

A-VIP: Anonymous Verification and Inference of Positions in Vehicular Networks

F. Malandrino C. Casetti C.F. Chiasserini M. Fiore R. S. Yokoyama C. Borgiattino
Politecnico di Torino Politecnico di Torino Politecnico di Torino INSA Lyon ICMC, Univers. São Paulo Politecnico di Torino

Abstract—Knowledge of the location of vehicles and tracking of the routes they follow are a requirement for a number of applications. However, public disclosure of the identity and position of drivers jeopardizes user privacy, and securing the tracking through asymmetric cryptography may have an exceedingly high computational cost. In this paper, we address all of the issues above by introducing A-VIP, a lightweight privacy-preserving framework for tracking of vehicles. A-VIP leverages anonymous position beacons from vehicles, and the cooperation of nearby cars collecting and reporting the beacons they hear. Such information allows an authority to verify the locations announced by vehicles, or to infer the actual ones if needed. We assess the effectiveness of A-VIP through testbed implementation results.

I. INTRODUCTION

Borrowing from a well-established communication pattern in wireless LANs, vehicular networks have adopted the term *beaconing* to indicate the periodic broadcasting of messages to neighboring vehicles or road-side units (RSUs). These messages, defined in both the SAE J2735 [1] and the ETSI ITS [2] specifications, carry vehicle identifier, timestamp and location information, and may be used for safety purposes as well as for cooperative awareness.

In a number of scenarios, location accountability is a requirement in order to provide services to the community or to drivers. Therefore, beacons used for vehicle position identification and tracking in such contexts need to be made secure. For instance, secure reporting of vehicle location can substantiate drivers' claims in case of accidents. At the same time, secure location verification by authorities can provide non-repudiation and liability for those involved.

The information carried by beacons may be secured through the use of an on-board tamper-proof Hardware Security Module (HSM) as well as signatures, cryptography and certificates [3]. However, these solutions suffer from three major drawbacks, concerning (i) users' privacy, (ii) computational costs of security and (iii) the system trust on user correctness. As for the first aspect, when not strictly required, public disclosure of the vehicle identity to all receiving devices in the proximity of a beaconer is an issue: vehicles can be tracked by adversaries, jeopardizing drivers' privacy and requiring complex pseudonym management [4]. There is thus a need for separating secure position identification by authorities and by peer users in the vehicular network, the latter representing an undesirable feature from a user privacy viewpoint. As for the second aspect, security mechanisms induce significant protocol overhead and computational complexity, up to the point where their use should

be largely dependent on applications and circumstances [5]. Finally, basic solutions cannot guarantee the correctness of the location information provided by a user who owns the required cryptographic material, but has a malfunctioning GPS receiver or is capable of tampering with GPS data before they are input to the HSM.

In this paper, we address the issues above by proposing A-VIP (Anonymous Verification and Inference of Positions), a framework that, unlike previous work, (i) allows an authority to securely verify the positions claimed by vehicles without resorting to asymmetric cryptography – as is done instead in the current IEEE 1609.2 standard [6]; (ii) safeguards drivers' privacy with respect to other vehicles participating in the network; (iii) does not require uninterrupted radio coverage from a roadside unit (RSU) communication infrastructure. Furthermore, in presence of unverified location claims, A-VIP allows the authority to infer the actual position of malfunctioning or misbehaving vehicles.

To that end, A-VIP leverages a simple, yet effective, mechanism based on *anonymous position beacons* broadcasted by vehicles, which prevent overhearing nodes from identifying or tracking their source.

The rest of the paper is organized as follows. We describe the system scenario and communication protocol in Sec. II. Sec. III details the location verification and inference procedures, while Sec. IV discusses the resilience of A-VIP to a number of attacks by adversarial vehicles. The performance of A-VIP in a real-road environment is evaluated in Sec. V. Finally, Sec. VI concludes the paper.

II. ANONYMOUS POSITIONING PROCEDURE

We consider a vehicular network composed of vehicles communicating with each other and, occasionally, with roadside units (RSUs). Vehicles may have also a 3G/LTE radio interface, through which they can access the cellular network that fully covers the road topology. Both RSUs and cellular base stations allow vehicles to contact a well-known Location Authority (LA).

Vehicles are equipped with GPS, thus, unless otherwise specified, they know their own position and share a common time reference. Each vehicle owns cryptographic material, i.e., a certified identity, and a long-term secret key, that can be used to establish a secure channel with the LA at any time, through either an RSU or the cellular infrastructure.

Vehicles that comply with the A-VIP mechanism are defined as *correct*, while the others may be: (i) *faulty*, that is, they follow the protocol but provide incorrect information

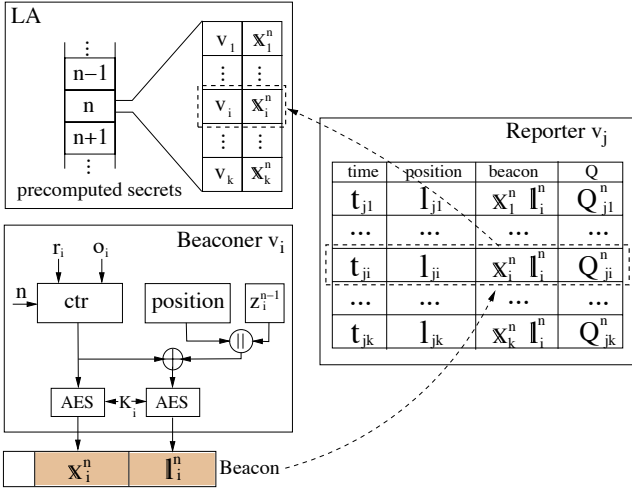


Fig. 1. Overview of A-VIP procedures by the beaconer, reporter and LA.

due to, e.g., GPS malfunctioning, or (ii) *adversarial*, i.e., their aim is to announce a fake position and have it verified, either to discredit nearby users, or to disrupt the A-VIP operation. To that end, adversarial nodes can either deviate from the A-VIP communication protocol procedure or comply with it, but injecting false information. In this work, we consider internal, independent adversaries, as colluding attacks are largely impractical in our vehicular scenario.

A. A-VIP goals

A-VIP aims at allowing an authority to (i) track vehicles, by verifying the positions they announce while guaranteeing their privacy with respect to other vehicles, and (ii) detect faulty or adversarial nodes and infer their actual locations. Such goals are achieved with low computational complexity.

B. Communication procedures

The procedures in the A-VIP protocol are described below, while a schematic overview is shown in Fig. 1.

Registration. The registration procedure takes place every time a vehicle is started, and is repeated after a registration validity time has expired. It is performed over the secure channel between the vehicle and the LA, established with the long-term key via the RSU infrastructure, if available, or through 3G/LTE, otherwise.

Let us assume that a generic vehicle v_i sends a registration request at time instant $t_{i,0}$. The LA records such an instant and returns to the vehicle a registration triplet (K_i, r_i, o_i) where K_i is a short-term 128-bit AES symmetric key, and r_i, o_i are random integers. The triplet is used to compute a time-dependent secret $\mathbb{x}_i(t)$, shared between the vehicle and the LA. As detailed later, when sent by v_i to the LA, $\mathbb{x}_i(t)$ allows the LA to verify the freshness of a beacon transmission and the identity of its originator. In order to compute it, the two entities initialize a counter to r_i and increment it by o_i every τ_b seconds, e.g., at every beacon transmission. The updated counter is then encrypted with K_i using AES in counter mode [7] (AES-CTR). Thus, in general, if $t_{i,0} + n\tau_b \leq t < t_{i,0} + (n+1)\tau_b$, then $\mathbb{x}_i(t) =$

$E_{K_i}\{r_i + no_i\} = \mathbb{x}_i^n$. Note that both r_i and o_i can be picked at random since the chances of collision among \mathbb{x}^n values, related to different vehicles at the same time, are negligible.

The LA is then in a position to precompute all the upcoming values of \mathbb{x}_i^n for a period that depends on the registration validity time.

Anonymous beaconing. When travelling, all correct vehicles broadcast a beacon every τ_b , as foreseen by current standards¹. Also, beacon transmissions occur at a power level common to all correct vehicles and at the basic data rate. We assume the beacon to be split into two parts: an encrypted one, for the purposes set forth in this paper and an unencrypted one, where plaintext content can be broadcast for such purposes as collision avoidance or cooperative applications. We assume however that the beacon is anonymous. When not transmitting, the vehicle listens to the channel, overhearing beacons broadcasted from other vehicles and collecting the information therein for later reporting to the LA.

The beacon content is assembled using the triplet assigned to a vehicle during the registration. Specifically, the n -th beacon issued by a vehicle v_i carries two pieces of information, as shown in the ‘‘Beaconer’’ box of Fig. 1:

- the time-dependent secret \mathbb{x}_i^n , which can be computed by v_i and by the LA, independently of each other;
- the encrypted current location announced by the vehicle $\mathbb{l}_i^n = E_{K_i}\{(l_i^n || z_i^{n-1}) \oplus (r_i + no_i)\}$, computed using the short-term pairwise key K_i from the triplet. The plaintext location l_i^n is concatenated with the one-bit flag z_i^{n-1} used to notify the LA whether the beacon issued at step $n-1$ was affected by a replay attack (as explained in Sec. IV). Such a string is then XOR’ed with the plaintext counter value $(r_i + no_i)$, to ensure freshness of the beacon positioning content and thwart partial-replay attacks (as also detailed in Sec. IV).

Reporting. When a beacon issued by a vehicle v_i is correctly received by a vehicle v_j , the latter is required to store the following entry in a *report table*, such as the one depicted in the ‘‘Reporter’’ box of Fig. 1:

- the time t_{ji} at which the beacon is received;
- its own position l_{ji} at the time the beacon was received;
- the secret \mathbb{x}_i^n carried in the beacon;
- the encrypted position \mathbb{l}_i^n of v_i carried in the beacon;
- a field Q_{ji}^n , indicating the received signal quality (e.g., the received signal power computed by the radio interface driver).

Every τ_r seconds (report interval), v_j generates a *report message* including the report table, populated with data collected from all newly overheard beacons. The report is transmitted to the LA, via the RSU or via the cellular network, ensuring authentication and integrity through standard procedures.

¹We acknowledge that there have been proposals to suppress beacons so as to reduce the channel contention [8]. We stress that our approach can easily keep that into account, by updating n at each τ_b , no matter whether a beacon is transmitted or not.

III. POSITION VERIFICATION AND INFERENCE

When the LA receives reports from vehicles, it processes them so as to (i) determine the locations announced by cars in the system, (ii) verify such locations and (iii) infer the actual positions of vehicles deemed to have advertised an incorrect location.

Let the LA divide the road topology into discretized spatial *tiles*, whose set is denoted by \mathcal{S} . Also, let \mathcal{V} be the set of vehicles that the LA has to verify. Upon receiving a report message from vehicle $v_j \in \mathcal{V}$, the LA processes one report table entry at a time, as follows:

- it extracts the time t_{ji} at which v_j received the beacon;
- for each $v_k \in \mathcal{V}$, it computes n such that $t_{k,0} + n\tau_b \leq t_{ji} < t_{k,0} + (n+1)\tau_b$, i.e., $n = \lfloor (t_{ji} - t_{k,0})/\tau_b \rfloor$, and it looks up the precomputed secret value \mathbb{x}_k^n that matches the \mathbb{x}_i^n in the report table entry (LA box in Fig. 1).

When a match is found, the LA identifies v_i as the vehicle that sent the beacon and retrieves the triplet associated to it. Then, the LA can take the following actions:

- it decrypts the location \mathbb{I}_i^n announced by v_i in the beacon reported by v_j ;
- it checks z_i^{n-1} ; if the bit is set, it discards the entry;
- else, if z_i^{n-1} is unset, it stores n , the position l_i^n included in the beacon by v_i , and the position l_j^n announced by v_j in the report table entry. The LA also stores the signal quality indicator, Q_{ji}^n , that v_j measured on the beacon received from v_i .

The LA leverages the information extracted from the report table entry to identify the possible tiles corresponding to a vehicle position. For the sake of clarity, in the remainder of this section we drop the time notation, thus assuming that all measures refer to the same beacon broadcast interval n .

We now briefly outline the approach adopted by the LA to identify the tiles corresponding to the position of the beaconer v_i , leveraging the signal quality Q_{ji} value of a beacon reception. For such an approach to be viable, the LA needs a model of the propagation conditions in the area where the broadcast transmission took place. Deterministic (e.g., ray-tracing), stochastic, or measurement-based models can be used to that end: the A-VIP procedure does not change and is performed as follows.

Let the propagation model be a function $h(s, t, Q_{ji}) : \mathcal{S}^2 \times \mathbb{R} \rightarrow [0, 1]$ that, for any pair of tiles (s, t) and any value of Q_{ji} , provides the probability $\mathbb{P}(R_t^{(j)} | B_s^{(i)}, Q_{ji})$ that a beacon sent by v_i from tile s can be received by v_j located in tile t , with the quality level Q_{ji} reported by v_j .

By applying Bayes' theorem, the LA can use such values to compute the probability $\mathbb{P}(B_s^{(i)} | R_t^{(j)}, Q_{ji})$ that the beaconer was in tile s , given that the beacon was heard by v_j in tile t , with a quality level Q_{ji} . Specifically,

$$\begin{aligned} p_{i,s}^{(j)} &= \mathbb{P}(B_s^{(i)} | R_t^{(j)}, Q_{ji}) \\ &= \frac{\mathbb{P}(R_t^{(j)} | B_s^{(i)}, Q_{ji}) \cdot \mathbb{P}(B_s^{(i)})}{\sum_{u \in \mathcal{S}} \mathbb{P}(R_t^{(j)} | B_u^{(i)}, Q_{ji}) \cdot \mathbb{P}(B_u^{(i)})} \end{aligned} \quad (1)$$

where $\mathbb{P}(B_x^{(i)})$, $x = s, u$, is the probability that the broadcasting vehicle v_i is in tile x at the moment of the transmission. This value depends on the vehicle density and the size of the considered area. For simplicity, we can assume $\mathbb{P}(B_s^{(i)}) = 1/|\mathcal{S}|$ for any v_i and any tile in $s \in \mathcal{S}$.

Upon receiving multiple reports, the LA can combine the above probabilities and compute the probability $P_{i,s}^{(n)}$ that v_i was in s while sending the n -th beacon, based upon the reports:

$$P_{i,s}^{(n)} = \frac{\prod_{j: v_j \in \mathcal{R}_i} p_{i,s}^{(j)}}{\sum_{u \in \mathcal{S}} \prod_{j: v_j \in \mathcal{R}_i} p_{i,u}^{(j)}} \quad \forall s \in \mathcal{S}, \quad (2)$$

where \mathcal{R}_i is the set of vehicles that reported v_i 's beacon.

An example is portrayed in Fig. 2, where two reporters, v_k and v_j , include different quality levels for a beacon received from v_i . For simplicity, in the figure we considered that the area corresponding to the value of Q , indicated by a reporter, maps onto an annulus comprised in its reception range. Then, the set of possible locations of the beaconer is the intersection of the two annuluses, shaded in Fig. 2(b).

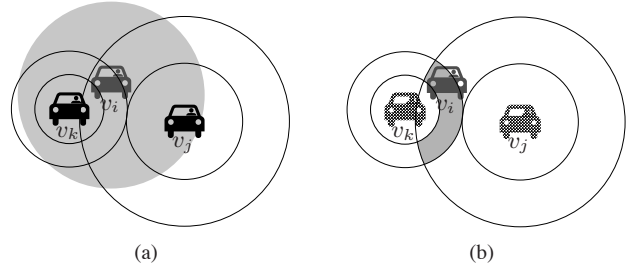


Fig. 2. The beacon broadcasted by v_i is reported by v_k and v_j . The shadowed area in (a) represents v_i 's transmission range. The annuluses denote the set of locations from which a beacon could be received by, respectively, v_k and v_j with the quality level indicated in their report. In (b), the intersection of the annuluses represents the possible positions of v_i .

IV. SECURITY ANALYSIS

Here, we discuss some possible attacks targeted at disrupting the position verification process described above.

Transmit-power attack. An attacker may maliciously increase or decrease its transmit power, thus affecting the Q -unaware and Q -aware approaches to the position verification and pretending to be closer or farther from the reporters than it actually is. However, while fooling a part of its neighbors, the attacker cannot help but appear inconsistent to the rest, since its announced position does not match the physical behavior of the transmission.

Replay attack. Such an attack has adversarial users replaying beacons from correct vehicles. Although the attacker can retransmit a copy of the beacon, it cannot tamper with its content, as both the secret \mathbb{x}_i^n and the beaconer position information are encrypted. We remark that encrypting the location together with the current counter value, as described in Sec. II-B, univocally ties it to \mathbb{x}_i^n . This prevents *partial replay attacks*, where the adversary only replays \mathbb{x}_i^n and modifies the position information \mathbb{I}_i^n .

Still, by performing a full replay at locations other than those of the original broadcast, the attacker could induce the LA to tag correct nodes as faulty. In such cases, the timing of the replay is of the essence:

- in case of a replay attack occurring more than τ_b seconds after the legitimate beacon was broadcast, the LA will no longer be able to match the secret in the beacon with any precomputed secret during that time frame, and the report table entry will be ignored;
- in case of a replay attack occurring less than τ_b seconds after the legitimate beacon was broadcast and reported multiple times by one or more witnesses, the LA will detect the duplicate entry and reject it.

In the latter case, the original beacon sender can detect the replay of its own beacon and report the misdeed by setting the z_i^{n-1} bit in its following beacon, as introduced in Sec. II-B. Recall that z_i^{n-1} can only be set by the original beacon sender, since it is encrypted along with the vehicle position and its freshness is ensured by the counter value. The LA will thus know that the beacon is invalid without affecting the vehicle credibility. The only result an attacker can achieve is thus to occasionally invalidate beacons from random vehicles. Jamming could yield the same effect with lower system complexity.

Wormhole attack. The replay attack can be combined with a wormhole attack, so that a full replay occurs less than τ_b seconds after the legitimate beacon was broadcast and in a different region (to avoid detection by the original sender). As a result, the replayed beacon will also be reported by witnesses other than the legitimate one. In this case, it is up to the LA to detect the inconsistency, by noting that the same beacon is heard by multiple witnesses farther apart than the nominal transmission range. The LA will thus be able to disregard both the original and replayed beacon entries. Additionally, the presence of a wormhole may be inferred, with the wormhole ends placed within the communication range of the reporting witness positions.

Phantom attack. An adversarial vehicle can run a phantom attack by never broadcasting beacons, nor reporting to the LA: such a vehicle would thus be completely transparent to the system. The advantages of such an adversarial strategy are however dubious. On the one hand, the attack could be used by a vehicle who is trying to escape liability after causing a car wreck. On the other, a phantom attacker who is falsely accused of being involved in an accident, would be unable to prove it was elsewhere.

Additionally, if the application is used for commercial purposes, such as to enforce e-toll, phantom attacks could pose a threat to the system. In such cases, the onboard devices are required to be tamper-resistant HSMs integrating the antenna apparatus, so that no vehicle can successfully disappear from the network.

Teleport attack. An adversarial user could impair local transmissions of its own beacons and use a wormhole to broadcast the same beacons at a location other than the one where it actually is. We refer to this as teleport attack,

enabling the adversary to, e.g., deny liability in any accident in which she is involved by beaconing at a distant, safe location. The same discussion as for the phantom attack applies here as well, and an integrated-antenna HSM is required to prevent teleport attacks when the goal is determining liability. Furthermore, A-VIP could be integrated with a dedicated solution for countering wormholes (see, e.g., [9]).

Sybil attack. It consists in a single vehicle employing multiple identities to corroborate its fake position advertisements. Through a sybil attack, an adversarial user could avoid broadcasting any beacon (i.e., perform a phantom attack), yet have multiple impersonated vehicles reciprocally (though falsely) report each other's beacons. Sybil vehicles can thus claim and mutually verify any possible position. However, in our system, identities cannot be fabricated but they must have been legitimately obtained, hence successively stolen by the adversary. Such a hurdle makes the sybil attack often infeasible, or only feasible for a short time before the identity theft is discovered.

Still, if one wants to provide an additional defense against sybil attacks, a practical solution would be to let the LA accept reports only if the following conditions are satisfied: (i) a vehicle sending the report has to exhibit at least a symmetric link with another vehicle, i.e., v_j has to report a beacon from v_i , which, in turn, reports a beacon from v_j ; (ii) a vehicle has to exhibit at least f different symmetric links over any back-to-back reporting intervals: the higher the f , the more the sybil identities required to successfully mount the attack.

In order to meet the above two conditions, an attacker has to transmit beacons, thus revealing its presence to others. It follows that, through the reports of other vehicles, the adversary may be detected.

Colluding attack. Colluders may only report each other's beacons, to corroborate their own false claims. The same discussion as for sybil attacks applies here as well. Colluding attackers have the further burden of continuous platooning to be successful, which makes this attack impractical.

V. TESTBED RESULTS

In this section, we aim at acquiring a better understanding of the position estimation provided by our framework. We implemented the A-VIP protocol on real hardware and tested it on up to five vehicles circulating on a 2-km loop on a public road, with a single deployed RSU. In order to at least alleviate the sparseness of the testbed scenario, we had vehicles travel within each other's range for most of the time. Vehicles and RSU are equipped with an Alix PC Engines motherboard and feature a IEEE 802.11h radio card. Vehicles carry one 5 dBi omnidirectional antenna on their rooftops, and transmit at an output power of 18 dBm. A GPS receiver provides localization information to each vehicle. The route traveled by the testbed vehicles is portrayed in the left image of Fig. 3, except for an additional final loop around the A point performed on a private road.

The propagation model, used for the computation of the

probabilities $p_{i,s}^{(j)}$, is derived from experimental measurements collected by the same vehicles used in the testbed. The corresponding propagation map is depicted in the right plot of Fig. 3, where, for clarity of presentation, the values of the received signal power have been discretized into high, medium and low signal quality bins.

In order to assess the quality of the position estimation in our experimental testbed, we defined a metric called *location error*, defined for a vehicle v_i at time n as:

$$e_i^n = \sum_{s \in \mathcal{S}} P_{i,s}^{(n)} d(\ell_i^n, s) \quad (3)$$

The location error is computed from the distances $d(\ell_i^n, s)$ between the actual location of v_i at time n , i.e., ℓ_i^n , and the centers of all tiles s . More precisely, such distances are averaged using as a weight the probability $P_{i,s}^{(n)}$ that v_i is within tile s at time n , inferred from reports coming in from nearby vehicles as per (2).

Since all vehicles are correct in our testbed, the uncertainty comes from (i) the RF signal propagation, which is time-varying and may induce errors in the estimation process, (ii) beacon and report message losses, impairing the verification process at the LA, and (iii) the beacon transmission frequency, which cannot capture the movement of vehicles within the beaconing interval τ_b .

In fact, we found the first two sources of errors, i.e., the propagation variability and the packet losses, to have a negligible impact on experimental results. It is instead the latter aspect, i.e., the beacon periodicity, that affects A-VIP operations in a more severe way. This is proven in Fig. 4, portraying the location error computed as in (3), versus the beaconing interval τ_b , with varying number of participating vehicles. The error for $\tau_b = 1$ s is in the order of the tile size (set to 10 m in these tests) which is the maximum spatial precision A-VIP can achieve. As the beaconing interval increases, the error significantly grows: cars are allowed to travel farther between back-to-back beacon transmissions, forcing the LA to estimate their intermediate locations with less data.

However, such an effect can be contrasted by a larger set of vehicles taking part to A-VIP operations. The LA then receives positioning information at a higher frequency, since vehicles transmit beacons in a de-synchronized way during τ_b and each beacon triggers reports by nearby cars.

Overall, the minor impact of physical layer phenomena on A-VIP performance is a positive factor. Most of the problems arise from the setting of τ_b ; however, this is a tweakable parameter, which our experimental evaluation suggests could be dynamically set, according to road traffic density.

VI. CONCLUSIONS AND FUTURE WORK

We presented A-VIP, a lightweight privacy-preserving framework for verification and inference of vehicle positions by a Location Authority. A-VIP leverages computationally-inexpensive symmetric cryptography and reporting of anonymized beacons by nearby vehicles. Through testbed measurements, we have shown A-VIP to achieve its goals

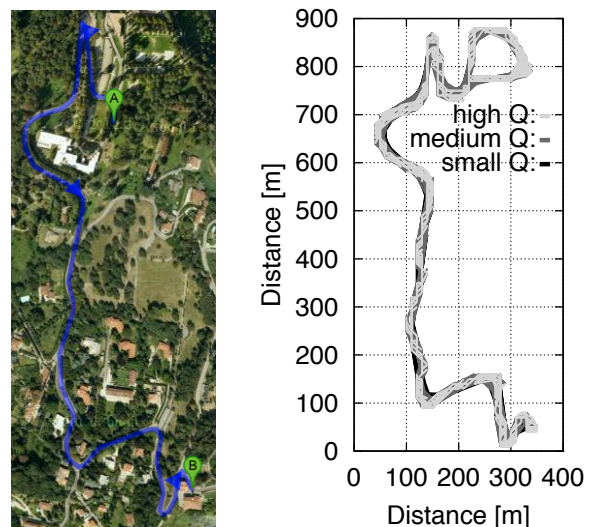


Fig. 3. Real-world (left) and signal propagation (right) maps.

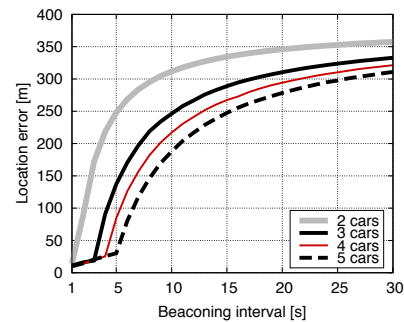


Fig. 4. Average location error versus the beaconing interval τ_b .

even in sparse vehicular settings with limited location error. Future work will address the evaluation of trust for vehicles involved in the cooperative location verification.

ACKNOWLEDGMENT

This work was supported by Regione Piemonte through the IoT_ToI project, and partially FAPESP and INCT-SEC.

REFERENCES

- [1] DSRC Message Set Dictionary, http://standards.sae.org/j2735_200911
- [2] ETSI TS 102 637-2: Intelligent Transport Systems (ITS) – Vehicular Communications – Basic Set of Applications – Part 2: Specification of Cooperative Awareness Basic Service, 2011.
- [3] Thales ISS, Thales, Jan. 2011 [Accessed on July 2012].
- [4] B. Wiedersheim *et al.*, “Privacy in Inter-Vehicular Networks: Why Simple Pseudonym Change Is not Enough.” *IEEE WONS*, 2010. Kranjska Gora, Slovenia, 2010.
- [5] E. Schoch, F. Kargl, “On the Efficiency of Secure Beaconing in VANETs,” *ACM WiSec*, 2010. 111-116
- [6] IEEE 1609.2 Draft Standard for Wireless Access in Vehicular Environments (WAVE) - Security Services for Applications and Management Messages, 2006.
- [7] W. Diffie, M. Hellman, “Privacy and Authentication: An Introduction to Cryptography,” *Proceedings of the IEEE*, vol. 67, no. 3, pp. 397–427, Mar. 1979.
- [8] M. Röckl, K. Frank, T. Strang, M. Kranz, J. Gacnik, J. Schomerus, “Hybrid Fusion Approach combining Autonomous and Cooperative Detection and Ranging methods for Situation-aware Driver Assistance Systems,” *IEEE PIMRC*, 2008.
- [9] J. Eriksson, S. Krishnamurthy, M. Faloutsos, “TrueLink: A Practical Countermeasure to the Wormhole Attack in Wireless Networks,” *IEEE ICNP*, 2006.