

Mix-Zones for Location Privacy in Vehicular Networks

Julien Freudiger, Maxim Raya, Márk Félegyházi,
Panos Papadimitratos and Jean-Pierre Hubaux
EPFL, Switzerland
firstname.lastname@epfl.ch

ABSTRACT

Vehicular Networks (VNs) seek to provide, among other applications, safer driving conditions. To do so, vehicles need to periodically broadcast safety messages providing precise position information to nearby vehicles. However, this frequent messaging (e.g., every 100 to 300ms per car) greatly facilitates the tracking of vehicles, as it suffices to eavesdrop the wireless medium. As a result, the drivers privacy is at stake. In order to mitigate this threat, while complying with the safety requirements of VNs, we suggest the creation of mix-zones at appropriate places of the VN. We propose to do so with the use of cryptography, and study analytically how the combination of mix-zones into mix-networks brings forth location privacy in VNs. Finally, we show by simulations that the proposed mix system is effective in various scenarios.

1. INTRODUCTION

Vehicular Networks (VNs) consist of vehicles and Road-Side Units (RSUs) equipped with radios. Using Vehicle-to-Vehicle (V2V) and Vehicle-to-Infrastructure (V2I) communications, vehicles share safety-related information and access location-based services. Initiatives in Europe [6] and the US [8] are evaluating VNs promises of safer driving conditions and more efficient traffic management. Envisioned safety-related applications require the vehicles to periodically broadcast their current position, speed and acceleration in authenticated *safety messages*. This messaging increases the awareness of vehicles about their neighbors' whereabouts and warns drivers off dangerous situations.

The nature of the wireless communications makes eavesdropping particularly easy. All an adversary needs to do is to deploy its devices across the area of the network that it wishes to monitor. At the same time, safety messages provide rich information on their senders, for example, the vehicle's location. This essentially allows automatic tracking of the vehicle whereabouts. Thus, it can reveal private information regarding the activities of the driver. The wide

availability of VN-compatible radios (802.11-based) makes such a threat even more credible.

The use of randomly changing identifiers (i.e., *pseudonyms*) has been proposed to decorrelate the identity of vehicles from their locations. The purpose of such a scheme is to achieve *unlinkability* between the vehicle and its pseudonyms in the long run. However, updating the pseudonym of a vehicle in a monitored region is ineffective, because the location information of safety messages can still be used for tracking. Therefore, changing pseudonyms is effective only within regions in which monitoring is impossible.

In this paper, we address the problem of achieving *location privacy* in VNs with randomly changing identifiers in the presence of a global passive adversary. We are fully aware of the need for strong security in VNs, along with privacy protection [23, 19, 22, 20, 18, 4]. In particular, in this context, pseudonyms are anonymized public keys. Our proposal fits in this framework of pseudonymous authentication. Our contribution is threefold. First, we propose a protocol to create cryptographic mix-zones at road intersections. This solution thwarts computationally-bounded eavesdroppers while preserving the functionality of safety messages. Second, we analyze the location privacy achieved by combining mix-zones into mix-networks in VNs. These so-called *vehicular mix-networks* leverage on the mobility of vehicles and the dynamics of road intersections to mix vehicle identifiers. Finally, we show by means of simulations the effectiveness of the proposed mix system.

The remainder of the paper is organized as follows. In Section 2, we present the VN system model and privacy-related adversary model. In Section 3, we introduce the CMIX protocol. In Section 4, we establish an analytical model of the amount of location privacy achievable in VNs with mix-zones. In Section 5, we evaluate the effectiveness of vehicular mix-networks by means of simulations. In Section 6, we discuss the global passive internal adversary threat model and alternative CMIX protocol designs. In Section 7, we present previous contributions on location privacy in mobile networks. Finally, we conclude the paper in Section 8.

2. SYSTEM AND THREAT MODEL

In this section, we briefly introduce VNs and describe the threat model we consider throughout the paper.

2.1 Vehicular Networks

Vehicular Networks (VNs) consist of vehicles, Road-Side Units (RSUs) and a collection of backbone servers accessible via the RSUs. We assume that a single VN Operator (VNO)

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WiN-ITS 2007 Vancouver, British Columbia, Canada
Copyright 200X ACM X-XXXXX-XX-X/XX/XX ...\$5.00.

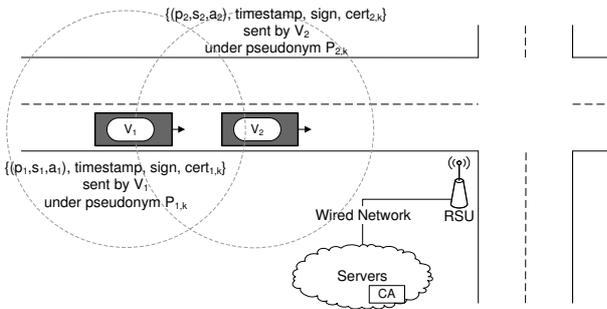


Figure 1: Vehicular Networks. Safety messages are emitted periodically (typically, every 100 ms to 300 ms) [8]. The Certification Authority (CA) is accessible through the RSUs. p_i , s_i and a_i are the vehicle i position, speed, and acceleration.

is responsible for deploying RSUs in the network. Due to their relevance to life-critical applications, VNs have to satisfy several strict requirements, namely *sender and data authenticity*, *availability*, *liability*, and *real-time delivery*. VNs are a very challenging application of ad-hoc networks as all the above prerequisites must be achieved under the stringent conditions created by a highly dynamic mobile environment.

To prevent accidents and inform each other of dangerous situations, vehicles periodically broadcast *safety messages* indicating their position, speed, acceleration, and possibly many other types of information (Fig. 1). We assume that each vehicle is equipped with a GPS device that provides accurate location information within an acceptable error margin.

Similarly to the work in [23, 19, 22, 18], we assume that a suitable public key infrastructure is available in VNs and that the messages are properly signed to ensure the liability of their sender in case of an accident. A certificate is attached to each message to enable other vehicles to verify the sender’s authenticity. Vehicles are equipped with *Tamper-Proof Devices* (TPDs) that guarantee the correct execution of cryptographic operations and the non-disclosure of private keying material. TPDs come with their own battery and clock. Prior to entering the network, each vehicle i has to register with a Certification Authority (CA) and preloads a large set of *pseudonyms* $P_{i,k}$, with $k = 1, \dots, F$, where F is the size of the pseudonym set. The CAs are fully *trusted* third parties and interoperable entities, operated by governmental organizations, that conform to privacy policies and keep the relation of the pseudonyms to the driver’s real identity secret. In case of liability issues, this relation can be made public by law enforcement. For each pseudonym $P_{i,k}$ the corresponding CA generates a unique public/private key pair $(K_{i,k}, K_{i,k}^{-1})$ and a corresponding certificate $Cert_{i,k}(K_{i,k})$.¹ Each vehicle sequentially updates its pseudonym at regular time intervals independently of other vehicles. Pseudonyms have a short validity period and cannot be reused.

2.2 Threat Model

VNs significantly facilitate the tracking of vehicles. In fact, the cost of tracking vehicles by radio eavesdroppers

¹The pseudonym can be the public key itself

is reduced compared to that of tracking vehicles with cameras. Similarly, the tracking efficiency is increased because an eavesdropper obtains identifiers, location and other information from safety messages.² Hence, the tracking granularity is high. Moreover, unlike mobile phones and laptop wireless adapters, vehicle transceivers cannot be switched off [20]. Consequently, vehicles’ whereabouts can be monitored at all times. This can provide significant information about their drivers’ activities, and can be exploited for targeted advertisement or surveillance. For example, an adversary can deduce from mobility traces the home and work place of a driver, and hence his real world identity [16].

The adversary is either internal or external, that is, utilizes devices that are legitimate members of the VN [23, 18], or any other radio transceiver [19]. In this paper, we are concerned with achieving location privacy against an external adversary.

An *external* adversary installs its own radio receivers near the road network and *passively* eavesdrops vehicle safety messages. Outside the range of its radio receivers, the adversary cannot overhear transmissions. Thus, its strength depends on the number of its eavesdropping devices. A *global* adversary has a complete view of the monitored network. Such an adversary can be put in place by exploiting already deployed 802.11 networks. For example, wireless social communities (e.g., FON [10]), or WiFi operators (e.g., Google) provide low cost wireless internet connectivity via WiFi networks in cities. With minor software or hardware modifications, this infrastructure can eavesdrop VN communications.

On the other hand, setting up a network of internal eavesdroppers would be much harder. The adversary would need to obtain legitimate devices, e.g., vehicles equipped with transceivers. The use of a TPD prevents adversaries from compromising cryptographic material. However, the VNO, which is a partially trusted third party, could be enticed to passively monitor the position of vehicles. We do not consider this type of adversary. We assume instead in this paper that the VNO assists in setting up privacy protection mechanisms.

3. CRYPTOGRAPHIC MIX-ZONES IN VEHICULAR NETWORKS

Anonymous systems, as described by Chaum [5], aim at increasing the adversary’s workload to uniquely identify the *author* of an *action*. In this section, we present a cryptographic technique to create anonymizing regions, that is, mix-zones [2] in VNs. The idea for mix-zones is to prevent the adversary from accessing the content of (safety) messages, including the vehicle’s signatures that are trivially linkable to the corresponding pseudonym, and thus be unable to connect two pseudonyms successively used by the same vehicle.

The effectiveness of anonymizing regions in providing location privacy depends on the density of vehicles and the unpredictability of their whereabouts. We propose to create mix-zones at predetermined locations and to force pseudonym changes to take place within those regions. Because the highest mixing of vehicles occurs at road intersections where

²We do not consider the monitoring of other type of vehicular communications (e.g., infotainment) that do not contain precise location information.

the speed and direction of vehicles change the most (i.e., it is an appropriate mix context [11]), we propose placing mix-zones at road intersections. We assume that all vehicles participate in the anonymization process at every road intersection.

3.1 The CMIX Protocol

Since the location of mix-zones is fixed, the adversary can identify them and thus could easily attempt to eavesdrop transmissions originating in the mix-zone area. To solve this problem, we introduce the *CMIX Protocol* to create Cryptographic MIX-zones (CMIXes): all legitimate vehicles within the mix-zone obtain a symmetric key from the road-side unit (RSU) of the mix-zone, and utilize this key to encrypt all their messages while within the zone. The symmetric key is obtained through a key establishment phase. To ensure the functionality of safety messages, this mix-zone key can be obtained by nodes approaching the mix-zone with the help of a key forwarding mechanism, and, finally, the RSU can swap to a new key through a key update mechanism.

3.1.1 CMIX Key Establishment

Vehicles rely on the presence of RSUs at road intersections to initiate a *Key Establishment* mechanism and establish a symmetric key. RSUs advertise their presence by periodically

$$\begin{aligned} v_i \rightarrow \text{RSU}: & \text{Request}, T_s, \text{Sign}_i(\text{Request}, T_s), \text{Cert}_{i,k} \\ \text{RSU} \rightarrow v_i: & E_{K_{i,k}}(v_i, SK, T_s, \text{Sign}_{\text{RSU}}(v_i, SK, T_s)), \text{Cert}_{\text{RSU}} \\ v_i \rightarrow \text{RSU}: & \text{Ack}, T_s, \text{Sign}_i(\text{Ack}, T_s), \text{Cert}_{i,k} \end{aligned}$$

Table 1: The Key Establishment protocol. T_s is a time stamp, $\text{Sign}()$ is the signature of the message, Cert is the certificate of the message sender.

cally broadcasting beacons. As soon as a vehicle v_i enters in the of transmission range of an RSU, R_{Beacon} , it initiates the key establishment protocol described in Table 1. As the vehicle knows its own location and the location of the RSU (announced in the beacon), it can estimate whether it is within the mix-zone, defined by a transmission range $R_{\text{CMIX}} < R_{\text{Beacon}}$. If so, the vehicle v_i broadcasts one or, if needed, several key request messages (first message in Table 1). The RSU replies with the symmetric key SK encrypted with the public key of vehicle v_i and a signature. v_i receives and decrypts the message. If the message is validated, v_i acknowledges it and SK can be used to encrypt all subsequent safety messages until v_i leaves the mix-zone. In case RSUs are co-located (i.e., their mix-zones overlap), vehicles are aware of all CMIX keys so that they can decrypt all messages. Alternatively, co-located RSUs could coordinate to use the same CMIX key.

3.1.2 Key Forwarding

Vehicles in the *extended mix-zone*, that is, at a distance d from the RSU where $R_{\text{CMIX}} < d < R_{\text{Beacon}}$ may be unable to obtain directly the key from the RSU; for example, they are beyond their transceiver’s range for bidirectional communication. Thus, they cannot decrypt safety messages coming out of the CMIX. As vehicles know they are within an RSU transmission range, when they receive encrypted safety messages, they issue one or, if needed, several key requests to obtain the SK key with the help of vehicles already in the mix-zone which are aware of it.

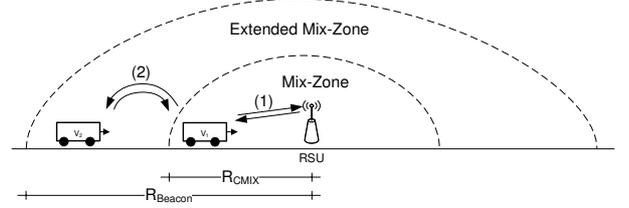


Figure 2: Extended Mix-zones. (1) v_1 uses the Key Establishment to learn the symmetric key. (2) v_2 uses the Key Forwarding protocol.

Consider the example of Figure 2: vehicle v_1 already knows the CMIX key and can forward it to v_2 . Hence, the RSU leverages on the vehicles in the mix-zone. In our example, when v_2 enters the extended mix-zone, as soon as it receives an encrypted (intelligible) message, it initiates the broadcast of one or, if needed, several key requests. v_1 eventually receives a key request from v_2 , and forwards it the symmetric key:

$$E_{K_{2,k}}(v_2, v_1, SK, T_s, \text{Sign}_{\text{RSU}}(v_1, SK, T_s))$$

The signature from the RSU, along with the time stamp, allows to validate the transmitted symmetric key. Note that vehicles in the extended region do not encrypt their safety messages with the CMIX key before entering the mix-zone (R_{CMIX}). The entire above message is in addition signed by v_1 .

3.1.3 Key Update

We propose a *Key Update* mechanism to renew or revoke CMIX symmetric keys. The RSU is responsible for such key updates and determines when to initiate the process. Key updates occur only when the mix-zone is empty and vehicles obtain new keys via the key transport and key forward protocols. The CA obtains the new symmetric key from the RSU over a secure channel, to satisfy the liability requirements (i.e., possibly, decrypt safety messages in the future). The robustness provided by the system is increased, if key updates are asynchronous across different base stations. But there is a trade off between security and cost, as frequent updates can incur additional overhead.

3.2 Analysis of the CMIX Protocol

The CMIX protocol requires the exchange of two messages. One or several key request messages are sent until either an RSU or a vehicle receive it. Such transmission overhead can be kept low: in a dense traffic scenario, one key request should suffice before receiving a reply, whereas in a low-density scenario the message overhead has low impact. When a key request is broadcasted, potentially every vehicle in the transmission range could send back a key reply. To avoid such reply flooding, a number of mechanisms can be used (e.g., random backoff mechanism); we will evaluate those in future work. Upon receipt of the mix-zone key, the vehicle sending the key request acknowledges the acquisition of the key, to prevent additional neighboring vehicles from forwarding again the key. Cryptographic overhead can also be easily sustained with relatively low delay [21].

In terms of security, the adversary is computationally bounded and unable to launch brute-force cryptanalytic attacks on the mix-zone encrypted messages. Because messages are au-

thenticated, an external adversary cannot impersonate vehicles, and replay attacks would not be successful neither thanks to time stamps. Similarly, the adversary cannot forge RSU messages or impersonate an RSU; as a result, cannot create fictitious mix-zones with symmetric keys it controls. By appending the node identifier to the key reply, an adversary is prevented from executing several instantiations of the key establishment protocol in parallel.

To impede the proper unfolding of the protocol, an external adversary can attempt to selectively jam messages from the RSU containing the symmetric key and force vehicles into transmitting their safety messages in clear. Still, the key forward protocol can increase the resilience. However, jamming is largely orthogonal to our problem. If nodes (detect they) are jammed, communication is impossible in the first place. In case of being unable to obtain the CMIX key because of denial of service, vehicles can always revert to regular operation, ignoring the mix-zone.

Furthermore, an internal adversary could simply purchase as many vehicles as the number of intersections it wishes to monitor. Then, every time a new update takes place, it would legitimately obtain from the RSU all mix-zone keys. Protection against such high-cost adversaries (recall: the TPD protects from illegitimate extraction of credentials and keys from the vehicle, and purchasing vehicles is costly) is part of future work.

Finally, an adversary can monitor vehicles entering and exiting the CMIX and compute the most probable mappings of events. We evaluate in the next section the amount of information this kind of adversary can gather. We show that the construction of mix-zones into mix-networks in VNs increases the robustness of the system against privacy degradation attacks by accumulating the anonymity over several mix-zones.

4. LOCATION PRIVACY OF VEHICULAR MIX-NETWORKS

In the context of VNs, the CMIX protocol creates mix-zones at road intersections. By connecting various mix-zones, we accumulate the anonymity achieved by each mix-zone and obtain a larger system referred to as a *mix-network*. Inspired by existing works on mix-zones, we derive an analytical model of the location privacy achieved by *vehicular mix-networks*.

We begin our analysis by developing a metric for location privacy in mix-zones and modeling two types of adversary: (i) a *weak adversary* and (ii) a *strong adversary*. The strength of the adversary depends on the information available to it. Then, we show that the mobility of vehicles and the physical properties of intersections affect the achievable unlinkability.

4.1 Vehicular Mix-Zones

A mix-zone is an anonymizing region that obfuscates the relation between entering and exiting vehicles. The adversary observes the *timing* and the *location* of the entering and exiting vehicles in order to derive a probability distribution over the possible mappings. In VNs, because we assume that mix-zones are located at road intersections, the timing of events depends on the *delay characteristics* of the intersection structure. Likewise, the location of entering and exiting vehicles depends on their *trajectory* in an intersection.

Beresford introduces a framework in [2] for the analysis of mix-zones, which was later adapted to VNs by Buttyan *et al.* in [3]. The event k of a vehicle entering the mix-zone at time τ on road n is denoted by $k = (n, \tau)$; similarly, the event l of a vehicle exiting the mix-zone at time τ' on road e is denoted by $l = (e, \tau')$. $p_{n,e}$ is the probability of exiting at e knowing that the vehicle entered at n . $q_{n,e}(t)$ is the probability that the time needed to enter at n and exit at e is equal to t . In a simple example with two vehicles, $p_{k \rightarrow l} = p_{n,e} q_{n,e}(t)$ is the probability of the mapping of an entry event k to an exit event l . For m vehicles, we use the derivation by Beresford in [1]. The location privacy of a vehicle corresponding to event l is the entropy of $p_{k \rightarrow l}$ ([25])

$$H(l) = - \sum_{k=1}^N p_{k \rightarrow l} \log_2(p_{k \rightarrow l}) \quad (1)$$

The entropy increases with two factors: (i) the number N of vehicles in the mix-zone and (ii) the similarity of the distribution of $p_{k \rightarrow l}$ to the uniform distribution. The precision of the probability distributions of $p_{n,e}$ and $q_{n,e}(t)$ relies on the information available to the adversary and is thus closely related to the threat model.

4.1.1 Weak Adversary

The Weak Adversary (WA) is aware of the set of vehicles entering/exiting the mix-zone but not of their timing and trajectories, e.g., an external passive adversary with a limited view of an intersection. Thus, to the adversary, vehicles spend a constant time C within a mix-zone, i.e., $q_{n,e}(t) = C$. Similarly, the trajectory of vehicles in intersections cannot be evaluated and $p_{n,e} = C$. Hence, $p_{k \rightarrow l}$ is a uniform distribution and the upperbound on the achievable location privacy is

$$H(l) \leq \log_2(N) \quad (2)$$

The maximum location privacy depends exclusively on the number N of vehicles in the mix-zone. The weak adversary provides an upperbound to the achievable unlinkability.

4.1.2 Strong Adversary

The Strong Adversary (SA) captures the sequentiality and location of events by gathering entering/exiting times and positions of vehicles. Hence, the effectiveness of a mix-zone relies on the delay characteristics of the intersection and on the trajectory of the vehicles in the road network. We assume that in a time interval Γ , N vehicles arrive at the mix-zone where N is determined by a Poisson distribution with parameter ρ . Hence, vehicles' arrival times τ_i for $i = 1 \dots N$ are distributed according to a uniform distribution, $\tau_i \sim \mathcal{U}(\Gamma)$.

Delay characteristics of VNs depend on the road intersection. The SA models road intersections with normal distributions. The delay is modeled by a normal distribution that depends on the trajectory of the vehicle in the intersection. For example, if d is the number of road segments that meet at an intersection, and we have $d = 4$; for vehicles arriving from n_1 , their delay characteristics is

$$q_{n_1, e_i}(t) \sim \mathcal{N}(\mu_{1,i}, \sigma_{1,i}) \quad (3)$$

where $i = 1 \dots d$ and e_1, e_2, e_3 and e_4 indicate (with respect to n_1) the directions *u-turn*, *left*, *straight* and *right*, respectively. By varying μ and σ , the SA models various types

of intersections. We assume that a strong attacker has an a-priori knowledge about the intersections (such as the intersection structure) to compute the delay parameters (μ, σ) .

The adversary models the trajectory of vehicles by choosing $p_{n,e} < 1$ according to an a-priori knowledge of the road topology and intersection types such that $\sum_{e=1}^d P_{n,e} = 1$ for every entering point n . For example, if $P_{n,e}$ takes a uniform distribution, then vehicles move according to a random walk.

4.2 Vehicular Mix-Networks

Typically, several mixes are used in chain to accumulate the unlinkability provided by each mix-zone. By connecting various independent mix-zones, we obtain a larger system that we call a *vehicular mix-network*. We approximate the location privacy of a vehicle v traversing a vehicular mix-network composed of L mix-zones as

$$H_{tot}(v, L) = \sum_{i=1}^L H_i(v) \quad (4)$$

where H_i is the location privacy gathered at each mix-zone. The vehicular network environment affects the efficiency of vehicular mix-networks in three aspects.

First, each mix-zone has its proper delay characteristics. In other words, the delay characteristics are very dependent on the intersection type. Hence, some intersections might provide very good anonymity while in others, vehicles will be easily traceable. The total unlinkability achieved by vehicle v then depends on the set $S_v \subset S$ and on the number $L_v = |S_v|$ of traversed mix-zones.

Second, the sparsely connected topology of VNs might facilitate the analysis of the network by the adversary by reducing the number of possible mappings. Yet, Danezis shows in [7] that sparsely connected mix-networks achieve the same anonymity as fully-connected mix-networks after $\mathcal{O}(\log|S|)$ steps, where S is the set of all mix-zones. In other words, when L is in the order of $\log|S|$, vehicular mix-networks achieve the same performance as fully connected networks. For example, in a vehicular mix-network with 10000 CMIXes, a vehicle must go through 4 intersections.

Third, the density of vehicles per mix-zone (intersection) changes over time and it decreases the average achievable anonymity per mix-zone. Indeed, the anonymity achieved with a fixed number of vehicles per mix-zone is higher than the anonymity achieved with a varying number of vehicles. Using Jensen's inequality with a concave function and (2), we have

$$E[\log_2(N)] \leq \log_2(E[N]) \quad (5)$$

Hence, density diversity in mix-zones decreases the achievable unlinkability. For these reasons, vehicular mix-networks are highly *dynamic*.

5. SIMULATIONS

We simulate a simple vehicular mix-network and evaluate the achievable location privacy under various network conditions.

5.1 Setup

The VN is modeled in Matlab as a 10×10 Manhattan network with $d = 4$, where d is the number of road segments that meet at each intersection. We assume M vehicles in

the network and an average number of vehicles per intersection $\rho = \frac{M}{10 \cdot 10}$. Each intersection is modeled by three parameters: (1) vehicle trajectories $p_{n,e}$, (2) vehicle arrival times Γ and (3) vehicle delays (μ, σ) . Vehicles move at each time step between adjacent intersections according to $p_{n,e}$. $P_{n,n}$, the probability of a u-turn, is set to 0 for every road direction. A $p_{n,e}$ vector is randomly defined for each intersection in the network. The vehicles' arrival times are uniformly distributed between 0 and Γ seconds. When Γ is small, vehicles enter the mix-zone almost simultaneously. This models a highly congested traffic. In a low congestion traffic scenario, Γ is large, and vehicle arrival times are easily distinguishable. Each intersection delay characteristics are modeled by randomly choosing parameters (μ, σ) . We assume that the adversary knows the system parameters and has therefore a perfect a-priori knowledge.

All results are the average of 20 simulations of the Manhattan network and are presented with 95% confidence intervals. We compute the achievable location privacy in the presence of a weak and strong adversary, as defined in the previous section.

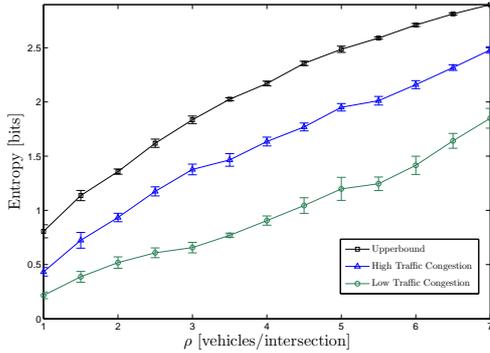
5.2 Results

Figure 3(a) presents the average location privacy obtained in an intersection for various vehicle densities. We show that the achieved location privacy varies with respect to the traffic congestion and the vehicle density. We observe that the less congested the traffic is, and the easier it is for the adversary to track vehicles based on their delay characteristics. The upperbound is the average of Equation 2 over the number of user per intersection in the simulation. Figure 3(b) presents the success probability of an adversary in tracking vehicles. The adversary success probability is the ratio of the number of successfully mapped vehicles to the total number of vehicles in a mix-zone, averaged over all mix-zones. As expected, the success ratio decreases as the entropy increases. Even with a high ρ and congested traffic, the adversary success ratio is relatively high. We show next that the combination of mix-zones into mix-networks significantly reduces this success ratio.

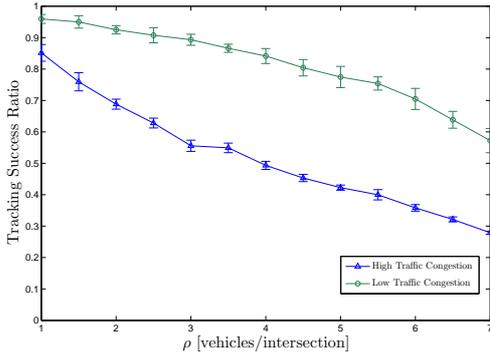
Figures 4(a) and 4(b) show the same metrics as before (entropy and adversary success ratio) in the case of a mix-network. As expected (Equation 4), the accumulation of anonymity over mix-zones increases almost linearly and the adversary success ratio in tracking vehicles correctly from source to destination becomes negligible. The achieved location privacy varies linearly in the number of bits, but exponentially in the number of vehicles. For example, a vehicle that traverses 10 mix-zones gathers on average 20 bits of anonymity (i.e., is indistinguishable among 2^{20} vehicles).

6. DISCUSSION

Information that links messages to the same vehicle can be obtained not only from the safety messages but also from other sources. At the physical layer, vehicles might be identifiable by fingerprinting their transceivers [9]; however, to the best of our knowledge, this has not been shown on VN radios, while in different contexts, successful identification was achieved but only with some significant probability. At the data link and network layers, using the same MAC address would render messages linkable. Thus, the proposed solution is to randomly change the MAC and network (IP, where applicable) address. Different aspects and methods



(a)



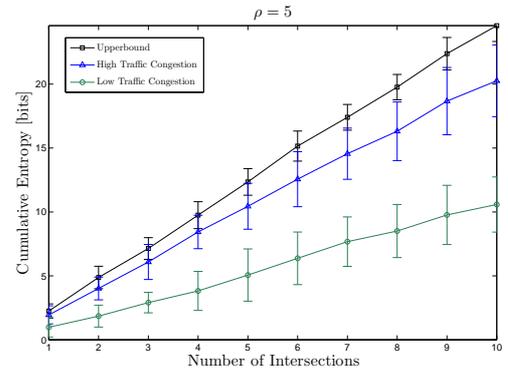
(b)

Figure 3: Vehicular Mix-zone simulation results. (a) Location privacy for various vehicle densities. (b) Adversary success probability.

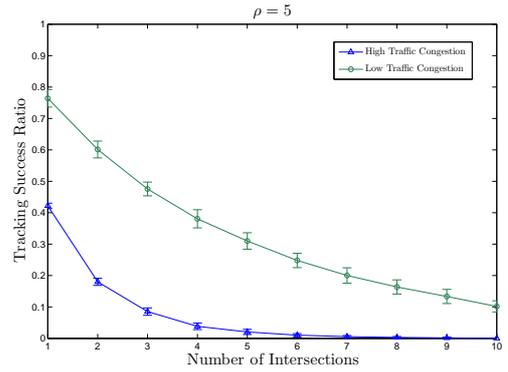
for data link and network layer address changes, with respect to the pseudonym changes, are identified in [18].

A VNO could use the deployed RSUs and act as a global, or more precisely within its area of coverage, internal adversary. Essentially, the VNO monitors the points in the network (intersections) that can provide valuable information to track vehicles. Buttyan *et al.* show that changing pseudonym in mix-zones between RSUs does not suffice to obtain location privacy [3] when an adversary monitors more than 50% of the intersections of the road network. Note also that location information in safety messages is very precise and facilitates tracking. Considering if an alternative way to ensure accident avoidance can be devised, as well as a form of *private positioning* that preserves the functionality of safety messages but achieves location privacy, are interesting directions of future work.

The efficiency of vehicular mix-networks is independent of the underlying CMIX protocol. Hence, several alternative CMIX protocols can be envisioned. The simplest CMIX protocol preloads one unique symmetric key in all vehicles throughout the network and the key is protected from external adversaries by the TPD. The robustness of such a protocol is very low, as the single key can be compromised. Instead of preloading one key, or letting the RSU decide which keys to use, an alternative solution is to let vehicles



(a)



(b)

Figure 4: Vehicular Mix-network simulation results with $\rho = 5$. (a) Location privacy for various vehicle densities. (b) Adversary success probability.

construct the symmetric key using group key agreement protocols. Each vehicle contributes its own secret to create the symmetric key. However, the rapid inclusion and exclusion of members would be hard. Furthermore, a protocol exploiting the network topology and vehicle mobility could potentially reduce the complexity of key establishment. In particular, a protocol using exclusively symmetric operations to create cryptographic mix-zones would be an attractive approach.

7. RELATED WORK

In order to achieve location privacy in a pervasive computing environment, Beresford and Stajano propose in [2] the concept of mix-zones where a natural mixing of mobile nodes occurs. Mix-zones are anonymized regions of the network wherein mobile nodes change their identifiers to obfuscate the relation between entering and exiting events. Numerous proposals have since followed to allow for the creation of mix-zones, improve their mixing performances or to measure their effectiveness.

Huang *et al.* [14] suggest using random *silent periods* wherein mobile nodes turn off their transceivers and update their identifiers. In [15], they further propose arranging silent periods into cascades and to enhance anonymity, coordinating silent periods. Gruteser and Grünwald propose in [12]

a network-centric solution in that direction. In [17], Li *et al.* suggest a user-centric solution letting mobile nodes coordinate their silent periods and decide whether to change pseudonyms. In [24], Sampigethaya *et al.* introduce CARAVAN, a VN location privacy scheme also based on silent periods. In a different direction, Hoh and Gruteser [13] present a path perturbation algorithm: nodes tamper with the broadcasted location data to reduce the spatial and/or temporal granularity of location information. This does not apply to VNs, because it would jeopardize vehicular safety applications. Finally, Buttyán *et al.* assume in [3] that vehicles exploit the separation of RSUs to mimic silent zones and create anonymizing regions. Their analysis shows that the effectiveness of this approach is limited as an adversary monitoring 50% of the intersections can successfully track 60% of the vehicles with very high probability because of the non-uniformity of the traffic.

8. CONCLUSION

In this paper, we study the problem of providing location privacy in vehicular networks. We introduce the CMIX protocol to create cryptographic mix-zones at road intersections wherein vehicles can change their pseudonyms. The combination of mix-zones into vehicular mix-networks permits the accumulation of unlinkability over the VN. Vehicular mix-networks rely on the mobility of vehicles to provide location privacy without jeopardizing the efficiency of safety messages. We model analytically and evaluate by means of simulations the unlinkability provided by vehicular mix-networks. Our results show that, although the unlinkability of individual mix-zones can be relatively low in some cases, the accumulated unlinkability of the mix-networks is generally very high.

9. REFERENCES

- [1] A. R. Beresford. Location privacy in ubiquitous computing. Technical Report 612, University of Cambridge, January 2005.
- [2] A. R. Beresford and F. Stajano. Mix-zones: User privacy in location-aware services. In *Proceedings of PerSec*, 2004.
- [3] L. Buttyán, T. Holczer, and I. Vajda. On the effectiveness of changing pseudonyms to provide location privacy in VANETs. In *Proceedings of ESAS*, 2007.
- [4] G. Calandriello, P. Papadimitratos, A. Lloy, and J.-P. Hubaux. Efficient and robust pseudonymous authentication in vanets. In *The Fourth ACM International Workshop on Vehicular Ad Hoc Networks (VANET)*, 2007.
- [5] D. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2):84–90, February 1981.
- [6] Car2Car Consortium. <http://www.car-to-car.org/>.
- [7] G. Danezis. Mix-networks with restricted routes. In *Proceedings of PET*, March 2003.
- [8] 5.9GHz DSRC. <http://grouper.ieee.org/groups/scc32/dsrc/index.html>.
- [9] K. J. Ellis and Nur Serinken. Characteristics of radio transmitter fingerprints, 2001. *Radio Science*, Volume 36, Issue 4, p. 585-598.
- [10] Fon. <http://www.fon.com>.
- [11] M. Gerlach and F. Güttler. Privacy in VANETs using changing pseudonyms - ideal and real. In *Proceedings of VTC*, April 2007.
- [12] M. Gruteser and D. Grunwald. Enhancing location privacy in wireless LAN through disposable interface identifiers: a quantitative analysis. *Mobile Networks and Applications*, 10(3):315 – 325, June 2005.
- [13] B. Hoh and M. Gruteser. Protecting location privacy through path confusion. In *Proceedings of SECURECOMM*, pages 194–205, 2005.
- [14] L. Huang, K. Matsuura, H. Yamane, and K. Sezaki. Towards modeling wireless location privacy. In *Proceedings of PET*, 2005.
- [15] L. Huang, K. Matsuura, H. Yamane, and K. Sezaki. Silent cascade: Enhancing location privacy without communication QoS degradation. In *Proceedings of SPC*, pages 165–180, 2006.
- [16] J. Krumm. Inference attacks on location tracks. In *Pervasive Computing*, May 2007.
- [17] M. Li, K. Sampigethaya, L. Huang, and R. Poovendran. Swing and Swap: User-centric approaches towards maximizing location privacy. In *Proceedings of WPES*, pages 19–28, 2006.
- [18] P. Papadimitratos, L. Buttyan, J.-P. Hubaux, F. Kargl, A. Kung, and M. Raya. Architecture for secure and private vehicular communications. In *Proceedings of 7th IEEE International Conference on ITS Telecommunications ITST*, 2007.
- [19] P. Papadimitratos, V. Gligor, and J.-P. Hubaux. Securing vehicular communications - assumptions, requirements, and principles. In *Proceedings of Workshop on Embedded Security in Cars (ESCAR)*, 2006.
- [20] P. Papadimitratos, A. Kung, J.-P. Hubaux, and F. Kargl. Privacy and identity management for vehicular communication systems: A position paper. In *Proceedings of Workshop on Standards for Privacy in User-Centric Identity Management*, 2006.
- [21] M. Raya and J.-P. Hubaux. The security of vehicular ad hoc networks. In *SASN*, 2005.
- [22] M. Raya and J.-P. Hubaux. Securing Vehicular Ad Hoc Networks. *Journal of Computer Security, Special Issue on Security of Ad Hoc and Sensor Networks*, 15(1):39 – 68, 2007.
- [23] M. Raya, P. Papadimitratos, and J.-P. Hubaux. Securing vehicular communications. In *IEEE Wireless Communications Magazine*, 2006.
- [24] 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.
- [25] A. Serjantov and G. Danezis. Towards an information theoretic metric for anonymity. In *Proceedings of PET*, 2002.