

On the Performance of Secure Vehicular Communication Systems

Giorgio Calandriello*, Panos Papadimitratos[†], Jean-Pierre Hubaux[†], Antonio Lioy*

*DAUIN, Politecnico di Torino, Italy

{giorgio.calandriello, antonio.lioy}@polito.it

[†]LCA, IC, EPFL, Switzerland

{panos.papadimitratos, jean-pierre.hubaux}@epfl.ch

Abstract—Vehicular communication (VC) systems are being developed primarily to enhance transportation safety and efficiency. Vehicle-to-vehicle communication, in particular, frequent cooperative awareness messages or safety beacons, has been considered over the past years as a main approach. Meanwhile, the need to provide security and to safeguard users’ privacy is well understood, and security architectures for VC systems have been proposed. Although technical approaches to secure VC have several commonalities and a consensus has formed, there are critical questions that have remained largely unanswered: Are the proposed security and privacy schemes practical? Can the secured VC systems support the VC-enabled applications as effectively as unsecured VC would? How should security be designed so that its integration into a VC system has a limited effect on the system’s performance? In this paper, we provide answers to these questions, investigating the joint effect of a set of system parameters and components. We consider the state-of-the-art approach in secure VC, and we evaluate analytically and through simulations the interdependencies among components and system characteristics. Overall, we identify key design choices for the deployment of efficient, effective, and secure VC systems.

Index Terms—Security and Protection, Wireless communication, Reliability, Availability, and Serviceability

I. INTRODUCTION

Vehicular communication (VC) systems will comprise vehicles and fixed road-side equipment (RSU) with wireless transceivers, and sensing and processing units. Vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) communication will enable a range of applications, with transportation safety playing a predominant role. Almost all research and development efforts converge to safety applications based on V2V communication, with vehicles frequently *beaconing* their status (e.g., position, speed, direction), along with warnings about potential dangers.

Nonetheless, VC systems can be vulnerable to attacks and jeopardize users’ privacy: An adversary could, for example, inject beacons with false information, or collect vehicles’ messages to track their locations and infer sensitive user data. Industry, academia, and authorities have recently understood that security and privacy protection are prerequisites for the deployment of VC systems. Security architectures were developed by the IEEE 1609.2 working group [1], the SeVeCom project [2], [3], following the earlier NoW project [4] and now the Car-to-Car Communication Consortium (C2C-CC) [5] and the eSafety eSecurity WG activities [6].

Across projects and working groups, secure VC systems rely on public key cryptography and digital signatures to protect V2V and V2I messages; *Certification Authorities* (CAs) manage credentials for legitimate participants (vehicles and RSUs). Pseudonymous authentication, with vehicles using short-lived credentials and public-private key pairs, provides protection of privacy along with security (authentication, integrity and non-repudiation as primary requirements). Security mechanisms protect all traffic sent across the 802.11p data link [7], including the safety beacons each vehicle transmits, typically every 100 to 300 ms.

Adding security for this high-rate communication would incur high overhead, both in terms of communication and processing. Consider, for example, a vehicle receiving digitally signed safety beacons from a hundred vehicles within range; it would need to validate a high percentage or almost all of those within a short delay in the order of a hundred milliseconds [7]. Even if VC is effective under such dense network conditions, the additional security overhead could cause failure in meeting the delay and reliability requirements of safety applications. This is especially so because the VC environment lacks abundant resources (bandwidth, computational power).

The following question naturally follows: Can secure VC systems be practical? Given the current system constraints and design approaches, could the addition of security and privacy mechanisms make VC systems ineffective? We address this problem in this paper, building on our previous work ([8], [9]). Based on broadly accepted approaches for secure and privacy-enhancing VC [1]–[4], we first outline how pseudonymous authentication is possible without repeated interactions with the CAs. Then, we present a proposal for reducing the security overhead without harming the effectiveness of the VC system, and we investigate how variants of secure VC instantiations affect the system performance. In particular, our main contribution is the first comprehensive evaluation of secure VC operations: (i) We evaluate the communication reliability, and then (ii) we determine if and how VC nodes can sustain the incurred processing load, providing an approximate analytical evaluation and closely matching simulation results. Having determined if VC nodes have sufficient processing power, (iii) we consider the overall system performance with respect to transportation safety and (iv) transportation efficiency, evaluating secure VC-enabled applications for a broad range of system configurations. We find that appropriately designed

security and privacy-enhancing VC systems can essentially support a safety application as effectively as unsecured VC systems can. Moreover, (v) we investigate revocation and propose a practical method for anonymous authentication schemes in VC, and (vi) we discuss additional technical issues and demonstrate how our model can assess processing resources needed for future systems.

We discuss the problem at hand and outline our investigation approach in Sec. II and present the representative secure and privacy-enhancing VC schemes we evaluate in Sec. III. The simulation setup, our analysis and experimental results follow in Sec. IV–Sec. VIII, along with the treatment of revocation in Sec. IX and other technical issues and extensions for future instantiations in Sec. X. We discuss related work in Sec. XI before we conclude.

II. PROBLEM AND APPROACH OVERVIEW

We want to determine whether the broadly accepted state of the art of secure VC is viable, especially considering how challenging VC environments are; because heavy-traffic scenarios (thus, dense network topologies) - with tens, one hundred or more vehicles (nodes) within range - can often occur. The traditional approach has been to analyze the protocol overhead and the network performance. However, in VC systems the objective is not to have a well-performing network *per se*, but to effectively support VC-specific applications. This is why we investigate the overall system performance, considering five dimensions: (i) *communication technology*, (ii) *system resources*, (iii) *network configuration and environmental factors*, (iv) *security protocols*, and (v) *supported applications*.

The technology commonly accepted for V2V and V2I communication is the IEEE 802.11p [10], which is incorporated in the Dedicated Short Range Communication (DSRC) - Wireless Access in a Vehicular Environment (WAVE) [11] and the Communication Access for Land Mobiles (CALM) [12] standards. Vehicles transmit periodic *safety beacons* on one dedicated channel, at a system-selectable beaconing rate. *Bandwidth*, one of the primary system resources, is determined by the standards, and it is considered fixed for this investigation. The second primary resource, *processing power*, can be adapted. Here, we take into consideration platforms that are currently used in VC prototypes, but any system should have sufficient processing power for its designated tasks. Thus, the system designer can always increase the processing power at the expense of increased cost.

The use of specific *cryptographic primitives* and other *protocol functionalities* determine the processing load for each node (vehicle). Here we consider the basic pseudonymous authentication approach, which has gained broad acceptance: It provides message authentication, integrity, non-repudiation and it makes it hard for two or more messages from the same sender to be linked¹. Given the large number of temporary identities (pseudonyms) in the system, pseudonymous

authentication can become cumbersome to manage. Therefore, we consider here a novel scheme, first presented in [8], [9], to alleviate this constraint, thanks to a more powerful but also more expensive anonymous authentication primitive. We describe these security protocols in Sec. III.

We consider *transportation safety and efficiency applications* as they are distinctive features of VC systems (compared to other mobile computing systems) and two main driving forces for the VC systems deployment. Moreover, they are, especially the safety ones, the most challenging among VC-enabled applications; their stringent time constraints and their critical nature can affect the well-being of the vehicle passengers. We focus here on one safety application, *Emergency Braking Notification (EBN)*, and one efficiency application, *Decentralized Floating Car Data (DFCD)*.

In order that the appropriate processing power can be determined and provisioned, we provide a framework to analyze the effect of a given *processing load* on the node performance. Then, we consider a system for which processing is not a bottleneck (otherwise, the system would certainly fail) and we evaluate the effectiveness of the applications. Conversely, given appropriate design choices (i.e., equipment with sufficient power), our investigation reveals the effect of other parameters and their interdependencies.

III. SECURE COMMUNICATION

Each node (vehicle) has a long-term, unique identity and corresponding credentials managed by a *Certification Authority (CA)*; without loss of generality, we assume there is a single CA, even though in reality a CA hierarchy would be present [13]. Instead of using their long-term credentials, vehicles obtain from the CA and use a set of short-lived certified public keys that do not identify the vehicle; then, they digitally sign messages with the corresponding private keys. As this is the widely used approach of *pseudonymous authentication* [1]–[5], we refer to it as the *Baseline Pseudonym (BP)* scheme, and define its operation in Sec. III-A. We consider only the vehicles, as the privacy of RSUs or other infrastructure does not need to be protected.

As the BP scheme requires numerous short-lived certificates and keys per vehicle, the stronger the protection of privacy is sought the higher the number of identities would be. For large-scale systems, this and the cost of periodically pre-loading vehicles with temporary keys and credentials can become a significant burden. To reduce the key management complexity and enhance the system usability and efficiency, we propose that nodes self-generate, i.e., self-certify, their own pseudonyms. With this method, first described in [8], [9], vehicles do not need to be side-lined or to compromise their user's privacy if a “fresh” pseudonym is no longer available; no “over-provisioning” in the supply of pseudonyms is necessary; and the cost of obtaining new pseudonyms over an “out-of-band” channel is avoided.

This can be achieved with the use of *anonymous authentication* primitives, notably *Group Signatures (GS)* we describe in Sec. III-B. As GS is hard to use for all VC messages, because of the GS processing and communication overhead,

¹More precisely, it allows that messages produced by a node over a protocol-selectable period of time, τ , be linked. But messages m_1, m_2 generated at times t_1, t_2 respectively, such that $t_2 > t_1 + \tau$, should not be linkable. The shorter τ is the fewer the linkable messages are and the harder tracking a node becomes.

in Sec. III-C we propose our *Hybrid Pseudonym (HP)* scheme that allows vehicles to generate on-the-fly their pseudonyms, by combining the BP and GS approaches. HP alleviates the management overhead of the BP, but in principle it is more costly than BP (due to HP's use of GS). To reduce the cost of HP to equal roughly that of BP and to increase the robustness of any pseudonymous approach, we propose a set of optimizations (Sec. III-D).

Concerning revocation (Sec. IX and X), all the approaches make use of *Revocation Lists* (RL), generated by the CA and distributed to vehicles primarily via the infrastructure [2], [13]. When a node validates a certificate, it checks whether the sender is revoked; if successful (i.e. the sender is not revoked), it proceeds with validating the message (signature(s)).

A. Baseline Pseudonym (BP) Scheme

Each node V is equipped with a set of *pseudonyms* that are certified *public keys* without any information that identifies V . More specifically, for the i -th pseudonym K_V^i for node V , the CA provides a certificate $Cert_{CA}(K_V^i)$, which is simply a CA signature on the public key K_V^i (unlike the common notion of certificate, for example the X.509 certificate). The node uses the private key k_V^i for the pseudonym K_V^i to digitally sign messages. To enable message validation, the pseudonym and the certificate of the signer are attached in each message. With $\sigma_{k_V^i}()$ denoting V 's signature under its i -th pseudonym and m the signed message payload, the message format is:

$$M1 : m, \sigma_{k_V^i}(m), K_V^i, Cert_{CA}(K_V^i)$$

Upon receipt of $M1$, a node, with the public key of the CA available, validates $Cert_{CA}(K_V^i)$, and then verifies the signature using K_V^i .

Each pseudonym is used at most for a period τ (referred in the rest of the paper as the *pseudonym lifetime*) and then discarded. We abstract away a number of possible implementation aspects, such as (i) the dynamic adaptation of the period of pseudonym usage, (ii) the number of pseudonyms (K_V^i and the corresponding k_V^i , $Cert_{CA}(K_V^i)$) that are pre-loaded to V , (iii) the frequency of pseudonym refills, and (iv) policies for pseudonym change, such as factors rendering a pseudonym change unnecessary (e.g., a TCP connection to an access point), and interactions of pseudonym changes with the network stack [14]. All these are important yet largely orthogonal to this investigation. The CA maintains a map from the long-term identity of V to the $\{K_V^i\}$ set of pseudonyms provided to a node. If presented with a message $M1$, the CA can perform the inverse mapping and identify the signer.

B. Group Signature (GS) Scheme

Each node V is equipped with a secret *group signing key* gsk_V , with the *group members* comprising all vehicles registered with the CA. A *group public key* GPK_{CA} allows for the validation (by any node) of any *group signature* $\Sigma_{CA,V}$ generated by a group member. Intuitively, a group signature scheme allows any node V to sign a message on behalf of the group, *without* V 's identity being revealed to the signature verifier. Moreover, it is impossible to link any two signatures

of a legitimate group member. Note that no public key or other credentials need to be attached to an anonymously authenticated message; the format is:

$$M2 : m, \Sigma_{CA,V}(m)$$

Group signatures, introduced by Chaum [15], are revisited in numerous works, e.g., [16]–[19], with formal definitions in [20], [21]. For the rest of the discussion, we assume and use the group signature scheme proposed in [22]. If the identification of a signer is necessary, the CA can perform an *Open* operation [20], [21] and reveal the signer's identity.

C. Hybrid Pseudonym (HP) Scheme

The combination of the BP and GS schemes is the basic element of our proposal [8], [9]. Each node V is equipped with a group signing key gsk_V and the group public key GPK_{CA} (recall that the group is the total of vehicles registered with the CA). Rather than generating group signatures to protect messages, a node generates its own set of pseudonyms $\{K_V^i\}$ (according to the BP public key cryptosystem). As for the BP scheme (Sec. III-A), a pseudonym is a public key without identification information, and $\{k_V^i\}$ is the set of corresponding private keys. For HP, the CA does not provide a certificate on K_V^i ; instead, V uses gsk_V to generate a group signature $\Sigma_{CA,V}()$ on each pseudonym K_V^i instead. In other words, it generates and “self-certifies” K_V^i on-the-fly, by producing $\Sigma_{CA,V}(K_V^i)$. Similarly to $M1$, V attaches $\Sigma_{CA,V}(K_V^i)$ to each message, and signs with the corresponding k_V^i :

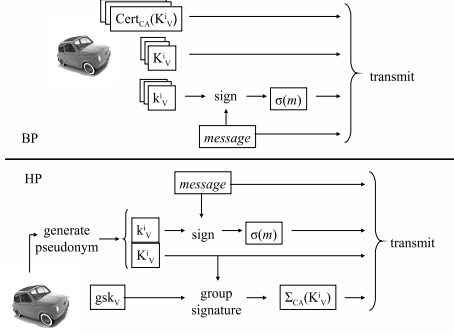
$$M3 : m, \sigma_{k_V^i}(m), K_V^i, \Sigma_{CA,V}(K_V^i)$$

When a node receives a message $M3$, the group signature $\Sigma_{CA,V}(K_V^i)$ is verified, using GPK_{CA} . If successful, the receiver infers that a legitimate system (group) member generated pseudonym K_V^i . We emphasize that, as per the properties of group signatures, the receiver/verifier of the certificate *cannot* identify V and *cannot* link this certificate and pseudonym to any prior pseudonym used by V . Once the legitimacy of the pseudonym is established, the validation of $\sigma_{k_V^i}(m)$ is identical to that for $M1$. To identify the message signer, an *Open* on the $\Sigma_{CA,V}(K_V^i)$ group signature is necessary; message m is bound to K_V^i via $\sigma_{k_V^i}(m)$, and K_V^i is bound to V via $\Sigma_{CA,V}(K_V^i)$. Fig. 1(a) compares the BP and HP.

D. Optimizations for the BP and HP Schemes

We describe optimizations to reduce overhead (Optimizations 1 and 2) and enhance robustness (Optimization 3). We employ the notation of the HP scheme, but the same considerations hold for BP too. Fig. 1(b) summarizes Optimizations 2 and 3.

Optimization 1: On the sender's side, $\Sigma_{CA,V}(K_V^i)$ is computed only once per K_V^i , because $\Sigma_{CA,V}(K_V^i)$ remains unchanged throughout the pseudonym lifetime τ . For the same reason, on the verifier's side the $\Sigma_{CA,V}(K_V^i)$ is validated upon the first reception and stored, even though the sender appends it to multiple (all) messages. For all subsequent receptions, if $\Sigma_{CA,V}(K_V^i)$ has already been seen, the verifier skips its validation. This optimization is useful because in practice $\tau \gg \gamma^{-1}$, where γ is defined as the *beacon frequency*.



(a) Illustration of the BP and HP schemes.


(b) Illustration of Optimizations 2 and 3, with $\alpha = 5$ and $\beta = 2$. One LONG message is sent every 5, and repeated 2 times after a pseudonym renewal.

Fig. 1: Illustration of the BP and HP security schemes and related optimizations.

Optimization 2: The sender appends its signature $\sigma_{k_V^i}(m)$ to all messages, but it appends the corresponding $K_V^i, \Sigma_{CA,V}(K_V^i)$ only once every α messages. We term such messages ($M1$ and $M3$) as *LONG*. $M4$ is defined as follows:

$$M4 : m, \sigma_{k_V^i}(m)$$

We denote $M4$ as *SHORT*, and α as the *Certificate Period*. $\alpha \in [1, \tau\gamma]$, where $\tau\gamma$ is the total number of transmissions during the pseudonym lifetime τ . To allow the user to choose the right K_V^i to verify an incoming *SHORT* message, all messages will carry a randomly generated 4-byte *keyID* field. This does not affect privacy as all *SHORT* messages signed under the same K_V^i can be trivially linked.

When a pseudonym change occurs, the new triplet $\sigma_{k_V^{i+1}}(m), K_V^{i+1}, \Sigma_{CA,V}(K_V^{i+1})$ must be computed and transmitted. V will sign messages with the new k_V^{i+1} corresponding to K_V^{i+1} from then on.

Optimization 2 can affect the protocol robustness, if the message that carries $K_V^{i+1}, \Sigma_{CA,V}(K_V^{i+1})$ is not received. Then, nodes in range of V must wait for α messages for the next pseudonym transmission, while being unable to validate any message from V . This can be dangerous if vehicles are close to each other and/or are moving at high relative speeds. Thus, we propose the following scheme to mitigate this problem.

Optimization 3: V repeats the transmission of $K_V^{i+1}, \Sigma_{CA,V}(K_V^{i+1})$ for β consecutive messages when K_V^{i+1} is issued, with β denoted as the *Push Counter*. After the β repetitions, with $\beta \in [0, \alpha - 1]$, the normal sequence 1 *LONG*, $\alpha - 1$ *SHORT* starts again.

IV. EVALUATION OVERVIEW

We analyze the system performance of secure VC along the dimensions presented in Sec. II. Given the complexity of the problem, we employ simulation as a primary tool of analysis and we provide analytical approximations. We evaluate the effectiveness of the EBN and the DFCD applications in a variety of setups, to gain insight into the role of each of

Parameter	Symbol	Range	Unit
Certificate Period	α	1,5,10,15,30,50	messages
Push Period	β	0 to 10	messages
Beacon Frequency	γ	3.33 and 10	beacons/s
Pseudonym Lifetime	τ	60	s
Number of Neighbors	N	160, 240, 320	vehicles
Packet Payload	m	200	bytes
Initial Vehicle Spacing	s	$20 \pm 1.5, 150, 200$	m
Average Vehicle Speed	v	65 and 80	Km/h
Road Setup	-	4,6,8	lanes
Security Schemes	-	BP and HP	-
Nominal Commun. Range	r	200	m

TABLE I: System parameters and values assigned for the evaluation.

the system parameters; indicative values for these are summarized in Table I. We study challenging or even extreme transportation conditions, because the system has to remain operational even under these conditions. We are interested in parameters that are not related to security (and of course in security protocol parameters) precisely because we want to capture their importance for the secure VC system.

We assume that only vehicles transmit, because RSUs would always be less numerous (each serving an area with tens or hundreds of vehicles), and often completely absent; thus, almost all of the safety-related data will be generated by vehicles. Finally, we also assume that all beacons carry relevant information for safety applications. We couple the *ns-2* simulator, which simulates V2V communication, with a custom module written in C, which simulates (i) the EBN and DFCD applications and its effect on vehicles movement and (ii) the security processing of messages. We choose such a combination because we could not find another publicly available simulation environment with security functionality integrated and with nodes adjusting their behavior according to the messages they receive.

First, we evaluate the cryptographic overhead, in terms of communication and processing, and we choose representative cryptographic primitives, security levels, and a reference platform (Sec. V) and then analyze the communication reliability (Sec. VI). Based on these two elements, we study the effect

Algorithm	Security level (bits)	Sign (ms)	Verify (ms)	Signature (bytes)	Public key (bytes)	Private key (bytes)
ECDSA-192	96	0.5	3	48	25	24
ECDSA-256	128	0.8	4.2	64	33	32
GS	128	53.7	49.3	225	800	64

TABLE II: Computation costs on a 1.5GHz Centrino processor and communication overhead for different signing algorithms: Elliptic Curve Digital Signature Algorithm (ECDSA) using different standardized elliptic curves, and a representative efficient Group Signature (GS) algorithm [22].

	Sign (ms)	Verify (ms)	Overhead (bytes)
BP LONG	1.3	7.2	141
HP LONG	54.2	52.3	302
SHORT	0.5	3	52

TABLE III: Processing delay (in ms) and communication overhead (in bytes) for different packet types.

	Packets per beacon period γ^{-1}
BP LONG	13.9
HP LONG	1.9
SHORT	33.3

TABLE IV: Maximum number of verifiable packets per γ^{-1} s, for $\gamma = 10$.

of the processing overhead on individual nodes (Sec. VII) and develop an approximate model. Assuming that nodes are provisioned with sufficient processing power, we perform a system-wide analysis for the considered applications (Sec. VIII). Finally, we address revocation, which can be costly (Sec. IX), and demonstrate how processing power provisioning can be done for future secure VC systems (Sec. X).

V. CRYPTOGRAPHIC OVERHEAD

We use EC-DNA as the basic signature algorithm [23], the group signature algorithm proposed by [22], and security level of $t = 96$ bits for message signatures and $t = 128$ bits for CA certificates in BP and for group signatures used in GS and HP. High security might not be necessary for the short-lived K_V^i , but it is required for the long-term keys and CA certificates.

Overhead: The $K_V^i, Cert_{CA}(K_V^i)$ is 89 bytes for BP, and with $\sigma_{K_V^i}(m)$ and $KeyID$ the overhead is 141 bytes per message. For GS, the overhead is $\Sigma_{CA,V}(m)$, thus 225 bytes per message. For HP, the overhead is $\sigma_{K_V^i}(m), K_V^i, \Sigma_{CA,V}(K_V^i), KeyID$, in total 302 bytes per message. For the $\alpha - 1$ *SHORT* messages, the overhead is $\sigma_{K_V^i}(m), KeyID$, thus 52 bytes. The effective overhead reduction depends on the value of α (see Sec. VI).

Computation: We make use of a Centrino machine with the clock speed set at 1.5 GHz, which is close to the CVIS (Cooperative Vehicle-Infrastructure System) vehicle PC, a rather powerful platform (compared to generally available embedded processors) adopted for the development of future VANET applications [24]. We obtain an EC-DNA benchmark on the platform through the OpenSSL standard test suite [25]. As for group signatures, a well-established implementation of the chosen algorithm [22] is not yet available. Thus, to

estimate the processing delay, we calculate the number of 32-bit word scalar multiplications required for GS signing and verifying; we extract the relevant data from [26] and [27] and we benchmark the scalar multiplication operation.

Table II shows the costs for signature, verification and overhead for the chosen algorithms. To obtain individual processing delays for a given type of message, it suffices to take the sum of the corresponding cryptographic primitive delays (M1, M3 and M4). Security levels are $t = 96$ for $\sigma_{K_V^i}(m)$, and $t = 128$ for $Cert_{CA}(K_V^i), \Sigma_{CA,V}(m)$ and thus $\Sigma_{CA,V}(K_V^i)$; we summarize results per message in Table III.

VI. COMMUNICATION RELIABILITY

The communication reliability is of central importance and depends on the channel properties and load; the more loaded the channel is, the more likely it is for a packet collision to occur at the wireless medium, which depends on the number of transmitters, N , the beacon frequency, γ , and the packet size (including the security overhead). We implement beacons with information on vehicle position, speed, and direction, with a timestamp, and safety warnings in a payload, m , of 200 bytes. The physical layer models realistic radio propagation [28], [29], with a nominal communication range of $r = 200$ m and a bandwidth of 6 Mb/s [7], [30], [31].

We estimate, with the help of detailed simulations, the average probability of successful reception at a receiving node at the center of a 200-meter radius disc that covers the entire width of a multi-lane highway and is filled with N uniformly spaced neighbors. We consider various settings, increasing the number of lanes and decreasing the vehicle density, varying the size of N from 8 to 160; a subset of these settings (four-, six- and eight- lanes with average spacing of 20 meters) is used in Sec. VIII. This metric is independent of the distance between the transmitter and the receiver. The reception probability as a function of the sender-receiver distance is presented in our previous work [9], and it is omitted here for simplicity and due to lack of space. The results are shown in Fig. 2 with 95% confidence intervals and are repeated 5 times, with each repetition lasting 60 s of simulated time.

We note that the 802.11p broadcast communication, a CSMA/CA protocol without acknowledgements, could be modeled and evaluated analytically, for example, in terms of the probability of successful reception and throughput, following numerous works for similar CSMA/CA protocols under various assumptions (on traffic conditions, presence or absence of channel errors, packet sizes, impact of hidden terminals, etc). An effort to derive a precise analytical model for VC is orthogonal to our investigation. What we need here




Fig. 2: Reliability of message reception for secure VC (beaconing), as a function of the neighborhood size, N .

γ (bcn./s.)	HP $\alpha = 1$	HP $\alpha = 10$	BP $\alpha = 1$	BP $\alpha = 10$
10	5020	2770	3410	2609
3.33	1671.66	922.41	1135.53	868.70
2	1004	554	682	521.8

TABLE V: Offered load per transmitter, in bytes/s, for different security schemes and settings.

is an accurate evaluation of the communication reliability, as a stepping stone for our security-related investigation; this can be obtained via detailed simulations.

With the fixed available bandwidth, specific for the communication technology, the communication reliability depends on the offered load; Table V summarizes the load for each of the scenarios in this paper. Fig. 2 shows the estimate of the probability of reception, P (i.e. the ratio of received messages over transmitted beacons), as a function of the number of transmitters, N , the beacon frequency, γ , and the protocol parameter α . We observe that the communication performance degrades fast with N when γ is high (Fig. 2(a) for HP), whereas the degradation is much slower as N increases for lower γ values. The effect of increasing α , thus reducing overhead is significant even when γ is not very high (e.g., $\gamma = 3.33$ beacons/s, Fig. 2(b) again for HP). Finally, as shown by Fig. 2(c) for $\gamma = 10$ and $\alpha = 10$, the BP and HP schemes perform almost identically.

These results show that γ turns out to be the most significant channel load factor. Choosing a smaller value for γ decreases the channel saturation and thus the processing overhead (fewer messages are sent); but it also affects the transportation safety, as we show in Sec. VIII. At the same time, the appropriate choice of BP and HP parameters can reduce security overhead (notably α , as it will become clear in Sec. VIII the needed β values incur very limited overhead). The almost identical P for BP and HP also show the benefit from the proposed optimizations, as both schemes have comparable overhead (with the advantages of HP).

VII. PROCESSING OVERHEAD

We want to answer the following questions: (i) How many packets does a given node V_R have to verify per time unit, in various VC settings? (ii) What is the additional message

verification delay introduced by security? We consider one *beacon period*, i.e. γ^{-1} seconds, as the time unit, as specified by transportation safety requirements.

The BP and HP schemes use two general message types, according to the induced security communication overhead: *SHORT* messages carrying a node signature, and *LONG* messages carrying a node signature and certificate. Each node transmits one *LONG* message every α *SHORT* messages, with β additional consecutive *LONG* messages sent upon a pseudonym change.

The processing load at some node V_R depends on the number of packets it needs to verify. This consists primarily of signature verifications for essentially all received beacons, as they carry safety-related information. In a given slot, if V_R has N neighbors in range, it should validate $O(N)$ messages per time unit. Due to Optimization 1, V_R needs to validate the certificate signature only the first time it receives it from each neighbor. In contrast, V_R generates only one signature per time unit, and for HP specifically it generates one group signature per pseudonym lifetime.

N nodes V_i , $i = 1, \dots, N$, produce messages at an aggregate rate λ , and V_R processes them at a rate μ . λ depends on the number of neighbors, N , the message generation rate, γ , the type of generated messages, and the reliability of communication across the wireless channel. μ depends on the choice of security primitives (and their security level) and the available on-board processing power. Thus, μ is constant for a given system configuration.

We view the system operation as a queue: We are interested in the system stability, which depends on λ , μ , and the queuing model. To identify an appropriate model, we characterize the arrival process and demonstrate that it can be closely approximated by a Poisson process. Then, we provide an analytical estimate for the average arrival rate, λ , and validate it through simulations. Finally, we apply queuing theory results, in order to answer the questions outlined above.

A. Characterization of the Arrival Process

We simulate the system and collect the message inter-arrival times at some V_R , for different setups. Then, we fit known distributions to the empirically obtained data and perform a χ^2 test to assess the quality of the fitting (p-value=0.05). We




Fig. 3: HP scheme: Comparison between analytical approximation (labeled as $M/D/1$) and simulations: Arrival rate, λ_1 , and processing delay W_1 , for *SHORT* messages, as a function of the neighborhood size N .

find that the exponential distribution fits well the empirical data; its memoryless property and the orderliness of packet reception (any node receives one packet at a time and no two or more arrivals occur simultaneously) suffice to approximate arrivals as a Poisson process. Fig. 3(a) shows the collected data and the fitted exponential distribution (superimposed the distribution with the λ approximated by Eq. 2 below): the fit is indeed good. This is valid for the type of traffic under consideration, i.e. safety beaconing, which is going to be the majority of the V2V traffic. If the type of traffic changes, the arrival process would in principle change.

We assume one processor at V_R with deterministic service times. We consider both *LONG* and *SHORT* packets in the same single queue, with no priority policy and no preemption². The queue is then a multi-class M/D/1, in this case with two classes. The average waiting time, W_i , is given by [32]:

$$W_i = t_i + \frac{\sum_{j=1}^r \lambda_j t_j^2}{2(1-\rho)} \quad (1)$$

where W_i , λ_i and t_i are the total time in queue, the arrival rate and the service time of the i -th (out of $r=2$) classes respectively, $\rho = \sum_{i=1}^r \rho_i$ and $\rho_i = \lambda_i t_i$. The queue length, L_i , can be derived from Eq. 1 and Little's law [33].

B. Estimation of λ

An estimate for λ_1 , the arrival rate for *SHORT* messages (derived in the Appendix) is:

$$\lambda_1 = NP(1 - (1 - P)^K) \quad (2)$$

with N the number of neighbors, P the average reception probability for messages (beacons), and $K = \frac{\lceil \tau \gamma / \alpha \rceil}{2}$. We focus on *SHORT* packets because they are the majority of the processing load, as explained above and in [9]. From the description of the BP and HP schemes (with Optimizations), the simulations, and the derivation, it appears that *LONG* messages have a limited impact.

In Fig. 3, we plot the analytical and simulation results, for N ranging from 4 to 48 vehicles and $\gamma = 10$ or $\gamma = 3.33$

beacons/s; we average over 1000 randomly seeded simulations. Fig. 3(b) shows how many packets V_R must process as a function of N , and that this relation is almost linear. As Table IV shows, 333 signature verifications (*SHORT* packets) per second is the maximum the node we consider here can handle: for $\lambda_1 \geq 333$ msg/s, the node would be unable to keep up and its queue of messages would grow fast. We observe that for $\gamma = 10$ (most frequent in the literature), λ_1 increases fast towards this threshold, and the message processing delay, W_1 , increases fast with N (Fig. 3(c)).

Consider an example to illustrate this: with 80 transmitting vehicles in range of V_R , $\alpha = 10$, $\beta = 0$, $\gamma = 10$, and $\tau = 60$, and the resultant $P = 0.655$ (Fig. 2(a)). We assume a highway scenario and a simple content-based optimization: V_R processes a beacon if it comes from a neighbor moving in the same direction (stream of traffic). For simplicity, if the two parts of the road are equally balanced, we consider $N = 40$ vehicles out of the 80 neighbors in range. From Eq. 2 we obtain $\lambda_1 = 264.3$ msg/s and from the simulation of the same scenario, the arrival rate would be 259.7 msg/s.

VIII. TRANSPORTATION SAFETY AND EFFICIENCY

We investigate how security affects transportation safety in *two settings*. First, we consider *pairs of vehicles*: one is in a dangerous condition transmitting an EBN message and the other is approaching, is previously unaware of the transmitting vehicle and must receive the EBN message. We analyze a fundamental metric, *the ability to be notified early*: We capture this as the distance at which the receiver is first able to validate the safety messages. The second setting we consider is more involved: We study the occurrence of collisions among vehicles in a *platoon of one hundred vehicles*, with and without the use of security; the latter serves clearly as a benchmark. We also investigate the impact of penetration rate of vehicular communication rate, to gain insight on how security affects the effectiveness of VC systems, and the EBN application in an urban setting. In terms of transportation efficiency, we investigate the performance of the DFCD application, notably the delay needed to float the information and the distance covered from the location of the event generation.

²We adopt this as a baseline approach; other policies and system models could be employed.

Here we integrate the results obtained in the previous sections; we assume that vehicles have sufficient processing power and are able to verify the signatures on all incoming packets. We average over 1000 randomly seeded simulations and present results with 95% confidence intervals. Recall that Table I summarizes parameter values.

A. Simulation Setup

We consider four-, six- and eight-lane scenarios, with vehicles placed in two opposing two-, three- and four-lane flows of traffic, respectively. This corresponds to a neighborhood N of 80, 120 and 160 vehicles respectively. Vehicles are 4-meters long and are initially uniformly randomly placed along each lane, with an average vehicle-to-vehicle distance of s meters. We focus on one lane of traffic within such a neighborhood of N vehicles, which changes mildly because of mobility. In the two-vehicle setting, there is a small initial “gap” in one lane, depending on the initial spacing of the pair of vehicles; e.g., when they are at 200 meters, there are initially 10 vehicles less present, or in other words 70 (110 or 150) vehicles instead of 80 (120 or 160). The vehicle velocities are initially random with an average v , unless stated otherwise; velocities are adapted according the VC system functionality and, in the platoon setting, upon visual contact with the preceding vehicle’s brake lights. Vehicles do not change lanes during the simulation, and they process messages originating from vehicles in the same traffic flow (i.e., with the same heading).

We consider an Emergency Braking Notification (EBN) application, with one vehicle in an emergency situation that brakes and starts the transmission of EBN messages. Braking has two effects: (i) it turns on the vehicle’s rear red lights that visually warn drivers within range of sight (which depends on the simulated weather conditions), and (ii) it triggers the transmission of EBN warning messages. Besides warning other vehicles, an EBN-warned vehicle warns its driver to start braking shortly afterwards. We model driver reaction times as a result of VC-enabled and visual warnings, with a random variable uniformly distributed between 0.75 and 1.5 s. We model weather conditions by setting vehicle braking capabilities and visibility conditions; for example, on a wet road, braking is possible at a rate of 4 m/s^2 and a driver can see up to 30 m. Our simulation conditions agree with related transportation engineering work, e.g., [34], [35].

We also consider a Decentralized Floating Car Data (DFCD) application [36], concerned with how effectively data generated by one vehicle can propagate to an area and a platoon. We consider a highway scenario (platoon of 100 vehicles, with average velocity of 110 Km/h and average spacing of 30 m), with the first (heading) vehicles in the platoon detecting a situation (e.g., traffic jam, road work, etc.) and then starting broadcasting a related notification. Each receiving vehicle will relay a notification with the same content as long as it receives it from at least three distinct vehicles. We consider the DFCD running in parallel with the safety beaconing, and piggy-back DFCD notifications on beacons. With 100 randomly seeded runs of 20 sec of simulated time (after the warm-up of 60 sec), we study the effect of security on three parameters: (i)

the time for the notification to reach 90% of the platoon, (ii) the time for it to “cover” a distance of 1000 m from the origin of the vehicle first detecting the situation (event), and (iii) the maximum distance the notification propagated.

Two-Vehicle Setup: We consider one transmitter, V_t , and one receiver, V_R , at an initial distance of d meters, with V_R always behind V_t in the same lane and with a velocity for V_R higher than that of V_t . V_R moves at a constant relative speed Δv with respect to V_t , without any other vehicle in between. For simplicity, we elect s to be such that it is less than or equal to the nominal communication range at the beginning of each simulation. We choose two setups, one with $\Delta v = 20 \text{ Km/h}$ and $s = 150 \text{ m}$, and the other with $\Delta v = 35 \text{ Km/h}$ and $s = 200 \text{ m}$. We evaluate how the optimization parameters, α and β , affect the distance, D , at which V_R receives the first $K_{V_t}^i, \Sigma_{CA,V}(K_V^i)$ from V_t . In this setting, we wish to test the ability of the secure VC system to deliver safety information, especially under challenging conditions; e.g., V_R is very close to V_t when the latter changes to a new pseudonym (and private key). Therefore, in order to evaluate the VC performance alone, we do not consider the rear red lights of V_t ; which, would naturally warn the driver of a V_R in line of sight.

Vehicle Platoon Setup: We focus on a platoon of one hundred cars along a single lane moving with similar velocities, denoted as V_1 to V_{100} , with V_1 for the vehicle at the front and V_{100} at the rear of the platoon. We utilize values of $s = 20 \text{ m}$, velocities on the average $v = 80 \text{ Km/h}$. We analyze how many collisions occur when the leading vehicle V_1 brakes for an emergency and starts sending EBN messages. Once some V_i , with $i > 1$ receives the warning, it starts sending EBN messages itself. As proposed in [37]–[39], when V_i receives a warning from a V_j with $j > i$, it stops transmitting warnings, because it assumes that at least one vehicle behind V_i has already been warned. In this setting, we consider rear red-light warnings. For the urban environment, we consider multi-lane settings and platoons of vehicles that move more slowly and are closer to each other than in the highway settings (average velocity of 45 Km/h and a vehicle distances of 10 m, on a wet road). We choose pseudonym lifetime $\tau = 60 \text{ s}$. We consider the first 60 s of the simulation time as a warm-up period, during which no emergency conditions arise. This approximates a realistic situation: When an emergency arises, vehicles have already validated (identified) some of their neighbors and can thus immediately accept their warnings. The simulation concludes when all vehicles in the platoon are immobile, with V_1 not moving further after its emergency braking.

First, we consider scenarios where all vehicles are equipped with VC systems. Intuitively, full VC deployment can lead to better safety, but it also corresponds to more strenuous conditions (processing and communication overhead). Nonetheless, VC will be deployed gradually, over a period of several years. Thus, we define the *penetration rate*, pr , as the fraction of VC-enabled vehicles, and we analyze the system behavior as a function of pr . Equipped vehicles behave as described above, whereas non-equipped vehicles rely only on visual means (the red lights of the preceding vehicle) to detect emergencies. We analyze this scenario with 4 lanes of traffic, α equal to 1 or




Fig. 4: Distance D , at the time of reception of the first certificate at the trailing vehicle, V_R , as a function of α , β .

10, β to 0, γ to 10 or 3.33 beacons/s, and pr from 0.05 to 1.

B. Two-Vehicle Simulation

Impact of pseudonym change on safety: In Fig. 5(a), the distance D at which V_R receives the first certificate $\Sigma_{CA,V}(K_V^i)$ decreases as the Certificate Period α increases: If a *LONG* message from V_t is missed, V_R has a chance to receive the next one only after α additional beacons from V_t . Nonetheless, we observe that messages from V_t can be validated in all cases before the distance becomes dangerously small.

Missing a new pseudonym could be dangerous if V_R (and in general any vehicle) is close to V_t at the time of pseudonym change and has high positive relative speed (i.e., approaches fast V_t). To capture such situations, we fix $s = 30$ m and we vary $\Delta v = 10, 20$ and 50 Km/h (Fig. 4(b)): The effect of α remains, but we also observe that with a higher Δv , the drop in the reception distance with α is faster. Overall, pseudonym switching can be risky if it happens when vehicles are close to (fast approach) each other.

Optimization 3, not used so far, can address this problem, as shown in Fig. 4(c), for $\alpha = 10$ and with β varying from 0 to 10. Even a single “pushed” message ($\beta = 1$) enables reception within roughly 2 meters after the pseudonym change, regardless of speed and initial distance (clearly, the actual reception distance depends on those parameters). Increasing redundancy, that is setting β beyond 3, does not improve robustness any further. We observe in Fig. 4(c) that even $\beta = 1$ is near-optimal.

Intuitively, this is because the probability of receiving a *LONG* message when V_R and V_t are very close to each other is relatively very high. If these two nodes are far apart, the probability of successful reception is relatively low, but there will be several opportunities (subsequent transmissions) for V_R to receive a *LONG* message from V_t . More precisely, let Z be a discrete random variable for the slot in which a *LONG* message is first received (during the lifetime of the pseudonym). If $E(Z)$ is the mean value, the average distance from V_t that the first *LONG* is received by V_R is $D = E(Z)\gamma^{-1}\Delta v$. V_t transmits *LONG* packets for β consecutive slots, and every $k\alpha$ slots. The probability that a *LONG* packet is first received at some j slot is simply the probability it is not received at any of the $i < j$ slots

and it is received at the j -th, where i, j take values from $I = 1, 2, \dots, \beta, \beta + 1, \dots, \beta + 1 + \alpha, \dots$:

$$P_j = P(j) \prod_{k \in I} (1 - P(k)) \quad (3)$$

The probabilities at each slot, $P(k)$, differ, as they depend on the distance of the two nodes (and their neighborhood more generally). Using values for different distances obtained experimentally (Sec. VI and [9]), we compare Eq. 3 to the simulation results, for $s = 30$ m, $\Delta v = 20$ Km/h, $\alpha = 10$, 8 lanes of traffic and $\gamma^{-1} = 100$ ms. Fig. 4(c) shows D calculated with the help of Eq. 3 and the experimental data.

C. Platoon Analysis

1) *Fully deployed VC:* Fig. 5(a) shows the percentage of vehicles, within a platoon of one hundred, which crash as a function of α . As a reference, we simulate the VC system and EBN application without security. In the absence of V2V communications, 80%-100% of vehicles crash; for the same scenarios, safety messaging reduces the number of crashes to approximately 10% of all vehicles. Then overall, as expected, security increases crashes compared to unsecured VCs: the increased network overhead and protocol restrictions on which alert messages can be validated delay the reception of valid EBN messages.³

However, we observe that the secure VC protocol parameters affect the safety application. We observe first a decrease in the average fraction of crashes, as α increases, and then a slow increase as α increases further. This is due to two competing factors: The increase of α reduces the channel load and thus increases the per-packet reception probability, but the authentication delay for a receiver missing a *LONG* packet also increases; e.g., for $\alpha = 50$, the authentication delay is (at least) 5 s. Optimization 3 reduces the number of crashes with respect to the non-optimized protocol, with the same value for α , as it adds negligible overhead but manages to reduce the authentication delay, as explained above.

Fig. 5(b) shows the percentage of vehicles that collide in an urban setting. Again, we observe that security has a limited

³For non-secured VC, the x-axis, the α , is not a parameter that affects its operation. This is why the corresponding curve is essentially flat, with minor variability due to the randomly seeded simulation scenarios.




Fig. 5: Transportation safety with secure VC: Collisions as a percentage of vehicles, in an emergency situation; varying road conditions, message rates, protocol parameters, and VC technology penetration rate; HP scheme, $\gamma = 10$ beacons/s.

effect on the non-secured EBN application performance. The difference from the highway scenario is the way security parameters affect performance. This is a dense VC network, but vehicles move with lower velocities and they are more densely packed. The lower inter-vehicle distances lead to higher reception probability, even though the channel remains relatively heavily loaded; thus, β does not give a significant advantage. Low inter-vehicle distances (high reception probability) result in an increase in vehicle collisions (crashes) only for high α (higher authentication delay). This is analogous to the highway observations; the difference is that low vehicle-distances (lower driver reaction margins) counter balance low velocities.

2) *Effect of secure VC penetration rate on safety:* Fig. 5(c) shows what happens in a scenario, with average vehicle spacing of 40 m, speed of 65 Km/h, braking capability of 6 m/s^2 and visual range of 70 m (modeling dry road conditions and good weather). First, we observe a percentage of crashes between 25-30% and none (lower than that for the earlier more challenging scenarios). The curves for different γ values are well separated, with more crashes for $\gamma = 3.33$, which indicates that this beacon frequency would not be sufficient to warn drivers (although the reception reliability is higher (Fig. 2(a)). Tuning α affects the number of crashes when pr is in the range of 40-80 %; for $\alpha = 10$ we observe an increase variable from 15 to 40 % compared to the case with $\alpha = 1$. Conversely, if VC has relatively low ($pr < 0.4$) or high ($pr > 0.8$) penetration, security optimizations have a limited impact.

3) *Effect of secure VC on efficiency applications:* Fig. 6 shows the communication performance for the DFCD application, operating on top of, thus being constrained by, the more resource-intensive safety beaconing.⁴ We observe that security parameters affect the performance. More so for α , which yields the best performance for all three metrics for values close to $\alpha = 10$. Increasing α reduces communication overhead, increases reliability, and reduces processing delays. For high α values (e.g., above 20), there is a mild deterioration performance due to the increase in authentication delay.

⁴We do not measure the effect the transportation efficiency application has (e.g., on average trip times) as this would require a global system consideration.

	BP	GS-I	GS-II
RL size	144R KB	32R B	32R B
Revocation check cost	1 μ s.	15R ms.	30 ms.
Hash table construction	0.1R s.	-	1.5R s.
Memory requirements	140R KB	-	38.4R KB

TABLE VI: Indicative values for revocation check costs as a function of R , for $C_p = 15$ ms, $k=100$, $\tau=60$ s.

Intuitively, the parameter values affect how the network can relay the DFCD notifications. If nodes do not have the credentials of the neighbors when they receive a notification, especially as they need at least three distinct valid notifications, they will be slow in relaying it. The effect of β relates to that: with less frequent retransmissions of LONG packets, loss upon a pseudonym change would incur further delay. $\beta = 1$ remedies this, keeping stable performance for all three metrics, even for high α .

Comparing the secured to the unsecured VC system DFCD performance, we see tangible deterioration because of security: The same maximum area is covered (with an approx. 1.5% deterioration only), but with approximately 42% higher delay to cover 1000 m behind the event (in the sense of the traffic flow), and approximately 30% higher delay to cover 90% of the platoon nodes. This is because the scenario is challenging (multi-lane, heavily loaded setting), exacerbated by the need to have 3 distinct copies of the DFCD notification. With the expected latency per hop varying per DFCD message type, but typically in the order of 500 ms, we see that delays of 1.2 sec for 1000 m (approx. five hops in our setting) for the secure VC system are satisfactory; so is the delay for the 90% platoon coverage (approx. 3km in our setting), and the maximum covered area (the total of the platoon vehicles).⁵

IX. REVOCATION

Here we discuss the revocation costs, based on the use of Revocation Lists (RLs). This is a largely orthogonal problem to

⁵We did not consider here propagation of the information by vehicles in the opposite traffic flow, or geocasting, techniques assumed in place in order to achieve expected areas of coverage in the order of 10 or 20km (depending on the nominal communication range).




Fig. 6: Decentralized Floating Car Data (DFCD); communication performance.

this investigation and out of the scope of this paper; moreover, there are several unknown parameters and factors in terms of the instantiation of a revocation solution. Nonetheless, in order to provide a complete picture, we consider the revocation overhead for each of the security schemes considered. We did not consider revocation in Sec. VII, but we do this in Sec. X: it suffices to add the revocation processing delay to the one for validating *LONG* messages (the first-received *LONG* per node and τ for BP and HP).

The basic difference between BP and HP (and GS) schemes is that the former deals with short-term keys and the latter with long-term ones. The number of vehicles that would be revoked is not currently known and it is hard to estimate, because it would depend on policy decisions, the size of the system in each region, among other currently unknown aspects. Here we denote the number of revoked vehicles as R . Then, for the HP (and GS) scheme, the size of the revocation list would be $|RL_{HP}| = R$. Whereas, for BP, $|RL_{BP}| = c \times R$, where c is essentially the number of temporary keys each vehicle holds at the time of its revocation.

c can be a large number, in the order of 10^4 to 10^6 ; but, again, its actual value depends on factors out of the scope of this paper: for example, the ability of vehicles to have frequent access to a trusted third party to obtain their short-term certificates, or the autonomy of vehicle policies mandate. We can consider that c corresponds approximately to the number of pseudonyms a vehicle obtains at a “*pseudonym refill*”, and we can assume that pseudonyms are valid only between two consecutive refills. Then, a revoked node running BP would be unable to obtain a new set of pseudonyms [14], and RL_{BP} would include only the pseudonyms granted at the last refill. Consider an example for c : with $\tau=60$ s, in one day a vehicle V would “consume” 1440 pseudonyms. Assuming that pseudonym refills take place once per month, then $c = 43200$; if the refills were made once per year, then $c = 518400$. In the rest of this paper, to provide illustrative examples, we assume that on the average $c \approx 10^4$.

The cost to verify whether a pseudonym is revoked is the cost of a lookup into the RL . This can be achieved in constant time, e.g. by using a hash table. In this case, the construction of the data structure is proportional to $|RL_{BP}|$, and it must be performed every time a new RL_{BP} is received. The required memory is also $\approx |RL_{BP}| \times E_{BP}$ bytes, where E_{BP} is the

size of one entry in the hash table. E_{BP} , composed by a serial number and a revocation date, is 14 bytes [1]. Thus, for each revoked vehicle (i.e. long-term identity) with c pseudonyms, at least 140 KB would be needed.

For GS and HP schemes, we extract relevant data from [22], [27]. Each entry in RL_{HP} is a *revocation token* of 32 bytes (Note: for simplicity, we use interchangeably the subscript GS for HP in terms of revocation.). Then, we consider two related revocation methods proposed in [22]: the first one, we term *GS-I*, incurs processing costs proportional to the RL size; the second one, *GS-II*, has fixed costs independent of the RL size, but it might allow the linking of some Group Signatures [22].

GS-I: The revocation tokens, each $E_{GS}=32$ bytes, are used directly for the revocation check process. The cost to verify one entry is $C_p \times |RL_{GS}|$, where C_p is the cost of computing one bilinear map. Group signatures of not-revoked nodes cannot be linked under any circumstances, but checking if a signer is revoked requires a traversal of the entire RL_{HP} (in other words, it is linear in the number of revoked vehicles).

GS-II: The basic difference from *GS-I* is the calculation of the Group Signatures, which include some intended recipient, S , a random positive integer r , now chosen by the signer to have a value less than a security parameter k . As it will be explained in further detail below, this construction allows S to pre-compute k revocation values and check the revocation status of the signer through a simple look-up; if, however, k were low, the signer might be forced to re-use r values, and these group signatures from that signer could be linked.

As the safety beacons are broadcasted, we need to adapt the scheme to the VC context: We redefine S to be $S = \langle G, T \rangle$, where G indicates a geographical area and T a time interval. Essentially, in a given area and time, every recipient can perform the fast revocation status check. For practical reasons, G and T can be coarsely defined, so that receiving nodes can easily determine the appropriate values (e.g., with the help of their on-board clock, GPS receiver, or other localization means with the help of terrestrial infrastructure).

Upon reception of a new RL_{GS} , a verifier V_R in S pre-computes and stores the k revocation values for each entry in the RL_{GS} , at a cost of $2C_p$ per entry. The cost to build this data structure is then $2C_p \times R \times k$. Upon receipt of a $\Sigma_{CA,V}(K_V^i)$, the verifier V_R performs a lookup into the table and if no match is not found (i.e. the signer is not revoked)




Fig. 7: Processing delay in future secure VC systems, with processing acceleration, as a function of network size.

it validates the signature. The cost to verify if the sender is revoked is $2C_p$, plus the lookup cost, which is negligible compared to C_p . The memory needed is approximately $|RL_{GS}| \times k \times E_{GS}$ bytes. E_{GS} is the result of one pairing computed from the revocation token, and its size is 384 bytes.

For a given $S = \langle G, T \rangle$, the value of k should be chosen such that a single sender is not forced to use the same r twice or more. Basically, it should be $k \geq \lceil |T|/\tau \rceil$. However, high k would increase the pre-computation costs, which also depends on how the RL_{GS} changes over time. Investigating trade-offs due to chosen values, e.g., τ , $|T|$, is left as future work. In order to provide a numeric example, we fix $k=100$, which corresponds to $|T|=1\text{h } 40\text{min}$, and we summarize the results in Table VI, assuming $C_p = 15\text{ms}$, $k=100$, $\tau=60\text{s}$, and the basic operations on the hash table, such as memory copy and data lookup, to be $1 \mu\text{s}$.

Clearly, the BP scheme incurs the minimum computational overhead but it has by far the longest RL. Moreover, the GS-I method could be cumbersome to apply, especially for the platform we considered here and for sizeable RLs. Then, the cost of GS-II is independent of R but it remains higher than that of BP. Nonetheless, GS-I and GS-II have the advantage of much smaller RLs, and thus require lower memory and bandwidth for storage and transmission.

Finally, recall that for the GS scheme, the revocation status check must be performed for each message; whereas for the HP scheme, the check is needed only once per previously unseen pseudonym. It is also important to note that the revocation of a node implies its anonymity is lost; then, any entity that has a transcript of its past transmissions in a given area, can use the corresponding revocation token and identify which messages in the transcript were sent by the revoked node.

X. DISCUSSION AND FUTURE WORK

The introduction of anonymous authentication, for infrequent operations, as in HP, offers the advantage of simpler administration (much lower number of credentials, on-the-fly generation), but it increases processing cost. This is so for each GS-related operation (LONG packets for HP). Anonymous authentication brings forth another challenge: a misbehaving node could abuse its anonymity, generate and sign multiple pseudonyms, and thus appear as multiple nodes, also known as a Sybil node behavior.

Each legitimate node that monitors its neighborhood and relies on its own sensor readings (e.g., infra-red or radar transceivers) could weed out such Sybil identities (pseudonyms), as pointed out in [40], or it could detect other inconsistencies, e.g., between the apparent density vehicles (inflated by a Sybil node) and the average velocity (in principle dropping with increased vehicle density). But there are proactive defenses to contain Sybil misbehavior. First, the invocation of the GS primitive can be regulated by trusted hardware, e.g., the Hardware Security Module (HSM) of the SeVeCom architecture [2]. A different approach is that of utilizing primitives that limit the number of GS generations per unit of time; e.g., [41] can be set to allow each node to generate one GS per pseudonym lifetime. If a Sybil node exceeds this, any such GS (and thus pseudonym) will be linkable by any receiving correct node.

The addition of new features, e.g., a time-limited GS generation, will have additional resource requirements. This is a more general concern. But it is natural to expect that on-board processing capabilities will also improve. For example, there are already secure VC boxes that achieve nearly a ten-fold increase (www.escrypt.com) in terms of processing (thanks to cryptographic hardware acceleration) compared to our benchmark platform.

In Fig. 7, we consider processing delay for various processing acceleration factors over our benchmark platform. A factor x decreases processing delay per cryptographic operation by $1/x$. We demonstrate how our model of Sec. VII can yield resource provisioning as the scale increases. Fig. 7 shows that in order to keep the per-safety-beacon delay flat for a large neighborhood (160 vehicles within range, verification from 80 ones), $x = 2$ is necessary. This is sufficient even for real-time revocation (i.e., 30ms delay for each LONG for the benchmark platform, and 15ms for $x = 2$), or time-limited GS scheme (e.g., for [41], we approximate that a time-limited GS-verification delay is approximately 2.5 times than of the “plain” GS). Note the progressive increase of the delay (y-axis) and recall that x speeds up not only the relatively infrequent LONG processing but also the majority of the lower-cost SHORT processing.

Future work will include additional characteristics of the transportation environment, e.g., traffic lights, change of routes other safety applications (corner-collision avoidance), alter-

native communication technologies, and any roles of infrastructure. Alternative security and privacy enhancing mechanisms is another future work direction, including alternative cryptographic primitives and functionality and further implementations and benchmarking. A different but related issue: general privacy considerations, across the data-handling “ecosystem”, beyond the mostly location-privacy related VC-centric considerations.

XI. RELATED WORK

The use of pseudonyms was first envisioned in [42] and more recent works considered their use in the context of VCs, e.g., [14], [43], [44]. More generally, several recent works are concerned with different aspects of security and privacy of vehicular networks: They either outline challenges [45], [46], describe particular attacks [47], [48] or more general attack overviews [49], [50], propose mechanisms [40], [44], [51], [1], and schemes for revocation [2], [13], [52]. [53], and the recent follow-up [54] that provides performance evaluation results, combine public and symmetric key cryptography to authenticate messages and is complementary to our work.

The idea of pseudonym self-generation for ubiquitous computing is proposed, independently of our work, in [55]. More recently, [56] applied that crypto-system to VANET. These works do not consider all the system-level issues we consider in this work, such as certificate distribution and application robustness. Our findings and mechanisms also apply to their work, complementing and extending it.

An alternate approach to reducing packet overhead and computation efforts is presented in [57]: it proposes that a signer attaches its certificate to messages only when it detects a change in its neighborhood, with such changes detected from beacons. In dense topologies, the results of [57], although obtained in less realistic conditions, are comparable to ours.

Two recent works propose to use bilinear pairings to provide privacy in VANET. The approach of [58] is similar to our GS scheme, thus it would be cumbersome if not impossible to apply for safety beaconing. [59] employs a mix of traditional public-key cryptography and bilinear pairings; this bears some resemblance to our HP scheme, but it is mainly limited by its strong reliance on the presence of RSUs, not envisioned to be densely present in most, if not all, deployments.

A few other papers [30], [37]–[39] propose and analyze VC-based algorithms for transportation safety, and transportation engineers studied how to reduce car collision chains [34], [35]. But the combined study of transportation safety applications enabled by VC and the effect of security overhead was not considered.

XII. CONCLUSION

We have analyzed the effect of security on the VC system effectiveness, notably a safety and an efficiency application. We have provided a framework to analyze the performance of secure VC systems, along with schemes that reduce the complexity and the overhead of security, we considered multiple system operation dimensions and identified interdependencies of various factors. We strongly believe that systematic

evaluation of the overall performance is critical, especially for pervasive computing systems that are tightly coupled to their users. As security and privacy are paramount for those systems, yet they incur significant overhead, designs should be validated, to show that the secured systems can be effective as envisioned and needed. We do this for a system as complex as VCs: This being the first work taking this approach, we aspire to produce and see further results for a technology that can be very widely deployed in the near future.

APPENDIX

We derive here the approximation for λ_1 , the average arrival rate for *SHORT* messages, i.e. Eq. 2 in Sec. VII. We consider a set of N transmitters V_i and one receiver V_R , all running the protocol with the same configuration (i.e. all V_i use the same beacon interval γ , the same pseudonym lifetime τ , and the same Certificate Period α , and $\beta = 0$ as a minor simplification due to the low effect of β on overhead for the values recommended by the findings in Sec. VIII). As discussed in Sec. VIII and Sec. VI, each message is received with a different probability depending not only on the overall setup but also on the (fast changing) distance between sender and receiver; obtaining these values is far from trivial.

Here, we make a few simplifying assumptions: We consider some slot, t , and assume that all N transmitters send a beacon, *SHORT* or *LONG* during that slot. Recall that each *LONG* beacon entails a calculation equivalent to a *SHORT* beacon (due to the verification of the ECDSA signature), plus an additional overhead when V_R receives a *LONG* one with a new pseudonym (due to the GS verification). We assume that each beacon is received with probability P , independently from all other $N - 1$ beacons. (The probability of reception is for example the average of the probabilities of reception at different distances for the given neighborhood, obtained from Sec. VI).

For some sending node V_i , its message will be verified by V_R with probability $p \times \text{Prob}\{V_i\text{'s } LONG \text{ was already received}\}$. This is determined by how many times the given V_i transmitted a *LONG* message (as we assume that the reception of *LONG* packets is also with probability P). By the definition of the scheme (HP or BP), during τ seconds, a pseudonym lifetime, V_i transmits $\lfloor \tau\gamma/\alpha \rfloor$ *LONG* packets.

The number of *LONG* packets from V_i that V_R observed (i.e. could potentially receive being in range) depends on various factors beyond this model (e.g., mobility, individual vehicle trajectories, road shape, communication obstructions). Here we make one final simplifying assumption: Each V_i is on the average “half-way through” its current τ , thus, it has on the average transmitted $K = \frac{\lfloor \tau\gamma/\alpha \rfloor}{2}$ *LONG* beacons.

Then, the $\text{Prob}\{V_i\text{'s } LONG \text{ was already received}\}$ is essentially the $\text{Prob}\{V_i\text{'s } LONG \text{ was received at least in one of } K \text{ tries}\}$; this is equal to $1 - \text{Prob}\{V_i\text{'s } LONG \text{ was received in none of } K \text{ tries}\} = 1 - (1 - P)^K$. As all of the N nodes sent out a packet, in each slot, on average V_R will receive

$$\lambda_1 = NP(1 - (1 - P)^K) \quad (4)$$

to process. This completes the derivation of Eq. 2.

REFERENCES

- [1] IEEE1609.2, "IEEE Trial-Use Standard for Wireless Access in Vehicular Environments - Security Services for Applications and Management Messages," Jul. 2006.
- [2] P. Papadimitratos, L. Buttyan, T. Holczer, E. Schoch, J. Freudiger, M. Raya, Z. Ma, F. Kargl, A. Kung, and J.-P. Hubaux, "Secure Vehicular Communication Systems: Design and Architecture," *IEEE Comm. Mag.*, vol. 46, no. 11, Nov. 2008.
- [3] F. Kargl, P. Papadimitratos, L. Buttyan, M. Müter, B. Wiedersheim, E. Schoch, T.-V. Thong, G. Calandriello, A. Held, A. Kung, and J.-P. Hubaux, "Secure Vehicular Communication Systems: Implementation, Performance, and Research Challenges," *IEEE Comm. Mag.*, vol. 46, no. 11, Nov. 2008.
- [4] "NoW: Network on Wheels," uRL: <http://www.network-on-wheels.de/>.
- [5] "The Car-to-Car Communication Consortium." [Online]. Available: <http://www.car-to-car.org>
- [6] "The eSafety eSecurity Working Group." [Online]. Available: http://www.esafesupport.org/en/esafety_activities/esafetyworking_groups/eseurity.htm
- [7] "DSRC: Dedicated Short Range Communications," <http://grouper.ieee.org/groups/scc32/dsrc/index.html>.
- [8] G. Calandriello, P. Papadimitratos, J.-P. Hubaux, and A. Liou, "Efficient and Robust Pseudonymous Authentication in VANET," in *VANET '07*, Montreal, Canada, Sept. 2007.
- [9] P. Papadimitratos, G. Calandriello, A. Liou, and J.-P. Hubaux, "Impact of Vehicular Communication Security on Transportation Safety," in *MOVE '08*, Phoenix, AZ, USA, Apr. 2008.
- [10] "IEEE P802.11p/D3.0, Draft Amendment for Wireless Access in Vehicular Environments (WAVE)," July 2007.
- [11] "Dedicated Short Range Communication at 5.9 GHz Standards Group." [Online]. Available: <http://www.iteris.com/itsarch/html/standard/dsrc5ghz-b.htm>
- [12] "ISO TC204 Working Group 16." [Online]. Available: <http://www.calm.hu/>
- [13] P. Papadimitratos, G. Mezzour, and J.-P. Hubaux, "Certificate Revocation List Distribution in Vehicular Communication Systems," in *VANET '08*, San Francisco, CA, USA, Sept. 2008.
- [14] P. Papadimitratos, L. Buttyan, J.-P. Hubaux, F. Kargl, A. Kung, and M. Raya, "Architecture for Secure and Private Vehicular Communications," in *ITST'07*, Sophia Antipolis, France, Jun. 2007.
- [15] D. Chaum and E. van Heyst, "Group Signatures," in *EUROCRYPT '91*, Brighton, UK, Apr. 1991.
- [16] G. Ateniese and G. Tsudik, "Group Signatures à la carte," in *SODA '99*, Baltimore, MD, USA, Jan. 1999.
- [17] P. Syverson and S. Stubblebine, "Group Principals and the Formalization of Anonymity," in *FM '99*, Toulouse, France, Sept. 1999.
- [18] E. Brickell, J. Camenisch, and L. Chen, "Direct Anonymous Attestation," in *CCS '04*, Washington DC, USA, October 2004.
- [19] D. Boneh, X. Boyen, and H. Shacham, "Short Group Signatures," in *Crypto '04*, Santa Barbara, CA, USA, Aug. 2004.
- [20] M. Bellare, D. Micciancio, and B. Warinschi, "Foundations of Group Signatures: Formal Definition, Simplified Requirements and a Construction based on Trapdoor Permutations," in *Adv. in Cryptology*, May 2003.
- [21] M. Bellare, H. Shi, and C. Zhang, "Foundations of Group Signatures: The Case of Dynamic Groups," in *CT-RSA '05*, San Francisco, CA, USA, Feb. 2005.
- [22] D. Boneh and H. Shacham, "Group Signatures with Verifier-Local Revocation," in *CCS '04*, Washington DC, USA, Oct. 2004.
- [23] IEEE 1363a-2004, "IEEE Standard Specifications for Public-Key Cryptography - Amendment 1: Additional Techniques," 2004.
- [24] "The CVIS project, <http://www.cvisproject.org/>."
- [25] "OpenSSL," <http://www.openssl.org>.
- [26] M. Brown, D. Hankerson, J. Lopez, and A. Menezes, "Software Implementation of the NIST Elliptic Curves Over Prime Fields," in *CT-RSA 2001*, San Francisco, CA, USA, April 2001.
- [27] N. Kobitz and A. Menezes, "Pairing-Based Cryptography at High Security Levels," Cryptology ePrint Archive, Report 2005/076, 2005.
- [28] M. Nakagami, "The m-distribution, a General Formula of Intensity Distribution of the Rapid Fading," in *Statistical Methods in Radio Wave Propagation*, W. G. Hoffman, Ed. Oxford: Pergamon, 1960.
- [29] Q. Chen, F. Schmidt-Eisenlohr, D. Jiang, M.-T. Moreno, L. Delgrossi, and H. Hartenstein, "Overhaul of IEEE 802.11 Modeling and Simulation in ns-2," in *MSWiM '07*, Chania, Greece, Oct. 2007.
- [30] Q. Xu, T. Mak, J. Ko, and R. Sengupta, "Vehicle-to-Vehicle Safety Messaging in DSRC," in *VANET '04*, Philadelphia, PA, USA, Oct. 2004.
- [31] M. Torrent-Moreno, D. Jiang, and H. Hartenstein, "Broadcast Reception Rates and Effects of Priority Access in 802.11-based VANETs," in *VANET '04*, Philadelphia, PA, USA, Oct. 2004.
- [32] M. Hazewinkel, Ed., *Encyclopaedia of Mathematics*. Springer-V., 2002.
- [33] W. Feller, *An Introduction to Probability Theory and Its Applications, Volume 1*. Wiley, 1968.
- [34] S. Diwan, B. Dalla Chiara, and F. Deflorio, "Effect of Vehicle to Vehicle and Vehicle to Infrastructure Communication Systems on Transportation Safety," in *AATT '06*, Chicago, IL, USA, Aug. 2006.
- [35] B. Dalla Chiara, F. Deflorio, and S. Diwan, "Communication Among Vehicles and Infrastructures: Evaluation of Outcomes on Road Safety," in *ITS '08*, Geneva, Switzerland, Jun. 2008.
- [36] "Intelligent Transport Systems (ITS), Vehicular Communications (VC), Basic Set of Applications, Definitions," ETSI TR 102 638 V1.1, Tech. Rep., Jun. 2009.
- [37] S. Biswas, R. Tatchikou, and F. Dion, "Vehicle-to-Vehicle Wireless Communication Protocols for Enhancing Highway Traffic Safety," *IEEE Communications Magazine*, vol. 44, Jan. 2006.
- [38] T. ElBatt, S. K. Goel, G. Holland, H. Krishnan, and J. Parikh, "Co-operative Collision Warning using Dedicated Short Range Wireless Communications," in *VANET '06*, Los Angeles, CA, USA, Sept. 2006.
- [39] X. Yang, J. Liu, F. Zhao, and N. Vaidya, "A Vehicle-to-Vehicle Communication Protocol for Cooperative Collision Warning," in *MobiQuitous '04*, Boston, MA, USA, Aug. 2004.
- [40] P. Golle, D. Greene, and J. Staddon, "Detecting and correcting malicious data in VANETs," in *VANET '04*, Philadelphia, PA, USA, October 2004.
- [41] J. Camenisch, S. Hohenberger, M. Kohlweiss, A. Lysyanskaya, and M. Meyerovich, "How to Win the Clone Wars: Efficient Periodic n-Times Anonymous Authentication," in *CCS*, Alexandria, VA, USA, Oct. 2006.
- [42] D. Chaum, "Security Without Identification: Transaction Systems to Make Big Brother Obsolete," *Comm. of the ACM*, vol. 28, no. 10, pp. 1030-1044, 1985.
- [43] M. Gerlach, A. Festag, T. Leinmuller, G. Goldacker, and C. Harsch, "Security architecture for vehicular communication," in *WIT 2007*, Hamburg, Germany, March 2007.
- [44] M. Raya and J. Hubaux, "Securing Vehicular Ad Hoc Networks," *Journal of Computer Security*, vol. 15, no. 1, pp. 39 - 68, 2007.
- [45] M. El Zarki, S. Mehrotra, G. Tsudik, and N. Venkatasubramanian, "Security Issues in a Future Vehicular Network," in *EW '02*, Florence, Italy, Feb. 2002.
- [46] B. Parno and A. Perrig, "Challenges in Securing Vehicular Networks," in *HotNets-IV*, College Park, MD, USA, Nov. 2005.
- [47] M. Jakobsson, X. Wang, and S. Wetzel, "Stealth Attacks in Vehicular Technologies," in *VTC-Fall '04*, Los Angeles, CA, USA, Sept. 2004.
- [48] J. Blum and A. Eskandarian, "The Threat of Intelligent Collisions," *IT Professional*, vol. Vol. 6, pp. 24-29, 2004.
- [49] A. Aijaz, B. Bochow, F. Dötzer, A. Festag, M. Gerlach, R. Kroh, and T. Leinmüller, "Attacks on Intervehicle Communication Systems - an Analysis," in *WIT '06*, Hamburg, Germany, Mar. 2006.
- [50] P. Papadimitratos, V. Gligor, and J.-P. Hubaux, "Securing Vehicular Communications - Assumptions, Requirements, and Principles," in *ES-CAR '06*, Berlin, Germany, Nov. 2006.
- [51] M. Li, R. Poovendran, K. Sampigethaya, and L. Huang, "CARAVAN: Providing Location Privacy for VANET," in *ESCAR '05*, Cologne, Germany, November 2005.
- [52] M. Raya, P. Papadimitratos, I. Aad, D. Jungels, and J.-P. Hubaux, "Eviction of Misbehaving and Faulty Nodes in Vehicular Networks," *IEEE JSAC*, Oct. 2007.
- [53] K. Laberteaux and Y.-C. Hu, "Strong VANET Security on a Budget," in *ESCAR '06*, Berlin, Germany, Nov. 2006.
- [54] J. Haas, Y.-C. Hu, and K. Laberteaux, "Real-World VANET Security Protocol Performance," in *Globecom*, Honolulu, HI, Nov. 2009.
- [55] K. Zeng, "Pseudonymous PKI for Ubiquitous Computing," in *EuroPKI*, Turin, Italy, Jun. 2006.
- [56] F. Armknecht, A. Festag, D. Westhoff, and K. Zeng, "Cross-layer Privacy Enhancement and Non-repudiation in Vehicular Communication," in *WMAN '07*, Bern, Switzerland, Mar. 2007.
- [57] F. Kargl, E. Schoch, B. Wiedersheim, and T. Leinmüller, "Secure and Efficient Beaconing for Vehicular Networks," in *VANET '08*, San Francisco, CA, USA, 2008.
- [58] X. Lin, X. Sun, P.-H. Ho, and X. Shen, "GSIS: A Secure and Privacy-Preserving Protocol for Vehicular Communications," *IEEE TVT*, vol. 56, no. 6, Nov. 2007.
- [59] R. Lu, X. Lin, H. Zhu, P.-H. Ho, and X. Shen, "ECPP: Efficient Conditional Privacy Preservation Protocol for Secure Vehicular Communications," in *INFOCOM '08*, Phoenix, AZ, USA, Apr. 2008.