Certificate Revocation in Vehicular Networks

Maxim Raya, Daniel Jungels, Panos Papadimitratos, Imad Aad and Jean-Pierre Hubaux

Laboratory for computer Communications and Applications (LCA)
School of Computer and Communication Sciences
EPFL, Switzerland
LCA-Report-2006-006

ABSTRACT

Among civilian communication systems, vehicular networks emerge as one of the most convincing and yet most challenging instantiations of the mobile ad hoc networking technology. Towards their deployment, security is a critical factor and significant challenge to be met. In this paper, we are concerned with the problem of certificate revocation in vehicular networks, a problem of central importance for any security architecture and particularly difficult for vehicular networks. We contribute a set of protocols for efficient and effective revocation, to evict illegitimate or faulty network nodes. Furthermore, we propose a protocol that enables nodes to collectively shield themselves against faulty of malicious operation of other nodes and contribute to their eviction. We show, by means of simulations, that our solution is feasible and achieves a sufficient level of robustness in spite of the unique challenges of the vehicular networking environment.

1. INTRODUCTION

A number of recent research initiatives supported by governmental and transnational organizations aim to enhance safety and efficiency of transportation systems. Vehicular communications (VC) or vehicular networks lie at the core of these efforts, which envision applications that, for example, provide warnings on environmental hazards (e.g., ice on the pavement), traffic and road conditions (e.g., emergency braking, congestion, or construction sites), and local (e.g., tourist) information.

To enable such applications, vehicles and road-side infrastructure units (RSUs), namely network nodes, will be equipped with on-board processing and wireless communication modules. Then, vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) (bidirectional) communication will be possible directly, when in range, or, in general, across multiple wireless links (hops), with nodes acting both as end points and routers. Instantiating the *mobile ad hoc networking* (MANET) technology can be the only means to realize safety and driving assistance applications through V2V communication, while, equally important, the deployment of an omnipresent infrastructure can be impractical and too costly.

A comprehensive set of security mechanisms integrated into the VC system is critical for their deployment, especially because of the life-critical nature of the vehicular network operation.

The security architecture will make use of a trusted third party or authority that manages the identities, credentials, and cryptographic keys of all network nodes. For the rest of the discussion we denote this as the certification authority (CA). This approach is deemed appropriate, instead of an ad-hoc or web-of-trust (PGP-like) method. Rigid identity and credential management processes for vehicles and drivers have been long in place, accountability and attribution of liability will continue to be crucial, and mechanisms for access control will be necessary. This is clearly reflected on the intent of the US DoT Intelligent Transportation Systems (ITS) initiative to base its security solutions on a Public Key Infrastructure (PKI), as stated by the IEEE 1609 family of standards for Wireless Access in Vehicular Environments (WAVE) [2].

Without the appropriate certificate, network nodes are essentially unable to participate in the system operation. Nevertheless, the certificate(s) cannot be valid for unlimited periods of time after their generation. Moreover, the CA reserves the right to revoke the certificate(s) of any node, and essentially evict the node from the network.

Timely access to revocation information is important for the robustness of its operation: messages from faulty, compromised, or otherwise illegitimate (e.g., from an administrative point of view), and overall potentially dangerous, nodes can be ignored. However, revocation will be a particularly hard problem in vehicular networks.

At first, the massive scale of these systems, with an increasing 750 million vehicles worldwide today, will impose a heavy administrative load. Furthermore, each node will be each registered with (and thus obtain its certificates from) one among a large number of CAs, as the currently large number of organizations responsible for transportation matters (e.g., Departments of Motor Vehicles (DMVs)) indicates. Meanwhile, vehicles will roam freely at high velocities, and the size of the network coverage (and secure communication) area could grow arbitrarily, to include any part of the drivable world. As a result, due to the related cost, assuming an omnipresent infrastructure (e.g., coverage of highways with wireless access points) is clearly a fallacy.

Under these conditions, the problem at hand is to how to design a system that can distribute revocation information. Clearly, the road-side infrastructure can act as

the gateway of the CA to the vehicular network. However, the challenge remains: what would be the course of action when such infrastructure is unavailable or unreachable? In such case, how can legitimate nodes be protected until they obtain the revocation information on faulty nodes?

Our contribution in this paper addresses exactly this challenging problem. We design a solution comprising three protocols tailored to the VC properties and efficiency constraints: two of the protocols rely on the availability of infrastructure, and the third one compensates for sporadic connectivity to the infrastructure. We first propose a protocol to efficiently distribute revocation information, even in adverse network coverage conditions. Second, we propose a protocol that leverages on the availability of an on-board *trusted* processing and storage device: the authority issues a command to the device to erase the stored credentials, thus directly evicts the targeted to-be-revoked node.

Finally, we propose a protocol that allows the neighbors of a faulty node to detect, or at least strongly suspect, its deviation from the implemented protocols. Nodes can immediately shield themselves, ignoring messages from such a node, and then collectively report their observations to the authority, once it becomes reachable, possibly triggering the attacker's credentials revocation. We emphasize however that no group of nodes (vehicles) has the power to revoke another node. The authority is the sole entity with the right to initiate a revocation protocol. This design choice is to ensure robustness to 'black-mail' attacks, retain accountability, and yet equip nodes with a rapid reaction and self-protection tool. Indeed, our performance evaluation results show that misbehavior alerts spread to a high percentage of the attacker's neighbors, despite the very short contact time between the protocol participants.

The paper is organized as follows. In Section 2, we survey related work. In Section 3, we describe the system model. In Section 4 we define the problem statement. In Section 5 we introduce our revocation protocols. We evaluate them in Section 6. Finally, we conclude in Section 7.

2. RELATED WORK

Existing works on vehicular network security [13, 15] propose the usage of a PKI and digital signatures but do not provide any mechanisms for certificate revocation, even though it is a required component of any PKI-based solution.

Revocation has been considered mostly in the context of the wireline Internet and the design of Public Key Infrastructure (PKI) services [9]. Certificate and certificate revocation list (CRLs) formats for individuals [17], methods for their storage [4], use of a permanent identifier [14] for each system entity, and attribute certificates [7] have been standardized. Delta-CRLs, that is, periodic dissemination of incremental revocation information, and as a substitute or in addition to the CRLs, a protocol for any client to obtain the revocation status of a specific certificate in a timely manner [12] have been proposed. These fundamental notions and mechanisms are clearly relevant to this work on revocation, even though their adaptation and augmentation to include new VANET-relevant attributes (e.g., geographical region) or identities would be necessary.

Nevertheless, the design of mechanisms to disseminate the revocation information across the VC system has not been considered in the wireline Internet context (for a survey and discussion of tradeoffs see [20, 18]). Due to the network volatility and scale, the overhead of querying a server to obtain timely revocation status, assuming the server is reachable, would be impractically high. Furthermore, for the same reasons, schemes such as [19], which distributes the load of a server to a set of participating clients redundantly forwarding revocation information, would not be practical for deployment within the vehicular network, but only meaningful behind the fixed infrastructure.

Our work addresses exactly these challenges, the network volatility, the network scale in numbers of nodes and size of the coverage area, and number of involved CAs. We propose targeted revocation information dissemination, to the most relevant region of the network, precisely aimed eviction of a network entity with a help of an on-board trusted computing device, efficient dissemination of CRLs in a compressed mode, use of broader coverage one-way, CA-to-clients, channels, and a self- or to be more precise collective defense mechanism (which can also contribute information towards revocation) before revocation of faulty nodes can be confirmed.

In the context of vehicular networks, the IEEE 1609.2 Draft Standard [2] is the only reference on certificate revocation. It proposes the distribution of CRLs and short-lived certificates, but does not elaborate how to achieve this. Short-lived certificates are also proposed in [10], which relies on Merkle tree constructions but does not consider revocation. Short lifetimes essentially trade off the need for revocation at the expense of vulnerability, but such an approach is not appropriate for VC environment.[1] Frequent refreshing of certificates, exactly because the vulnerability window must be very low due to the critical nature of VC communications, will create an overwhelming load at the side CA, as well as on the network, for communication with the relevant CA.

Instantiating a CA in the context of mobile ad hoc networks has been investigated, with the distribution of its functionality to a number of servers jointly acting with the help of threshold cryptography enhancing robustness and reachability of the CA [21]. However, these schemes have not considered the problem of revocation, and more so within the VC environment. On the other hand, instantiation of the CA functionality (or part thereof) by impromptu coalitions of network nodes (e.g., [11, 6]) cannot be applicable in VC systems. Al-

[1] An exception can be context-specific credentials, allocated, for example, to a vehicle entering a highway segment and 'purchasing' access to a service. However, this is orthogonal to the problem we are considering here.

lowing any ad hoc and in general small subset of adversarial nodes to maliciously accuse and evict legitimate nodes or even admit at will nodes to the network would be an unacceptable breach of the VC system security, which mandates accountability and liability.

In the direction of the misbehavior detection, [8] describes a scheme based on statistical methods to mitigate attacks injecting false data in vehicular network. To exclude malicious nodes, the maintenance of node reputation metrics has been proposed, e.g., in [5] or [6]. We discuss further the misbehavior detection, which is orthogonal to our problem, in Section 5.1.

In contrast, we propose a distributed system for neighbor warning essentially as a "fall-back" solution, which nonetheless has added value. In case of CA unavailability (e.g., due to lack of connectivity), nodes can individually (although possibly using input from other nodes) detect faulty participants, protect themselves, warn their neighbors, and report their related findings when possible to the CA. Yet, they are not empowered to revoke the certificate of any other participant. Such an action can be taken only by the CA.

3. SYSTEM MODEL

We classify the VC entities as *nodes* and *authorities*. Focusing on the network operation, we view authorities as network entities rather than public agencies or corporations with administrative powers. Furthermore, we do not distinguish users, i.e., individuals operating vehicles, from the vehicles.

3.1 Communication model

Messages are transmitted either periodically, e.g., every 0.3 s for *safety messages*, or triggered by in-vehicle or network events. Unicast communication is clearly possible, yet a large fraction of the traffic is broadcasted across the network, with restrictions, e.g., based on the location of the sender and receivers, determining its propagation. We do not dwell on the network and application protocols, as they clearly depend on the specific instantiations.

At the data link layer, the Dedicated Short Range Communications (DSRC) protocol, based on IEEE 802.11 technology and currently towards being standardized as IEEE 802.11p [1], provides omnidirectional transmission ranges of typically 300 to 1000 m. In the rest of the paper, we assume that communication takes places over 802.11p, unless noted otherwise.

Beyond DSRC, vehicular networks can leverage on other wireless communication technologies, such as the (licensed-frequency) existing cellular networks, broadband wireless (e.g., WiMax), or the low-speed radio broadcast system used nowadays for traffic information.

We denote a subset of the network nodes as the *infrastructure*; this comprises the RSUs, i.e., short-range DSRC *base stations*, and *mobile* units. The latter can be public safety (highway assistance, fire-fighting) or police vehicles, aerial vehicles (e.g., police helicopters), or public transport vehicles (e.g., buses, trams). These nodes are assumed to be trustworthy.

3.2 Authorities

Drawing from the analogy with existing administrative processes and automotive authorities (e.g., city or state transit authorities), a large number of CAs will exist. Each of them is responsible for the identity management of all vehicles registered in its *region* (national territory, district, county, etc.). Each vehicle (and node in general) that wishes to be part of the vehicular network is registered with exactly one CA. We note that organizations such as closed highway authorities may distribute their own special-purpose credentials to vehicles (acting in a sense as CAs). However, in this work, we do not consider such cases. Finally, we note that reachability to a CA from the vehicular network is not assumed guaranteed. Infrastructure nodes serve as the gateway of the CA to/from the vehicular network. We assume that the connection of the CA to the static infrastructure nodes is over wireline secure links.

3.3 Trusted Components

Many vehicles are already equipped with hardware components and firmware that regulate or record information on their operation; examples are speed limiters, tachographs, and *event data recorders* (*EDRs*). These components are considered to be critical by manufacturers and legislators. We assume that nodes are equipped with trusted components (*TCs*), i.e., tamper-resistant hardware and firmware performing cryptographic operations and providing storage.

The TC stores the node credentials, prevents their exposure to the on-board computer, and provides an interface for the latter to request cryptographic operations on any input to the TC. We also assume that each TC stores a special-purpose public/private key dedicated for communication with the certification authority. We emphasize that this special-purpose keying material is not used for operation on any other messages and it is not related to other cryptographic keys the node may utilize. Yet, again, all cryptographic material and operations are stored and performed respectively by the TC.

3.4 Adversary model

This paper deals with revoking certificates of detected attackers. As will be mentioned later, the design of the detection system itself is out of scope of this paper. Hence, the only assumptions on the attacker that we make in the following are those that affect the functionality of the proposed revocation protocols. Otherwise, the attacker fits the general model described in other works [13, 15].

An essential assumption we make is *the existence of an honest majority in the attacker's neighborhood* (defined in Section 5.4.2). As we will show later, this allows vehicles to rely on quantity (number of honest neighbors) in order to establish sufficient trust level in the messages of other vehicles. Although it seems to be limiting, this assumption is reasonable if we look at the existing vehicular transportation systems. The actual percentage of malicious attackers is very low and the real threats come from malfunctioning devices.

4. PROBLEM STATEMENT

Valid credentials are a prerequisite for the participation in the VC system. We emphasize, however, that the possession of credentials does not necessarily imply the correct operation of the nodes. The problem at hand, given the system model described in Section 3, is to enable *revocation* of any of the credentials held by any of the nodes.

Revocation is the operation, undertaken by the organization and in this context network party issuing the certificate, that seeks to invalidate a certificate before the end of its lifetime or validity period.

The reasons for revocation are largely orthogonal to the operation itself. For example, it may be necessary for administrative reasons (e.g., change of registration domain), or because the related cryptographic material have been compromised (e.g., a private key has been detectably disclosed), or because the the registration terms or obligations have been violated, or, within the latter category, the holder of the credentials exhibits faulty behavior with respect to the implemented protocols.

Revocation is performed through distribution of the revocation information to all system entities. The appropriate mechanism to do so, an active topic of research, depends on the targeted system. The distribution points of the revocation information can be multiple and either operate in a redundant way or have distinct functionality (e.g., distributing information for specific types of certificates). In this work, we are not concerned with the organization of the system in that sense, but rather the way revocation information is made available at the wireless mobile vehicular network.

Our ultimate goal is to design a system capable of evicting illegitimate and faulty nodes, with their (progressive) isolation as a first step. Revocation is the primary means to achieve this. Our objective is to design a scheme tailored to characteristics of the considered environment, in order to achieve *timely* and *efficient* distribution of revocation information across the vehicular network. At the same time, we are after a scheme that can *robustly* and *efficiently* achieve isolation of illegitimate and faulty nodes as well as assist to their eventual eviction.

5. REVOCATION PROTOCOLS

In this section, we introduce our solutions to the revocation problem in VANETs. Given the diversity of possible scenarios (e.g., with or without infrastructure availability) in which revocation may be necessary, we have developed three novel certificate revocation protocols that together cover the whole spectrum of such scenarios. The following sections present them in more detail.

5.1 Solution Overview

As discussed earlier, the major drawback of CRLs in VANETs is their lack of scalability. But still this simple solution remains one of the most attractive choices. To adapt it to the VANET scale, we have opted for compressing CRLs, which turned out to be feasible using Bloom filters. Hence, the first revocation protocol is naturally called **RC^{2}RL** (**R**evocation using **C**ompressed **C**ertificate **R**evocation **L**ists). Just like traditional CRLs, C^{2}RLs are distributed to vehicles with information on revoked certificates. But C^{2}RLs are much smaller than CRLs.

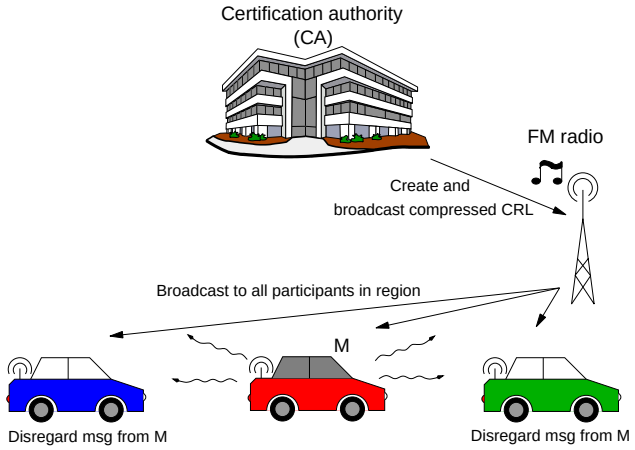

Figure 1: RC^{2}RL

In some special cases, all the certificates of a given vehicle need to be revoked. Because of the large number of certificates a vehicle carries, a more efficient approach to revoke these certificates is to prevent the TPD of the vehicle from using them. To accomplish this, the CA can instruct this TPD to stop all security functions, thus practically killing it. The protocol that kills the TPD of a vehicle is called **RTPD** (**R**evocation of the **TPD**).

We had also to consider the rather frequent VANET scenario, at least in the early years of deployment, where an infrastructure coverage is not available. In this case, any revocation attempt by the CA may not reach the destination area (the neighborhood of an attacker). To circumvent this problem, we designed **DRP** (**D**istributed **R**evocation **P**rotocol) that allows a *temporary* revocation of an attacker until connection to the CA is established.

It should be noted here that an essential enabler of any of the following revocation protocols, and especially DRP, is an *attacker detection system*. This may not necessarily mean a single mechanism but rather a set of mechanisms, each suitable for a specific scenario. For example, vehicles could use a local decision making algorithm, while the CA can leverage on its global knowledge of the network and run a more sophisticated, and also trusted, detection system. The orthogonality of the design intricacies of a such a system, as well as the lack of space, allow us leave this subject for future work. A good starting point on the subject in VANET literature is the system proposed in [8].

5.2 RC^{2}RL

CRLs need to be compressed, considering the limited available bandwidth. Since they contain very little re-

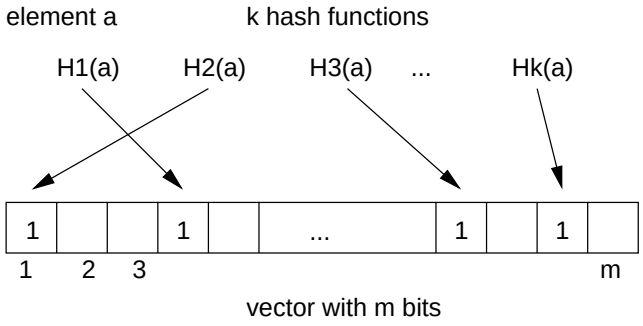

Figure 2: Bloom filter

dundancy, they cannot be efficiently compressed using normal lossless methods. We therefore use Bloom filters that are a special form of lossy compression [3]. They are a probabilistic data structure used to test whether an element is a member of a set. A main characteristic of Bloom filters is that they return a configurable rate of false positives (elements, which are claimed to be in a set, but actually are not). However, there are no false negatives: if the Bloom filter claims that an element is not in the set, we are sure that it is not.

A Bloom filter (Fig. 2) consists of a sequence of m bits, initially all set to zero (the vector). A key or element can be included in the filter by hashing it with a specific number k of independent hash-functions (each with a range 1, . . . , m), and by setting to 1 the bits that the hash-functions point to in the vector. By adding several keys to the filter, it is possible that one bit is set to 1 multiple times. To check afterwards if an element is contained in the filter, the element is hashed and the statuses of the corresponding bits are checked in the filter. If at least one bit that should be one is zero, one can surely affirm that the element is not contained in the filter. On the other hand, if all necessary bits equal one, with high probability the element is included. It may however also be possible that the bits were set to 1 by a combination of several other keys. Therefore, the more elements are added to the set, the larger the probability of false positives.

When receiving safety messages, vehicles verify the signing keys with the currently valid Bloom filter. We will discuss quantitative aspects of Bloom filters in Section 6.1. The CA broadcasts the C^{2}RLs as follows:

$$CA \rightarrow * : C^2RL, Sig_{CA}[C^2RL]$$

where $Sig_{CA}()$ denotes the digital signature of CA.

5.3 RTPD

There are situations in which all the keys of a vehicle have to be revoked. The most notable examples are when the vehicle is identified as an attacker or when one of its VANET components is malfunctioning. In this case, periodically creating a C^2RL containing all the keys of the vehicle and distributing it to the whole VANET is an overkill. Instead, there is a simpler and much more efficient solution, namely *Revocation of the Tamper-Proof Device* (RTPD), illustrated in Fig. 3.

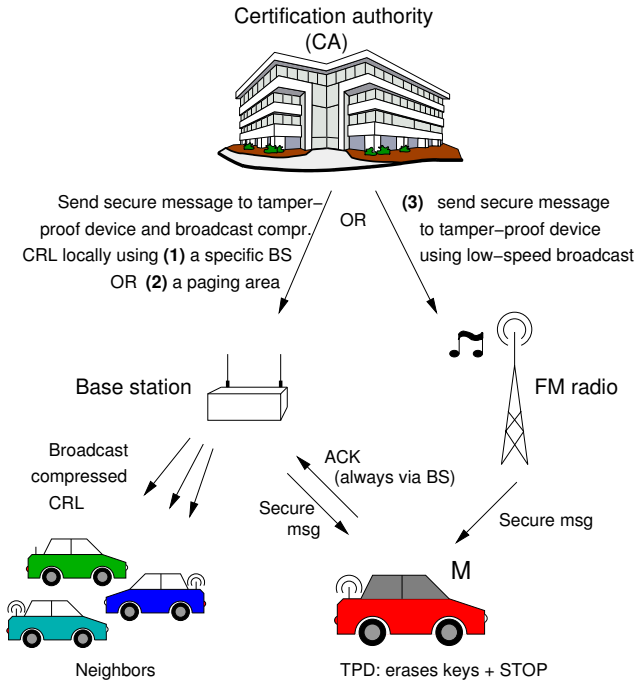

Figure 3: RTPD

The CA generates a revocation message for the TPD of vehicle M; this message contains M's *Electronic License Plate* (ELP_M) and a timestamp T. The message is first encrypted with M's public key PuK_M to prevent other vehicles that can eavesdrop the message from uncovering the identity of the revoked vehicle. Then the CA signs the encrypted message, which allows the base stations and the TPD to verify the authenticity of the message. If it is correctly verified, the base stations forward the message to the TPD. Since only the TPD is able to decrypt the message, it is difficult for attackers to specifically block messages destined to their own TPD. The CA's message will have the following content:

$$CA \rightarrow TPD_M : Sig_{CA}[E_{PuK_M}(ELP_M|T)]$$

where $E_{PuK_M}()$ denotes encryption with public key PuK_M.

There are several options for channeling this message to the targeted TPD. The first and most obvious choice would be to route it to the BS closest to the concerned vehicle, if its location is known to the CA. Otherwise the CA defines a paging area, consisting of several base stations in the region of the vehicle's most recent locations (trajectory extrapolation based on the vehicle's expected speed and acceleration can be useful in determining the paging area). If all else fails, the CA can recur to other distribution media mentioned in Section 3.1, such as low-speed radio broadcast.

When the TPD of the vehicle receives the message, it immediately erases the keys and stops signing VANET messages. It sends back a timestamped and signed ac-

knowledgment, as soon as it comes into reach of a base station:

$$TPD_M \overset{BS}{\rightarrow} CA : ACK, T, Sig_M[ACK|T]$$

Last but not least, the CA needs to warn the neighbors of the attacker that they may have received bogus safety messages signed with revoked keys. Therefore, it creates a C^{2}RL containing all keys of the revoked vehicle with currently valid validity windows. This C^{2}RL can be especially useful if the attacker tries to block the reception of messages from the CA. Although using the full-blown version of RC^{2}RL creates unnecessary overhead, locally sent C^{2}RLs do not create the problem of managing a very high number of different revocation lists at the concerned vehicles, since they are only of temporary nature.

5.4 DRP

Both RC^{2}RL and RTPD are revocation protocols initiated by the CA, which implies that the CA should be aware of the attacker vehicle and go through the process of generating and sending the revocation list or message. This raises two questions: (i) how does the CA learn about the attacker in the first place? and (ii) what happens if the C^{2}RL or the TPD-killer message are not delivered in time to their respective destinations, thus creating a vulnerability window? The answer to these two questions is the *Distributed Revocation Protocol* (DRP). The main principle of DRP is simple: the neighbors of the attacker vehicle take care of detecting and *temporarily* revoking it. The temporariness of revocation in this case comes from the fact that, in contrast to RC^{2}RL and RTPD, DRP is not a real revocation protocol but rather a *warning system against attackers*. As the CA is the only authority capable of revoking keys (given all the related administrative responsibilities and costs), it is normal that the neighbors of an attacker should, in the best case, be able to identify it and inform other neighbors, as well as the CA. Hence, any so-called "revocation" by the neighbors of a vehicle is temporally limited to the duration of contact between the vehicle and these specific neighbors. But once enough evidence against the attacker is gathered, the CA can initiate one of the previously described revocation protocols.

As mentioned in Section 5.1, each vehicle should run an attacker detection system, not covered in this paper. Being *aware* of the attacker allows the vehicle to ignore any messages sent by the attacker. But it also requires the vehicle to observe the attacker over some time before detecting its misbehavior. This is where VANET properties, especially high mobility and the broadcast nature of communications, come into the picture. On one hand, the ephemeral nature of contacts between vehicles (e.g., assuming a transmission range of 300 m and two vehicles moving in opposite directions on the same highway, each at an average speed of 100 km/h, their contact time is merely around 11 s) requires the detection system on any vehicle to continuously run in order to judge frequently changing neighbors. This suggests

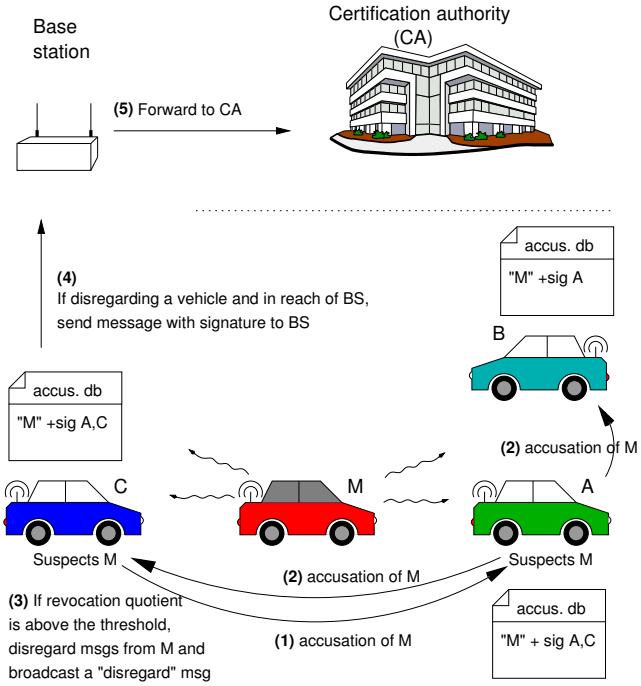

Figure 4: Distributed Revocation Protocol (DRP). Vehicle C has reached the revocation threshold for vehicle M, and broadcasts a disregard message. B is only in the transmission range of A, and gets the information from A when A reaches the threshold and sends a disregard message.

that a standalone detection system is not efficient. But on the other hand, high mobility and broadcast communications mean that vehicles come also very fast in reach of an attacker's future neighbors. Hence, the latter can benefit from the judgments of the first. In other words, sharing information between the standalone detection systems can highly improve their efficiency. This is the key concept behind DRP, shown in Fig. 4.

More precisely, vehicles that detect an attacker will broadcast *warning* messages to all vehicles in range. The latter can use this information as input to their respective detection systems. In this paper, we consider the case of vehicles that receive warning messages before being able to make any observations of the attacker and hence totally rely on these messages. The final step is to report the attackers or defective devices to the CA as soon as possible (i.e., when in reach of a base station or public vehicle).

5.4.1 Neighbor Warning System for DRP

As all private vehicles have initially the same trust level, a message sent by one vehicle warning that one of its neighbors is an attacker should not be considered by the receiver since both the *accuser* and the *accused* are attackers with the same probability, from the receiver's point of view. Hence a warning system that can be useful in DRP should rely on the collective informa-

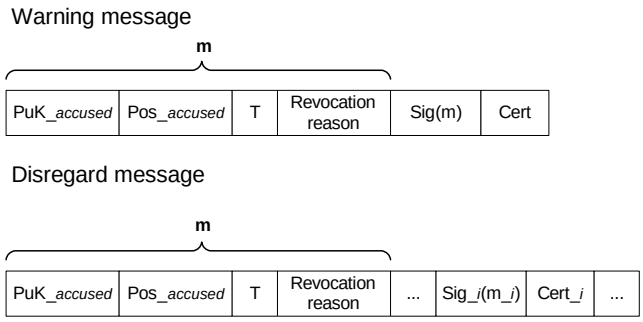

Figure 5: Formats of warning and disregard messages.

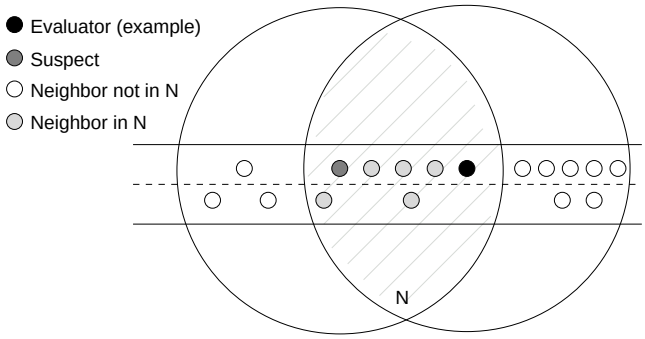

Figure 6: Definition of the attacker's neighborhood.

tion gathered from a vehicle's neighborhood. Again, as all vehicles can be attackers with the same probability, the exchanged warning messages may contain correct or wrong *accusations*. Given the limited amount of evidence available to vehicles receiving this information, they have to rely on the assumption of honest majority and crosscheck all the received accusations.

In this paper, we use a simple algorithm for summing accusations with an additional feature inspired by [6]: an accusation issued by a node has a lower weight when this node is accused already by other participants. For a given vehicle, if the sum of weighted accusations (the *revocation quotient*) against it exceeds a defined threshold, it is locally revoked by DRP. More precisely, warning messages are transformed into *disregard* messages that indicate to all the neighbors of the attacker to ignore its message.

It should be stressed here that this algorithm is just an example and other accusation aggregation systems can be devised. Our choice landed on this rather simple system because it requires no setup overhead, like incentive systems, nor long observation periods, like reputation systems. In addition, it involves no interactive mechanisms, like group agreement protocols; this prevents our system from being dependent on specific participants. As stated earlier, the only requirement for the proposed neighbor warning system is the existence of an honest majority.

The difference between warning and disregard messages, the formats of which are shown in Fig. 5, is that a specific number of *supporting signatures* are included by the sender in the latter. This increases the credibility of the message under the assumption of an honest majority.

5.4.2 Definition of the Attacker's Neighborhood

The attacker detection system, introduced in the previous section, relies on the neighbors of a *suspect* vehicle to accuse it in case of misbehavior and warn other vehicles. Once these vehicles have enough information about the suspect, they can evaluate whether it is misbehaving; we refer to them as *evaluators*. Hence, it is important to define the suspect's *neighborhood* N. New vehicles coming into the range of a suspected vehicle will use the information provided by their neighbors to make a decision on whether to trust it. It is essential here to avoid the *hidden vehicle problem* whereby a vehicle that is not in the suspect's range, and hence cannot evaluate it, is considered when making a decision (Fig. 6). Therefore, the neighborhood N of a suspect vehicle also depends on the vehicle evaluating it. We define $N(s,e)$[2], where s and e refer to suspect and evaluator respectively, as the intersection of the coverage areas (defined by the transmission range) of both the suspected vehicle and the vehicle evaluating it. In this way, we make sure that all vehicles in N can hear the suspect, and thus potentially accuse it, and at the same time be able to report their accusations to the evaluators.

It should be noted that all vehicles broadcast their position information, hence an evaluator vehicle can approximately determine which vehicles are in its own and in the suspect's transmission ranges simultaneously and thus determine N. The suspect's position is reported in the warning messages received by the evaluator.

5.4.3 Computation of the Revocation Quotient

The parameters that are used to calculate the revocation quotient for key j at each vehicle (with accusations from different keys i) are the following:

- A_i is the total number of accusations (issued by different public keys) heard by the vehicle against public key i. A_i is used to lower the weight of the accusations made by key i if it was already accused (i.e., it has lower credibility).
- P_i is the accumulated sum of $|N_i|$, the number of neighbors of i (as explained in Section 5.4.2) over time (while key i is used).
- α_i is the normalized value of the total number of accusations with respect to the total number of neighbors of key i during its usage ($0 \leq \alpha_i \leq 1$). This value is computed as follows: $\alpha_i = \frac{A_i}{P_i}$.
- ω_i is the weight of any accusation made with public key i. This weight depends on the number of

[2]In the following, we will use N instead of $N(s,e)$ unless there is a possibility of confusion of definitions.

accusations made against key i: $\omega_i = 1 - \alpha_i$, giving $0 \leq \omega_i \leq 1$. Therefore, the weight of a node against which there are no accusations equals 1. For trusted nodes and the accuser itself this number is always fixed to one.

- R_j is the revocation quotient defining whether the certificate for public key j should be revoked. It is computed as follows: $R_j = \frac{1}{P_j}(\sum_{i=1}^{P_j} \sigma_{ij}\omega_i\mu_{ij})$, where $\sigma_{ij} = 1$ if there is an accusation against j issued with public key i, otherwise 0. $\mu_{ij} \in [1, S]$ is a parameter corresponding to the severity class of the accusation from i against j and S is the scale of the severity class preloaded in the TPD.

The revocation quotient threshold (R_T) is a configurable parameter. A typical value would be 0.5 (majority vote). If $R_j > R_T$, key j becomes untrustworthy, i.e., messages signed with this public key are disregarded. A vehicle that has accumulated enough accusations against an attacker to reach the threshold and disregard messages from it is called hereafter a *warned vehicle*. An *initially warned vehicle* is a warned vehicle that disregards all bogus messages, because it has already reached the warned state before receiving the first message from M (because it has received enough accusations to reach the revocation threshold).

6. EVALUATION

In this section, we evaluate the performance of the different revocation protocols introduced in this paper. Notably, the size of Bloom filters and the performance of distributed revocation under stringent VANET conditions are key to determining the feasibility of the proposed solutions. The following results show that VANET certificate revocation can be efficiently done.

6.1 Bloom Filter Size

To reduce the communication overhead, we have chosen a Bloom filter size of few tens of kbytes. To reduce the probability of false positives with a reasonable computation overhead, we have chosen to use 10 hash functions, which is close to the optimal number for a sufficiently high number of keys as Fig. 7 illustrates. When introducing fewer keys (and thus having a lower probability of false positives) the result is sub-optimal for this filter-size. To stay closer to the optimal result, it would be necessary to use more hash-functions, which also means to have more computational burden on the processor of the vehicle, and more delays. Therefore this is a good tradeoff between the size, the probability of error, and the computation cost.

6.2 Simulation of DRP

As DRP relies on the ad hoc operation of vehicles within short time delays, we have simulated it using *ns-2*. We have considered three different parameters in our evaluations. The first one is the traffic model: we used a freeway (FW), a city (WU), and a mixed (freeway/city) (AO) scenarios; WU and AO are realistic scenarios taken from [16]. The second factor is the density of vehicles,

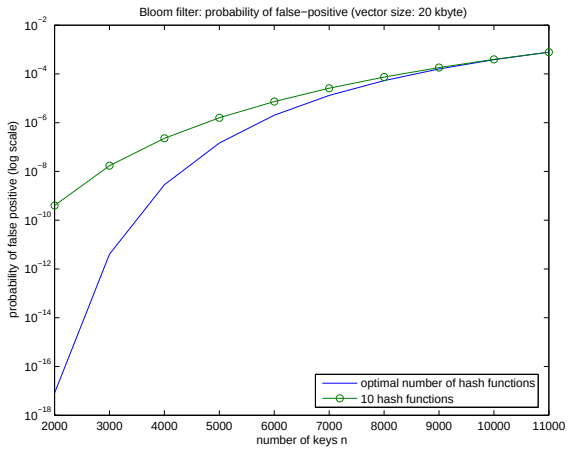

Figure 7: Bloom filter size.

which reflects the time of the day (e.g., density is much higher at peak hours). The last factor is the vehicle speed. The presented results are the average of 50 simulation runs.

6.2.1 Vehicle Density

In the WU and AO scenarios, we can see that for a very low vehicle density the percentage of *initially warned vehicles* (Section 5.4.3) is low (Fig. 8(a)). In this case it is not possible to pass information from one vehicle to another in a reliable way and to warn other vehicles beforehand. However, this value rapidly grows and stabilizes between 80 and 90%. For those vehicles not initially in the warned state, the average time to reach the threshold to disregard messages stabilizes at a value lower than 2 s (Fig. 8(b)). For the freeway scenario, there is a slight decrease in performance for very high densities. This can be explained by the fact that the number of packet collisions increases when the density increases. Because of the hidden node problem, even though we are not yet at channel capacity, the number of colliding packets increases and creates some loss. Since in this scenario speeds are higher (and thus contacts shorter), this may result in the fact that some accusations are not received by every neighbor.

From this we can conclude that even with a relatively low density, which is also the case with a low market penetration at the beginning of VANET deployment, our revocation system is still able to accumulate enough information to perform successfully within comfortable delays.

6.2.2 Average Speed

We evaluate the same metrics for different average speeds in the scenario (Fig. 8(d) and 8(e)). In the urban areas (AO and WU), we can see that higher speeds give better results: again, since more participants can be contacted in the same time interval, it is possible to accumulate more accusations against the attacker and to warn other vehicles (in these two cases, the maximum

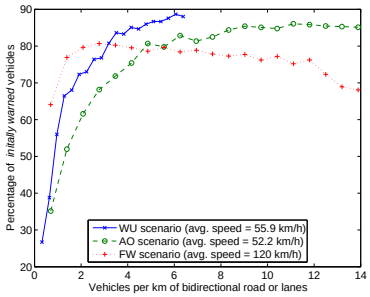

(a) Percentage of initially warned vehicles as a function of the vehicle density (in the FW scenario, the density is per lane)

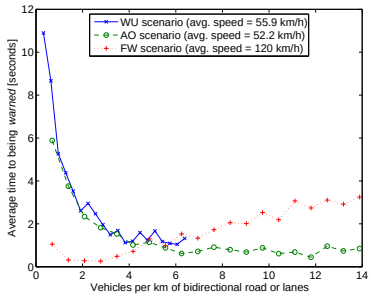

(b) Average time to being warned as a function of the vehicle density

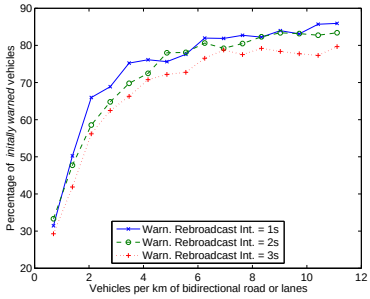

(c) Percentage of initially warned vehicles as a function of the vehicle density (AO Scenario)

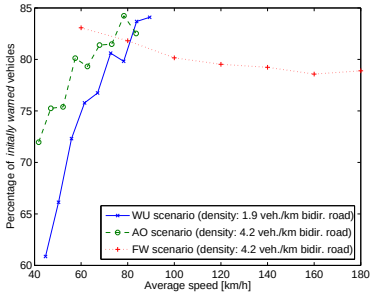

(d) Percentage of initially warned vehicles as a function of the average vehicle speed

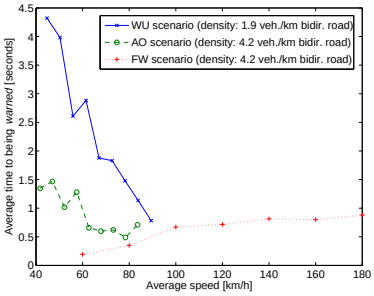

(e) Average time to being warned as a function of the average vehicle speed

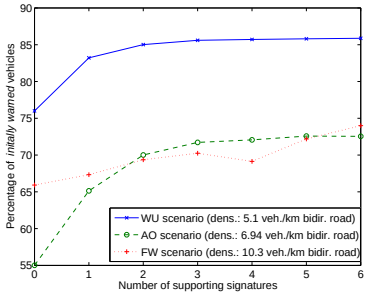

(f) Percentage of initially warned vehicles as a function of the number of supporting signatures

Figure 8: System performance vs. density of vehicles, average speed and number of supporting signatures

average speed is 90 km/h, higher values are rare). In the freeway scenario, the average speed is much higher, and performance decreases slightly for very high speeds: this can be explained by the fact that the contact time becomes very short (approaching to the same order of magnitude as Accusation message sending). In this case, some messages may be lost, thus resulting in a slightly higher time to reach the necessary threshold.

Still, we can see that the distributed revocation can execute within the delay bounds imposed by short contact times and cover a considerable percentage of concerned vehicles.

6.2.3 *Effect of Warning Rebroadcast Interval*

If a vehicle continues to receive bogus messages from the attacker signed with the same public key, it does not send additional warning messages, unless the difference between the current time and the timestamp of its last sent warning message (against this key) is larger than the *Warning Rebroadcast Interval* (WRI) parameter. The WRI is used to prevent vehicles from flooding the channel by sending an accusation message every time they hear a bogus message; this prevents a possible DoS attack based on excessive channel load or by creating too much computation overhead for signature creations. However, if a vehicle sends a warning message only once and then stops participating in the warning process against a potential attacker, newly arriving vehicles will not be able to accumulate enough accusation information and may trust messages from the malicious vehicle, although all other neighbors agreed to disregard it. Hence, the parameter WRI is a trade-off between sending too often the same warnings, and quickly informing new neighbors about misbehaving vehicles.

Changing the WRI has only a small impact on the percentage of initially warned nodes in the case of low vehicle density (Fig. 8(c)). In higher density situations, we can however observe a slight increase of warned vehicles with smaller WRI, when sending accusations more frequently, even though this also increases channel load. As the results show, WRI can be kept large without much degradation in the performance of DRP, but with a considerable reduction in channel load.

6.2.4 *Number of Supporting Signatures*

The *number of supporting signatures* (Section 5.4.1) is a parameter of DRP. It may be influenced by the cryptographic system that is used: if signature sizes are large, the number of supporting signatures may have to be kept small, to maintain a reasonable packet size. But a minimum number of signatures is still required to as-

sure a sufficient level of credibility to disregard messages. From Fig. 8(f) it can be seen that the number of supporting signatures has an impact on the number of warned vehicles for low numbers of supporting signatures. But starting from only 4 signatures, there is little change in the results. This can be explained by the fact that at some point, the number of vehicles that broadcast exactly the same supporting signatures (and thus give no new information) becomes higher. Hence, sufficient credibility in disregard message contents can be achieved with relatively small overhead.

7. CONCLUSION

In this paper, we addressed the certificate revocation problem in vehicular networks. To take into account the specific properties of these networks, we designed three novel revocation protocols, each one adapted to a specific scenario. We have also devised a method for efficiently compressing CRLs using Bloom filters. Last but not least, we have evaluated the proposed solutions, notably the ad hoc protocol (DRP) and shown that it is suitable for highly mobile networks.

Given the broad scope of the subject tackled in this paper, there is ample space for future work. We will especially focus on the attacker detection system needed in DRP.

8. REFERENCES

[1] 5.9 GHz Dedicated short range communications (DSRC). http://grouper.ieee.org/groups/scc32/dsrc/index.html.

[2] IEEE P1609.2 Version 1 - Standard for Wireless Access in Vehicular Environments - Security Services for Applications and Management Messages.

[3] Burton H. Bloom. Space/time trade-offs in hash coding with allowable errors. *Commun. ACM*, 13(7):422–426, 1970.

[4] S. Boeyen, T. Howes, and P. Richard. Internet X.509 Public Key Infrastructure LDAPv2 Schema. RFC 2587, 1999.

[5] S. Buchegger. *Coping With Misbehavior in Mobile Ad-hoc Networks*. PhD thesis, Swiss Federal Institute of Technology (EPFL), April 2004.

[6] Claude Crépeau and Carlton R. Davis. A certificate revocation scheme for wireless ad hoc networks. In *SASN '03: Proceedings of the 1st ACM workshop on Security of ad hoc and sensor networks*, pages 54–61, New York, NY, USA, 2003. ACM Press.

[7] S. Farrell and R. Housley. An Internet Attribute Certificate Profile for Authorization. RFC 3281, 2002.

[8] Philippe Golle, Dan Greene, and Jessica Staddon. Detecting and correcting malicious data in VANETs. In *VANET '04: Proceedings of the first ACM workshop on Vehicular ad hoc networks*, pages 29–37, New York, NY, USA, 2004. ACM Press.

[9] R. Housley, W. Polk, W. Ford, and D. Solo. Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile. RFC 3280, 2002.

[10] Markus Jakobsson and Susanne Wetzel. Efficient attribute authentication with applications to ad hoc networks. In *VANET'04*, pages 38–46, New York, NY, USA, 2004. ACM Press.

[11] Jiejun Kong, Haiyun Luo, Kaixin Xu, Daniel Lihui Gu, Mario Gerla, and Songwu Lu. Adaptive security for multilevel ad hoc networks. *Wireless Communications and Mobile Computing*, 2(5):533–547, 2002.

[12] M. Myers, R. Ankney, A. Malpani, S. Galperin, and C. Adams. X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP. RFC 2560, 1999.

[13] B. Parno and A. Perrig. Challenges in securing vehicular networks. In *Proceedings of HotNets-IV*, 2005.

[14] D. Pinkas and T. Gindin. Internet X.509 Public Key Infrastructure Permanent Identifier. RFC 4043, 2005.

[15] M. Raya and J. P. Hubaux. The security of vehicular ad hoc networks. In *Proceedings of SASN'05*, Alexandria, VA, USA, November 2005.

[16] Amit Kumar Saha and David B. Johnson. Modeling mobility for vehicular ad-hoc networks. In *VANET '04: Proceedings of the 1st ACM international workshop on Vehicular ad hoc networks*, pages 91–92, New York, NY, USA, 2004. ACM Press.

[17] S. Santesson, M. Nystrom, and T. Polk. Internet X.509 Public Key Infrastructure: Qualified Certificates Profile. RFC 3739, 2004.

[18] Petra Wohlmacher. Digital certificates: a survey of revocation methods. In *MULTIMEDIA '00: Proceedings of the 2000 ACM workshops on Multimedia*, pages 111–114, New York, NY, USA, 2000. ACM Press.

[19] Rebecca N. Wright, Patrick D. Lincoln, and Jonathan K. Millen. Efficient fault-tolerant certificate revocation. In *CCS '00: Proceedings of the 7th ACM conference on Computer and communications security*, pages 19–24, New York, NY, USA, 2000. ACM Press.

[20] Peifang Zheng. Tradeoffs in certificate revocation schemes. *SIGCOMM Comput. Commun. Rev.*, 33(2):103–112, 2003.

[21] Lidong Zhou and Zygmunt J. Haas. Securing ad hoc networks. *IEEE Network*, 13(6):24–30, 1999.