

This is the peer reviewed version of the following article:

Hardware limitations to secure C-ITS: experimental evaluation and solutions / Pollicino, F.; Stabili, D.; Ferretti, L.; Marchetti, M.. - In: IEEE TRANSACTIONS ON VEHICULAR TECHNOLOGY. - ISSN 0018-9545. - 70:12(2021), pp. 12946-12959. [10.1109/TVT.2021.3122333]

Terms of use:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

11/01/2026 08:49

Hardware limitations to secure C-ITS: experimental evaluation and solutions

Francesco Pollicino, Dario Stabili, Luca Ferretti and Mirco Marchetti

University of Modena and Reggio Emilia, Italy

Email: {francesco.pollicino, dario.stabili, luca.ferretti, mirco.marchetti}@unimore.it

Abstract—Cooperative Intelligent Transportation Systems (C-ITS) improve driving experience and safety through secure Vehicular Ad-hoc NETWORKS (VANETs) that satisfy strict security and performance constraints. Relevant standards, such as the IEEE 1609.2, prescribe network-efficient cryptographic protocols to reduce communication latencies through a combination of the Elliptic Curve Qu-Vanstone (ECQV) implicit certificate scheme and the Elliptic Curve Digital Signature Algorithm (ECDSA). However, literature lacks open implementations and performance evaluations for vehicular systems.

This paper assesses the applicability of IEEE 1609.2 and of ECQV and ECDSA schemes to C-ITSs. We release an open implementation of the standard ECQV scheme to benchmark its execution time on automotive-grade hardware. Moreover, we evaluate its performance in real road and traffic scenarios and show that compliance with strict latency requirements defined for C-ITS requires computational resources that are not met by many automotive-grade embedded hardware platforms. As a final contribution, we propose and evaluate novel heuristics to reduce the number of signatures to be verified in real C-ITS scenarios.

Index Terms—VANETs, ECQV, C-ITS, V2V.

I. INTRODUCTION

Cooperative Intelligent Transportation Systems (C-ITS) [33] improve the driving experience by adopting communications among roadside infrastructure, road users and vehicles. Green wave systems, smart traffic lights and smart parking are just a few representative examples of the future smart city features designed to decrease traffic, energy consumption, pollution and optimize commuting, with direct consequences for both driving experience and quality of life. However, designing and implementing these complex systems is a challenging task. As an example, the novel communication networks must support a highly heterogeneous environment, that comprises many vehicles and board manufacturers, and comply with the strict C-ITS constraints in terms of low latency, security, and dynamic network configuration. Standards for C-ITS communications have been developed in the United States, Europe, and Japan [4], [16], [19]. While each standard defines its own data structures and supported security mechanisms, they are all built upon the IEEE 802.11p standard, hence the management of the PHY and MAC layers is the same regardless of the regional standard. In this paper, we investigate security solutions for Vehicular Ad-hoc NETWORKS (VANETs) and focus on the integrity and authenticity guarantees of vehicle communications. To this aim, we consider the security protocols described in the IEEE 1609.2 standard defining the secure message formats for WAVE, policies for the management of

the security certificates, and the supported digital signature and encryption algorithms.

This paper proposes four main contributions to the state of the art. First, we present an implementation of the Elliptic Curve Qu-Vanstone (ECQV) implicit certificate scheme and the Elliptic Curve Digital Signature Algorithm (ECDSA) that is compliant with the IEEE 1609.2 standard and evaluate its deployment on automotive-grade boards. To the best of our knowledge, this is the first open implementation of implicit certificates for resource-constrained devices in terms of computational power and memory. Second, we investigate the feasibility of the implicit certificate scheme in multiple C-ITS scenarios characterized by different latency constraints identified by the National Highway Traffic Safety Administration (NHTSA) for safety-critical communications between vehicles [29]. Third, we analyze the applicability of the automotive boards against realistic scenarios. The realistic scenarios are built on different areas of the city of Modena (Italy) and simulated using real traffic data provided by the municipality. These experiments highlight the limitations imposed by the constraints of several automotive-grade embedded platforms. The last contribution of this paper is the proposal of a prioritization strategy to improve the applicability of automotive-grade boards to real traffic scenarios. The proposed strategy heuristics adopt position-based heuristics based on information included in V2V messages. We assess the effectiveness of the proposal by using the simulations based on the proposed traffic scenarios.

The rest of the paper is organized as follows. Related work is presented in Section II, while the base knowledge required for the understanding of this paper is presented in Section III. Section IV presents our prototype implementation of the ECQV and ECDSA schemes for low-power devices, as well as its applicability to representative automotive-grade boards. The description of the realistic traffic scenarios and the results of our simulations are presented in Section V. Section VI proposes the novel heuristics for the prioritization strategy. Finally, conclusions and future research directions are presented in Section VII.

II. RELATED WORK

Securing C-ITS implies protecting connected vehicles from cyber-attacks. These attacks are aimed to disrupt both communications between the microcontrollers composing the internal vehicular network (intra-vehicular communications) and communications between the vehicles and other external devices

(inter-vehicular communications). Research efforts on intra-vehicular communications propose different prevention [40], [52], detection [13], [31], [53] and reaction [27], [56] approaches against a variety of cyber-attacks to the main communication buses supporting intra-vehicular communications [26], [34]. However, the focus of this paper is on security approaches for inter-vehicular communications that enable C-ITS.

Several security schemes applied to inter-vehicular communications have already been proposed to guarantee the integrity and authenticity of V2V communications [21], [32], [39]. These solutions are the adaptation of secure standard protocols designed for the IT domain and fail to consider the requirements of the automotive environment.

Other works considering the peculiarities of V2V communications focus on privacy [24], [43] and safety requirements [58]. However, these works are only based on theoretical analyses, whereas this paper also leverages experimental evaluations based on realistic simulations that demonstrate the limitations of general purpose automotive-grade boards in supporting V2V communications in a real city and for real traffic scenarios. Many papers already published in the literature that leverage simulations are based on the LuST scenario [12], developed with the mobility data and road structure of the city of Luxembourg. However, the traffic in LuST is extremely sparse and spans a wide area, while our goal is to analyze the applicability of security solutions for V2V communications in dense scenarios including rush hours in which many vehicles are concentrated in a small area. For these reasons, the traffic scenarios analyzed in this paper comprise four different representative areas of the city of Modena (Italy) and different traffic conditions.

Analyses in the context of realistic traffic scenarios have been proposed to compare the effects of obstacle shadowing during transmission of messages using both IEEE 802.11p and ARIB STD-T109 PHY layers [17], but it did not consider the full protocol stack nor the timings required to process received messages and to verify digital signatures. In this paper, we consider the full IEEE 1609 protocol stack and we focus on the computational requirements of the boards to process received messages and verify their authenticity.

Although the authors of [9] analyzed the energy consumption required for ECDSA signature verification in inter-vehicular communications, they did not consider the specifications of the IEEE 1609 standard including implicit certificates. On the other hand, this paper focuses on the timings of the cryptographic operations of both ECQV and ECDSA when deployed on automotive-grade embedded systems, and evaluates their applicability in real traffic scenarios.

The authors of [5] analyze if VANET communications security based on ECQV certificates can satisfy latency constraints that are needed to guarantee safety in case of a car crash. Although the proposed objectives and results are interesting, the proposed testbed is based on a laptop computer and does not evaluate timings by using a proper implementation on hardware that typically characterizes automotive boards. Moreover, the evaluation is based on a simplified scenario based on synthetic traffic conditions composed of a fixed number of

vehicles on a straight road traveling at constant speeds. On the other hand, the evaluation presented in this paper is based on multiple and different traffic conditions characterizing various city zones, with no fixed number of vehicles and variable travel speeds, hence better representing realistic communications between connected vehicles. The evaluation proposed in this paper focuses on real automotive-grade boards and is based on the first open prototype implementation for ECQV and ECDSA. The prototype implementation is optimized for low-power devices and allows to better analyze the performance of the implemented protocols in real-world scenarios.

One of the challenges of implementing cryptographic protocols in embedded hardware is to guarantee adequate performance to verify digital signatures. A possible solution is to use specialized hardware implementations. As an example, the authors of [28] adopted a FPGA that allows to verify up to 250 ECDSA signatures every 100 ms. Further improvements have been proposed in [48] to reduce the power consumption. However, both implementations use elliptic curves at a 81-bit security level (the size of the cryptographic prime groups is about 163-bits), while in this paper we focus on a 128-bit security level (we use the *secp256r1* curve as required by the standards). Another approach is the one proposed by [22], which designs an application-specified integrated circuit implementation for ECDSA verification that can achieve a verification throughput of up to 2700 signatures every 100 ms. Other commercial solutions based on dedicated hardware exist, such as RoadLINK SAF5400 by the NXP Semiconductors [47] that claim a throughput of up to 200 signatures every 100 ms. However, implementations of commercial boards are typically closed source and cannot be easily audited by researchers or updated in case of security vulnerabilities [44]. In this paper we propose an open source implementation of the ECQV/ECDSA stack for general purpose microcontrollers that is accessible and auditable.

This work is also motivated by the lack of shared and established workloads to assess the computational requirements of microcontrollers responsible for V2V communications. As an example, IEEE 1609.2 prescribes that V2V messages are sent every 100 ms. However, the standard does not provide any minimum requirement in terms of the number of messages that a recipient should be able to manage in a given time frame. The time required to handle a single message is dominated by the verification of its attached digital signature, hence constraints related to the number of messages that can be received in a fixed time frame directly translate to limitations to the number of the surrounding vehicles or the hardware requirement of microcontrollers. We argue that the number of the surrounding vehicles is an independent variable and that urban environments lead to a concentration of vehicles (and hence to a number of received messages) far higher than the scenarios already considered in previous studies [5], [12]. We improve related works by proposing analyses based on a wide range of requirements in terms of allowed latency and throughput, that are compliant with standards for safety-critical communications. Moreover, we propose novel heuristics for prioritizing signature validations that allow constrained automotive-grade boards to selectively verify the authenticity

of a subset of more relevant V2V messages.

We observe that although the proposed analyses are based on the established Dedicated Short Range Communication (DSRC) protocol stack (see Section III-A), they are also applicable to emerging Cellular-V2X (C-V2X) due to the similarities of the security services. C-V2X is a promising alternative or complementary vehicular communication technology that is being developed as part of the overall Third Generation Partnership Project (3GPP) process to advance cellular systems from 4G to 5G technologies. It has been introduced in Release 14 of 3GPP [45], [46]. The primary goal of C-V2X technologies is to supersede the IEEE 802.11p/DSRC stack. We observe that C-V2X adopts communication channels with higher bandwidth with regard to DSRC in the context of V2I communications. However, the bandwidth of C-V2X in the context of V2V communications is only slightly better or comparable [6], [54]. Finally, we highlight that C-V2X replaces the PHY and the MAC layers, and leverages all the existing standards in the applications, message/facilities, security services, and the Transport/Networking layers defined by SAE International, ETSI, and IEEE [38]. Since this paper focuses only on the characteristics of these higher layers, the proposed methodologies are also applicable to C-V2X. Due to the comparable performance in V2V communications, quantitative results based on simulations should also apply to C-V2X. We leave further evaluations as future work.

III. BASE KNOWLEDGE

We describe base knowledge regarding the Dedicated Short Range Communication (DSRC) stack (Section III-A) and the IEEE 1609.2 standard (Section III-B)

A. Dedicated Short Range Communication

The Dedicated Short Range Communication (DSRC) stack is the current standard for V2V and V2I wireless communications. The PHY and MAC layers of DSRC are based on IEEE 802.11p [20], the network layer adopts the IEEE 1609 Wireless Access in Vehicular Environments (WAVE) protocol [19], and the transport layer uses the well-known Internet protocols (UDP and TCP). At the application layer, the DSRC stack includes 15 different types of messages [41]. As an example, the *Map* message provides intersection and roadway lane geometry data for one or more locations, and the *TravelerInformationMessage* contains a variety of traffic conditions such as weather, local or regional emergencies and speed warnings. The most used schema is the Basic Safety Message (BSM), which is designed for low latency and safety-relevant applications. The default size of each safety message is 254-byte and includes two types of data. The first part of a BSM message is mandatory and contains the core information about the vehicle (i.e. its size) and its status (i.e. speed, position, and accelerations). The second part is optional and adds a variable number of event-related data, such as notification about the activation of safety-related subsystems within the vehicle (e.g., the activation of the ABS).

B. IEEE 1609.2

The IEEE 1609.2 [18] standard includes (I) the *security services* that must be supported by all the devices, (II) the schema of the messages, and (III) the adopted cryptographic schemes and protocols.

1) *Security services*: IEEE 1609.2 defines two types of security services: the WAVE Internal Security Services and the WAVE Higher Layer Security Services. The services composing the WAVE Internal Security Services are the *Secure Data Service* (SDS) and the *security management*. These two services are used to convert Protocol Data Units (PDUs, unsecured by definition) into Secured Protocol Data Units (SPDUs) and vice-versa.

The security services defined in the WAVE Higher Layer Security Services are the *Certificate Revocation List Verification Entity* (CRLVE) and the *Peer-to-Peer Certificate Distribution Entity* (P2PCDE). The former service is used to validate incoming *Certificate Revocation List* (CRL), forwarding the revocation lists to the *Security Services Management Entity* (SSME); while the latter service is used to enable Peer-to-Peer certificate distribution.

2) *Message schemas*: The IEEE 1609.2 standard supports three types of SPDUs: unsecured, signed, and encrypted. Unsecured SPDUs do not provide any security guarantees. Signed SPDUs guarantee authenticity and non-repudiation and can be used for authorization mechanisms and Basic Safety Messages (which are not encrypted). Encrypted SPDUs guarantee confidentiality. It is possible to combine Signed and Encrypted SPDUs into a single SPDU through encapsulation to guarantee authenticity, non-repudiation and confidentiality. We remark that signature verification is performed for each received message, and in a V2V communication protocol the number of messages received by any given vehicle is greatly larger than the number of messages sent within the network. Due to these characteristics, one of the most computationally expensive operations operated by vehicles is the verification of the digital signatures attached to Signed SPDUs.

3) *Cryptographic schemes*: The IEEE 1609.2 standard requires adopting the Elliptic Curve Digital Signature Algorithm (ECDSA) instantiated with one out of three elliptic curves identified as *NIST-P256*, *BrainpoolP256r1* and *BrainpoolP384r1*. Moreover, IEEE 1609.2 requires a Public Key Infrastructure (PKI) to distribute public keys, where trusted *Certification Authorities* (CA) bind the identity information of communicating parties to their public keys within *certificates*. Two types of certificates are supported: traditional *certificate chains* and *implicit certificates* [11].

A certificate chain includes the public keys of the sender of the message and of all the intermediate CAs, hence the size of the certificate chain is proportional to the number of intermediate CAs. Verifying the certificate chain requires to verify the digital signatures attached by all the intermediate CAs up to the Root Certificate.

For implicit certificates, the IEEE 1609.2 standard specifies usage of the Elliptic Curve Qu-Vanstone (ECQV) scheme. Intuitively, an ECQV certificate does not include the public key of the sender but allows a recipient to re-compute it by using the certificate of the sender and the public key of the CA, thus

Algorithm 1 Certificate Sign Request

```

1: function CSRGEN( $ID$ )
2:    $sk, PK = \text{KEYGEN}()$ 
3:   return  $CSR = (ID \parallel PK)$ 

```

Algorithm 2 Certificate generation

```

1: function CRTGEN( $sk_{CA}, PK_{CA}, CSR$ )
2:    $ID, PK \leftarrow CSR$ 
3:   do
4:      $sk', PK' = \text{KEYGEN}()$ 
5:      $P = PK + PK'$ 
6:      $CRT := (P, ID)$ 
7:      $x = \mathcal{H}_\ell(P \parallel ID)$ 
8:     while  $(x \cdot P + PK_{CA}) = \mathcal{O}$ 
9:        $r = (x \cdot sk' + sk_{CA}) \bmod \ell$ 
10:    return  $(CRT, r)$ 

```

saving network usage. Network savings of implicit certificates can be considered negligible in most Web scenarios because communications are mostly high-bandwidth and can tolerate latency. Moreover, certificates are only used once during the secure channels handshakes. However, vehicular networks are characterized by datagram-oriented communications that include small data packets, each including a digital signature and a certificate (especially safety-critical packets, see Section IV). Moreover, vehicular networks, especially vehicle-to-vehicle communications, are deployed on low-rate wireless networks and have tighter latency requirements. Hence, ECQV is the best fit for this scenario.

The ECQV operations framework includes four routines: *certificate sign request*, *certificate generation*, *certificate reception* and *public key extraction*. These operations, combined with the *signature* and *verification* procedures of ECDSA, allow guaranteeing message integrity and authenticity by using trusted CAs as trusted anchors. In the following, we describe the flow of the operations from the generation of an ECQV certificate to the signature verification by referring to Figures 1.

The actors are the *requester client*, the *certification authority* (CA) and the *receiver client*. The requester client generates a *Certificate Sign Request* (CSR) via the *CSR generation* function (*CSRGen*). The *CSRGen* function requires the *ID* of the entity requesting the certificate and produces an output composed by the *CSR* (that includes *ID* and *PK*, which is the intermediate public key of the requester) and the intermediate private key *sk* (Algorithm 1).

The *CSR* is sent to the CA that produces the corresponding certificate *CRT* and the private key contribution *r* by using the *CRT generation* function (*CRTGen*, Algorithm 2).

After having received the *CRT* from the CA, the requester validates it with the *CRT reception* function (*CRTReception*). Upon verification, the requester generates the final key pair (sk_U, PK_U) by using *CRT* and *r* (Algorithm 3).

The private key sk_U is used to generate the signature σ of the messages *m* with the ECDSA *signing* function (*Sign*). The client then broadcasts the message *m*, the signature σ and its own certificate *CRT* (Algorithm 4).

Each receiver client uses the *public key extraction* function (*Extract*) to extract the public key of the sender client from the *CRT* (Algorithm 5), which is later used for the verification of

Algorithm 3 Certificate reception

```

1: function CRTRECEPTION( $PK_{CA}, CRT, r$ )
2:    $P, ID \leftarrow CRT$ 
3:    $x = \mathcal{H}_\ell(P \parallel ID)$ 
4:    $PK_U = sk_U \cdot B$ 
5:    $PK'_U = \text{EXTRACT}(PK_{CA}, CRT)$ 
6:   if  $PK_U \neq PK'_U$  then
7:     return InvalidCertificate
8:   return  $(sk_U, PK_U)$ 

```

Algorithm 4 Message Sign

```

1: function SIGN( $m$ )
2:    $h = \mathcal{H}(m)$ 
3:    $z = h[:L_n]$ 
4:   do
5:     do
6:        $k \xleftarrow{\$} \{0, 1\}^{n-1}$ 
7:        $(x_1, y_1) = k \cdot B$ 
8:        $r = x_1 \bmod n$ 
9:       while  $r \neq 0$ 
10:       $s = k^{-1}(z \cdot sk_U) \bmod n$ 
11:      while  $s \neq 0$ 
12:    return  $\sigma = (r, s)$ 

```

the signature of the message with the *verify* function (*Verify*, Algorithm 6).

IV. PROTOTYPE IMPLEMENTATION AND MICROBENCHMARKS

This section describes the microbenchmarks of ECDSA and ECQV on representative automotive-grade boards. Results have been obtained by using our prototype implementation that supports the NIST P256 curve (*secp256r1*) as required by the IEEE 1609.2 standard. The implementation depends on the *microECC* [25] library, that provides efficient elliptic curve operations on constrained devices thanks to optimized ASM code for ARM platforms.

The implementation complies with the following best practices:

- cryptographic operations comprising secret information adopt time-constant code to avoid timing side-channels (e.g., no conditional branches, time-constant scalar point multiplication [7]);
- no variables are allocated by using dynamic memory allocation;
- random numbers are generated with a hardware TRNG when available, otherwise we use the software pseudo-random number generator provided by the board library;
- although the implementation supports multiple elliptic curves, the code selects only the required curve at compile time to avoid wasting storage by including useless code;
- elliptic curve points are always transferred by using the compressed format defined in the standard SEC 1 [10] to reduce the network overhead, while cryptographic operations are computed by using the uncompressed representation (x, y) .

The source code of the prototype implementation is open for reuse and inspection by researchers and industry practitioners¹.

¹<https://weblab.ing.unimore.it/resources/uECQV.zip>

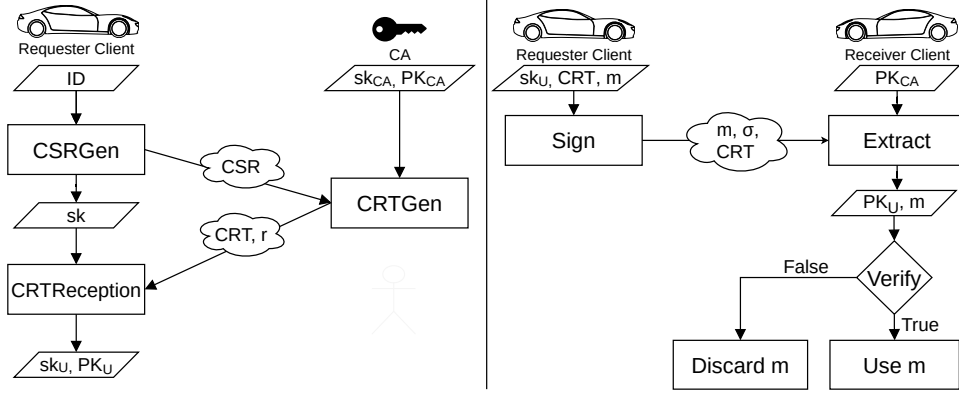


Fig. 1: Operation flow for the generation of the ECQV certificate and ECDSA signature based on ECQV certificates

Algorithm 5 Public key extraction

```

1: function EXTRACT( $PK_{CA}, CRT$ )
2:    $P, ID \leftarrow CRT$ 
3:    $x = \mathcal{H}_\ell(P \parallel ID)$ 
4:    $PK_U = x \cdot P + PK_{CA}$ 
5:   if  $PK_U = \mathcal{O}$  then
6:     return InvalidCertificate
7:   return  $PK_U$ 

```

Algorithm 6 Signature verification

```

1: function VERIFY( $m, \sigma, PK_U$ )
2:    $r, s \leftarrow \sigma$ 
3:    $h = \mathcal{H}(m)$ 
4:    $z = h[: L_n]$ 
5:    $w = s^{-1} \bmod n$ 
6:    $u_1 = z \cdot w \bmod n$ 
7:    $u_2 = r \cdot w \bmod n$ 
8:    $(x_1, x_2) = u_1 \cdot B + u_2 \cdot PK_U$ 
9:   if  $r \equiv x_1 \bmod n$  then
10:    return ValidSignature
11:  else
12:    return InvalidSignature

```

The testbed of the evaluation included the following automotive-grade boards:

- **A72**: based on 64-bit ARM Cortex-A72 quad-core CPU operating at 1.5GHz.
- **A53**: based on 64-bit ARM Cortex-A53 quad-core CPU operating at 1.2GHz.
- **ARM11**: based on ARM1176JZFS single-core CPU operating at 700MHz.
- **M4**: based on 32-bit Cortex-M4 CPU operating at 80MHz, with 1MB of flash memory and 128KB of RAM memory.
- **M3**: based on 32-bit ARM Cortex-M3 CPU operating at 84MHz, with 512KB of flash memory and 96KB of RAM memory.

These boards are similar to automotive microcontrollers produced by *STMicroelectronics* [1]. For completeness, we also include results of an **x86_64** architecture: a modern laptop

	x86_64	A72	A53	ARM11	M4	M3
KeyGen	0.32	1.33	2.82	14.67	109.67	147
CsrGen	0.33	1.37	2.79	14.21	109.73	148
CrtGen	0.62	2.85	5.83	30.02	233.45	317
CRTReception	0.60	2.79	5.62	29.29	231.47	313
Extract	0.34	1.44	3.03	15.38	120.13	162
Sign	0.32	1.44	3.05	15.38	118.00	161
Verify	0.35	1.58	3.36	16.57	133.47	182

TABLE I: Timings of the operations on the *secp256r1* curve [ms]

with an Intel Core i7-9750H.

We observe that the implementation is single-threaded. This is a typical design choice for most cryptographic algorithms that cannot be easily parallelized without the risk to introduce security vulnerabilities. This is not a limitation since low-power embedded devices are based on single-core architectures, and more powerful architectures can verify concurrently multiple signatures on different cores.

The columns of Table I represent the platforms used for the evaluation, while the rows represent the cryptographic operations of ECDSA and ECQV: key generation (*KeyGen*), certificate request generation (*CsrGen*), certificate generation (*CrtGen*), certificate reception (*CRTReception*), extraction (*Extract*), signature (*Sign*) and verification (*Verify*). All results are expressed in milliseconds. We observe that the most powerful ARM architectures (A72 and A53) can execute all operations in a few milliseconds, while cheaper ARM architectures (ARM11) are an order of magnitude slower. Moreover, ultra-low power architectures (M4 and M3) are two orders of magnitude slower, thus requiring a few hundred milliseconds to execute each operation. We observe that although the certificate generation operation (*CrtGen*) can usually be deployed on a dedicated server machine, the implemented library would also be able to generate novel certificates with low-power devices.

We now analyze the applicability of our implementation deployed on automotive-grade boards for securing V2V communications. To this aim, we consider communication requirements defined by the National Highway Traffic Safety Administration (NHTSA) [30], which considers multiple ve-

hicle communication scenarios and defines the constraints that must be satisfied to guarantee safety. Among these constraints, we are interested in the *allowable latency*, which is the maximum latency allowed for end-to-end communication and processing of data, and the *communication range*, which is the maximum distance for which the communication is of interest for the receiver. The report identifies 8 high-priority and safety-critical scenarios, and for each of them defines the associated allowable latency and communication range:

- Pre-Crash Sensing: 20ms and $\sim 50m$;
- Traffic Signal Violation Warning: 100ms and $\sim 250m$;
- Curve Speed Warning: 1s and $\sim 200m$;
- Emergency Electronic Brake Light: 100ms and $\sim 300m$;
- Cooperative Forward Collision Warning: 100ms and $\sim 150m$;
- Left Turn Assistant: 100ms and $\sim 300m$;
- Lane changing Warning: 100ms and $\sim 150m$;
- Stop Sign Movement Assistance: 100ms and $\sim 300m$;

We evaluate the timings required to compute all the due cryptographic operations for *sending* and *receiving* an authenticated message on the automotive-grade boards. We define three device roles: **full sender (FS)**, **direct sender (DS)**, and **receiver (R)**. The device operating as *FS* generates a novel certificate for each communication, as used in the Butterfly protocol [49], while a *DS* device requests a valid certificate offline and uses the same certificate for multiple communications. The operations required by a *FS* are *CsrGen*, *CRTReception*, and *Sign*, while the only required operation for a *DS* is *Sign*. The *R* devices must extract the public key from the certificate and verify the validity of the signature, thus requiring the execution of *Extract* and *Verify*. Table II shows the timings required for the cryptographic operations executed by each role.

	Full Sender	Direct Sender	Receiver
x86_64	1.25	0.32	0.69
A72	5.61	1.44	3.02
A53	11.46	3.05	6.39
ARM11	58.88	15.38	31.94
M4	459.20	118.00	253.60
M3	622.00	161.00	344.00

TABLE II: Timings for the cryptographic operations of the roles of the device on different platforms [ms]

Table III summarizes the applicability of the considered automotive-grade boards for secure V2V communications in the scenarios outlined by the NHTSA. The rows of Tables III represent the three most strict allowable latencies (20ms, 100ms, and 1000ms), while the columns represent the different platforms and roles. The *FS* and *DS* sub-columns are used to present the applicability of each board as a sender device with regard to send rates, and the *R* sub-column shows the maximum number of messages that the platform can validate as a receiver within the allowed latency. We remark that the actual number of received messages depends on the number of vehicles within the communication range, hence it might exceed the validation capability of a board. We investigate the capabilities of the boards in real traffic scenarios in Section V.

In the following we summarize the main results. With an allowable latency of 20 ms, it is possible to deploy the *x86_64*, the *A72* and the *A53* boards in the *FS* role, while in the *DS* role it is possible to also deploy the *ARM11* board. When using the boards as *R* role, only the *x86_64*, *A72* and *A53* systems are able to verify the signatures of 29, 6 and 3 incoming messages within the allowable latency. In safety-critical applications with allowable latency of 100 ms, it is possible to use the *x86_64*, *A72*, *A53*, and *ARM11* boards in all roles, while the *M4* nor *M3* boards cannot be deployed for any role. The maximum number of messages that the boards are able to verify within the 100ms allowable latency in the *R* role are 144, 33, 15, and 3 for the *x86_64*, *A72*, *A53*, and *ARM11* boards, respectively. Moreover, we highlight that both the *M4* and *M3* boards are not able to send any message nor to verify any signature within allowable latencies of 20 and 100 ms due to their limited computational power. Finally, with an allowable latency of 1000 ms, it is possible to deploy all the boards for all the roles, but the *M4* and *M3* board can only verify 3 and 2 signatures respectively. Finally, the *M4* and *M3* boards can still be deployed for non-safety-critical scenarios with higher allowable latencies, as demonstrated in [36].

V. SIMULATION IN REALISTIC TRAFFIC SCENARIOS

We assess the applicability of the automotive-grade boards to realistic traffic scenarios in simulated environments. We describe the characteristics of the traffic scenarios in Section V-A and the results of the simulations in Section V-B, and we analyze the applicability of the boards in Section V-C. In Section V-D we adopt the results of the simulations to assess the advantages of implicit certificates over certificate chains in V2V networks.

A. Realistic traffic scenarios

To recreate realistic simulations, we use data regarding the city of Modena for multiple road and traffic scenarios. We characterize each scenario by using multiple parameters, including the average number of vehicles per hour, the distribution of the road vehicles per hour, and the average number of vehicles in a kilometer for different roads. For clarity of exposure, we describe the traffic scenarios by the average number of vehicles per kilometer (traffic index Ti) during rush hour in descending order.

- **Roundabout (Grapes)** [$Ti = 27.33$]: represents the roundabout located in the southwest part of the city, nearby the Department of Engineering “Enzo Ferrari” of the University of Modena and Reggio Emilia. The Grapes scenario connects 2 segments of a 6-lanes expressway (two road segments in each direction that become three in the proximity of the roundabout) with one segment of urban road and one segment of extra-urban road, for a total of 12 road kilometers. The number of simulated vehicles during normal traffic and rush hour are 167 and 328 respectively.
- **Highway (H-Way)** [$Ti = 14.41$]: represents the highway junction located between the “Campogalliano” and the

	x86_64			A72			A53			ARM11			M4			M3		
	FS	DS	R	FS	DS	R	FS	DS	R	FS	DS	R	FS	DS	R	FS	DS	R
20 ms	✓	✓	29	✓	✓	6	✓	✓	3	×	✓	×	×	×	×	×	×	×
100 ms	✓	✓	144	✓	✓	33	✓	✓	15	✓	✓	3	×	×	×	×	×	×
1000 ms	✓	✓	1449	✓	✓	331	✓	✓	156	✓	✓	31	✓	✓	3	✓	✓	2

TABLE III: Analysis of the constraints of the boards. Applicability for both sender roles and maximum number of received messages for the receiver role

“Modena Nord” highway toll booths. The Highway scenario includes two highway segments connected with a T-junction. The first highway segment is a 4-lanes highway (two lanes for each direction), while the second highway is a 6-lanes highway (three lanes for each direction). Vehicles can only enter or exit the highway through entry or exit lanes. The total road length for this scenario is 54 kilometers. The number of simulated vehicles during normal traffic and rush hour are 131 and 778 respectively.

- **University campus (Campus)** [$T_i = 6.58$]: represents the area surrounding the Department of Engineering “Enzo Ferrari”, located in the southwest part of the city, near the main city hospital. This scenario is composed of multiple urban roads connected through plain intersections. The Campus scenario includes a residential area, a shopping center, the hospital and the engineering campus, for a total of 52 kilometers of roads. The number of simulated vehicles during normal traffic and rush hour are 159 and 342 respectively.
- **Modena Automotive Smart Area (MASA)** [$T_i = 4.24$]: represents the Modena Automotive Smart Area [2], located behind the main railway station of Modena. The MASA scenario is a roughly rectangular residential area, connecting 4 perimeter urban roads through roundabouts, for a total of 34 kilometers. The number of simulated vehicles during normal traffic and rush hour are 77 and 144 respectively.

B. Simulations results

We simulated the considered scenarios by using VEINS [51], that is an open source framework for vehicular network simulations based on OMNeT++ [57] for the simulation of networks, and on SUMO [23] (Simulation of Urban MObility) for the simulation of the road traffic. In our simulation, we adopted the IEEE 802.11p, IEEE 1609.4 DSRC/WAVE [15], Physical Layer [8], Obstacle Shadowing [50] and Antenna Patterns [14] modules. The first two modules (IEEE 802.11p and IEEE 1609.4 DSRC) extend the VEINS capabilities by enabling the DSRC/WAVE stack, including Quality-of-Service channel access, the Wave Short Message (WSM) management and periodic beaconing of BSMs. For details on the protocols please refer to Section III. The other modules simulate the propagation and attenuation of the wireless signals to recreate proper signal coverage of messages in urban environments.

The simulations aim to provide meaningful insights about the actual number of messages received by all the vehicles in a particular road scenario, hence we recreated the four scenarios

based on the city of Modena. Each scenario is used twice in our simulations, once to simulate the communication between vehicles during the rush hour as a high traffic condition (*Rush*), and once to simulate the communications during normal traffic conditions (*Normal*).

We recreated the scenarios within the simulated environment by using the road and polygonal maps of the buildings available at Open Street Map [55]. The vehicles simulated in our scenarios are programmed to send beacon messages with a frequency of $10Hz$ (i.e. one message each 100 milliseconds), as recommended by the SAE J2945-201712 [42] standard. The four scenarios were simulated for five minutes for both *Rush* and *Normal* traffic conditions. Results are summarized in Table IV, including the total number of vehicles and the total number of messages exchanged in the simulated VANETs.

In the following, we analyze the results of the simulations that are of interest with regard to the specifications included in the NHTSA report for 20 ms and 100 ms safety-critical messages respectively (see Section IV).

1) *20 ms allowable latency*: The results of the analysis on the number of messages received by the vehicles in the realistic scenarios with allowable latency of 20 ms and a distance of 50 m are presented in Table V. The columns represent the simulated areas with different traffic conditions and the rows represent the maximum, minimum and average (rounded to the lowest integer) number of messages received by a single vehicle within the allowable latency. The table shows that the maximum number of messages received by the vehicles ranges from 2 to 5 and 5 to 12 in *Normal* and *Rush* traffic conditions. In all scenarios and conditions the minimum number of messages is equal to 1.

2) *100 ms allowable latency*: The results of the analysis on the number of messages received by the vehicles in the realistic scenarios with allowable latency of 100 ms and within all distances are presented in Table VI. The table shows that the maximum number of messages received by the vehicles ranges from 27 to 83 and 61 to 175 in *Normal* and *Rush* traffic conditions. In all scenarios and conditions the minimum number of messages is equal to 1.

C. Applicability of the automotive-grade boards

To determine the applicability of the automotive-grade boards we analytically evaluate the time required by each board to verify the authenticity of the messages received by a vehicle in each traffic scenario. To this aim, we multiply the maximum number of messages received by a vehicle by the time required by the boards to extract the implicit certificate and verify the digital signature (see the *Receiver* timings in Table II of Section IV).

	Highway		MASA		Campus		Grapes	
	<i>Rush</i>	<i>Normal</i>	<i>Rush</i>	<i>Normal</i>	<i>Rush</i>	<i>Normal</i>	<i>Rush</i>	<i>Normal</i>
# cars	778	131	144	77	342	159	328	167
# messages	21,227,443	830,210	1,106,882	262,669	4,645,905	980,121	15,462,204	3,427,977

TABLE IV: Overview of the simulation results on the different scenarios in both rush hour and normal traffic conditions

	Grapes		Highway		Campus		MASA	
	<i>Rush</i>	<i>Normal</i>	<i>Rush</i>	<i>Normal</i>	<i>Rush</i>	<i>Normal</i>	<i>Rush</i>	<i>Normal</i>
# max of msgs	12	5	6	2	5	3	5	2
# min of msgs	1	1	1	1	1	1	1	1
# avg of msgs	4	2	2	1	2	1	2	1

TABLE V: Results of the simulations with allowable latency of 20 ms and a distance of 50 m

	Grapes		Highway		Campus		MASA	
	<i>Rush</i>	<i>Normal</i>	<i>Rush</i>	<i>Normal</i>	<i>Rush</i>	<i>Normal</i>	<i>Rush</i>	<i>Normal</i>
# max of msgs	175	83	89	38	69	34	61	27
# min of msgs	1	1	1	1	1	1	1	1
# avg of msgs	85	42	33	12	25	12	20	9

TABLE VI: Results of the simulations with a maximum allowable latency of 100ms (all distances)

Table VII shows the results for 20 ms allowable latency. The rows include the boards and the columns include the traffic scenarios. We highlight in gray the cells of the table that refer to the boards that do not satisfy the latency requirement. The results show that no automotive-grade board can verify all the received messages within 20 ms in the worst case traffic scenario (*Rush*) by using a single core. Only the *A72* board is able to satisfy all *Normal* traffic scenarios. Moreover, the reference *x86_64* architecture can verify the signatures of all the received messages in all scenarios.

Table VIII shows the results for 100 ms allowable latency. We omit the boards with timings greater than 4000 ms. The results show that no automotive-grade board can verify all the received messages within 20 ms in the worst case traffic scenario (*Rush*) by using a single core. Only the *A72* board is able to satisfy all *Normal* traffic scenarios. Moreover, the *x86_64* reference architecture can verify the signatures of all the received messages in all scenarios. The results show that no automotive-grade board can verify all the received messages in both traffic scenarios. Only the *x86_64* reference architecture is able to verify all the signatures in all scenarios, except for the *Grapes* scenario in the *Rush* traffic condition.

The results highlight that even modern vehicles will fail in verifying all safety-related messages in real urban scenarios. As a result, a few safety messages will either be completely ignored, or accepted without verification. Both solutions are unacceptable since they expose drivers and road users to safety risks and cyberattacks. However, we also remark that improved results could be achieved if we consider the concurrent verification of digital signatures by multiple cores of a single board. To this aim, it is possible to operate an approximated analytical evaluation by dividing the timings reported in Tables VII and VIII by the number of cores of the considered board.

D. Comparison with certificate chains

We compare the effectiveness of implicit certificates and explicit certificate chains.

We analyze the sizes of implicit and explicit certificates for ECDSA instantiated over a 256 bit curve (*secp256r1*), that is the recommendation of IEEE 1609.2 standards for vehicular communications. Each BSM message transmitted by using the DSRC stack over V2V networks is authenticated with implicit certificates and includes 64 bytes for the ECDSA digital signature, 93–180 bytes for the implicit certificate (the size depends on the attached metadata and on padding), and a minimum of 254 bytes for the payload (see Section III-A). Hence, the average size of a full message ranges from 410 to 500 bytes.

Since the size of an ECDSA digital signature is 64 bytes, a message authenticated with explicit certificates (with a certificate chain with no intermediate CAs) is about 64 bytes longer (the actual size may vary depending on padding). The average size of a full message ranges from 474 to 565 bytes. By considering these estimations, using ECQV implicit certificates reduces network usage ranging from 10% to 15%.

The NHTSA report [29] defines variable data rates and distance ranges for V2V communications, which nominally range from 3 Mb/s to 27 Mb/s and from 300 m to 1000 m. The maximum data rate suggested by the NHTSA report for the control channel related to safety applications is 6 Mb/s, while the 27 Mb/s data rate should only be used for bursts. Increasing the data rate from 6 Mb/s to 27 Mb/s can cause higher packet losses and reduced communication ranges. By considering a bandwidth of 6 Mb/s, a latency of 100 ms, an average packet size of 460 bytes for ECQV implicit certificates and of 520 bytes for explicit certificates, the maximum number of messages supported by the network is 163 and 144 messages respectively.

The validation of an explicit certificate requires two digital signature verification operations: the first to validate the signature of the explicit certificate, and the second to validate the signature of the BSM. On the other hand, the ECQV implicit certificate scheme requires only one signature verification, that

	Grapes		Highway		Campus		MASA	
	<i>Rush</i>	<i>Normal</i>	<i>Rush</i>	<i>Normal</i>	<i>Rush</i>	<i>Normal</i>	<i>Rush</i>	<i>Normal</i>
x86_64	8	4	4	2	4	2	4	2
A72	36	15	18	6	15	9	15	6
A53	77	32	38	13	32	19	32	13
ARM11	383	159	191	64	160	96	160	64
m4	3043	1268	1521	507	1268	761	1268	507
m3	4128	1720	2064	688	1720	1032	1720	688

TABLE VII: Time [*ms*] required by each board to verify all the signatures with a maximum allowable latency of 20ms using a single core.

	Grapes		Highway		Campus		MASA	
	<i>Rush</i>	<i>Normal</i>	<i>Rush</i>	<i>Normal</i>	<i>Rush</i>	<i>Normal</i>	<i>Rush</i>	<i>Normal</i>
x86_64	121	57	61	26	48	23	42	19
A72	529	251	269	115	208	103	184	82
A53	1118	530	569	243	441	217	390	173
ARM11	5590	2651	2843	1214	2214	1086	1948	862

TABLE VIII: Time [*ms*] required by each board to verify all the signatures with a maximum allowable latency of 100 milliseconds using a single core

is, to validate the signature of the received BSM. However, the receiver must first extract the sender's public key by using the ECQV *public key extraction* operation. An analytic estimation based on the timings of the operations obtained by using the proposed implementation on the automotive-grade boards (see Table I) let us conclude that implicit certificates are comparable to explicit certificates in terms of computational costs.

Our analysis confirms that the use of implicit certificates as indicated by the IEEE 1609.2 standard is advantageous to reduce network overhead (and to increase the throughput of wireless vehicular networks).

VI. PRIORITIZATION STRATEGY BASED ON SENDERS POSITIONS

The results of the evaluation of the applicability in realistic scenarios (Section V) show that the computational constraints of the boards may prevent verification of some messages within safety-critical timings. In this section, we propose a scheduling strategy to identify a subset of messages that should be verified with a higher priority. The strategy adopts heuristics based on the relative positions of the vehicles. Upon reception of a message, the receiving vehicle extracts the geographical coordinates of the sender. If the coordinates are within a certain area surrounding the receiver, they are verified and analyzed before the other messages. Two rationales guide this approach. First, messages sent by vehicles that are relatively far from the receiver do not convey meaningful information for immediate reaction, hence it is safe to ignore their content, thus decreasing the number of signature validations. Second, messages sent by vehicles that are equally distanced from the receiver may have different importance depending on their relative positions. In the following, we analyze the impact of the proposed strategy in the scenarios considered in Section V by considering three different areas: circular, elliptical, and forward-facing. We consider the same scenarios of the simulation presented in Section V-A and analyze the number of messages sent within the areas. For each area and

for each scenario, we evaluate the average time required by the boards to analyze all the prioritized messages.

A. Circular area heuristic

We consider a distance-only approach where the receiving vehicle controls whether the sender vehicles are within a circle with radius r . The values of the radius r in our analysis are 150, 250 and 300 meters, that are the distance requirements proposed by the NHTSA report for the 100ms latency. As a comparison, we also report results for a radius of 1000 meters, that is the typical maximum communication range of DSRC. Table IX shows the average number of messages sent within the circle of radius r . The columns represent the scenarios used in our evaluation and the rows represent the values of the radius r . The table shows that on average the 26%, 59% and 66% of the messages are sent by vehicles within a maximum distance of 150, 250 and 300 meters. All messages are within the maximum distance of 1000 meters.

Table X shows the timings required for each board to verify messages authenticity within the circular area within 100ms. Cells of the table with a gray background identify the scenarios that cannot be deployed due to high timings. We observe that it is possible to deploy the A72 board to satisfy almost all scenarios with a radius of 150 meters (the only exception is Grapes with heavy traffic conditions). We observe that the A53 platform is not able to satisfy almost all scenarios, and the ARM11 is not able to satisfy any scenarios. For this reason, we omit the m4 and m3 boards that have even worse performance.

B. Elliptical area heuristic

The elliptical area considers the surroundings of the vehicle by building an ellipse centered on the vehicle. This type of area gives more priority to vehicles that are behind or in front of the receiver, and reduces the priority of side vehicles. This heuristic is useful in particular in all those scenarios that do not require information from the side vehicles, as

		Grapes		Highway		Campus		MASA	
		Rush	Normal	Rush	Normal	Rush	Normal	Rush	Normal
100 ms	< 150m	54	23	21	8	18	9	19	8
	< 250m	113	52	48	20	40	20	39	18
	< 300m	125	59	55	22	47	23	44	20
	≤ 1000m	175	83	89	38	69	34	61	27

TABLE IX: Number of messages in different ranges according to the distance of the sender vehicle using the circular area

	<i>r</i>	Grapes		Highway		Campus		MASA	
		Rush	Normal	Rush	Normal	Rush	Normal	Rush	Normal
x86_64	150m	37	16	14	6	12	6	13	6
	250m	78	36	33	14	28	14	27	12
	300m	86	41	38	15	32	16	30	14
	1000m	121	57	61	26	48	23	42	19
A72	150m	163	69	63	24	54	27	57	24
	250m	341	157	145	60	121	60	118	54
	300m	378	178	166	66	142	69	133	60
	1000m	529	251	269	115	208	103	184	82
A53	150m	345	147	134	51	115	58	121	51
	250m	722	332	307	128	256	128	249	115
	300m	799	377	351	141	300	147	281	128
	1000m	1118	530	569	243	441	217	390	173
ARM11	150m	1725	735	671	256	575	287	607	256
	250m	3609	1661	1533	639	1278	639	1246	575
	300m	3993	1884	1757	703	1501	735	1405	639
	1000m	5590	2651	2843	1214	2204	1086	1948	862

TABLE X: Timing [*ms*] of the automotive-grade boards with circular heuristic for different radius values with a latency of 100*ms*.

the Pre-Crash Sensing in one-way road (e.g. Highway), Lane Change Warning that provides a warning to the driver if an intended lane change may cause a crash with nearby vehicles, Cooperative forward collision warning that is designed to aid the driver in avoiding or mitigating collisions with the rear-end of vehicles, Emergency Electronic Brake light application that sends a message to other vehicles following behind, and Left Turn Assistant with regard to vehicles coming from the opposite direction.

We model the elliptical area by using the equation $\frac{x^2}{a^2} + \frac{y^2}{b^2} = 1$, where *a* denotes the side distance and *b* denotes the distance from the front and rear vehicles. We consider a single value for the parameter *a* of 25 meters. Instead, we consider multiple values for the parameter *b* of 150, 250, 300 and 1000 meters in case of 100*ms* allowed latency, and a single value of *b* of 50 meters for the 20*ms* allowed latency. As a result, we consider a total of five different ellipses.

Table XI shows the average number of messages sent within the different ellipses. The columns represent the different scenarios and the rows represent the values of the two parameters *a* and *b*. The table shows that the number of messages received within 20*ms* allowed latency ranges from 6% to 10% for the different scenarios, and on average the 17%, 18% and 20% of messages are received within 100*ms* for values of *b* equal to 150*m*, 250*m* and 300*m*, respectively.

Table XII shows the timings required for each board to verify messages authenticity within 20*ms* and 100*ms* allowed latencies. We observe that it is possible to deploy the A72 and A53 boards to satisfy almost all scenarios within 20*ms*

allowed latency (the only exception is Grapes with heavy traffic conditions). Instead, ARM11 and less powerful boards are not able to satisfy any scenario within 20*ms* allowed latency. For 100*ms* allowed latency, the Grapes scenario with heavy traffic conditions still cannot be satisfied by any board (with the exception of the reference x86_64 architecture) and the Grapes scenario with normal traffic conditions can be satisfied only by the A72 board. The ARM11 board can satisfy only the Highway with normal traffic conditions. Less performing boards (that are omitted from the table) cannot satisfy any scenarios.

C. Forward-facing area heuristic

The forward-facing area considers a circular arc in front of the vehicle. This type of area helps to further improve all the scenarios in which priority should be given to front vehicles like the Pre-Crash Sensing, Left Turn Assistant, and Stop Sign Movement.

We denote the angular aperture of the circular arc as α , for which we consider values of 90, 135, and 180 degrees. Moreover, we consider values for the radius of the arc equal to 50 and 300 meters for the 20*ms* and 100*ms* allowed latencies, respectively. These values are compliant with recommendations by the NHTSA for the considered scenarios.

Table XIII shows the average number of messages sent within the different forward-facing areas. The columns represent the different scenarios and the rows represent the values of parameter α . The table shows that the average number of received messages are 3%, 4% and 6% for 20*ms* allowed

<i>a</i>	<i>b</i>	Grapes		Highway		Campus		MASA	
		<i>Rush</i>	<i>Normal</i>	<i>Rush</i>	<i>Normal</i>	<i>Rush</i>	<i>Normal</i>	<i>Rush</i>	<i>Normal</i>
<i>20 ms</i>									
<i>25m</i>	<i>50m</i>	8	3	3	1	3	2	3	2
<i>100 ms</i>									
<i>25m</i>	<i>150m</i>	40	14	7	2	9	4	11	4
<i>25m</i>	<i>250m</i>	43	17	8	3	10	5	13	6
<i>25m</i>	<i>300m</i>	45	18	8	3	11	5	14	6
<i>25m</i>	<i>1000m</i>	51	22	9	3	12	6	17	7

TABLE XI: Number of messages originated within an elliptical area centered in the receiving vehicle

			Grapes		Highway		Campus		MASA	
<i>a</i>	<i>b</i>		<i>Rush</i>	<i>Normal</i>	<i>Rush</i>	<i>Normal</i>	<i>Rush</i>	<i>Normal</i>	<i>Rush</i>	<i>Normal</i>
20 ms										
x86_64	25m	50m	6	2	2	1	2	1	2	1
A72	25m	50m	24	9	9	3	9	6	9	6
A53	25m	50m	51	19	19	6	19	13	19	13
ARM11	25m	50m	256	96	96	32	96	64	96	64
100 ms										
x86_64	25m	150m	28	10	5	1	6	3	8	3
	25m	250m	30	12	6	2	7	3	9	4
	25m	300m	31	12	6	2	8	3	10	4
	25m	1000m	35	15	6	2	8	4	12	5
A72	25m	150m	121	42	21	6	27	12	33	12
	25m	250m	130	51	24	9	30	15	39	18
	25m	300m	136	54	24	9	33	15	42	18
	25m	1000m	154	66	27	9	36	18	51	21
A53	25m	150m	256	89	45	13	58	26	70	26
	25m	250m	275	109	51	19	64	32	83	38
	25m	300m	288	115	51	19	70	32	89	38
	25m	1000m	326	141	58	19	77	38	109	45
ARM11	25m	150m	1278	447	224	64	287	128	351	128
	25m	250m	1373	543	256	96	319	160	415	192
	25m	300m	1437	575	256	96	351	160	447	192
	25m	1000m	1629	703	278	96	383	192	543	224

TABLE XII: Timing [*ms*] of the automotive-grade boards with an elliptical heuristic for different parameters.

latency, and 6%, 8% and 36% for 100ms allowed latency, for 90, 135 and 180 degrees, respectively.

Table XIV shows the timings required for each board to verify messages authenticity within 20ms and 100ms allowed latencies. We observe that it is possible to deploy the A72 and A53 boards to satisfy almost all scenarios within 20ms allowed latency. Instead, ARM11 and less powerful boards are not able to satisfy any scenario within 20ms allowed latency. For 100ms allowed latency, the A72 board satisfies almost all scenarios. The A53 board satisfies all scenarios for α equal to 90 and 135 degrees, but it can only satisfy a few scenarios for 180 degrees. The ARM11 board can satisfy only very low traffic scenarios.

D. Evaluation of the prioritization strategy

Our results demonstrate that prioritizing messages based on the position of the sender is a viable solution to validate all relevant safety messages in real traffic conditions by adopting general purpose automotive-grade boards. By using the circular area heuristic described in Section VI-A we confirmed that the distance-prioritization scheme implicitly specified in the

NHTSA requirements help to reduce the number of signatures by up to 75% with a distance of 150m, but this heuristic is not sufficient with greater distances (e.g. 300m) and heavy traffic scenarios as shown in Table X. By introducing more complex heuristics like the elliptical area described in Section VI-B we can improve the applicability of the boards enabling low power boards like the A53 to be suitable in most scenarios as shown in Table XII. Finally, with the Forward-facing area heuristic VI-C, we can further improve the applicability of the boards by reducing by up to 90% the number of the messages in some specific scenarios (see Table XIV).

A question that is worth analyzing is whether the proposed prioritization strategy affects the security guarantees of the protocol. In particular, we show that the strategy does not improve nor decreases security guarantees. We analyze the security guarantees with regard to the IEEE 1609.2 standard and the NHTSA specifications. The proposed strategy does not conflict with safety requirements specified by the NHTSA (Section IV) and can be potentially integrated with other prioritization strategies based on the classes of messages. Moreover, the proposed strategy requires that the authenticity

	α	Grapes		Highway		Campus		MASA	
		Rush	Normal	Rush	Normal	Rush	Normal	Rush	Normal
20 ms	90°	1	1	1	1	1	1	1	1
	135°	1	1	1	1	1	1	1	1
	180°	5	2	2	1	2	1	2	1
100 ms	90°	13	6	4	1	6	3	3	1
	135°	15	7	6	2	6	3	5	2
	180°	69	31	24	11	24	12	22	12

TABLE XIII: Number of messages in different ranges according to the distance of the sender vehicle using the Forward-facing Area

		Grapes		Highway		Campus		MASA	
		Rush	Normal	Rush	Normal	Rush	Normal	Rush	Normal
20 ms and 50 m									
x86_64	90°	1	1	1	1	1	1	1	1
	135°	1	1	1	1	1	1	1	1
	180°	4	2	1	1	2	1	2	1
A72	90°	3	3	3	3	3	3	3	3
	135°	4	3	3	3	3	3	3	3
	180°	15	7	6	3	7	4	7	3
A53	90°	7	6	6	6	6	6	6	6
	135°	9	6	6	6	6	6	6	6
	180°	33	15	13	6	15	9	14	6
ARM11	90°	37	32	32	32	32	32	32	32
	135°	44	32	32	32	32	32	32	32
	180°	164	74	66	32	75	45	71	32
100 ms and 300 m									
x86_64	90°	9	4	3	1	4	2	2	1
	135°	10	5	4	1	4	2	3	1
	180°	48	21	16	8	17	8	15	8
A72	90°	40	18	12	3	18	9	10	4
	135°	45	22	18	5	18	9	15	6
	180°	208	93	71	34	72	36	68	37
A53	90°	85	37	25	7	37	18	21	8
	135°	95	47	39	10	37	18	32	12
	180°	440	197	151	72	153	77	143	78
ARM11	90°	424	186	127	33	186	91	106	40
	135°	473	233	195	51	187	92	158	60
	180°	2201	982	754	361	765	386	716	390

TABLE XIV: Timing [ms] of the automotive-grade boards with forward-facing heuristic for different angular apertures.

of all the messages must be verified by using the ECQV scheme as specified by the IEEE 1609.2 standard.

We observe that the proposed strategy is not meant to defend receiving vehicles against Denial of Service (DoS) attacks by adversaries that are within the communication range of the receiver. First, attackers could transmit any number of messages to saturate the wireless communication channel and thus completely jam the wireless network [3]. This type of attack is outside the scope of the paper because it regards the physical layer of the communication protocols stack. Second, attackers could send messages with fake sender position either with legitimate or illegitimate digital signatures. The first case is outside the scope of the paper because we consider that legitimate senders never send false information. The effects of false information sent by legitimate vehicles are further analyzed in [35]. In the second case, messages with illegitimate digital signatures are discarded by receivers.

However, the proposed prioritization strategy is based on senders positions before verifying message authenticity. A potential attack could envision sending a high number of messages with fake sender positions that satisfy the heuristic that could saturate the verification throughput of the receiver, thus possibly causing a Denial of Service. Despite this type of attack, the proposed strategy does not introduce security issues. First, since the prioritization heuristics are modeled after the NHTSA heuristics, the messages that are discarded due to the attacks are those that have also a lower priority for safety recommendations. Second, an adversary with the same capabilities (e.g., sending a message within the communication range of a certain vehicle) can operate an even more powerful DoS attack by simply jamming the physical layer of the wireless network, as mentioned above.

VII. CONCLUSIONS

This paper presents an experimental evaluation of ECQV performance on automotive-grade boards for their application in VANETs communications. As a first contribution, we propose an open implementation of ECQV and ECDSA cryptographic operations, available at [37]. This implementation is then used to test the performance of four different automotive-grade boards with different roles, in terms of the maximum number of messages that these platforms can verify in a given time window. This second contribution allows to estimate the highest possible workload that each board can sustain, and demonstrates their applicability in both normal and safety-critical scenarios identified by the NHTSA in V2V communications. As a third contribution, we performed an experimental evaluation of the applicability of the boards in realistic scenarios, representing different portions of a real city (Modena, Italy) characterized by different traffic conditions. These experiments allow us to define several reference workloads that are representative of the number of messages that a single car has to receive and validate in an urban scenario. These workloads demonstrate that even powerful boards are not able to verify the signatures of all the incoming messages within the maximum allowed time. To mitigate this issue and improve the applicability of constrained hardware platforms to the V2V context, as a final contribution we propose and evaluate different heuristics to prioritize signature validation.

ACKNOWLEDGMENT

This research has received funding from COSCA (*Conceptualising Secure CARs*), a project supported by the European Union's Horizon 2020 research and innovation programme under the NGL_TRUST grant agreement no. 825618.

REFERENCES

- [1] STMicroelectronics website. https://www.st.com/content/st_com/en.html, Last visited April 2021.
- [2] MASA: Modena Automotive Smart Area. <https://www.automotivesmartarea.it/?lang=en>, Last visited Sept. 2020.
- [3] AL-KAHTANI, M. S. Survey on security attacks in vehicular ad hoc networks (vanets). In *2012 6th international conference on signal processing and communication systems* (2012), IEEE.
- [4] ASSOCIATION OF RADIO INDUSTRIES AND BUSINESS. 700 mhz band intelligent transport systems. ARIB STD-T109, 2012.
- [5] BAE, M. A. R., SIMPSON, L., FOO, E., AND PIEPRZYK, J. Broadcast authentication in latency-critical applications: On the efficiency of ieee 1609.2. *IEEE Trans. Vehicular Technology* (2019).
- [6] BAZZI, A., CECCHINI, G., MENARINI, M., MASINI, B. M., AND ZANELLA, A. Survey and perspectives of vehicular wi-fi versus sidelink cellular-v2x in the 5g era. *Future Internet* 11, 6 (2019), 122.
- [7] BRIER, E., AND JOYE, M. Weierstraß elliptic curves and side-channel attacks. In *Proc. Int'l Work. public key cryptography* (Feb. 2002).
- [8] BRONNER, F., AND SOMMER, C. Efficient multi-channel simulation of wireless communications. In *2018 IEEE Vehicular Networking Conference (VNC)* (2018).
- [9] BUCHER, H., KLIMM, A., SANDER, O., AND BECKER, J. Power estimation of an ecdsa core applied in v2x scenarios using heterogeneous distributed simulation. In *2015 IEEE/ACM 19th International Symposium on Distributed Simulation and Real Time Applications (DS-RT)* (2015).
- [10] CERTICOM RESEARCH. Standards for EfficientCryptography 1. SEC 1, 2009.
- [11] CERTICOM RESEARCH. Sec 4: Elliptic curve qu-vanstone implicit certificate scheme, standards for efficient cryptography group. version 1.0.
- [12] CODECA, L., FRANK, R., AND ENGEL, T. Luxembourg sumo traffic (lust) scenario: 24 hours of mobility for vehicular networking research. In *2015 IEEE Vehicular Networking Conference (VNC)* (2015), IEEE.
- [13] DUPONT, G., DEN HARTOG, J., ETALLE, S., AND LEKIDIS, A. A survey of network intrusion detection systems for controller area network. In *Proc. IEEE Int'l Conf. Vehicular Electronics and Safety* (Sep. 2019).
- [14] ECKHOFF, D., BRUMMER, A., AND SOMMER, C. On the impact of antenna patterns on vanet simulation. In *2016 IEEE Vehicular Networking Conference (VNC)* (2016).
- [15] ECKHOFF, D., SOMMER, C., AND DRESSLER, F. On the necessity of accurate IEEE 802.11 p models for IVC protocol simulation. In *2012 IEEE 75th Vehicular Technology Conference (VTC Spring)* (2012), IEEE, pp. 1–5.
- [16] ETSI. ITS-G5 access layer specification for intelligent transport systems operating in the 5 ghz frequency band. EN 302 663, 2019.
- [17] HEINOVSKI, J., KLINGLER, F., DRESSLER, F., AND SOMMER, C. Performance comparison of ieee 802.11 p and arib std-t109. In *2016 IEEE Vehicular Networking Conference (VNC)* (2016).
- [18] IEEE. Standard for Wireless Access in Vehicular Environments—Security Services for Applications and Management Messages. Std. 1609.2a-2017.
- [19] IEEE. Guide for wireless Access in vehicular environments (WAVE) architecture. *IEEE: Piscataway, NJ, USA* (2013), 1609.
- [20] IEEE 802.11 WORKING GROUP AND OTHERS. Ieee standard for information technology—telecommunications and information exchange between systems—local and metropolitan area networks—specific requirements—part 11: Wireless lan medium access control (mac) and physical layer (phy) specifications amendment 6: Wireless access in vehicular environments. *IEEE Std 802.11*, 11 (2010).
- [21] KARGL, F., PAPADIMITRATOS, P., BUTTYAN, L., MÜTER, M., SCHOCH, E., WIEDERSHEIM, B., THONG, T.-V., CALANDRIELLO, G., HELD, A., KUNG, A., ET AL. Secure vehicular communication systems: implementation, performance, and research challenges. *IEEE Comm. Magazine* 46, 11 (2008).
- [22] KNEŽEVIĆ, M., NIKOV, V., AND ROMBOUTS, P. Low-latency ecdsa signature verification—a road toward safer traffic. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* (2016).
- [23] LOPEZ, P. A., BEHRISCH, M., BIEKER-WALZ, L., ERDMANN, J., FLÖTTERÖD, Y.-P., HILBRICH, R., LÜCKEN, L., RUMMEL, J., WAGNER, P., AND WIEBNER, E. Microscopic traffic simulation using sumo. In *The 21st IEEE International Conference on Intelligent Transportation Systems* (2018).
- [24] LU, R., LIN, X., LUAN, T. H., LIANG, X., AND SHEN, X. Pseudonym changing at social spots: An effective strategy for location privacy in vanets. *IEEE transactions on vehicular technology* (2011).
- [25] MACKAY, K. MicroECC. <https://github.com/kmackay/micro-ecc>, Last visited Mar. 2020.
- [26] MARCHETTI, M., AND STABILI, D. READ: Reverse Engineering of Automotive Data Frames. *IEEE Trans. Information Forensics and Security* 14, 4 (2019).
- [27] MATSUMOTO, T., HATA, M., TANABE, M., YOSHIOKA, K., AND OISHI, K. A method of preventing unauthorized data transmission in controller area network. In *2012 IEEE 75th Vehicular Technology Conference (VTC Spring)*.
- [28] NABIL, G., NAZIHA, K., LAMIA, F., AND LOTFI, K. Hardware implementation of elliptic curve digital signature algorithm (ecdsa) on koblitz curves. In *2012 8th International Symposium on Communication Systems, Networks & Digital Signal Processing (CSNDSP)* (2012), IEEE.
- [29] NATIONAL HIGHWAY TRAFFIC SAFETY ADMINISTRATION. Vehicle safety communication projet – task 3 final report. DOT HS 809 859, March 2005.
- [30] NATIONAL HIGHWAY TRAFFIC SAFETY ADMINISTRATION. Vehicle safety communication projet – final report. DOT HS 810 591, Apr. 2006.
- [31] NOWDEHI, N., AOUDI, W., ALMGREN, M., AND OLOVSSON, T. CASAD: CAN-Aware Stealthy-Attack Detection for In-Vehicle Networks. arXiv 1909.08407, 2019.
- [32] PAPADIMITRATOS, P., BUTTYAN, L., HOLCZER, T., SCHOCH, E., FREUDIGER, J., RAYA, M., MA, Z., KARGL, F., KUNG, A., AND HUBAUX, J.-P. Secure vehicular communication systems: design and architecture. *IEEE Comm. Magazine* 46, 11 (2008).
- [33] PAUL, A., CHILAMKURTI, N., DANIEL, A., AND RHO, S. Chapter 2 - intelligent transportation systems. In *Intelligent Vehicular Networks and Communications* (2017), Elsevier.
- [34] PESÉ, M. D., STACER, T., CAMPOS, C. A., NEWBERRY, E., CHEN, D., AND SHIN, K. G. LibreCAN: Automated CAN message translator. In

- Proc. 2019 ACM SIGSAC Conf. Computer and Communications Security* (2019).
- [35] POLLICINO, F., STABILI, D., BELLA, G., AND MARCHETTI, M. Six-Pack: Abusing ABS to avoid Misbehavior detection in VANETs. In *2021 IEEE 93rd Vehicular Technology Conference (VTC2021-Spring)* (2021), IEEE, pp. 1–6.
 - [36] POLLICINO, F., STABILI, D., FERRETTI, L., AND MARCHETTI, M. An experimental analysis of ecqv implicit certificates performance in VANETs. In *Proc. IEEE 92nd Vehicular Technology Conf.* (Nov. 2020).
 - [37] POLLICINO, F., STABILI, D., FERRETTI, L., AND MARCHETTI, M. Implementation of the ECVQ implicit certificate scheme for low-power devices. <https://weblab.ing.unimore.it/resources/uECQV.zip>, Last visited Sept. 2020.
 - [38] QUALCOMM. ITS Stack 80-PE732-64 REV A. <https://www.qualcomm.com/media/documents/files/c-v2x-its-stack.pdf>, Last visited Aug. 2021.
 - [39] RAYA, M., AND HUBAUX, J.-P. The security of vehicular ad-hoc networks. In *Proc. 3rd ACM Work. Security of ad-hoc and sensor networks* (Jul. 2005).
 - [40] ROSENSTATTER, T., SANDBERG, C., AND OLOVSSON, T. Extending AUTOSAR's Counter-Based Solution for Freshness of Authenticated Messages in Vehicles. In *Proc. IEEE 24th Pacific Rim Int'l Symp. Dependable Computing* (Dec 2019).
 - [41] SAE INTERNATIONAL. Dedicated short range communications (dsrc) message set dictionary. *SAE International* (2016).
 - [42] SAE INTERNATIONAL. Dedicated short range communications (dsrc) message set dictionary. *SAE International* (2016).
 - [43] SAMPIGETHAYA, K., LI, M., HUANG, L., AND POOVENDRAN, R. Amoeba: Robust location privacy scheme for vanet. *IEEE Journal on Selected Areas in communications* (2007).
 - [44] SCHINK, M., WAGNER, A., UNTERSTEIN, F., AND HEYSZL, J. Security and trust in open source security tokens. *IACR Transactions on Cryptographic Hardware and Embedded Systems* (2021), 176–201.
 - [45] SEMICONDUCTORS, N. 3GPP TS 36.213, v 14.2.0. <https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=2427>, Last visited Aug. 2021.
 - [46] SEMICONDUCTORS, N. 3GPP TS 36.300, v 14.4.0. <https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=2430>, Last visited Aug. 2021.
 - [47] SEMICONDUCTORS, N. RoadLINK SAF5400. <https://www.nxp.com/products/wireless/dsrc-safety-modem/roadlink-saf5400-single-chip-modem-for-v2x:SAF5400>, Last visited Aug. 2021.
 - [48] SGHAIER, A., ZEGHID, M., AND MACHHOUT, M. Fast hardware implementation of ecdsa signature scheme. In *2016 International Symposium on Signal, Image, Video and Communications (ISIVC)* (2016), IEEE.
 - [49] SIMPLICIO, M. A., COMINETTI, E. L., PATIL, H. K., RICARDINI, J. E., AND SILVA, M. V. M. The unified butterfly effect: Efficient security credential management system for vehicular communications. In *Proc. 2018 IEEE Vehicular Networking Conf.* (Dec 2018).
 - [50] SOMMER, C., ECKHOFF, D., GERMAN, R., AND DRESSLER, F. A computationally inexpensive empirical model of ieee 802.11 p radio shadowing in urban environments. In *Eighth international conference on wireless on-demand network systems and services*.
 - [51] SOMMER, C., GERMAN, R., AND DRESSLER, F. Bidirectionally coupled network and road traffic simulation for improved ivc analysis. *IEEE Transactions on Mobile Computing* (2011).
 - [52] STABILI, D., FERRETTI, L., AND MARCHETTI, M. Analyses of secure automotive communication protocols and their impact on vehicles life-cycle. In *Proc. IEEE Int'l Conf. Smart Computing* (June 2018).
 - [53] STABILI, D., AND MARCHETTI, M. Detection of missing can messages through inter-arrival time analysis. In *Proc. IEEE 90th Vehicular Technology Conf.* (Sep. 2019).
 - [54] STORCK, C. R., AND DUARTE-FIGUEIREDO, F. A survey of 5g technology evolution, standards, and infrastructure associated with vehicle-to-everything communications by internet of vehicles. *IEEE Access* 8 (2020), 117593–117614.
 - [55] SUMO. OSMWebWizard. <https://sumo.dlr.de/docs/Tutorials/OSMWebWizard.html>, Last visited Sept. 2020.
 - [56] TSVIKA, D., MONTVELISKY, Y., MARCHETTI, M., STABILI, D., COLAJANNI, M., AND WOOL, A. Vehicle Safe-Mode, Concept to Practice. Limp-Mode in the Service of Cybersecurity. *SAE International Journal of Transportation Cybersecurity and Privacy* 3, 1 (2020).
 - [57] VARGA, A., AND HORNIG, R. An overview of the omnet++ simulation environment. In *Proceedings of the 1st international conference on Simulation tools and techniques for communications, networks and systems & workshops*, ICST (Institute for Computer Sciences, Social-Informatics and ...
 - [58] XU, Q., MAK, T., KO, J., AND SENGUPTA, R. Vehicle-to-vehicle safety messaging in dsrc. In *Proceedings of the 1st ACM international workshop on Vehicular ad hoc networks* (2004).