

A Comprehensive Solution for Securing Connected and Autonomous Vehicles

Mohsin Kamal^{*}, Christos Kyrkou^{*}, Nikos Piperigkos[†], Andreas Papandreou[†], Andreas Kloukiniotis[†], Jordi Casademont[‡],
Natlia Porras Mateu[§], Daniel Baos Castillo[§], Rodrigo Diaz Rodriguez[¶], Nicola Gregorio Durante[¶],
Peter Hofmann^{||}, Petros Kapsalas^{**}, Aris S. Lalos^{†,††}, Konstantinos Moustakas[†],
Christos Laoudias^{*}, Theocharis Theocharides^{*,o}, Georgios Ellinas^{*,o}

^{*}*KIOS Research and Innovation Center of Excellence, University of Cyprus, Nicosia, 1678, Cyprus*

[†]*Department of Electrical and Computer Engineering, University of Patras, Greece*

[‡]*Universitat Politecnica de Catalunya and Fundacio i2CAT, Barcelona, Spain*

[§]*Nextium by Idneo, Barcelona, Spain*

[¶]*Atos IT Solutions and Services Iberia S.L., Madrid, Spain*

^{||}*Deutsche Telekom Security GmbH, T-Systems, Berlin, Germany*

^{**}*Panasonic Automotive, Langen, Germany*

^{††}*Industrial Systems Institute, Athena Research Center, 26502, Platani, Patras, Greece*

^o*Department of Electrical and Computer Engineering, University of Cyprus, Nicosia, 1678, Cyprus*

Abstract—With the advent of Connected and Autonomous Vehicles (CAVs) comes the very real risk that these vehicles will be exposed to cyber-attacks by exploiting various vulnerabilities. This paper gives a technical overview of the H2020 CARMEL project (currently in the intermediate stage) in which Artificial Intelligent (AI)-based cybersecurity for CAVs is the main goal. Most of the possible scenarios are considered, by which an adversary can generate attacks on CAVs, such as attacks on camera sensors, GPS location, Vehicle to Everything (V2X) message transmission, the vehicle’s On-Board Unit (OBU), etc. The counter-measures to these attacks and vulnerabilities are presented via the current results in the CARMEL project achieved by implementing the designed security algorithms.

I. INTRODUCTION

The global market for Connected and Autonomous Vehicles (CAVs) is expected to reach \$7 trillion by 2050 [1]. With the introduction of new standards for vehicular communications (e.g., IEEE 802.11p and LTE-PC5 that operate in infrastructure-free mode and unlicensed frequency bands and LTE-Uu that uses licensed bands and requires an operator’s infrastructure), the risk of cyber-threats has increased dramatically [2]. The *CARMEL* project, currently in the intermediate stage, aims to provide innovative intrusion detection and prevention solutions for the automobile industry by applying advanced AI and Machine Learning (ML) techniques.

The *CARMEL* cyber defense solutions tackle cyber-attacks targeting the CAV’s: (i) camera that uses AI to

understand the scene and take vital decisions, (ii) message exchange with nearby CAVs, (iii) OBU, and (iv) GPS location. Additionally, V2X communication is secured using a Public Key Infrastructure (PKI), based on the European Telecommunications Standards Institute (ETSI) Intelligent Transportation Systems (ITS) standards, enhanced with an efficient and scalable Certificate Revocation Lists (CRL) distribution system.

Specifically, one of the objectives in *CARMEL* is to demonstrate AI/ML-based techniques for detecting and mitigating cyber-attacks on the camera system/data in automated driving systems. In many ITS applications, the patterns of interest cannot be reliably classified by explicit programming due to the diversity of the data streams. In such cases, if sufficiently large amounts of example (training) data are available or feasible to obtain, then ML approaches are applied (typically based on neural networks). Another objective is to protect the connected vehicles by designing algorithms for detecting anomalies on the communication interface, such as the exchange of messages and GPS location spoofing attacks.

CARMEL addresses cybersecurity challenges by introducing two new hardware elements where ML algorithms can be executed. A Multi-access Edge Computing (MEC) module is proposed, connected to the fixed radio infrastructure, which receives all messages transmitted by CAVs and performs interoperability and security functions. In addition, an Anti-Hacking Device (AHD), combined with the OBU is introduced, which is deployed in the vehicle, aiming to tighten the security measures and make it very hard to hack the vehicle’s systems by providing algorithms for secure boot, secure firmware update, etc. The *CARMEL* ecosystem and cybersecurity solution is shown in Fig. 1.

The key contributions of this paper are: (i) providing insights on vulnerabilities and attacks on different elements of CAVs, and (ii) presenting an overview of cybersecurity solu-

This work was supported by the European Union’s H2020 research and innovation programme under the *CARMEL* project (Grant agreement No 833611). The work of M.K., C.K., C.L., T.T., and G.E. has also been supported by the European Union’s Horizon 2020 research and innovation programme under grant agreement No 739551 (KIOS CoE - TEAMING) and from the Republic of Cyprus through the Deputy Ministry of Research, Innovation and Digital Policy. The work of J.C. has also been supported by grants PID2019-106808RA-I00 and PID2020-112675RB funded by MCIN/AEI/10.13039/501100011033.

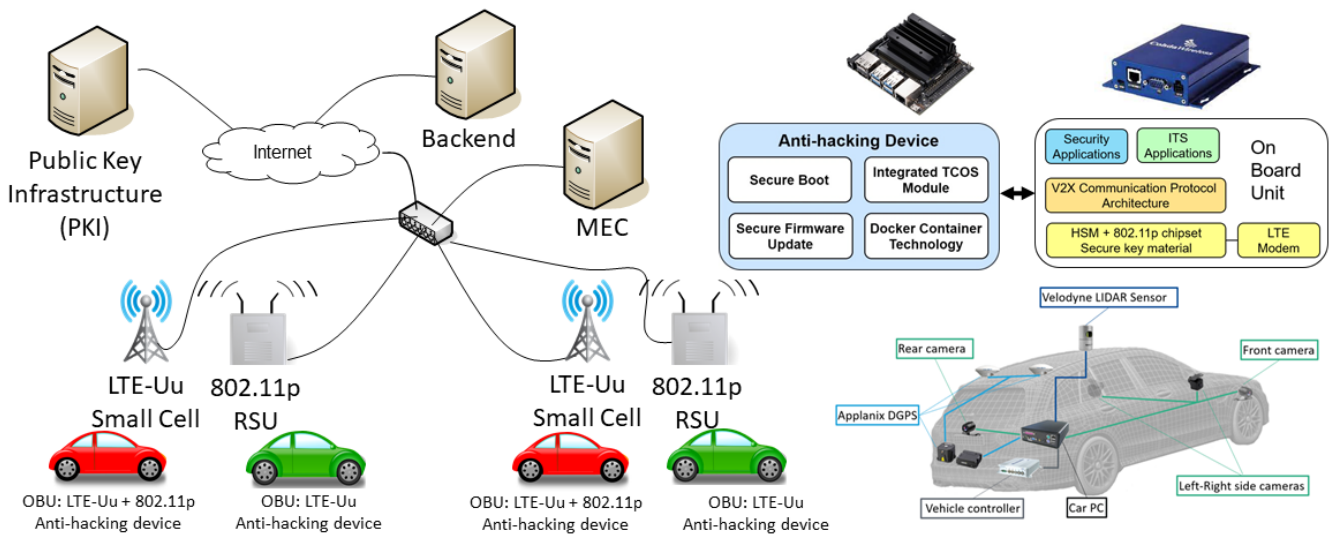


Fig. 1: CARMEL cybersecurity solution: Connectivity infrastructure (left), the CAV (bottom-right), and the coupled OBU and AHD (top-right) integrated into the CAV.

tions for CAVs and the results achieved so far by CARMEL.

II. SECURING AUTONOMOUS VEHICLES

The autopilot feature in autonomous vehicles depends heavily on computer vision and AI algorithms. Sensor data can be manipulated directly to elicit false algorithmic inferences. In camera sensor attacks, an attacker manages to access the vehicle’s critical systems and install and activate malicious software that distorts the captured camera data. Adversarial attacks seek small perturbations of the input guided by the perception module, causing large errors in the estimation by the perception modality. Image deterioration attacks aim to alter the input image in order to lead the vehicle perception modules to fail. In contrast to adversarial attacks, deterioration attacks are not guided by a target label or model behavior; rather they arbitrarily cause a general drop in the quality of the image (by adding noise or artifacts) so that the perception module’s output becomes erroneous (Fig. 2). Herein, mitigation strategies for both attack classes are described.

A. Adversarial Attack on Camera Sensor

A multi-modal fusion technique is implemented to robustify scene analysis in autonomous vehicles against adversarial attacks. A holistic approach has been proposed by integrating a denoising module atop an image segmentation network to




Fig. 2: Cyber-attack causing wrong inference.




Fig. 3: Attack scenario: 3D object detection (left), 2D robust segmentation (bottom) and multi-modal output (top).

improve its robustness to adversarial inputs. The denoising block, added before the segmentation model, handles images with arbitrary degradation levels and generates smooth restoration effects without artifacts. As a second layer of defense, a late fusion method has been applied, by extracting direct features from 3D point clouds and projecting them onto a 2D segmented image for identifying inconsistencies. As a result, anomaly detection on the camera sensor can be detected, leading to more secure perception systems.

Specifically, in the proposed solution, the models of AdaFM [3] and Deeplab [4] were utilized for the denoising and segmentation operations accordingly. Both of the previous models were trained separately using data from the KITTI dataset [5]. As it regards the impact of the denoiser on attacked images, some meaningful results were generated for multiple adversarial attacks. For instance, by applying a Projected Gradient Descent (PGD) attack to the clean images, the segmentation result was dropped to 0.69/0.53/0.17% Intersection over Union (IoU) for 2/4/8 levels of perturbations. By applying the denoiser, the performance was increased, reaching 0.74/0.72/0.72% IoU for the same levels of perturbations. Even though the denoiser achieves valuable results, a second layer of defense has also been added in the case where the adversarial noise could not be alleviated successfully.

Hence, PointRCNN [6] was integrated into the pipeline that achieves state-of-art results by reaching 78.70/54.41/72.11% recall for the classes of car/pedestrian/cyclist in a two-stage 3D object detection framework. Finally, by correlating outputs of different perception modules with additional sensor readings, it is possible to provide improved situational awareness for the driver. This is illustrated in Fig. 3, where the final decision of the perception engine is ascertained only from the lidar sensor, by projecting the 3D outputs to the image plane, as the 2D attacked image cannot be entirely restored.

B. Image Deterioration Attack on Camera Sensor

1) *Variational Techniques for Noise Suppression*: Well-examined approaches of noise suppression have mainly used tools of linear systems. However, linear methods of noise suppression are not efficient, since the nature of most image processing problems is often ill-posed. To alleviate the ill-posedness, a regularization process must be adopted, reformulating the ill-posed problem so that an appropriate solution can be found. Such a solution is usually sufficient to approximate the original problem at an admissible level. Two variational techniques examined are (i) total variation that models image restoration as a solution of the Total Variation Partial Differential Equation (PDE). Such an approach is efficient to restore images corrupted by multiple types of noise; (ii) context-aware operators based on anisotropy in diffusion. The Anisotropic PDE favors the evolution of level sets in some directions and prevents it in others. This is specified by local eigenvectors and eigenvalues of the diffusion tensor field [7].

2) *Deep-Learning-Based Techniques*: The problem at hand is image distortion due to added noise, as well as image manipulation, that can result in not detecting structural elements of a scene (e.g., traffic signs, pedestrians, etc.). Such attacks are easy to be performed, as the attacker does not need to have any information regarding the underlying algorithm and models and can result in an erroneous output of a perception module. Existing pre-processing approaches, such as filtering (bilateral or Gaussian), even though effective against specific attack types, may fail when artifacts are completely removed [8]. Hence, a solution called *DriveGuard* (Fig. 4) is developed, that uses convolutional autoencoder models to improve the robustness of the image segmentation models and proactively mitigate the effects of deteriorated image quality [9]. The proposed approach is based on a lightweight spatio-temporal autoencoder, utilizing separable convolutions, as an image reconstruction tool to robustify semantic segmentation for




Fig. 4: DriveGuard deep learning architecture.

autonomous vehicle applications. This approach is applied on images from the CARLA simulator [10] and Cityscapes dataset [11] and manages to mitigate attacks and restore image segmentation outcomes even for heavily distorted images to within 5-6% of the original model.

A detection process is also incorporated, operating in parallel to attack mitigation, that uses the input image and the generated output to detect whether an attack has occurred (by measuring their differences). Consequently, when the detector is triggered, the DriveGuard model outputs the mitigated generated image in case of an attack, or the unaffected input image, otherwise. Overall, the proposed approach is capable of recognizing an attack with an accuracy higher than 97.5% when tested on a dataset of attacked images, 500 from the CARLA simulator and 500 from the Cityscapes dataset.

III. SECURING CONNECTED VEHICLES

Attacks on OBU, V2X message transmission, and GPS location are discussed below, as well as the defense mechanisms against a detected malicious actor considering a delay in certificate revocation due to large message overhead. The attack response workflow is shown in Fig. 5, where a flag is raised after each attack resulting in the certificate revocation of the entity that is thought to be compromised.

A. Detection of Tamper Attack on Vehicle's OBU

1) *The Vehicle's OBU*: Standard cooperative cars are equipped with an OBU which provides secure communication functionalities. CARMEL introduced the necessity for the development of an enhanced OBU to comply with the current security regulations and to combine it with the AHD which is able to detect hacking attempts and functional misbehaviors using ML-based algorithms. The OBU architecture, includes the following main elements, as shown in Fig. 1 (top-right):

- Hardware Security Module (HSM), serving as a repository for private key data (for authentication and encryption), and as a cryptographic processor for sensitive operations. Hardware security functions are used to protect the OBU against

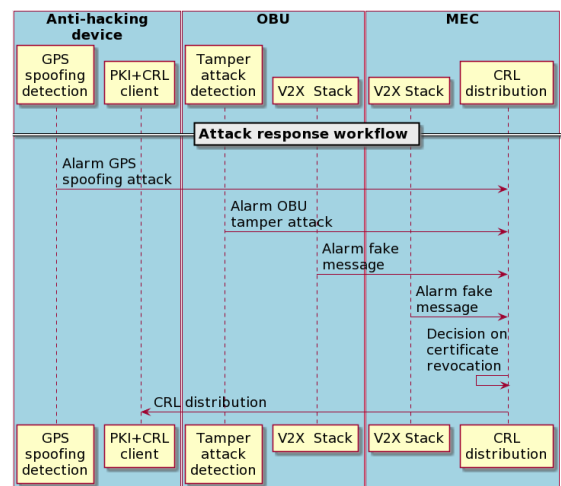


Fig. 5: Attack response workflow.

tamper attacks, through box opening detection, active hardware protection of susceptible signals, and environmental sensors to prevent fault injection attacks.

- Radio interfaces (IEEE 802.11p and LTE-Uu) for connecting to the PKI servers in order to obtain the pseudonymous Authorization Tickets (ATs) before being able to transmit ITS messages, and for real-time management of certificates.
- V2X Communication Protocol Architecture, that contains the software package enabling the OBU to generate messages that are properly extended to perform all security- and privacy-related functionalities.
- Security Applications module, that contains the software functions to interact with the PKI and manage the registration and authorization procedures, as well as to obtain the pseudonymous ATs and store them in the HSM.
- ITS Applications module, that represents any ITS application running on the vehicle, e.g., for exchanging CAMs
- AHD, that extends the OBU's security capabilities through advanced techniques for detecting misbehaviors and attacks.

2) *Attack Scenarios*: Attacks considered include both hardware and software attacks. Hardware attacks concern any physical manipulation (i) at the environmental conditions such as temperature and voltage, or (ii) directly at the physical OBU. Countermeasures introduced range from environmental sensors to wire-meshes or switches to detect any manual manipulation. Regarding software attacks, these include all attacks focused on the verification or encryption processes, thus countermeasures can rely on mutual authentication, data encryption, and secure boot. In case the OBU processes any system alarm in some way, it triggers secure-state actions to protect confidential information, e.g., zeroization of the private keys and any other confidential data.

B. Detection of Attack on V2X Message Transmission

The first step to establish V2X security is the PKI architecture, comprised five different servers (Root Certification Authority, Online Certification Authority, Enrollment Authority, Authorization Authority, Validation Authority), that supply ATs to the vehicles (after the enrollment and authorization phases) for digitally signing the transmitted packets. The private key of each AT is stored securely in the OBU's HSM so as not to be compromised. With this mechanism, every vehicle receiving a V2X message is able to detect if the message has been transmitted by a non-compliant vehicle or if it is supplanting its role. To prevent a spoofer from keeping track of a target vehicle, the OBU changes the AT from time to time. Thus, the challenge is to identify the optimum moment for this change. An ML-based ATs' scheduler is developed, that decides the best time to change the AT, by evaluating how easily the vehicle can be tracked considering the V2X messages in the area, how many remaining ATs are still available in the vehicle, and the time remaining to obtain new ATs. The algorithm, first selects a set of messages that are candidates to belong to the same vehicle. Then, using a Random Forest learning method, the system obtains a tracking score value representing the probability that two messages

were sent by the same vehicle. Finally, using an optimal stopping model, the AT change decision algorithm decides when is the best moment to change the AT depending on the number of vehicle's non-used ATs and the previous tracking score.

The other component that enhances performance and security is the MEC which processes all V2X messages received by the 802.11p RSUs and LTE small cells. These messages are checked and, using policies based on region of interest and type of vehicles, they are forwarded to vehicles using other V2X radio technologies. The MEC is also responsible to collect alarms triggered in the different elements of the system, evaluate their level of risk and, eventually, to revoke the certificates of the attacked vehicles. In case a decision has been reached that the certificates of the considered vehicles have to be revoked, they are included in the CRL and distributed to the rest of the vehicles (Fig. 5). In order to make this distribution scalable, CRLs are used based on a Bloom Filter (BF), which is a probabilistic data structure allowing to efficiently compress information using the output of multiple hash functions, representing a set in a space-efficient way.

C. Detection of a GPS Location Spoofing Attack

GPS location spoofing is a crucial security threat, which needs to be tackled in order to guarantee safety for CAVs.

1) *Cooperative Detection Solution*: When vehicles need to cooperate, the same reference system is often required. For that purpose, GPS is more attractive although it suffers from inaccuracies. This approach makes use of the star V2V topology formed by each ego vehicle (center of the star) and its neighbors in the form of the Laplacian matrix. Note that the centralized methods [12] and [13] which motivated us, exploited the overall Laplacian matrix of a V2V cluster of vehicles. The ego vehicle computes its differential coordinates with respect to its nearby vehicles by extracting the range measurements from visual sensors. Nearby vehicles also send their range measurements towards the ego as well as their GPS positions (potentially spoofed). Under the assumption that only a small number, i.e., 5-15%, of vehicles of the star topology may have been attacked, the same sparse optimization problem as in [12], [13] is formulated, though now focusing on a new distributed implementation method (star Laplacian matrix, GPS positions, and range measurements of the ego vehicle and its neighbors), called Robust Distributed Localization (RDL).




Fig. 6: Positioning accuracy CDF (RDL, intact, spoofed GPS).

The CARLA simulator is used during the experiments and evaluation. GPS noise is modeled as a Gaussian distribution, while to simulate spoofing a fixed outlier to the position is added. It is assumed that vehicles communicate amongst themselves if their distance is $\leq 30\text{m}$. In the worst case, 10% of the vehicles of the star topology (including the ego) are spoofed during the simulation. Fig. 6 shows the Cumulative Distribution Function (CDF) of positioning accuracy, demonstrating the feasibility of the proposed distributed approach (omitting results for the centralized approaches in [12] and [13], since the new approach is completely distributed). Even when the spoofed vehicles are 10%, RDL provides high enough accuracy to the ego vehicle resulting in 90% reduction of the spoofed GPS error. RDL proves to be much more accurate than the theoretically intact GPS, achieving maximum error lower than 8m, instead of 10m with GPS.

2) *In-Vehicle Detection Solution:* The in-vehicle solution against GPS location spoofing attacks fuses multi-source data, readily available from the CAV's on-board sensors.

Process Outline: In particular, the proposed detection solution is based on a GPS integrity check [14]. In the *prediction* phase, multi-sensory data collected through the OBU and/or the CAN bus are used to compute the CAV's predicted location in the next time step given the previous CAV location estimate. This is achieved by projecting the location ahead in time using the sensor readings and a CAV mobility model. In the *update* phase, the CAV location measurements, provided by a GPS-free localization algorithm (e.g., based on cellular networks), are used to update the predicted location by means of Bayesian filtering and derive a refined location estimate. Finally, in the *attack detection* phase, the CAV location provided by the GPS receiver is compared to the refined location. If their distance d (e.g., Euclidean or Bhattacharyya) exceeds a threshold T_d , then an attack is detected. In case of attack, as a possible mitigation measure the CAV may discard the GPS readings and rely instead on the refined locations.

Online Adaptive Threshold: In our previous work, the detection threshold T_d was determined *offline* [12]–[14]. Specifically, a dataset of GPS and refined locations was collected by a CAV moving in an attack-free scenario. The series of distance values was filtered (i.e., averaging with a sliding window of length w) and T_d was selected at the 95th percentile of the Empirical CDF. However, the uncertainty in the GPS location may vary depending on the environment, e.g., it is higher in urban areas compared to open-sky rural areas. Thus, a T_d selected when the CAV was moving in one environment may lead to false alarms or misdetections when the CAV moves in a different setting. This necessitates a time-varying *online* threshold that can adapt to dynamically changing conditions. A baseline online adaptive threshold approach is presented, where batch i that contains the last n data samples is used to calculate the distance values, filter them, compute the threshold T_d^i , and then detect possible attacks in the subsequent samples similarly to the *offline* approach. Only the attack-free distance values are used to compute T_d^i . Preliminary results indicate that detection accuracy can be improved, as the F1 score is

0.9711 compared to 0.9693 attained by the *offline* approach.

D. Latency in Certificate Revocation

In this scenario, a malicious actor is detected by the MEC but the updated CRL has not yet reached the related vehicles due to its size (and consequent transmission time), allowing the attacker to act with impunity until the CRL has been broadcasted completely. To fulfill the privacy requirements on the vehicle's identity, a single vehicle is assigned a high number of pseudonym temporal identities (ATs), that are revoked when the vehicle is invalidated, resulting in an n -fold increase of the CRL size when the number of invalidated vehicles increases.

To minimize the CRL distribution time, the revoked certificates are represented using a BF. Concerning its implementation and alignment with existing standards, the ETSI TS 102 941 [15] standard defines a sequence of HashedId8 to identify the ATs in the CRL. Each AT in the sequence is represented with an 8 bytes hash, and to optimize the size of the CRL to be broadcasted, the sequence of ATs was replaced with a serialized BF. Fig. 7 depicts the size of the standard CRL against the BF implementation of the CRL for 10^7 revoked ATs, considering the additional backup certificates assigned to the vehicles. As the BF bit-array pre-allocates the space to store the ATs, an adaptive approach was deployed, gradually increasing the BF array space with the growth of revoked AT.

A BF query on whether a certificate belongs to the set will never result in a false negative, i.e., revoked certificate detected as valid; but occasionally it may return a false positive (FP), i.e., a valid certificate detected as revoked. The False Positive Probability (FPP) can be tuned with the size of the filter (m), the number of elements in the set (n), and the number of hash functions used (k) as $[1 - (1 - 1/m)^{kn}]^k$.

To address the FP occurrence, the sender vehicle checks the AT it is about to use to sign the message to be transmitted against the CRL stored in its memory. If it is detected as a FP, the sender vehicle discards it and tries a different AT. To counterbalance the discarded ATs, vehicles are equipped with backup certificates which replace the ones detected as FPs.




Fig. 7: Adaptive Bloom Filter CRL for different values of False Positive Probability (FPP) and compression ratio (R).

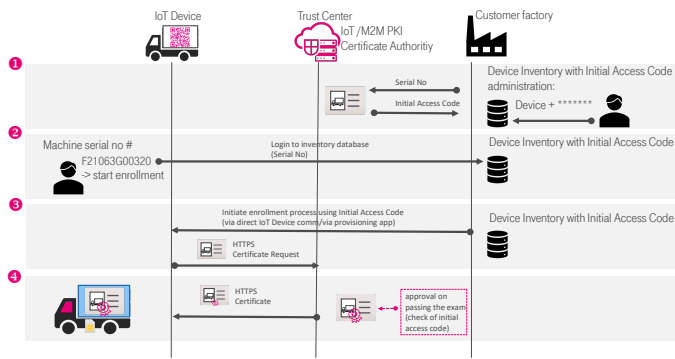


Fig. 8: Certificate provisioning.

IV. ANTI-HACKING DEVICE

Several security features are implemented in the AHD (Fig. 1) to harden the software against foreseen attacks:

- **Secure Boot:** The AHD hardware has fuses (write-once programmable storage locations) that contain the public keys of acceptable boot loader signatures.
- **Secure Firmware Update:** The AHD allows updating the firmware of the Internet (e.g., over the vehicle's communication controller via LTE/5G), accepting only firmware update files that are properly signed by the AHD vendor.
- **Docker technology:** The AHD encapsulates the detection algorithms and some system services into Docker containers.
- **Integrated HSM module:** The AHD contains the HSM in the form of a Telekom Card Operating System (TCOS) security chip. The TCOS module offers secure storage of private key materials and certificates and the ability to run sensitive cryptographic operations securely on chip.

In CARMEL, the HSM is utilized in two different scenarios: First, the PKI auto-enrollment process interacting with the DT-Sec PKI is implemented to provide a certificate for the embedded HSM. Second, it is shown how this can be used on the AHD to run a "threat event signing proxy" Docker container. The signing of threat event reports mitigates so-called "slander attacks", where the backend misbehavior detection system is flooded with false threat detection reports.

1) *HSM Integration and Auto-Enrollment Certificate Authority:* One of the main challenges in managing a distributed fleet of CAVs deployed over a geographically dispersed area is the secure and efficient management of the initial certificate provisioning and update. Expired certificates can cause service outages and device malfunctions and require costly manual maintenance, since the device cannot authenticate against cloud-based services for remote troubleshooting.

To counter these problems, the Auto-enrollment trust center was established, providing a PKI solution for the AHD by offering an interface using EST (Enrollment over Secure Transport standard, IETF RFC 7030 [16]) that runs over HTTPS. This allows devices to contact the trust center even behind firewalls or application-level HTTP proxies that might block other, proprietary non-HTTP-based protocols. The process is presented in Fig. 8. For the initial deployment of the certificates (e.g., in the AHD vendor's factory) the following

steps were implemented for this demonstration: (i) in the trust center a client account is set-up, generating a list of initial access codes that can later be used to request one certificate per code. The AHD vendor receives these access codes and stores them securely in a device inventory database; (ii) When the AHD is initialized, a new private key is created and loaded into the TCOS module; (iii) A certificate request is then created for the corresponding public key; (iv) The configured access code or secret is used to request a certificate from the trust center using the EST protocol; (v) The trust center sends the certificate, that is in turn stored by the AHD.

V. CONCLUSION

CARMEL presents cybersecurity solutions for CAVs addressing most of the possible attacks. An AHD is operated with secure protocols to safeguard a vehicle from any adversary. Further, deep multi-modal data analysis, noise suppression, and deep learning are used to protect autonomous vehicles against camera sensor attacks. For connected vehicles, algorithms are developed for secure V2X message transmission and GPS location spoofing attack detection and mitigation. The project's results show that all solutions provided perform effectively for the attack scenarios considered. In the final stage of CARMEL, all algorithms will be integrated in the AHD for real-time cybersecurity operations.

REFERENCES

- [1] X. Sun *et al.*, "A survey on cyber-security of connected and autonomous vehicles (CAVs)," *IEEE Trans. Intell. Transp. Syst.*, pp. 1–20, 2021.
- [2] M. Singh, "Cybersecurity in vehicular communication," in *Information Security of Intelligent Vehicles Communication*. Springer, 2021.
- [3] J. He *et al.*, "Modulating image restoration with continual levels via adaptive feature modification layers," in *Proc. IEEE CVPR*, 2019.
- [4] L.-C. Chen *et al.*, "Rethinking atrous convolution for semantic image segmentation," *arXiv preprint arXiv:1706.05587[cs.CV]*, 2017.
- [5] A. Geiger *et al.*, "Vision meets robotics: The KITTI dataset," *The Int. J. Rob. Res.*, 32(11):1231–1237, 2013.
- [6] S. Shi *et al.*, "PointRCNN: 3D object proposal generation and detection from point cloud," in *Proc. IEEE CVPR*, 2019.
- [7] P. Perona and J. Malik, "Scale-space and edge detection using anisotropic diffusion," *IEEE Trans. PAMI*, 12(7):629–639, 1990.
- [8] C. Kyrkou *et al.*, "Towards artificial-intelligence-based cybersecurity for robustifying automated driving systems against camera sensor attacks," in *Proc. IEEE ISVLSI*, 2020.
- [9] A. Papachristodoulou *et al.*, "DriveGuard: Robustification of automated driving systems with deep spatio-temporal convolutional autoencoder," in *Proc. IEEE WACV Workshops*, 2021.
- [10] C. Cornelius *et al.*, "Talk proposal: Towards the realistic evaluation of evasion attacks using CARLA," *arXiv preprint arXiv:1904.12622[cs.CV]*, 2019.
- [11] M. Cordts *et al.*, "The Cityscapes dataset for semantic urban scene understanding," in *Proc. IEEE CVPR*, 2016.
- [12] C. Vitale, C. Laoudias *et al.*, "The CARMEL project: A secure architecture for connected and autonomous vehicles," in *Proc. EuCNC*, 2020.
- [13] C. Vitale, N. Piperigkos *et al.*, "CARMEL: Results on a secure architecture for connected and autonomous vehicles detecting GPS spoofing attacks," *EURASIP J. Wirel. Commun. Netw.*, 2021(1):1–28, 2021.
- [14] M. Kamal *et al.*, "GPS location spoofing attack detection for enhancing the security of autonomous vehicles," in *Proc. IEEE VTC-Fall*, 2021.
- [15] ETSI, "Intelligent transport systems (ITS); Security; Trust and privacy management," *Technical Specification 102 941 V1.1.1 (2012-06)*.
- [16] M. Pritikin *et al.*, "Enrollment over secure transport," *RFC 7030*, 2013.