/S/77265, inside POSDRU Romania.

REFERENCES

[1] S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, S. Savage, K. Koscher, A. Czeskis, F. Roesner, and T. Kohno. Comprehensive experimental analyses of automotive attack surfaces. *Proceedings of the 2011 Usenix Security*, 2011.

[2] R. M. Gerdes, M. Mina, S. F. Russell, and T. E. Daniels. Physical-layer identification of wired ethernet devices. *IEEE Transactions on Information Forensics and Security*, 7(4):1339–1353, 2012.

[3] J. Hall, M. Barbeau, and E. Kranakis. Radio frequency fingerprinting for intrusion detection in wireless networks. *IEEE Transactions on Dependable and Secure Computing*, 2005.

[4] T. Hoppe, S. Kiltz, and J. Dittmann. Security threats to automotive CAN networks: Practical examples and selected short-term countermeasures. *Reliability Engineering & System Safety*, 96(1):11 – 25, 2011.

[5] N. Romero-Zurita, D. McLernon, M. Ghogho, and A. Swami. PHY Layer Security Based on Protected Zone and Artificial Noise. *IEEE Signal Processing Letters*, 20(5):487–490, 2013.

[6] I. Rouf, R. Miller, H. Mustafa, T. Taylor, S. Oh, W. Xu, M. Gruteser, W. Trappe, and I. Seskar. Security and privacy vulnerabilities of in-car wireless networks: A tire pressure monitoring system case study. In *Proceedings of USENIX Security Symposium*, 2010.

[7] C. Szilagyi and P. Koopman. Low cost multicast authentication via validity voting in time-triggered embedded control networks. In *Proceedings of the 5th Workshop on Embedded Systems Security*, page 10. ACM, 2010.

[8] A. Van Herrewege, D. Singelee, and I. Verbauwhede. CANAuth-a simple, backward compatible broadcast authentication protocol for CAN bus. In *9-th Embedded Security in Cars Conference*, 2011.