SLOW: A Practical Pseudonym Changing Scheme for Location Privacy in VANETs

Levente Buttyán[1] Tamás Holczer[1] André Weimerskirch[2] William Whyte[3]

[1] Laboratory of Cryptography and Systems Security (CrySyS)
Budapest University of Technology and Economics, Hungary
{buttyan, holczer}@crysys.hu

[2] escrypt Inc., USA
aweimerskirch@escrypt.com

[3] NTRU Cryptosystems, USA
wwhyte@ntru.com

Abstract

Untraceability of vehicles is an important requirement in future vehicle communications systems. Unfortunately, heartbeat messages used by many safety applications provide a constant stream of location data, and without any protection measures, they make tracking of vehicles easy even for a passive eavesdropper. One commonly known solution is to transmit heartbeats under pseudonyms that are changed regularly in order to obfuscate the trajectory of vehicles. However, this approach is effective only if some silent period is kept during the pseudonym change and several vehicles change their pseudonyms nearly at the same time and at the same location. Unlike previous works that proposed explicit synchronization between a group of vehicles and/or required pseudonym change in a designated physical area (i.e., a static mix zone), we propose a much simpler approach that does not need any explicit cooperation between vehicles and any infrastructure support. Our basic idea is that vehicles should not transmit heartbeat messages when their speed drops below a given threshold, say 30 km/h, and they should change pseudonym during each such silent period. This ensures that vehicles stopping at traffic lights or moving slowly in a traffic jam will all refrain from transmitting heartbeats and change their pseudonyms nearly at the same time and location. Thus, our scheme ensures both silent periods and synchronized pseudonym change in time and space, but it does so in an implicit way. We also argue that the risk of a fatal accident at a slow speed is low, and therefore, our scheme does not seriously impact safety-of-life. In addition, refraining from sending heartbeat messages when moving at low speed also relieves vehicles of the burden of verifying a potentially large amount of digital signatures, and thus, makes it possible to implement vehicle communications with less expensive equipments.

1. Introduction

Security in Vehicular Ad Hoc Networks (VANETs) is a topic of increasing theoretical and practical interest. European and American projects to implement vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) communications and thereby improve safety-of-life are progressing steadily [1, 2, 3, 4]. As deployment decision points for these projects draw nearer, the provision of adequate security mechanisms will be an important consideration for policy-makers.

In addition to the usual security requirements of confidentiality, authentication and integrity, VANET security typically presents an additional requirement, that of *privacy*. Informally, the privacy requirement represents a user's expectation that only appropriately authorized parties will be able to determine where he or she was at a given time. This informal definition may be formalized in many ways, and the definition of appropriately authorized parties may vary according to the circumstances and from jurisdiction to jurisdiction (or a

user may expect that no entity can track them at all).

As messages sent by the vehicles within the VANET may contain meta-information that endangers the privacy of the drivers, vehicle communication systems must satisfy the following two properties: *pseudonymity* and *unlinkability*. Pseudonymity means that identifiers in a message do not directly refer to the sender of the message, so an eavesdropper cannot easily determine the real identity of the sender. Unlinkability means that it is made difficult for an attacker to determine that two messages have come from the same vehicle. This second property is necessary to preserve privacy in the sense of our informal statement above because a physical observation of a vehicle at point A, and the ability to link its transmissions at A to transmissions at B, would allow an attacker to determine that the vehicle had also been at point B. Note that we do not address short-term linkability which is required in order to implement vehicle safety applications.

Making the security subsystem designer's job more complicated, most proposed V2X communications systems make use of an additional type of highly privacy-threatening message, known as the *heartbeat* (in America) or *beacon* (in Europe) message (see [5] for an example). This message is sent with a high frequency (10 Hz is often recommended) and contains the vehicle's current position and velocity, in order to improve the information that other drivers have about the traffic conditions in their immediate vicinity. An attacker can therefore attempt to trace a vehicle, and thereby break its location privacy, by "joining the dots" between two heartbeat messages with the same identifier (which we call *syntactic linking*) or by constructing a trajectory through a consistent series of (position, velocity) pairs (which we call *semantic linking*). The challenge to the security subsystem designer is to make these kinds of tracing as hard as possible.

This paper makes two main contributions: First, in Section 2, we provide a breakdown of the requirements that a system must address in order to provide privacy. The aim is to provide an analytical framework that future researchers can use to concisely state which aspects of privacy a new proposal does or does not address.

Second, we propose an approach for implementing mix zones that does neither require extensive RSU support nor complex communication between vehicles, and that does not endanger safety-of-life to any significant extent, while providing both *syntactic mixing* and *semantic mixing* (in the language of Section 2). To our knowledge, this is the first proposal that provides for semantic mixing while at the same time addressing the safety-of-life concerns that naturally arise when a vehicle tries to obscure its path. The key insights are simply that vehicles traveling at a low speed are less likely to cause fatal accidents, and that vehicles will be traveling at a low speed at natural mix-points such as signalled intersections. The main body of experimental work in this paper is therefore an investigation of the consequences for the untraceability of vehicles if they stop sending heartbeat messages when their speed drops below a certain threshold and change all their identifiers after such silent periods. We call our scheme SLOW, which stands for silence at low speeds. (We note that of course SLOW is not a full solution to untraceability, as it does not cover the safe use of silent periods at high speeds; other techniques will need to be used to give untraceability in this case).

Vehicles may of course choose to send heartbeats if necessary for safety-of-life reasons, for example if they sense an impending collision with a vehicle traveling above the threshold speed. Still, a large number of all vehicle interactions at intersections are non-life-threatening, therefore, assuming that exception cases can be properly defined and implemented, intersections (especially signalled intersections) seem to be a natural choice as practical "zones of silence" where large number of vehicles can mix. Hence, our dynamic mix zone that is automatically created around vehicles stopped at an intersection is likely to be maintainable in the great majority of cases.

Our work is inspired by the same insights as the work of [6]. However, [6] only addresses syntactic mixing, not semantic mixing, and requires the use of significant infrastructure. By replacing [6]'s cryptographic mix zones with zones of silence we address semantic mixing and infrastructure requirements simultaneously.

This paper has the following structure: We start by introducing our overall analytical framework in Section 2. We then survey previous work in Section 3. Next, in Section 4, we introduce our attacker model and our proposed solution, and in Section 5, we present the results of our experiments showing that our approach does indeed make tracing of vehicles hard for the attacker, and that it is usable in the real world. Finally, Section 6 presents conclusions and suggestions for further research.

2. Framework

Any system that aims to provide privacy for vehicles must address the following areas:

Syntactic privacy. In brief, all vehicles that use pseudonyms must change those pseudonyms from time to time. This area includes:

N1 *Pseudonymity*: An identifier that is available to an eavesdropper must not be directly linkable to the

vehicle (for example, it must not contain the VIN, the driver's name, or anything else an eavesdropper might know).

N2 *Change of identifiers*: Identifiers must change with some frequency[1].

N3 *Local synchronization of change of identifiers*: All identifiers, up and down the network stack, must change simultaneously. (This is not a communications issue as such, but a local engineering issue; however, it must be addressed).

N4 *Cooperative synchronization of change of identifiers* or *syntactic mixing*: A vehicle in an observed area must change its identifier at the same time as at least one other vehicle and the two (or more) changing vehicles must do so in a way that allows semantic privacy as defined below[2].

N5 *Pseudonym use*: This covers two intermingled areas:

N5.1 *Pseudonym format*: What cryptographic mechanism is used by psuedonym owners to authenticate that they are valid units within the system?

N5.2 *Pseudonym issuance and renewal*: How are pseudonyms issued? How does a vehicle avoid running out of them? (The answer to this may involve the identifier change frequency, N2.) What assumptions are necessary about the infrastructure to ensure that a vehicle is not left without pseudonyms?

Semantic privacy. This captures the idea that vehicles must not be traceable by reconstructing the trajectories implied by their heartbeat messages. This area includes:

M1 *Semantic unlinkability*: A vehicle's stream of heartbeat messages must be interrupted at some frequency for some period of time.

M2 *Semantic mixing*: Semantic unlinkability is valuable mainly in so far as it creates ambiguity for an attacker about whether a resumed stream of heartbeats comes from vehicle A or vehicle B.

Robust privacy. This captures how misbehaving entities within the system may affect privacy and security. This area includes:

R1 *Privacy-preserving bad-actor removal*: How is a misbehaving entity removed? Does this removal affect the privacy of its transmissions before it began to misbehave? Does its removal affect the privacy of other entities in the system?

R2 *Privacy against insider attacks*: How is privacy protected against bad actors in Law Enforcement or at a Certificate Authority (CA)?

This paper explicitly contributes in the area of syntactic mixing (N4), semantic mixing (M2), and semantic unlinkability (M1). Our results are based on the assumption that pseudonyms are changed whenever our criteria are met. This will be fairly frequent, on the order of once every few minutes for urban driving, implicitly addressing N2. An identifier change frequency this high may require frequent reissuance of pseudonyms, limiting the choices possible in areas N5.1 and N5.2. To the best of our understanding, our proposal is compatible with any reasonable solution for N1, N3, R1, or R2.

3. Related Work

There are a number of studies of pseudonym changes to assist syntactic unlinkability (N2). In [7], a periodic change of certificates is proposed based on the vehicle's driving and DSRC properties such as speed, transmission range, and transmission rate. The authors determine in their setting on a highway an appropriate time period for a certificate change of around one minute. Further approaches suggest changing pseudonyms once the best opportunity is identified. In [8], a vehicle first assesses its environment and determines how much uncertainty a pseudonym change at a given time would cause to the attacker. Once the level of expected uncertainty reaches a given threshold, a pseudonym change is triggered. A thorough analysis of the effectiveness of changing pseudonyms was performed in [9], where the authors show that even if pseudonyms are always successfully changed in an unobserved zone, the adversary is still able to trace vehicles with reasonable probability. The paper suggests that the success probability of the attacker saturates at around 0.6 for a strong adversary that observes more than 50% of the road network due to the non-uniformity of traffic. An interesting result is that the success rate mainly depends on the attacker's capabilities rather than on traffic density.

In [6], the authors suggest to construct mix zones for vehicles by cryptographic means. They propose to install such cryptographic mix-zones by deploying a special RSU at places with high traffic density such

[1]The frequency of change that provides privacy to the level expected by a user will in practice often depend on local regulation.

[2]Otherwise, an attacker who sees, for instance, identifiers (A,B,C,D) at time t and (A,B,C,E) at time $t+1$ will know that D and E refer to the same vehicle.

as crossroads. Once a vehicle enters a cryptographic mix-zone, they obtain a symmetric key from the RSU. While the vehicle is inside of the cryptographic mix-zone, all communication is encrypted and therefore an adversary cannot read-out useful information (including meta-information) from its messages. Vehicles in the mix-zone forward the symmetric key to vehicles that are in direct transmission range outside of the mix-zone such that these vehicles are also able to decrypt messages. Vehicles then change pseudonyms while being inside of the mix-zone. Another approach against global attackers but without infrastructure support was presented in [10, 11]. These papers suggest grouping vehicles together (for a few seconds) and introducing silent periods. Each vehicle group has a group leader that broadcasts information while the other vehicles are silent. Also, when vehicles change pseudonyms, they introduce a period of silence in order to reduce the available information for an attacker.

Another proposed approach provides multiple certificates in vehicles based on the combination of group signatures and multiple self-issued certificates [12, 13]. The disadvantage is that On Board Units (OBUs) need to perform expensive group signature verification operations, and that OBUs are empowered to mount Sibyl attacks. [14] uses group signatures to request temporary certificates from a CA in an anonymous manner without the disadvantages of the previous scheme, but at the cost of an available connection to the CA. Our solution suggested in the next section accounts for a global attacker without the support of the RSU infrastructure.

4. Attacker Model and Proposed Solution

We assume a global attacker that can get mass coverage. Conceptually, the attacker might be the RSU network operator that has access to messages received by all RSUs, or the attacker might have set up a network covering an entire city[3]. This is clearly an extremely powerful attack model, perhaps too powerful to be plausible, but we use this because if the system is secure in the face of this attacker it will be secure in the face of other, weaker attackers too.

The attacker can use two basic mechanisms to link transmissions from a vehicle: (1) linking pseudonyms or other identifiers between heartbeat messages (syntactic linking), and (2) using the position and velocity information in the heartbeat messages to reconstruct the trajectory of the vehicle (semantic linking).

We assume no supporting infrastructure in terms of an RSU network, therefore, vehicles must have a strategy to create their own mix zones, and that strategy must work even in the case where the attacker has 100% coverage. The defender's mechanism is to turn off radio transmissions (to make semantic linking difficult) and change pseudonyms (to make syntactic linking difficult) while the radio is turned off without endangering safety of life.

More precisely, the proposed solution, which we call SLOW for silence at low speeds, works as follows. We choose a threshold speed v_T, say $v_T = 30$ km/h. A vehicle will *not* broadcast any heartbeat message, or any other message containing location or trajectory data in the clear, if it is traveling below speed v_T, unless this is necessary for safety- of-life reasons. If the vehicle has not sent a message for a certain period of time, then it changes pseudonyms (identifiers at all layer of the network stack and related certificates) before the next transmission. Traffic signals in a crowded urban area seem like an ideal location for such a pseudonym change: whenever a crowd of vehicles stop at a traffic signal, they may go into one of several lanes, they may choose to turn or not turn, and so on. Thus, we create mix-zones at the point where there is maximum uncertainty about exactly where a vehicle is and exactly what it is going to do next. This is also a safe set of circumstances under which to stop transmitting. Only 5% of pedestrians struck by a vehicle at 20 km/h die [15] while at 50 km/h the figure is 40%. Presumably, vehicle-to-vehicle collisions where both cars are traveling at 30 km/h result in even fewer fatalities. Situations can be defined as exceptions. For instance, if vehicle A is stopped at a signal, but vehicle B coming up behind it emits a heartbeat that lets vehicle A know that there is a risk of a collision, then vehicle A can send out a heartbeat to warn vehicle B to brake. We note that our simulations do not include this exception case, because in practice these cases come up only rarely. Future research based on SLOW will investigate this exception case in greater detail. We also note that an attacker can abuse exception cases to break the silent period, but this attacker (unless it is an inside attacker) can be tracked down by standard methods and revoked.

Besides being very simple to implement, SLOW has other advantages. Traffic jams and slow traffic leads to a large amount of vehicles in transmission range and therefore requires extensive processing power to verify the digital signatures of all incoming heartbeat messages. By refraining from sending heartbeat messages, SLOW avoids the necessity of extensive signature verifications in traffic jams and slow traffic, and thus, reduces hardware cost. A more detailed analysis of the impact on computation complexity, as well as the level

[3] Fraunhofer Institute has established that the hardware cost (ignoring the backhaul connections) to set up receivers covering all 900 km^2 of Berlin is about 250,000 Euros.

of privacy and safety provided by our scheme will be presented in the next section.

5. Analysis

5.1. Privacy

It must be intuitively clear that a vehicle frequently sending out heartbeat messages is easy to trace, but to the best of our knowledge, no accurate experiment confirms this statement in VANET settings. As field experiments cannot be done due to the lack of envisioned VANET infrastructure, we carried out simulations to measure the level of traceability in an urban setting. We used the SUMO [16] simulation environment, as it is a realistic, microscopic urban traffic simulator. SUMO was set to use a 100 Hz frequency for internal update of vehicle position and velocities, and every Nth position (N depending on the heartbeat frequency) was considered to be available to the attacker as a heartbeat.

Note that tracing vehicles in an urban setting is essentially a multitarget tracking problem, which has an extensive literature, however, mostly related to radar development in the fields of aviation and sailing [17]. Yet, the following tracking approach, consisting of three steps, can be adopted to the vehicular setting too: First, the actual position and speed of the targets are recorded by eavesdropping the heartbeat messages. Based on the position and speed information, a predicted new position is calculated, which can be further refined by the help of side information such as the layout of the streets, lanes *etc.* At the next heartbeat, the new positions are eavesdropped and matched with the predicted positions.

We implemented an attacker that tracked the vehicles in the SUMO output based on the tracking approach described above. The attacker uses the last two heartbeat information to calculate the acceleration of the vehicles making the prediction of the next position more accurate. The vehicles are tracked from their departure to their destination. Tracking is considered successful, if the attacker has not lost a target through its entire journey.

The results of the tracking of 50 vehicles are shown in Figure 1. As we can see, if the beaconing frequency is 5-10 Hz, which is needed for most of the safety applications, then 75-80% of the vehicles are tracked successfully. By evaluating the unsuccessful cases, we observed that the target vehicles were lost at their destinations. More precisely, in the vast majority of the unsuccessful cases, when the target vehicle V_1 arrived to its destination and stopped sending more messages, if an other vehicle V_2 was in its vicinity, then the attacker continued tracking V_2 as if it was V_1. We counted this as unsuccessful case, because the attacker erroneously determined the destination of the target vehicle (i.e., it concluded that the destination of V_1 was that of V_2, and those two destinations have virtually never been the same). However, during the movement of the target vehicles (i.e., before they reached their destination), the attacker was able to track them with a remarkable 99% success rate. This confirms that semantic linking is a real problem.

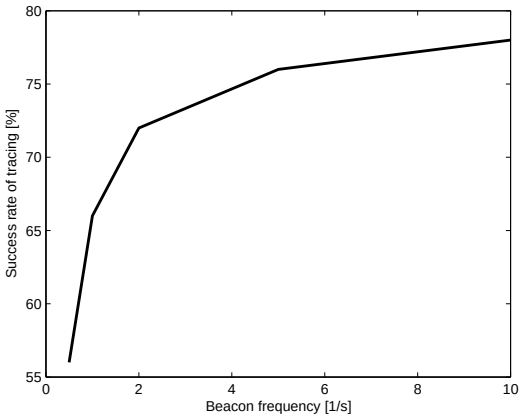

Figure 1. Success rate of an attacker performing vehicle tracking by semantic linking of heartbeat messages when no defense mechanisms are in use.

In any case, from a privacy point of view, a system where the users are traceable with probability 0.75-0.8 is not acceptable. Our proposed silent period scheme, where the vehicles stop sending heartbeat message below a given speed, mitigates this problem. It must be clear that the tracking algorithm described above does not work when the vehicles stop sending heartbeats regularly. Yet, the attacker may use other side information, such as the probability of turning to a given direction in an intersection, to improve the success probability of tracking despite the absence of the heartbeats. Thus, we need a new attacker model that also accounts for such side knowledge of the attacker.

We formalize the knowledge of the attacker as follows (for a summary of notations the reader is referred to Table 1): First, each intersection is modeled with a binary matrix J, where each row corresponds to an ingress lane and each column corresponds to an egress lane of the intersection, and J_{ij} (the entry in the i-th row and j-th column) is 1 if it is possible to traverse the intersection by arriving in ingress lane i and leaving in egress lane j. As an example, consider the intersection shown in Figure 2 and its corresponding matrix J defined in

Table 1. Notation

v_T	threshold speed
J	junction descriptor matrix
m	number of lanes towards the junction
n	number of lanes from the junction
T	probability distribution of the target's lanes
W	number of waiting vehicles per lanes
w	number of waiting vehicles in the junction
L	list of egress events
l_D	decision of the attacker
$\hat{l}$	the target's real egress event
L_S	list of suspect events

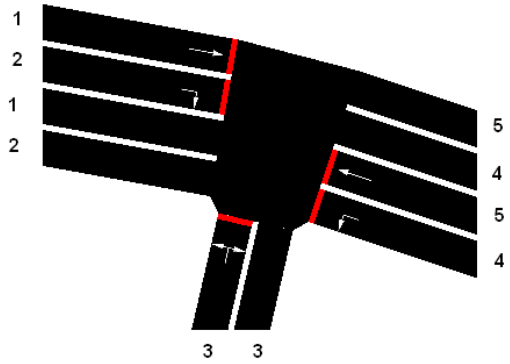

Figure 2. An example intersection, the corresponding matrix is given in (1)

(1).

$$J = \begin{pmatrix} 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 \end{pmatrix} \quad (1)$$

Second, we assume that the accuracy of GPS receivers does not permit to decide with certainty which lane of a road a given vehicle is using. Therefore, we also assume that the attacker knows on which road a target vehicle enters the intersection, but it does not know which ingress lane it is using. Nevertheless, the attacker may have some a priori knowledge on the probability of an incoming vehicle choosing a given ingress lane on a given road in a given intersection; such knowledge may be acquired by visually observing the traffic in that intersection for some time. These probabilities can be arranged in an m dimensional vector T, where the i-th element T_i is the probability of choosing ingress lane i when entering the intersection on the road that contains ingress lane i. As an example, consider the intersection in Figure 2, and the vector

$$T = (0.6, 0.4, 1, 0.8, 0.2)$$

This would mean that vehicles arriving to the intersection on the road that contains ingress lanes 1 and 2 choose lane 1 with probability 0.6 and lane 2 with probability 0.4. Note that vehicles arriving on the road that contains only ingress lane 3 have no choice, hence T_3 in this example is 1.

Third, when multiple possible egress lanes correspond to a given ingress lane (i.e., there are more than one 1s in a given row of matrix J), we assume that vehicles choose any of those egress lanes uniformly at random. For example, a vehicle arriving in ingress lane 1 of the intersection in Figure 2 can leave the intersection in egress lane 4 or 5 with equal probability.

Finally, when the target vehicle arrives at an intersection, there may already be some other vehicles waiting or moving below the threshold speed in that intersection. The number of such silent vehicles in ingress lane i is denoted by W_i, and the m dimensional vector containing all W_i values is denoted by W. Note that due to our previous assumption that the attacker is not always able to precisely determine the ingress lane used by an incoming vehicle, it is also unable to determine the exact values of all W_i's; nevertheless, it can use its experimental knowledge on the probabilities of choosing a given lane, represented by vector T, to at least estimate the W_i values.

Let us denote by L the list of vehicles that leave the intersection (and thus restart sending heartbeats) after the target entered the intersection (and thus stopped sending more heartbeats). More precisely, each element L_k of list L is a (timestamp, road) pair (t,r) that represents a vehicle reappearing on road r at time t. The objective of the attacker is to decide which L_k corresponds to the target vehicle. Let us denote by ℓ the list element chosen by the attacker, and let ℓ^* be the list element that really corresponds to the target vehicle. The attacker is successful if and only if $\ell = \ell^*$.

In theory, the optimal decision is the following:

$$\ell = \arg\max_k \Pr(L_k | J, T, W, L)$$

where $\Pr(L_k|J,T,W,L)$ is the probability of L_k being the right decision given all the knowledge of the attacker. However, it seems to be difficult to calculate (or estimate) all these conditional probabilities, as they have to be determined for every possible intersection (J), number of awaiting vehicles in the intersection (W), and observation of egress events (L).

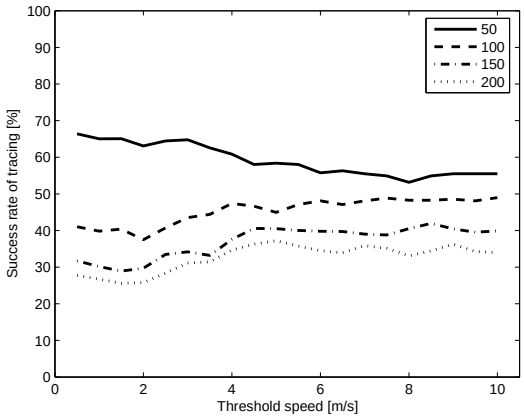

Figure 3. Success rate of our simple attacker in a single intersection. Different curves belong to different experiments with the total number of vehicles given in the legend.

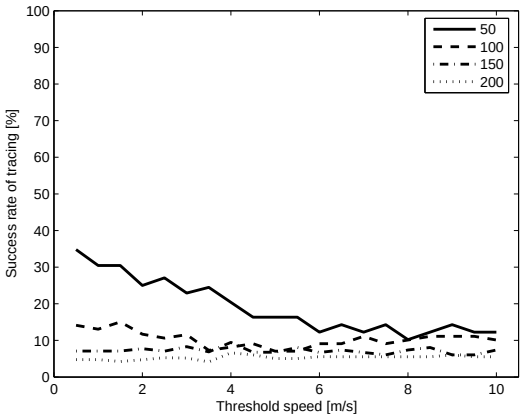

Figure 4. Success rate of our simple attacker in the general case, when the target traverses multiple intersections between its starting and destination points. Different curves belong to different experiments with the total number of vehicles given in the legend.

Hence, we assume a more simplistic attacker that uses the following tracking algorithm: Let us denote by w the total number of silent vehicles in the intersection when the target vehicle arrives and stops sending heartbeats. The attacker decides on the w-th element of L, unless that entry surely cannot correspond to the target (e.g., it is not possible to leave the intersection on the road in the w-th element of L given the road on which the target arrived to the intersection). When the w-th element of L must be excluded, the attacker chooses the element that is the closest in the list L and that cannot be excluded.

Our simple attacker model essentially assumes that traffic at an intersection follows the FIFO (First In First Out) principle. While this is clearly not the case in practice, our attacker still achieves a reasonable success rate in a single intersection as shown in Figure 3. One can see, for instance, that when the total number of vehicles is 100, the attacker can still track a target vehicle through a single intersection with probability around $\frac{1}{2}$.

Figure 4 shows the success rate of the attacker in the general case, when the target traverses multiple intersections between its starting and destination points. As expected, the tracking capabilities of the attacker in this case are worse than in the single intersection case. The quantitative results of our simulation experiments suggest that only around 10% of the vehicles can be tracked fully by the attacker when the threshold speed is larger than 22 km/h (approximately 6 m/s).

The effectiveness of the attacker depends on the v_T threshold speed and the density of the vehicles. In general the higher the threshold speed at which vehicles stop sending heartbeats, the higher the chance that the attacker loses the target (i.e., the lower the chance of successful tracking). Moreover, in a dense network, it is more difficult to track vehicles. Note, however, that there is an important difference in practice between the traffic density and the threshold speed, namely, that the threshold speed can be influenced by the owner of the vehicle, while the traffic density cannot be.

5.2. Effects on safety

The main objective of vehicular communications is to increase road safety. However, refraining from sending heartbeat messages may seem to be in contradiction with this objective. Note, however, that we propose to refrain from sending heartbeats only below a given threshold speed, and we argue below that this may not endanger the objective of road safety.

According to [15], only 5% of pedestrians struck by a vehicle at 20 km/h die, while this figure is 40% at 50 km/h. In [18], it is shown that in a 60 km/h speed limit area, the risk of involvement in a casualty crash doubles with each 5 km/h increase in traveling speed above 60 km/h. In [19], it is shown that 1 km/h change in speed can influence the probability of an accident by 3.45%.

The statistical figures above show that at lower speed the probability of an accident is lower too. This is because usually vehicles go at lower speed in areas where the drivers need to be more careful (hence the

speed limit). Thus, it makes sense to rely more on the awareness of the drivers to avoid accidents at lower speeds. On the other hand, at higher speeds, accidents can be more severe, and warning from the vehicular safety communication system can play a crucial role in avoiding fatalities.

5.3. Effects on computation complexity

A great challenge in V2V communication deployment is the processing power of the vehicles [20]. The most demanding task of the On Board Unit (OBU) is the verification of the signatures on the received heartbeat messages. This problem can be partially handled by not attaching certificates to every heartbeat message [12], but it does not solve the problem of verifying the signatures on the messages.

In principle, the heavier the traffic, the more vehicles are in each others communication range. More vehicles send more heartbeats overwhelming each other. The number of vehicles in communication range depends on the average speed of the traffic, assuming that the vehicles keep a safety distance between each other depending on their speed.

In Figure 5, the results of some simple calculations can be seen showing the number of signature verifications performed as a function of the average speed. In this calculation, we assumed that vehicles follow each other within 2 seconds. The communication range is assumed to be 100 m and the heartbeat frequency is 10 Hz. It can be seen in the figure that, in a traffic jam on an 8-lane road, each vehicle must verify as many as approximately 8,000 signatures per second. If SLOW is used with a threshold speed of around 30 km/h (approximately 8 m/s), then the vehicles never need to verify more than 1,000 signatures per second. (assuming all other parameters are the same as before). This approach also works well in combination with congestion control where the transmission power is reduced in high density traffic scenarios. Our approach therefore makes the hardware requirements of the OBU much lower and enables the use of less expensive devices.

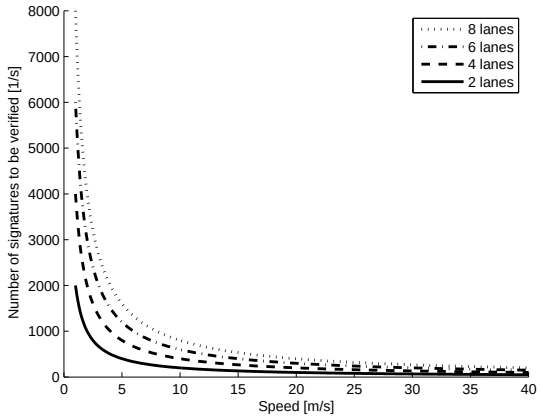

Figure 5. Number of signatures to be verified as a function of the average speed. The communication range is 100 m, and the heartbeat frequency is 10 Hz. Safety distance between the vehicles depends on their speed.

6. Conclusion and Future Research

In this paper, we proposed a simple and effective privacy preserving scheme, called SLOW, for VANETs. SLOW requires vehicles to stop sending heartbeat messages below a given threshold speed (this explains the name SLOW that stands for "silence at low speeds") and to change all their identifiers (pseudonyms) after each such silent period. By using SLOW, the vicinity of intersections and traffic lights become dynamically created mix zones, as there are usually many vehicles moving slowly at these places at a given moment in time. In other words, SLOW implicitly ensures a synchronized silent period and pseudonym change for many vehicles both in time and space, and this makes it effective as a location privacy enhancing scheme. Yet, SLOW is remarkably simple, and it has further advantages. For instance, it relieves vehicles of the burden of verifying a potentially large amount of digital signatures when the vehicle density is large, as this usually happens when the vehicles move slowly in a traffic jam or stop at intersections. Finally, the risk of a fatal accident at a slow speed is low, and therefore, SLOW does not seriously impact safety-of-life.

We evaluated SLOW in a specific attacker model that seems to be realistic, and it proved to be effective in this model, reducing the success rate of tracking a target vehicle from its starting point to its destination down to the range of 10–30%. A possible future extension of our work would be to investigate further attacker models and to study other metrics of privacy beyond the one we used that is based on the success probability of an attacker that attempts to track vehicles.

Possibilities for future research include the following:

1. Reducing heartbeat rates as the vehicle's speed reduces, rather than eliminating them altogether.

2. Further consider what the threshold speed should be, and what the rules governing exceptions should be, taking into account real-world data about in-

tersection collisions. For example, although collisions at 30 kmh are only occasionally fatal, a head-on collision between two vehicles travelling at 30 kmh each is effectively at 60 kmh.

References

[1] IntelliDrive Project, "www.intellidriveusa.org/."

[2] CVIS Project, "http://www.cvisproject.org/."

[3] SAFESPOT Project, "http://www.safespot-eu.org/."

[4] SeVeCom Project, "http://www.sevecom.org/."

[5] Society of Automotive Engineers, "Dedicated short range message set (dsrc) dictionary," SAE, Tech. Rep. Standard J2735, 2006.

[6] J. Freudiger, M. Raya, M. Felegyházi, P. Papadimitratos, and J.-P. Hubaux, "Mix-zones for location privacy in vehicular networks," in *Proceedings of the 1st International Workshop on Wireless Networking for Intelligent Transportation Systems (WiN-ITS 07)*, 2007.

[7] M. Raya, "The security of vehicular ad hoc networks," in *Proceedings of the 3rd ACM workshop on Security of ad hoc and sensor networks*. ACM New York, NY, USA, 2005, pp. 11–21.

[8] M. Gerlach, "Assessing and improving privacy in vanets," in *Proceedings of the Workshop on Embedded Security in Cars (escar06)*, 2006.

[9] L. Buttyán, T. Holczer, and I. Vajda, "On the effectiveness of changing pseudonyms to provide location privacy in vanets," in *In Proceedings of the Fourth European Workshop on Security and Privacy in Ad hoc and Sensor Networks (ESAS2007)*. Springer, 2007.

[10] K. Sampigethaya, L. Huang, M. Li, R. Poovendran, K. Matsuura, and K. Sezaki, "CARAVAN: Providing location privacy for VANET," in *Proceedings of Embedded Security in Cars (ESCAR 2005)*, 2005.

[11] K. Sampigethaya, M. Li, L. Huang, and R. Poovendran, "Amoeba: Robust location privacy scheme for vanet," *IEEE Journal on Selected Areas in Communications*, vol. 25, no. 8, pp. 1569–1589, 2007.

[12] G. Calandriello, P. Papadimitratos, J.-P. Hubaux, and A. Lioy, "Efficient and robust pseudonymous authentication in vanet," in *VANET '07: Proceedings of the fourth ACM international workshop on Vehicular ad hoc networks*. New York, NY, USA: ACM, 2007, pp. 19–28.

[13] F. Armknecht, A. Festag, D. Westhoff, and K. Zeng, "Cross-layer privacy enhancement and non-repudiation in vehicular communication," in *4th Workshop on Mobile Ad-Hoc Networks (WMAN)*, 2007.

[14] A. Studer, E. Shi, F. Bai, and A. Perrig, "TACKing Together Efficient Authentication, Revocation, and Privacy in VANETs," Carnegie Mellon CyLab, Tech. Rep., 2008.

[15] W. Leaf and D. Preusser, "Literature review on vehicle travel speeds and pedestrian injuries," National Highway Traffic Safety Administration, http://www.nhtsa.dot.gov/people/injury/research/pub/HS809012.html, October 1999.

[16] D. Krajzewicz, G. Hertkorn, C. Rössel, and P. Wagner, "Sumo (simulation of urban mobility); an open-source traffic simulation," in *Proceedings of the 4th Middle East Symposium on Simulation and Modelling (MESM2002)*, A. Al-Akaidi, Ed. Sharjah, United Arab Emirates: SCS European Publishing House, September 2002, pp. 183–187.

[17] M. Gruteser and B. Hoh, "On the anonymity of periodic location samples," in *Proceedings of the Second International Conference on Security in Pervasive Computing*. Springer, 2005, pp. 179–192.

[18] C. Kloeden, A. McLean, V. Moore, and G. Ponte, "Travelling speed and the risk of crash involvement," NHMRC Road Accident Research Unit, The University of Adelaide, 1997.

[19] A. Baruya, "Speed-accident relationship on different kinds of european roads," MASTER Deliverable 7, September 1998.

[20] F. Kargl, A. Kung, A. Held, G. Calandriello, T. V. Thong, B. Wiedersheim, E. Schoch, M. Müter, L. Buttyán, P. Papadimitratos, and J.-P. Hubaux, "Secure vehicular communication systems: implementation, performance, and research challenges," *IEEE Communications Magazine*, vol. 46, no. 11, pp. 110–118, 2008.