

Scaling VANET Security Through Cooperative Message Verification

Hongyu Jin and Panos Papadimitratos

Networked Systems Security Group, KTH Royal Institute of Technology, Sweden

{hongyuj, papadim}@kth.se

www.ee.kth.se/nss

Abstract—VANET security introduces significant processing overhead for resource-constrained On-Board Units (OBUs). Here, we propose a novel scheme that allows secure Vehicular Communication (VC) systems to scale well beyond network densities for which existing optimization approaches could be workable, without compromising security (and privacy).

Index Terms—Security, performance, scalability

I. INTRODUCTION

Vehicular Communication (VC) systems, notably Vehicle-to-Vehicle (V2V) and Vehicle-to-Infrastructure (V2I) communication, entail high-rate transmissions; typically, for safety applications, On-Board Units (OBUs) transmit at a rate of 10 messages (safety beacons) per second. Granted, there are methods that adapt the beaconing rate, but the challenge is clear: often, especially as VC systems get progressively widely deployed, each vehicle will have to process safety beacons, along with other traffic, from several tens of other vehicles within its OBU range. For example, 200 (300) messages/second for 20 (30) neighboring vehicles.

The provision of security and privacy protection aggravates the situation, adding communication overhead (digital signatures and certificates attached, thus longer messages) as well as computational overhead (digital signature verifications mostly, and signature calculations). This problem has been investigated in the literature, with a number of improvements (e.g., [3], [5]) compatible with the standardized pseudonymous authentication approach. For example, the certificates of the authenticating sender can be omitted periodically or based on the sender's context [5]; at the receiver side, they need to be validated once [3].

These approaches provide significant improvements and show how one can dimension processing power [3], but they are somewhat conservative: they assume each node validates all received messages it receives and deems relevant. Indeed, this is the straightforward approach. An alternative, adaptive, reactive approach has been considered in [10], but only for multi-hop messages.

It is important to realize that for safety beacon transmitted, there are N validations that take place; where N is the number of receiving neighbors. This is significant redundant effort, especially when most of the transmissions, termed Cooperative Awareness Messages (CAMs), are important yet not of high priority. Intuitively, if most nodes (vehicles) are benign, each

validation of a message they perform could serve their peers in the vicinity.

This is exactly the idea we promote in this short paper: Each node can notify its neighbors about its successful verification of some recently received beacons; each neighbor that validates such an augmented message can leverage this additional information and avoid verifying itself the corresponding beacons. Ideally, this could reduce significantly the overhead, as one costly cryptographic verification provides information for multiple messages. But this is a double-edged sword: if not done carefully, it leaves space for abuse by intelligent internal adversaries.

Our results show that the network density, which raises the scalability problem, can also provide a remedy. Our scheme trades a tiny window of vulnerability, allowing a miniscule fraction of beacons that could be mistaken as valid before an adversarial node be evicted.¹ At the same time, we get extensive improvement in terms of message validation delay, allowing, in fact, the network to scale to neighborhoods that are 50% to 100% larger than the ones for which optimized yet more conservative approaches would be saturated and unworkable.

In the rest of the paper, we present in detail our cooperative validation scheme (Sec. II), we provide a security analysis (Sec. III), and present a body of simulation results (Sec. IV) before some concluding remarks (Sec. V).

II. OUR SCHEME

Overview: Our scheme extends the traditional V2V message verification, leveraging neighboring peers to reduce validation delays without compromising the achieved security (and privacy). The basic idea is to augment each (safety) message with brief identifiers of previously validated messages. These identifiers indicate the corresponding messages have been verified by the sender. This is exactly where nodes can benefit from each other: accepting a message can help verifying the messages (received and queued) the identifiers in this message point to. In addition, to counter misbehavior, each node probabilistically selects a subset of the received identifiers and verifies by itself the signatures of the corresponding messages. Revocation would be triggered if any misbehavior is identified. Table I summarizes notation used in this paper.

¹We emphasize that if a message is critical, the receiver can always validate it in a traditional way, in addition to our optimistic approach.

TABLE I: Notation

N	Number of vehicles
$\{msg\}_\sigma$	Signed message
Pr_{check}	Probability of checking each peer-provided verification result
α	Number of verification results in a CAM
γ	CAM frequency
τ	Average message verification delay
$H()/H$	Hash function/Hash value
b	1 bit value, indicating the message is selected for checking

Message Generation and Reception: The format of a signed CAM in our scheme is changed into:

$$\{M\}_\sigma = \{CAM\ Fields, H_1..H_\alpha\}_\sigma. \quad (1)$$

Except the hashes, $H_1..H_\alpha$, we assume the rest of the fields are as defined in the standard [4]. $H_1..H_\alpha$ are the hashes of latest verified CAMs (based on the timestamps of the CAMs, not the times of reception or verification). For example, a vehicle, V , caches locally the hash values of the latest verified CAMs (for which V performed signature verifications, not cooperatively verified as described below) and includes them in its own (sent) CAMs.

Everytime a node receives a CAM, it generates a job based on the CAM. Here, we consider the processing of a received CAM as a job. The format of the job is defined as follows:

$$\{\{M\}_\sigma, H(\{M\}_\sigma), b = 0 \text{ or } 1\}. \quad (2)$$

The field b indicates the CAMs are selected for probabilistic checking. In case b is set to 1 for a job, the corresponding CAM cannot be verified through cooperative verification (peer-provided hashes): the signature of this CAM must be verified. For each new job, b is set to 0.

We assume a single thread for cryptographic verification in each OBU: a CAM received when the thread is busy needs to be queued. To increase efficiency, we *randomly select* the inserted position in the queue for each new job. This way, we reduce the probability that nearby receivers verify the same CAM roughly simultaneously. The verification of CAMs that sent from a node does not need to follow the sending order, as long as they are verified before they expire.

Cooperative Verification: Queue processing at a node is done according to Algorithm 1. When the queue is not empty, the node pops the first job from the queue, and accepts the CAM if the signature is valid. The hashes, $H_1..H_\alpha$, are used to verify the CAMs (for which $b = 0$) in the queue. For each cooperatively verified CAM, there is a probability Pr_{check} , that the CAM will be checked by validating the signature. If so, b is set to 1 and it is inserted after the last job with $b = 1$ (i.e., before the first job that $b = 0$). Otherwise, the CAM is accepted and removed from the queue.

The probabilistic checking of the claimed verifications (hashes) counters abuse, due to (i) the density of the neighborhood and the (extensive, most often) majority of benign nodes present, and (ii) the ability to locally contain/ignore misbehaving nodes and then globally evict them. The latter is

easily enabled by our scheme (simple cryptographic validation in lieu of misbehavior detection), yet the exact way to identify locally the wrongdoer and evict it is orthogonal and can be done by schemes proposed in the literature, e.g., [8].

Algorithm 1 Cooperative verification

```

1: while Queue is not empty do
2:   Pop a job,  $\{\{M\}_\sigma, H, b\}$ , from the head of Queue
3:    $M = \{CAM\ Fields, H_1..H_\alpha\}$ 
4:   if The signature of  $\{M\}_\sigma$  is valid then
5:     Accept  $M$ 
6:     for Each  $H_i$  in  $H_1..H_\alpha$  of  $M$  do
7:       if  $H_i$  is found in Queue, and  $b$  of  $M_i$  is 0 then
8:         Chooses 1 with probability  $Pr_{check}$ ,
9:         or 0 with probability  $1 - Pr_{check}$ 
10:        if Chooses 1 then
11:          Insert  $\{\{M\}_\sigma, H(M), b = 1\}$  into Queue,
12:          right after the last job, for which  $b$  is 1
13:        else
14:          Accept  $M_i$ ,
15:          and remove  $\{\{M_i\}_\sigma, H_i, b\}$  from Queue
16:        end if
17:      end if
18:    end for
19:  end if
20: end while

```

III. SECURITY ANALYSIS

Adversary Model: We consider internal adversaries seeking to abuse the system, in particular the cooperative verification scheme. Without loss of generality, let an adversary controlling one OBU injecting bogus, i.e., *not properly signed* messages. Then, let the adversary use a compromised OBU (equipped with the appropriate credentials) that transmits augmented messages falsely claiming previously transmitted bogus messages as verified. In the hope that those bogus messages received by benign nodes would be accepted without verification/checking. Such adversarial behavior is actually relevant to the optimistic cooperative validation approach we advocate here. An attacker could try to generate malicious CAMs given overheard hashes. However, given the properties of hash functions and the length of hash values (80 *bit*, as considered in Sec. IV), it is very straightforward that finding a message based on the hash values is very hard.

Analysis: Next, we discuss exactly how this misbehavior, specific to our scheme, is countered. We emphasize we are not concerned here with the validity of the message content, e.g., the correctness of a location or an alert about emergency braking; those are orthogonal and can be addressed by relevant consistency checking ([6], [7]) and data-centric security schemes [9]. Here, we are concerned with incorrectly signed (with arbitrary content) messages, and the attempt to legitimize them by improper, adversarial use of our scheme. The detection of any such false claim is straightforward for any legitimate receiver, as long as it cryptographically validates the purported as verified message (signature).

The use of pseudonymous authentication, as per the standards under development, guarantees *non-repudiation* and

message integrity and authentication; as long as the receiving node performs the cryptographic validation itself (message signature and attached pseudonym validation).

Revealing False Claims: To increase the probability of detecting such misbehaviors, the reasonable amount of peer-provided verification results should be checked. However, cooperation among nodes within a neighborhood could significantly increase this probability.

Consider a single adversary case, transmitting messages at a rate γ_{adv} . In the worst case (broadcasting aggressively bogus messages, hoping benign nodes receive as many bogus messages as possible), the adversary could broadcast bogus messages at a rate $\frac{\alpha}{\alpha+1} \cdot \gamma_{adv}$ and broadcast valid messages that “validate” those bogus messages at a rate $\frac{1}{\alpha+1} \cdot \gamma_{adv}$. More specifically, broadcast α bogus messages and broadcast the $(\alpha + 1)$ -th, as a valid one that includes hashes of the earlier α bogus messages. We seek to detect the adversary that provided such faulty claims and revoke its credentials. For any message, the probability of all α included hashes not being checked by a receiver is:

$$Pr_{skip} = (1 - Pr_{check})^\alpha. \quad (3)$$

Then, assuming v votes (misbehavior reports) are needed to cooperatively reveal a misbehavior, the probability of revealing such a message with N benign nodes in the neighborhood (assuming they have all received the bogus “validating” message) can be estimated as:

$$Pr_{reveal} = 1 - \sum_{i=0}^{v-1} \binom{N}{i} (Pr_{skip})^{N-i} (1 - Pr_{skip})^i, \quad (4)$$

where Pr_{skip} is calculated with Eq. (3). Pr_{reveal} increases with neighbor density; high neighbor density (thus, high message reception rate) environments are the ones our scheme fits best. Note that Pr_{reveal} is only the probability of revealing a single malicious message. The probability of revealing misbehavior after n malicious messages are sent out would be even higher: $1 - (1 - Pr_{reveal})^n$.

For example, Pr_{reveal} is around 0.80 in a network with $\alpha = 5$, $N = 15$, $Pr_{check} = 0.1$ and $v = 5$. This means the adversary would be revealed with high probability even after its first transmission of a false claim. A more intelligent adversary could include only a reduced number of bogus message hashes in each valid CAM, to reduce the probability of getting detected. However, this significantly weakens the adversary, which would still get revealed after several rounds.

Benign receivers may consume only a few bogus messages for the short period, causing minimal harm while ensuring integrity and authentication of the vast majority of messages in an efficient way; considering most of the CAMs are sent out with low priority in a usual environment. In addition, a node can choose to verify the CAMs sent out with high priority immediately, ignoring the queue size and our cooperative verification protocol.

Privacy: On top of security, we note that privacy is not weakened by our scheme. The hash values in a CAM do not

link transmissions of any other node, beyond what one can infer from geographical information included in the CAMs.

IV. EVALUATION

We use OMNeT++ [1] and the IEEE 802.11p module from Veins [2] to simulate our scheme and analyze the system performance. Let *waiting time* be the total time a received message waits in the queue before its verification. Our scheme achieves low delays, which would not have been possible for the standard approach (First-Come First-Served with verification of all messages, referred next as *baseline*).

TABLE II: System Parameters (**Bold** for Default Setting)

N	15, 20, 25, 30 , 35, 40
Pr_{check}	0.2 , 0.4, 0.6
α	3, 4, 5 , 6
τ	3, 5 , 7 <i>ms</i>
γ	10 <i>Hz</i>

Simulation Settings: Table II shows the system parameters and values used in the simulation. We consider a $200\text{ m} \times 200\text{ m}$ square area. The node we evaluate is placed at the center of the area and its neighbors (the rest of the nodes) are uniformly placed in the area. We consider a bit rate of 6 *Mbps*, with 300 *bytes* length for each CAM (including a signature and a certificate). In addition, we increase the payload length based on the number of hashes (80 *bit* for each) included, thus communication overhead of $80 \cdot \alpha$ *bits*. This amounts to extra communication delay in the order of 0.1 *ms* per CAM, considering the values we use for α . Certificate omission [3] can be used to decrease message verification delay. However, we do not explicitly address it in our simulation; rather, we assume the message verification delay, on average, has a deterministic value τ (including all operations, such as hash computation and queue search). For each simulation setting, we perform 5 randomly seeded experiments of 2 *min* and average over these 5 runs. The bold values in Table. II are the default ones used in our simulation. For example, N is the parameter we examine in Fig.1a, with the rest of the parameters having the default values: $Pr_{check} = 0.2$, $\alpha = 5$, $\tau = 5\text{ ms}$ and $\gamma = 10\text{ Hz}$. The default values of τ and γ are typical values based on the literature (e.g., [3], [4]).

Simulation Results: Fig. 1a shows the waiting time CDF as a function of N . We also evaluate the baseline scheme for $N = 15$: maximum sustainable neighbor size (with such queue) is around 20, as only 200 *msg/s* can be verified with $\tau = 5\text{ ms}$. Fig. 1a shows the baseline scheme ($N = 15$) performs similarly to our scheme with $N = 20$. In addition, around 90% of the messages have waiting times less than 0.3 *s* with $N = 40, 45$, which is a significant improvement considering the queue would not even be stable for the baseline scheme.

Fig. 1b illustrates waiting times for the default setting over the simulation time. We see spikes, but overall, waiting times are stable and do not increase as simulated time progresses. In Fig. 1c, we see that as the number of neighbors increases (thus, higher message reception rate), more messages need

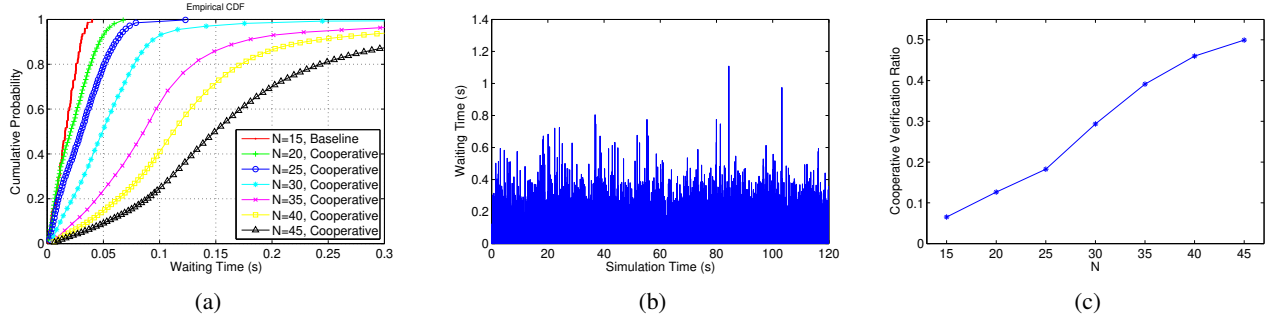


Fig. 1: (a) Waiting time CDF as a function of N . (b) Waiting time as simulated time progresses. (Default setting) (c) Cooperative message verification ratio as a function of N .

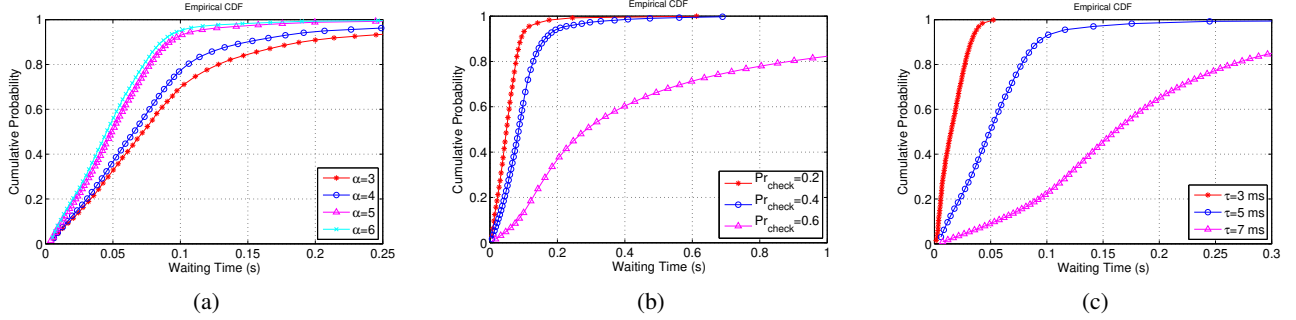


Fig. 2: Waiting time CDF as a function of (a) α , (b) Pr_{check} , and (c) τ .

to be verified based on the peer-provided verification results. This implies a higher risk of accepting bogus messages in a malicious environment. However, on the other hand, the probability of revealing (thus revoking) malicious nodes also increases as the neighbor density increases (Sec. III).

We evaluate the waiting time for different α values (Fig. 2a). We discover a threshold for α , above which peer-provided verification results cannot help: with $\alpha = 7$, the waiting time CDF almost overlaps with that for $\alpha = 6$. However, we can expect a higher threshold for α with larger N . Pr_{check} has an impact on the probability of bogus message (thus, malicious node) detection and the waiting time distribution: more signatures need to be verified for higher Pr_{check} . As shown in Fig. 2b, almost 80% of the messages have waiting times less than 0.8 s even with $Pr_{check} = 0.6$. Under this setting, there is even a high probability that malicious nodes can be detected locally/independently, since more than half of the messages are checked. Fig. 2c shows the impact of message verification delay on waiting time. In our simulation, we found with $\tau = 8$ ms, the queue is not stable anymore and the queue size goes to infinity. However, it is already a significant improvement considering only around 15 neighbors can be sustained for the baseline scheme with $\tau = 7$ ms.

We do not consider deadlines for CAMs here, but we can infer those; e.g., for the default setting, with a 0.1 s deadline, we can expect more than 90% of the received CAMs to be verified (Fig. 1a). Again, this would not have been possible for the baseline scheme, which can sustain around 20 neighbors in the same setting.

V. CONCLUSIONS

We demonstrated how our cooperative message verification scheme could enable secure VC at network densities even double compared to those prior approaches could be workable for. Though this is achieved by trading off a tiny vulnerability window, we showed this can be harmless. In addition, our scheme is orthogonal to all prior optimizations and could complement them.

REFERENCES

- [1] OMNeT++. <https://omnetpp.org/>.
- [2] Veins. <http://veins.car2x.org/>.
- [3] G. Calandriello, P. Papadimitratos, J.-P. Hubaux, and A. Liou. On the performance of secure vehicular communication systems. *IEEE TDSC*, 8(6), 2011.
- [4] ETSI EN 302 637-2. Intelligent transport systems; vehicular communications; basic set of applications; part 2: Specification of cooperative awareness basic service, Nov. 2014.
- [5] M. Feiri, J. Petit, and F. Kargl. Formal model of certificate omission schemes in vanet. In *IEEE VNC*, Paderborn, Germany, Dec. 2014.
- [6] A. Festag, P. Papadimitratos, and T. Tielert. Design and performance of secure geocast for vehicular communication. *IEEE TVT*, 59(5), 2010.
- [7] T. Leinmüller, C. Maihöfer, E. Schoch, and F. Kargl. Improved security in geographic ad hoc routing through autonomous position verification. In *International Workshop on VANET*, Los Angeles, CA, Sept. 2006.
- [8] T. Moore, M. Raya, J. Clulow, P. Papadimitratos, R. J. Anderson, and J. Hubaux. Fast exclusion of errant devices from vehicular networks. In *IEEE SECON*, San Francisco, CA, June 2008.
- [9] M. Raya, P. Papadimitratos, V. D. Gligor, and J.-P. Hubaux. On data-centric trust establishment in ephemeral ad hoc networks. In *IEEE INFOCOM*, Phoenix, AZ, Apr. 2008.
- [10] N. Ristanovic, P. Papadimitratos, G. Theodorakopoulos, J.-P. Hubaux, and J.-Y. L. Boudec. Adaptive message authentication for multi-hop networks. In *IEEE/IFIP WONS*, Bardonecchia, Italy, Jan. 2011.