Readjusting the Privacy Goals in Vehicular Ad-hoc Networks: A Safety-preserving Solution Using Non-Overlapping Time-slotted Pseudonym Pools

David Eckhoff[a], Christoph Sommer[b]

[a] *Robotics and Embedded Systems, Department of Computer Science, Technische Universität München, Germany*
[b] *Cooperative Mobile Systems, Heinz Nixdorf Institute and Dept. of Computer Science, Paderborn University, Germany*

Abstract

Current proposals for privacy measures in vehicular networking commonly suffer from either of three limitations: Many are so drastic that they compromise road traffic safety, a primary goal of vehicular networks. Others are more compliant, but (despite adding substantial overhead) are ineffective at protecting users' privacy against state-of-the-art attackers. The last group suffers from neither limitation, but undermine accountability and are thus in conflict with the requirements of future vehicular networks. As a consequence, workable privacy protection is often thought unattainable and privacy protection simply disregarded in the many field experiments, proposals, and standardization documents to date. In this work, we start fresh with a readjusted view on privacy goals and the capacities of a state-of-the-art attacker in mind, taking a structured approach to deriving a holistic solution for location privacy protection in Vehicular Ad-Hoc Networks (VANETs): We show that local privacy protection cannot be attained without compromising road traffic safety – nor does it have to be. Our approach is based on synchronized time-slotted pseudonym pools and the local announcing of pseudonym changes. By this, we overcome the privacy–safety problem while at the same time increasing privacy for all users. Our system is fully compatible with the requirements of vehicular networks and current standardization efforts.

1. Introduction

Vehicular networking, the wireless exchange of data between vehicles, is a key component of future intelligent transportation systems. Wirelessly sharing information among vehicles promises to improve road traffic safety and, as a pleasant side effect, can enable novel business concepts. Many vehicles in Japan can already rely on wireless short range communication technology and both the American IEEE and the European ETSI are finalizing standardization documents for the operation of future vehicular networks. Indeed, the US DOT has even announced its intent to make ad-hoc communication systems mandatory for new cars to reap the full benefits for road traffic safety that such a system can bring when deployed universally [1].

One of the key features targeted by the proposed systems is cooperative awareness, that is, vehicles communicating with each other to establish a virtual view of their surroundings for the sake of improving traffic safety. This is envisioned to be achieved with the help of periodic beacons, broadcast transmissions sent by a vehicle which include its current state. Beacons are commonly sent with a frequency of 1 Hz to 10 Hz. In IEEE WAVE and ETSI ITS-G5, each vehicle wirelessly informs all cars in its vicinity of state information that includes its current position, current heading, and current velocity.

Email addresses: david.eckhoff@in.tum.de (David Eckhoff), sommer@ccs-labs.org (Christoph Sommer)

As these transmissions are (and need to be) decodable by the general public, however, they can be received not just by other vehicles but by anyone with a receiver physically close enough to the sender (even at distances as far as 300 m to 800 m [2]). An adversary could exploit this and track a vehicle simply by linking consecutive transmissions. This can lead to a violation of location privacy of drivers, and through that also other types of privacy [3, 4].

This privacy problem in vehicular networks has been understood from the very beginning [5]. The consensus in vehicular network privacy research is to use changing short-term identifiers, that is, *pseudonyms*, instead of static ones to complicate tracking for any eavesdropping adversary. An important challenge is to employ a suitable pseudonym change strategy, i.e., when (or where) a vehicle should change its pseudonym to maximize its location privacy. A broad range of these strategies [6], many of them aligned with the general privacy framework proposed in [7], has been proposed in the time since – and some of them even also considered in various field trials [8, 9], albeit not always with the focus necessary to pave the way for a concrete strategy to become part of standardization documents.

A major obstacle in finding a suitable pseudonym changing strategy is the fact that there seems to be no agreement in the parameters [10]: Strategies differ with regard to the adversary against which they are protecting, how they influence other applications such as traffic safety, and their compatibility with other system requirements, such as accountability or computational complexity. We believe that

Accepted manuscript, share-able under CC BY-NC-ND. Published in Elsevier Computer Communications, DOI 10.1016/j.comcom.2018.03.006

to make privacy protection a fundamental part of future vehicular networks, they have to take into account all these constraints and requirements. For example, safety applications rely on receiving and linking periodic messages; any privacy protection mechanism interfering with these applications is thus unlikely to be deployed.

In this work, we take a realistic look at the requirements of envisioned intelligent transportation systems and propose a holistic pseudonym-based solution that increases privacy without sacrificing safety:

- We make use of non-overlapping time-slotted pseudonyms to increase the overall privacy protection.
- At the same time, we advocate putting an end to chasing the goal of confusing eavesdropping adversaries, as this is completely opposite to the primary purpose of vehicular networks: allowing vehicles to track other nearby vehicles to avoid collisions.
- We support this claim with a detailed simulation study based on synthetic mobility and on real-world traces to show that confusing local adversaries is not possible without also affecting traffic safety.
- We present the underlying model of a state-of-the-art attacker using a multi-target tracking algorithm.

Our results give insights into the limitations of pseudonym changing strategies and consequently allow us to effectively tackle the privacy–safety trade-off. In addition, our solution is insusceptible to Sybil attacks and allows for efficient and privacy-preserving certificate revocation. It is also fully compatible with the upcoming North American IEEE and European ETSI families of standards.

This manuscript constitutes an extended version of our previous work [11], now also including the proposed multi-target tracking algorithm used to model the attacker in our computer simulation studies and an in-depth description of the simulation setup.

The remainder of this manuscript is structured as follows: In Section 2 we describe the status quo of vehicular network privacy systems as envisioned in IEEE WAVE and ETSI ITS-G5. Here, and throughout the remainder of the manuscript, we will also refer to and discuss related work. Section 3 discusses the privacy threats of vehicular networks; Section 4 explains the constraints that privacy protection mechanisms must work within. In Section 5 we present our solution which we believe is a viable approach to coping with the location privacy challenges in VANETs without negatively impacting traffic safety. We back up central claims that motivated the construction of our solution using a novel model of a state-of-the-art attacker (Section 6) and investigating its leverage against state-of-the-art privacy solutions in a computer simulation study (Section 7).

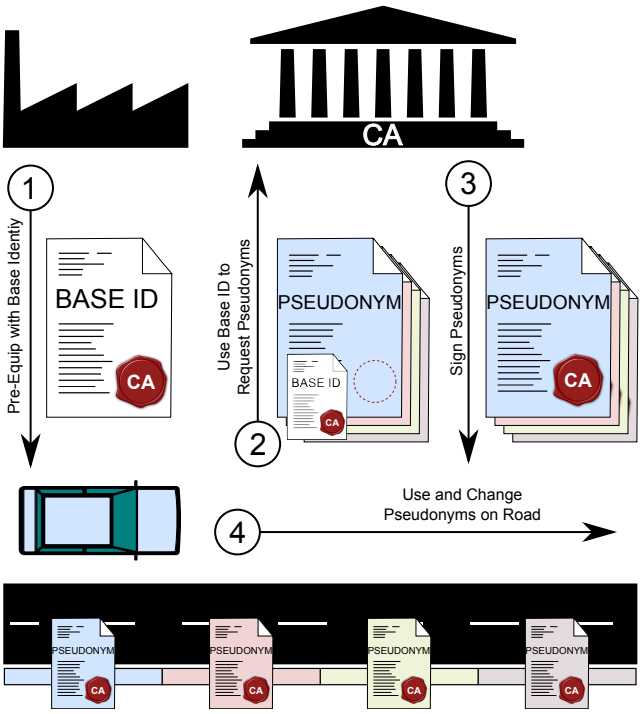

Figure 1: A simplified vehicular Public Key Infrastructure (PKI).

2. Status Quo: Vehicular PKI

Authenticity and integrity are essential security requirements in vehicular networks. Only authorized devices should be able to participate in the network and it must be guaranteed that forged messages can be detected as such. These security goals can be achieved by means of a Public Key Infrastructure (PKI) as described in standards of IEEE (1609.2-2016) and ETSI (102 941). In addition, this PKI is also the basis for privacy protection through the use of authenticated pseudonymous identifiers.

A (slightly simplified) explanation of the system is shown in Figure 1. Vehicles are equipped with a base identity (or long-term identifier), consisting of a certificate and public-private key pair (Step 1). This identity is unique to a certain vehicle and must therefore never be used for car-to-car communication. It is only used to generate or request pseudonyms (in the form of pseudonymous certificates) from a Certificate Authority (CA) trusted by all vehicles (Step 2). If the identity is valid (as indicated by a signature of the CA) and the information in the pseudonym request is correct, the CA signs the pseudonyms and sends them back to the vehicle (Step 3). Each vehicle maintains a pool of pseudonyms and uses a selected pseudonym as its visible address, that is, to sign and send messages over the wireless channel (Step 4). Other vehicles will only consider received messages if signed with a valid pseudonym.

It is, however, unclear how these pseudonym pools are organized and how vehicles should select which pseudonym to use for which transmission. For example, it was discussed that multiple (or even all) pseudonyms are valid at the same

time and that the On-Board Unit (OBU) of the vehicle can choose freely or randomly which pseudonym to use. This introduces the problem of Sybil attacks [12, 13], that is, one vehicle pretending to be many at the same time, thus subverting consensus-based approaches to credibility checks. Other vehicles would have no trivial method of identifying such an attack, as they cannot link different pseudonyms to the same vehicle. In earlier work, we have suggested the use of non-overlapping pseudonyms to avoid this problem [14].

Other proposals for privacy protection include the use of silent periods, that is, not transmitting beacons after a pseudonym change [15] or the use of group cryptography [16] to prevent eavesdropping. However, both of these proposals are not compatible with the upcoming standards as they interfere with traffic safety or conflict with the unencrypted transmissions of periodic beacons. Gerlach et al. have proposed to consider the context of a vehicle to determine when a pseudonym change can be effective [17], Freudiger et al. presented their concept of mix-zones, that is, geographic areas for pseudonym changes [18]. The results we present in Section 7 show that these proposals are not sufficient to protect the privacy of drivers.

As of today, the IEEE and ETSI family of standards do not recommend a specific pseudonym changing strategy, nor do they discuss existing solutions. The documents only mention the need to "use a pseudonym that cannot be linked to [...] the user's true identity" (ETSI 102 893-v1.1.1) and suggest to change it frequently "[...] to avoid simple correlation between the pseudonym and the vehicle" (ETSI 102 940-v1.1.1).

Similarly, it is still unclear how pseudonym pools have to be configured to work efficiently with certificate revocation, given the potentially large number of pseudonymous certificates each vehicle carries. Certificate revocation is the process of invalidating pseudonyms, e.g., when a vehicle is found to transmit faulty messages. ETSI ITS-G5 does not consider revocation of vehicular OBUs. Instead, it is argued that pseudonym pools should be small and the exclusion of certain vehicles can be achieved by simply not signing new pseudonym requests from them. The IEEE 1609.2-2016 standard supports a linkage-based revocation method, which we will discuss in detail in Section 5.3, where we also explain how it benefits from our proposal.

In conclusion it can be said that, while currently envisioned systems provide a solid basis for the deployment of privacy-enhancing technologies, there is a need for concrete recommendations when it comes to the usage of pseudonymous identifiers. We contribute to finding these recommendations by first identifying the exact requirements and constraints and then presenting a proposal which we believe satisfies these requirements.

3. Understanding the Privacy Challenge in Vehicular Ad-Hoc Networks (VANETs)

In order to properly address the privacy issues in Vehicular Ad-Hoc Networks (VANETs) we have to be clear about the exact nature of these issues. This includes the type and property of privacy at risk, the potential adversary, and the attack channels. Only if privacy risks are exactly defined can a privacy protection mechanism be designed.

There exist different taxonomies to categorize different types and properties of privacy. Finn et al. [19] divide privacy into seven types, namely privacy of person, behavior, communication, data, thoughts, location, and association. The lines between the types are blurred, and, through correlation, violation of one type can lead to the violation of other types as well. For example, correlation about the location of two persons can imply information about their association. We focus on *location privacy*, as this is the primary privacy type endangered by the periodic broadcast messages transmitted by intelligent vehicles.

Pfitzmann and Hansen have defined different properties of privacy [20]. These include anonymity, unlinkability, undetectability, unobservability, and pseudonymity. While all five are affected by vehicular networks, we concentrate on the unlinkability property, that is, the inability to link two messages. We will also show that unobservability, besides pseudonymity and anonymity, is a fundamental requirement to prevent tracking.

Looking at the vast literature on privacy protection in VANETs, it can be observed that there is no general agreement on who the primary adversary in these systems is [10]. Adversaries can be defined among different orthogonal dimensions: local vs. global, internal vs. external, passive vs. active, static vs. adaptive, and the amount of prior knowledge. There seems to be a tendency toward focusing on an external global passive adversary [10], that is, an adversary that can listen to all unencrypted communication in the network. In the case of vehicular networks, this includes all transmitted beacons. We will show that when considering the primary goal of VANETs, that is, improving traffic safety, it is unwise to try and defend against a global attacker: not because of issues of technical realization or cost, but because there can be no effective privacy protection against an omnipresent observer, as traffic safety and confusing nearby receivers are opposing goals.

The adversary and their strength have to be chosen carefully. Defending against attackers who eavesdrop on car-to-car communication to specifically target certain individuals might be infeasible, as these attackers might as well physically follow the car in question. Privacy protection in vehicular networks should therefore focus on the prevention of new attacks and not on precluding the ones that could be executed anyway. The chosen adversary model should account for this. We therefore focus on a local and passive adversary, who sets up one or multiple receivers (possibly in strategic positions) to eavesdrop on transmitted beacon messages. The goal of the adversary is

to track all vehicles through the network to create detailed mobility traces. How these traces are then processed, e.g., by correlating them with home and work addresses [21], is not directly relevant, because the goal of the privacy protection mechanism is to prevent the creation of these traces in the first place. Considering unencrypted beacons, there needs to be no differentiation between internal (i.e., other users or service provider) and external adversaries, as long as the attack is of a passive nature.

Lastly, it has to be defined through which channel adversaries obtain sensitive data. Possible channels are observable data, published data, re-purposed data, and leaked data [22]. We primarily consider observable data, that is, overheard beacon messages. Privacy mechanisms protecting this channel will then implicitly also protect attacks based on re-purposed data and leaked data, as they affect the possibility to collect overheard messages. When talking about the privacy implications of certificate revocation, we also account for attacks based on published data.

4. Requirements for VANETs Privacy Protection

Before privacy protection mechanisms can be proposed, it needs to be clear which use-case specific restrictions apply. Past field operational tests have shown that only privacy protections that do not negatively impact other objectives of the vehicular network have a chance of being deployed without being severely degraded (to the point of not providing privacy at all).

4.1. Accountability and non-repudiation

The possibility for an authoritative entity to resolve pseudonyms to base identities, that is, accountability, has been identified as an important requirement for vehicular networks. IEEE 1609.2-2013 already notes that methods to allow fully anonymous identifiers "[. . .] might conflict with other goals such as removing bad actors and supporting law enforcement access under appropriate circumstances". In addition, fully anonymous identifiers would also enable vehicles to plausibly deny having sent certain messages.

Misbehavior by an authority therefore cannot be made technically impossible, but has to be tackled legally or by policy. This means that all privacy protections interfering with accountability and non-repudiation are unlikely to be deployed in a real system. One promising approach to address this issue is separation of knowledge, as already stated in ETSI 102 941-v1.1.1: it requires multiple entities to collude in order to resolve a pseudonym.

4.2. Privacy–Safety Trade-off

One of the biggest challenges in privacy protection of vehicular networks is the so called privacy–safety trade-off. Improved traffic safety is one of the primary goals of intelligent transportation systems. Vehicles receive broadcast messages from other cars; based on the content of these messages (e.g., speed, location, and heading), OBUs can warn the driver and (semi-)autonomous vehicles can brake. To enable OBUs to reliably run these collision avoidance systems and other safety applications, they need to have an exact virtual representation of the vehicle's surroundings. This representation can be based on sensor readings such as radar or computer vision, but also on car-to-car communication. In the latter case, the goal of the receiving OBU is the same as for an eavesdropping attacker: the tracking of vehicles in the vicinity, albeit with different motives. Confusing a tracking adversary therefore also means potentially confusing the OBUs of other vehicles. Additionally, confusing an adversary by changing pseudonyms is rather difficult, as we will show in Section 5.2.

Never must privacy protection in vehicular networks cause a traffic accident or, even worse, injury or death. The fact that many people will likely value safety much higher than privacy in a potentially critical situation has to be accounted for when developing and deploying pseudonym changing strategies. This fact disqualifies a large number of proposed pseudonym changing strategies, most prominently, approaches incorporating silent times, that is, the omission of beacons for a certain period after a pseudonym change. Although remaining silent benefits privacy [15], it was shown that the effectiveness of traffic safety application is significantly reduced during these intervals [23].

We believe that the goal of pseudonym changing strategies has to be reconsidered. Privacy has to yield to traffic safety, and therefore confusing nearby receivers – other cars and adversaries alike – cannot be the goal. The pseudonym changing strategy must be designed in a way that it creates maximum confusion for adversaries outside of the transmission range with zero impact on traffic safety.

4.3. Storage and Computational Restrictions

The tasks envisioned for a vehicle's OBUs will be demanding in terms of computational power and storage capacity. The validity of each incoming message must be checked: this includes verifying the attached cryptographic signature, checking whether the used public key is on the stored certificate revocation list, and even whether the contents of the message are plausible. Applications such as collision avoidance consume additional computation power. With potentially hundreds of messages arriving and up to ten beacon messages generated every second, the resulting computational effort could be challenging. The vehicle's pseudonym pool (including the corresponding private keys) should be stored in costly tamper-proof storage and, if revocation is supported, all revoked public keys need to be stored as well. Lastly, if pseudonyms need to be re-filled at runtime, the cost associated with their download needs to be limited. Privacy protection mechanisms should therefore account for these limitations and refrain from computationally intensive and storage-heavy tasks if possible. With the expected increasing computational power of OBUs, this requirement will possibly be less relevant in the future.

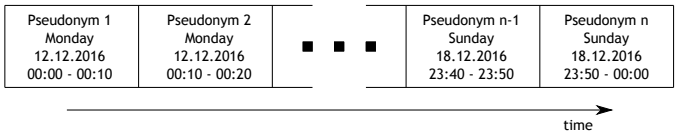

Figure 2: Time slotted pseudonym pool of 1 week length and 10 minute pseudonym validity.

4.4. Security Implications

The deployed privacy protection mechanisms should not interfere with the security of the system by opening new attack vectors. The compromise of one or a few OBUs should not affect the security of the entire system. If each vehicle maintains a pool of simultaneously valid pseudonyms, the physical compromise of one OBU enables a malicious vehicle to pretend to be multiple cars [13]. This does not affect the security of the entire system, because the individual attacker could be identified, but opens an attack vector that would not exist without the use of pseudonyms. In contrast, fully anonymous privacy protection mechanisms that do not provide accountability could potentially break system security when malicious users cannot be identified.

5. A Proposal for Privacy Protection

Based on the above considerations, we present our proposal for holistic location privacy protection in VANETs, with a particular focus on fulfilling the requirements of future intelligent transportation systems. Our solution aims at minimizing the privacy implications caused by the periodic transmissions of pseudonymous status beacons. We do not address issues of other layers potentially jeopardizing user privacy, such as applications that include very specific content (e.g., the vehicle dimensions) in beacon messages. Our proposal combines and adapts three existing building blocks to yield a workable solution, providing a basis for higher layer privacy protection as it secures the basic function of future VANETs, that is, cooperative awareness.

5.1. Non-Overlapping Time-Slotted Pseudonym Pools

The first and most important building block of our proposal is the use of non-overlapping time-slotted pseudonym pools [14, 24, 7]. Each vehicle maintains a pool of chronologically ordered pseudonyms as shown in Figure 2. The configuration relies only on two parameters, the length of the pseudonym pool and the validity duration of each pseudonym. We argue that for optimal privacy protection the pseudonym pools for all vehicles are synchronized, i.e., they use the same parameters for length and validity. These parameters also implicitly control the level of privacy protection and storage requirements.

Assuming synchronized clocks (e.g., via GPS), all vehicles change their pseudonym at exactly the same time (e.g., at 0:10 AM, following the example in Figure 2). Therefore, the use of time-slotted pseudonym pools also dictates the pseudonym changing strategy, and the time of validity controls the frequency of pseudonym changes. This also implies that vehicles that are not under adversary surveillance at that point in time will be using a pseudonym unknown to the adversary when they re-enter the adversary's transmission range. The impact of time synchronization is negligible, as clocks need not be more tightly synchronized than the interval in which the old pseudonym is still part of the messages. Since this interval is usually in the order of seconds, GPS clock synchronization is more than sufficient.

Another advantage of non-overlapping pseudonyms is that Sybil attacks are no longer possible, as for each point in time, a vehicle only has one valid pseudonym. The physical compromise of an OBU therefore does not introduce new attack vectors caused by the privacy protection mechanism.

Time-slotted pools also allow for easy estimation and control of the storage required on the OBU. They further allow for privacy-preserving and efficient certificate revocation, as we will show in Section 5.3.

Suitable settings for the validity duration and pseudonym pool lengths need to be based on average trip durations and capacities of designated OBUs. The more often a vehicle changes its pseudonym, the higher the likelihood that one of these changes was not overheard by an adversary. It is therefore desirable to reduce the slot time as much as possible with respect to storage requirements and pseudonym requesting overhead. This can be done without affecting traffic safety, as we will show in the following section.

An alternative modus operandi of time-slotted pseudonyms is to re-use them after a certain period, i.e., employing a circular pseudonym pool [14] (e.g., with a pool length of 100 minutes, 20 pseudonyms and a pool validity of 1 week [25]). This can considerably reduce the number of pseudonyms that need to be stored on an OBU but introduces another problem: Adversaries are potentially able to link seamlessly unrelated transmissions (e.g., days apart) that were done using the same pseudonym. Overhearing a pseudonym change then also means that the adversary is able to link two pseudonyms in the past and the future. With every iteration of the circular pool, an adversary could collect more information about a vehicle's pseudonyms and violate a driver's privacy not only in retrospect but also until the entire pseudonym pool is replaced. Additionally, pseudonym revocation would then require the entire pool to be revoked, disclosing a vehicle's past (cf. Section 5.3). We therefore do not recommend re-using pseudonyms as the privacy implications in these scenarios are difficult to assess.

5.2. Solving the Privacy-Safety Trade-Off

We identified the privacy-safety trade-off as one of the most important factors to consider when developing privacy protection in vehicular networks. With time-based pseudonyms, we take away the ability for vehicles to control when pseudonyms are changed. They do no longer have the option to postpone a pseudonym change until after a

critical traffic situation as their currently used pseudonym is no longer valid. With synchronized pools between all vehicles, the situation becomes even more critical. Imagine a busy traffic circle or intersection where dozens of cars change their identifier at exactly the same time. In a worst-case scenario, this could confuse safety applications and potentially lead to an accident that could have been prevented using properly functioning car-to-car technology. Even if safety applications are only very rarely confused by a pseudonym change, e.g., once in ten thousand critical situations, the sheer number of vehicles on the street will lead to cases where privacy protection caused a traffic accident.

Therefore we advocate to surrender privacy to nearby vehicles by advertising pseudonym changes, that is, temporarily adding the last used pseudonymous identifier to new messages (and sign the message with both old and new pseudonym [26]). Thus, even if a vehicle changes pseudonyms right before sending a critical message, as the old pseudonym is included in the message for some time, vehicles nearby can obtain at least as much information as they have in a no-privacy scenario. Other proposals have discussed the possibility of using overlapping time-slots and allowing a sender to keep signing with the old pseudonym, should the car be in a safety-critical situation. However, this requires the sending car (and not the receiver) to always fully understand whether a situation is critical, which is bound to be error-prone.

As for our proposal, we further claim that it has almost no negative impact on privacy, as it is almost impossible to confuse eavesdropping attackers. In Section 7, we want to back up this strong claim by an extensive simulation study.

5.3. Pseudonym Revocation

Revocation is the process of the CA excluding certain vehicles from the vehicular network by distributing a so called Certificate Revocation List (CRL) containing their valid pseudonyms. The reasons for revoking a vehicle's pseudonyms are diverse, but one of the main reasons is (intentional or unintentional) transmission of false messages. Revocation is a challenging process with regard to both efficiency and privacy [27]. Revoking a large number of vehicles results in long CRLs, and putting all pseudonyms of a vehicle on a list allows others to link these pseudonyms.

Assume an adversary operating several access points throughout a larger area. For every overheard message, he stores the location and the used pseudonym. While at the time of receiving a message the adversary might be unable to link pseudonyms heard at different locations, publishing a list of all past pseudonyms of a vehicle will allow him to do so in retrospect. This enables him to link different locations and possibly to de-anonymize the driver, e.g., by analyzing patterns in the recorded data such as daily commutes. Therefore, the published CRL must not include past pseudonyms that are no longer valid to preserve backward privacy of revoked vehicles.

Our proposed pseudonym strategy works well with a mechanism called linkage values, first introduced in [27], refined by [28], and later by [29]. The idea is to not simply publish a list of revoked pseudonyms, but to enable vehicles to compute this list based on publishing a secret key and the number of revoked pseudonyms. To this end, each pseudonym certificate is attached a linkage value C_v^i, where i is the certificate number and v is the vehicle. These values are linked by a known cryptographic hash function $h(\cdot)$ and a vehicle-specific secret key κ_v only known by the CA. The linkage value C_v^i can be computed by encrypting κ_v^i using the known symmetric encryption function $e(\cdot)$ as follows:

$$\begin{array}{ccccccc} \kappa_v & \xrightarrow{h(\kappa_v)} & \kappa_v^2 & \xrightarrow{h(\kappa_v^2)} & \kappa_v^3 & \xrightarrow{h(\kappa_v^3)} & \cdots \\ \downarrow e(\kappa_v) & & \downarrow e(\kappa_v^2) & & \downarrow e(\kappa_v^3) & & \\ C_v^1 & & C_v^2 & & C_v^3 & & \cdots \end{array} \tag{1}$$

Assume a vehicle v holds n pseudonyms, and the CA wishes to revoke this vehicle's future pseudonyms from certificate j on. It computes κ_v^j by hashing the stored secret κ_v repeatedly $j-1$ times. By publishing κ_v^j and the number of revoked pseudonyms $n-j$, each vehicle can compute all revoked linkage values $C_v^j \dots C_v^n$ and store them internally on the OBU. Because certificates contain the linkage value, cars can then check each received message against the stored CRL and discard the message if it was sent using a revoked pseudonym. Due to the irreversibility of hash function $h(\cdot)$, linkage values of older pseudonyms before j cannot be computed, thus preserving backward privacy for the revoked vehicles.

This mechanism benefits from the use of time-slotted pseudonym pools, as they introduce a chronological order and a clear partition into past, current, and future pseudonyms. It is then trivial to identify which pseudonyms have to be revoked and it can be guaranteed that at the time of revocation only one revoked pseudonym could have been already used by the revoked car. It therefore offers both efficiency and backwards privacy, clearly outperforming traditional CRL approaches.

The use of linkage values has been adapted recently in the IEEE 1609.2-2016 standard. To further reduce overhead, only deltas instead of the entire CRL can be distributed.

6. Simulation Study and Evaluation Framework

First indications of the difficulty to confuse local adversaries were given in [30] and [31]. The underlying scenarios, however, were simplistic and purely synthetic.

In our simulation study, we make use of both real-world traces and synthetic scenarios. To this end, we implemented a modern multi-target tracking algorithm within our Veins simulation framework [32], and evaluated how successful an adversary can track vehicles.

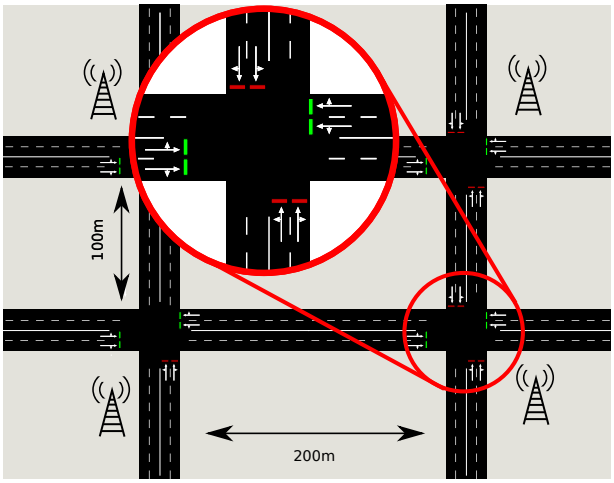

Figure 3: A grid scenario consisting of four intersections. Road segments leading to the intersections are 400 meters long. An adversary has placed 4 antennas and is able to fully monitor the scenario.

In total, we investigated three different scenarios. A simulated combination of four intersections (Figure 3), real-world mobility traces recorded in the NGSIM project [33] on highway Route 101 (Figure 4), and a stretch of simulated highway that is only partially covered by an adversary (Figure 5). We investigate different traffic volumes in the synthetic scenarios, ranging from nearly empty to almost clogged roads. This is complemented by the scenario utilizing a realistic mobility trace, which features 15 min of versatile traffic, including jams, traffic shock-waves, and free-flowing traffic.

In all scenarios, the adversary set up access points and tries to track all vehicles based on the received beacon messages. The adversary has exactly one chance to guess which vehicle was which when it leaves the simulation, and the vehicle counts as tracked if the adversary is correct.

We intentionally focus on a seemingly best case for privacy protection, that is, low beacon frequencies (including values lower than the recommended minimum in the standards [25, Table III-1],[34]), short pseudonym validity times (e.g., new pseudonyms for each message) to illustrate that even under these optimal conditions, privacy is extremely difficult to achieve.

The key simulation parameters are gathered in Table 1.

6.1. Tracking Framework

The idea of our evaluation is to assume the role of an adversary who tries to track vehicles based on their transmitted periodic beacon messages. To this end, we implemented a modern tracking system, the full background of which can be found in [35]. In this section only a brief summary of its theoretical background is reproduced.

There exists a large number of different tracking systems [36], many of them designed for specific purposes. In the field of vehicle tracking, it is common to use a tracking system design as depicted in Figure 6. The starting point for each tracking system is always a set of observations $O = \{o_1, \ldots, o_n\}$ made by an adversary. An observation can be obtained in various ways and it can consist of an arbitrary amount of information. For example, observations made by an adversary who set up a camera system on an intersection would consist of timestamps, positions, colors, and object dimensions, while observations obtained using a radio receiver would include information contained in the received message and other information the adversary can correlate. In the context of vehicular networks this applies to all information contained in the periodic Basic Safety Messages (BSMs) (or Cooperative Awareness Messages (CAMs), respectively) sent by all vehicles.

The goal of the adversary is to create a track for each vehicle. A track T_i is a finite sequence of observations, e.g., sent messages, that the adversary believes belong to the same vehicle. The problem of tracking can now be defined as finding the correct observation that belongs to an existing track. This is illustrated in Figure 7: Assume an adversary has already successfully tracked three vehicles using observations made at time $t = 1$ and $t = 2$. At time $t = 3$ the set of observations O includes three received broadcast messages. All tracks and observations are used as the input for the tracking algorithm to assign an observation $o_i \in O$ to a track T_j, or if not possible, to start a new track or end an existing track.

6.1.1. Filtering and Prediction

The first step in a tracking system is to filter the collected observations. Observations made with the help of sensors are usually subject to noise and are therefore inaccurate to some degree. Depending on the type of sensor and noise, there exist different filter mechanisms to adjust the readings and thereby increase their accuracy. In the context of position data, this is usually done by help of a Kalman filter [37]. In a vehicular network, the position data received by vehicles does not necessarily require filtering. Ideally, the transmitting vehicle itself should already transmit filtered position information as they have direct access to all sensors. The need for accurate position information is particularly relevant considering that transmitted position information is an important input for the safety applications of receiving vehicles. Also, the OBUs of the vehicles are expected to have limited processing power, making it possibly infeasible to run extensive filtering techniques for each neighboring vehicle. If vehicles transmit unfiltered position information, the adversary can apply filters using the included information in BSM and CAM broadcasts.

Once the observations have been filtered, the adversary predicts (or extrapolates) the next expected observation of each track. One method to arrived a prediction is to simply use the velocity and position of the latest observation in a track. However, observations may include much more than only position information and every piece of information can be used by the adversary to track a vehicle. For example, steering wheel angles, turn signals and other information included in the periodic safety messages can be exploited by the adversary to estimate a vehicle's next position. At

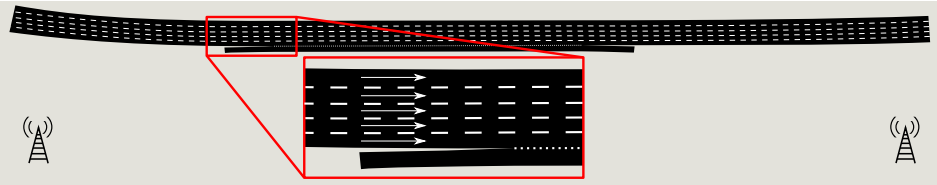

Figure 4: Highway 101 Scenario. Node movement is based on 900 s of real traffic on a 640 m stretch of Hollywood Fwy near Universal City Plaza, Los Angeles, CA. Five lanes are running in the same direction, temporarily joined by a sixth lane in the middle. Traffic was recorded by eight video cameras and post-processed to derive trace files [33]. An adversary is able to fully monitor the scenario.

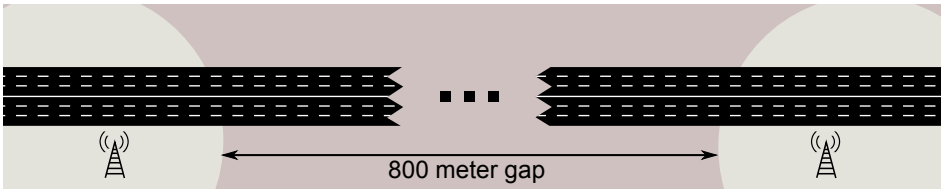

Figure 5: Freeway blind spot scenario: an adversary set up two access points on a highway, partially monitoring the scenario, unable to receive messages from vehicles in the 800 m blind spot.

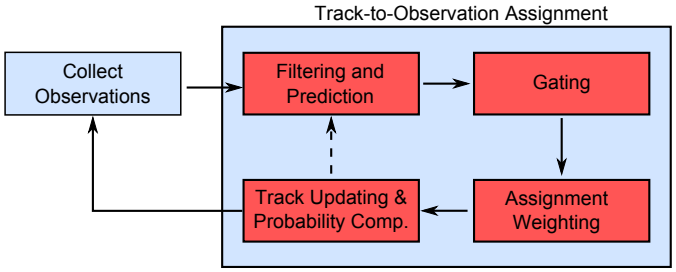

Figure 6: Architecture of a vehicle tracking system.

the end of the prediction phase, there exists exactly one estimated successor state for each track.

6.1.2. Gating

Gating is the process of eliminating all unlikely successors for each track to increase the performance of the tracking system by decreasing the overall number of required comparisons between observations and predictions. It is a per-track operation that identifies all $o_i \in O$ that cannot be used (or have a likeliness below a certain threshold) to continue a track T_j. Assume again the situation illustrated in Figure 7: When finding the possible successor for each track, some observations may be neglected because it might have been physically impossible for the vehicle associated with a track to reach the given position.

Gating is not limited to geographic areas, but can be extended to all kinds of information. For example, if vehicles transmitted their (most likely rounded) vehicle dimensions, all observations (that is, received broadcast messages) containing different vehicle dimensions could be disregarded. In the case of pseudonyms, and assuming pseudonyms are unique and not exchanged between vehicles, an attacker could discard observations with pseudonyms that are already associated with other tracks, eliminating the effect of isolated, non-coordinated pseudonym changes already at the gating stage. In general, it can be said that the gating process is often dependent on the Privacy-Enhancing Technology (PET) or the privacy vulnerability itself. Knowledge about the PET can be used to reduce the number of possible observations and some privacy vulnerabilities may even lead to a situation where an attacker can exclude all observations but one, e.g., when they are able to predict a certain state and only one observation matches their prediction or when a vehicle has not changed its pseudonym in the last obversation interval.

In this study, we use a gating mechanism which tries to eliminate all observations that cannot be physically reached from a given track endpoint. In addition to the distance from the observation to the track endpoint and the difference in speed, we also analyse the maximum yaw rate. The maximum yaw rate of a vehicle, given in degrees per second, depends on the current velocity of a vehicle, typical maximum values ranging from 75 °/s at very low speeds to 5 °/s on freeways. It has to be noted that the feasibility of this approach heavily relies on the underlying vehicle simulation model: For example, some simulators allow vehicles to turn around almost instantly, possibly leading to a failure of tracking when the maximum yaw rate is used for gating. To avoid over-fitting our tracking system with regard to a certain simulator, we also used real-world traces to test our hypothesis that local privacy cannot be achieved with standard-compliant beacon frequencies.

6.1.3. Assignment Weighting

After all unlikely observations are discarded for a certain track, the tracking algorithm estimates the likelihood of all remaining observations to continue the track. For that, a rating mechanism is needed. The most obvious rating is to assign each observation in the gating area the same probability [16], regardless of its distance or difference compared to a predicted position $\bar{e}_i$ of the track T_i. In most cases this will lead to a false sense of privacy, maximizing

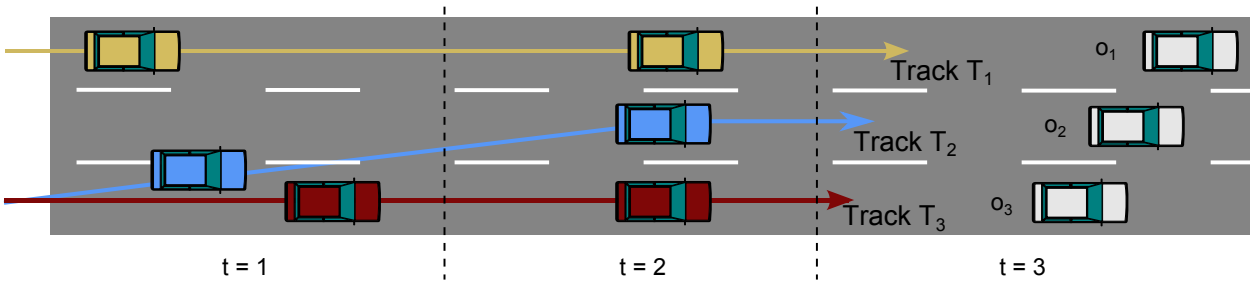

Figure 7: Tracking can be seen as the problem of assigning a new observation o_i to a track T_j.

metrics like entropy and reducing the maximum tracking time. This mechanism therefore corresponds to a weak adversary that cannot make use of the information included in the observations.

Reducing observations to only their position for the computation of the assignment weights is not preferable as all information contained in an observation can be used by a potential adversary. Therefore [36] and others propose the use of the Mahalanobis distance [38] to incorporate all possible dimensions of the target's state $X_t \in \mathbb{R}^n$. It is defined as $\sqrt{(\bar{e}-o)^\intercal \cdot S^{-1} \cdot (\bar{e}-o)}$ with S being the covariance matrix. As a worst case assumption, we assume the covariance matrix to be diagonal, that is, the variance σ^2 of each dimension to be uncorrelated; then, the Mahalanobis distance d_m between the estimated state $\bar{e}$ and an observation o regarding K dimensions of the state can be given in the form of:

$$d_m(\bar{e}, o) = \sqrt{\sum_{i=1}^{K} \frac{(\bar{e}_{[i]} - o_{[i]})^2}{\sigma_{[i]}^2}} \tag{2}$$

To expand on the principle of the Mahalanobis distance (also known as *squared statistical distance* or *normalized Euclidean distance*), assume the assignment weight depends on the actual positions p, the velocities v, and the heading ϕ of the estimated state $\bar{e}$ and an observation o. Then the distance becomes:

$$d_m(\bar{e}, o) = \sqrt{\frac{(p_{\bar{e}} - p_o)^2}{\sigma_p^2} + \frac{(v_{\bar{e}} - v_o)^2}{\sigma_v^2} + \frac{(\phi_{\bar{e}} - \phi_o)^2}{\sigma_\phi^2}} \tag{3}$$

The variances can be seen as a weighting mechanism for each of the terms, as they reflect the uncertainty of the prediction. Using this distance, we can create a ranking of observation-to-track assignments for every track.

6.1.4. Track Updating & Probability Computations

After the distance for each possible track-to-observation assignment has been calculated, the tracking algorithm has to decide which observation continues which track. This solution has to be unambiguous, that is, one observation must be used to continue only one track. Also, the tracking system can determine whether a new track has started (e.g., if an observation could not be assigned to a track) or a track has ended (if no suitable observation has been found to continue the track for a given time interval). This overall solution is referred to as a *hypothesis* and reflects one of the adversary's possible views of the system.

This problem can be mapped to the auction house problem [36]: First, an $n \times m$ track-to-observation assignment matrix with n tracks, m observations, and the corresponding entries a_{ij} to be values indicating the quality of the assignment of T_i to o_j is created. Finding the global optimum would then be the selection of ≤ 1 entries per row so that the sum of all selections becomes maximal. An easier to implement alternative leading to the same result is given by [39], which suggests that the track-to-observation assignment be converted to a graph $G = (V, E)$ with $V = \dot{V} \cup \tilde{V}$ and $\dot{V}$ being the track endpoints and $\tilde{V}$ being all observations. Edges E are a subset of the Cartesian product $E \subset \dot{V} \times \tilde{V}$, making G a directed graph with edges only from track endpoints to observations. Each edge e is assigned a cost c_e depending on the statistical distance between the track endpoint and the observations. The following method requires these costs to be high for small distances and vice versa, for example, by assigning each edge the negative value of the statistical distance [36]. Tracks are only connected with observations in their gating area. The goal is to find a solution $E_s \subset E$ that satisfies the properties that each $v \in \dot{V}$ has an outdegree ≤ 1 and each $v \in \tilde{V}$ has an indegree ≤ 1 while maximizing the sum of all costs $\sum_{e \in E_s} c_e$. The maximum matching problem can then be solved by the Edmonds algorithm [40] in $O(n \cdot m \cdot log(n))$ time as implemented in the Lemon template library [41].

After having computed a solution to the auction house problem, it is necessary to be able to assign probabilities to track-to-observation assignments and consequently to hypotheses. This probability is an important input for metrics such as the entropy or success rates. For that, we deploy the Joint Probabilistic Data Association (JPDA) method as described in [36].

For every assignment of track i to observation j a Gaussian likelihood value g_{ij} is computed. This is done using the statistical distance (cf. Equation (2)), the number of dimensions M, and the covariance matrix S_{ij} of these dimensions.

$$g_{ij} = \frac{e^{-d_m(\bar{e}_i, j)^2/2}}{(2\pi)^{M/2} \cdot \sqrt{|S_{ij}|}} \tag{4}$$

Parameter	Value
Simulation	extended Veins
Adversary Model	external, local, passive, static, domain-specific
Privacy Property	unlinkability
Data Source	observable information
Metrics	tracking fail rate
Scenario	grid, highway 101, blind spot freeway
Technology	IEEE WAVE
Beacon Frequency	0.25 Hz to 10 Hz
No. of Vehicles	25-600
Transmission Power	10 mW
Bitrate	6 Mbit/s
Radio Sensitivity	89 dBm
Thermal Noise	−110 dBm
Tracking Parameters	for 0.5s, 1s, 2s, 5s beacon int.
σ_p	1, 1.4, 4.38, 80
σ_v	1, 2, 2.5, 40
σ_ϕ	0.12, 0.26, 3.44, 15
Track Timeout	4s, 4s, 6.5s, 11s

Table 1: Setup and parameters used in the evaluation.

Furthermore assume the set of all selected assignments to be G and P_D to be the probability of successfully detecting an observation; in the context of wireless networks, this can be related to the packet loss rate (as a worst case assumption, we will assume this number is known) and the probability of a track actually ending (on a highway, this probability can be assumed to be zero). Then the unnormalized probability of a hypothesis $p'(H_k)$ can be computed using the extraneous return density β (in this case, the density of new tracks in the gating areas, a term that can be ignored in the scenarios considered in this study), the number of continued tracks m, the number of discontinued tracks e, the number of unassigned observations u, and the product of all assignments $g \in G$.

$$p'(H_k) = (1 - P_D)^e \cdot (P_D)^m \cdot \beta^u \cdot \prod_{g \in G} g \qquad (5)$$

This probability can be used to rank all hypotheses in terms of likelihood. It can also be normalized and be used to compute the probability for each observation-to-track assignment [35]. In this simulation study, while we compute various hypotheses for each observation period, we only store and continue with the most likely one. This hypothesis can be seen as the adversary's view on the current state of the system.

6.2. Metrics

There are many ways to measure the level of privacy in a system. Many metrics have been proposed [10], evaluating various aspects of privacy the system. As we assume that vehicles are pseudonymous (and anonymous) at the beginning of the simulation, we mainly focus on the linkability property of the system, that is, how well an adversary is able to link two messages. The drawback of many proposed metrics that can be used to measure linkability (e.g., the entropy) is that that their numeric values are difficult to comprehend and interpret not only for laypeople [10]. We therefore chose an easy-to-understand metric to present our results: We chose the tracking fail rate which can also be interpreted as the adversary failure rate. It is the chance of a vehicle evading tracking with a higher fail rate representing a higher level of privacy.

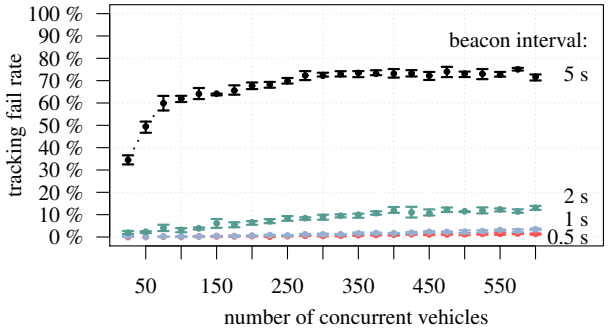

Figure 8: Chance of evading tracking in the synthetic grid scenario. Plotted are the averages over all simulation runs. Error bars extend from the 25 % to the 75 % quantiles.

7. Results

7.1. Urban scenario

As a first step, we investigated the effect of beacon intervals and traffic density on the adversary's capability to track vehicles. To support our claim that local privacy cannot be achieved without affecting traffic safety, we configured the scenario in the best possible way for privacy. Each beacon was sent with a new pseudonym, completely eliminating the possibility to link two messages based on the sender address. Further, the adversary was only allowed to utilize position, speed, and heading information in the beacons. In a real-world scenario, information such as the state of the turn signals or the steering wheel angle would allow much easier tracking [3]. We introduced a position noise of about 4 m, making it hard for the attacker to detect the lane on which a vehicle is driving.

Figure 8 shows our results for the synthetic grid scenario. Looking at beacon frequencies of 1 Hz and 2 Hz, we observe that the chance of not being tracked is lower than 5 %. We investigated how certain vehicles evaded tracking and found the primary cause to be packet loss, rendering these vehicles invisible to the attacker. With beacon frequencies below the specified minimum frequency of 1 Hz in the IEEE (see SAE J2945/1-2.2) and ETSI standards (see ETSI 302 637-2-V1.3.0), the level of privacy improved significantly. With only one beacon every 5 s, the adversary was no longer able to reliably track vehicles. The reason for this is that within

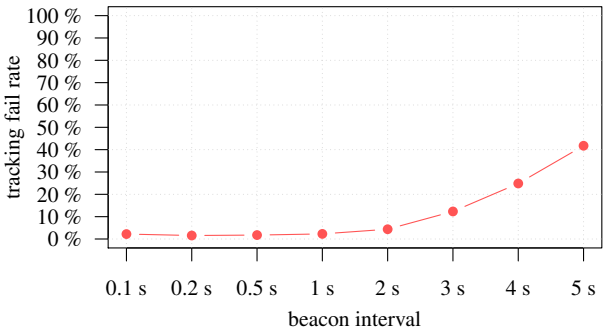

Figure 9: Chance of evading tracking in the (fully deterministic) real-world highway scenario.

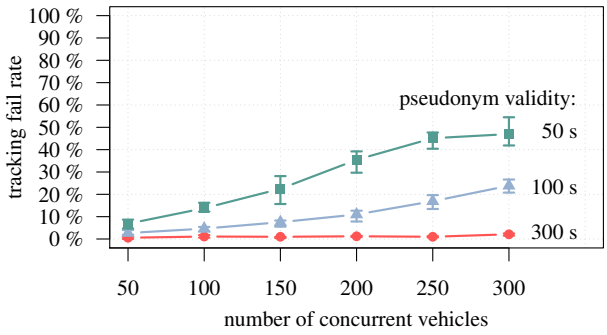

Figure 10: Chance of evading tracking in the blind-spot freeway scenario (beacon rate = 1Hz). Plotted are the averages over all simulation runs. Error bars extend from the 25 % to the 75 % quantiles.

5 s vehicles could perform complete turning maneuvers. This confused the adversary, additionally leading to error propagation in vehicle assignment. It has to be noted that these beacon frequencies are far beyond the safety requirement of vehicular networks [42] and are therefore not a viable configuration.

The synthetic nature of the intersection scenario could lead to a false sense of privacy protection. We therefore investigated a real-world trace recorded during the NGSIM project on the US American Highway Route 101 [33]. The trace contains vehicle information at 10 Hz resolution, including vehicle position and velocity, but not heading, which we added by computing position difference between two data points. We artificially created lower beacon frequencies by equidistantly sampling vehicle information with the desired frequency. Again, the adversary does not make use of identifiers, tracking solely using position, velocity, and heading.

Results are shown in Figure 9. Vehicles were unable to confuse our tracking algorithm for beacon intervals of 1 s and lower. In fact, tracking 'real' vehicles turned out to be easier than tracking simulated ones. A possible explanation is that in the simulation environment, vehicles sometimes behave unnaturally, disregarding the laws of physics, e.g., by suddenly turning around or instantly changing lanes. The results further indicate that at safety compliant beacon frequencies, confusing an eavesdropping attacker is nearly impossible. This means that all pseudonym changing strategies that do not alter the beacon frequency beyond a safety limit will be ineffective.

7.2. Freeway with radio blind-spot

The first results confirm our approach to not try and pursue privacy to nearby vehicles and instead fully concentrate on privacy protection when and where an adversary is not eavesdropping. The time in which a vehicle's transmissions cannot be overheard by an adversary must then be used effectively to increase the level of location privacy. To confuse an attacker, not only must a vehicle change its own pseudonym before re-entering an area covered by an adversary, but, ideally, many vehicles will have done the same to increase confusion for the adversary. We illustrate this effect by investigating a synthetic freeway scenario, where an adversary set up two receiver stations with an 800 m wide radio blind spot in-between. In this scenario, vehicles will use pseudonyms for more than one message and the adversary will exploit this by linking messages based on the used identifier. Pseudonym changes are not synchronized, i.e., vehicles change pseudonyms independently.

Figure 10 shows our results, comparing different pseudonym validities. With a validity time of 300s as recommended in the notice of proposed rulemaking by the National Highway Traffic Safety Administration [25], a gap of 800m was not enough to confuse the adversary noticeably. Please note that we are not arguing towards shorter pseudonym validity times as the effect on the tracking success clearly depends on the size of the radio blind spot. We rather aim to illustrate that only when an attacker is unable to overhear consecutive messages, can tracking be avoided. Shorter times for the pseudonym validity considerably increased the tracking fail rate in our scenario. Not being able to monitor lane changes and overtaking maneuvers in the blind spot makes it observably difficult for the adversary to re-identify vehicles that changed their identifier. At the highest traffic volume, a 50 s validity caused about 80 % of all vehicles to change their pseudonym in the blind spot. More than half of these vehicles could not be properly re-identified by the attacker, emphasizing the need for synchronous pseudonym changing to cause attacker confusion.

7.3. Synchronized vs non-synchronized changing

Lastly, we compared the impact of synchronized pseudonym changing versus non-synchronized pseudonym changing (e.g., [25]). Both approaches can make use of time-slots, pseudonym change advertising and efficient pseudonym revocation; they only differ in one aspect: In the synchronized setting, all vehicles have the exact same time-slot boundaries, causing every car to change its pseudonym at the same time. In the second approach, every car has in fact the same time-slot length, however, the boundaries of the slots are arbitrarily offset.

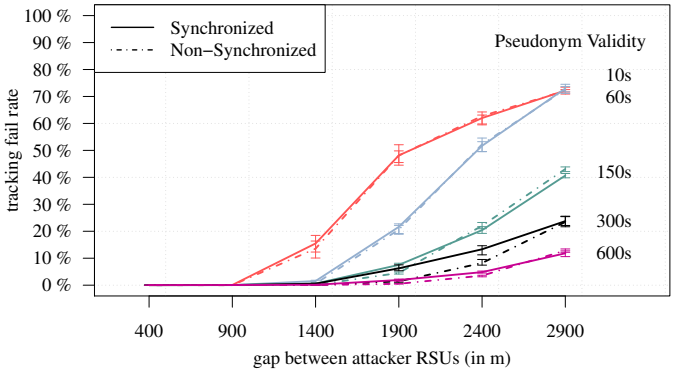

(a) Fixed traffic volume of 250 concurrent vehicles with a varying distance between attacker RSUs and different pseudonym validity times

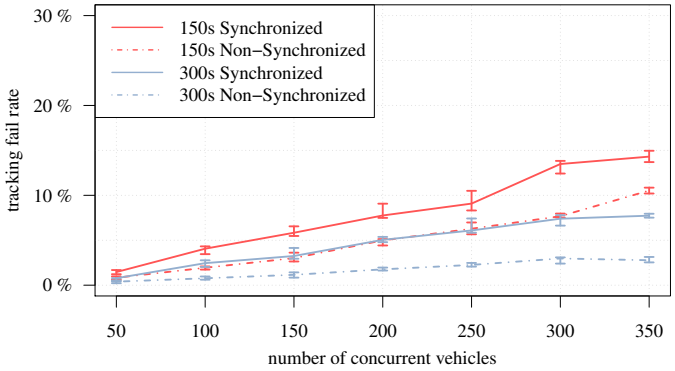

(b) Fixed distance of attacker RSUs of 2000m with a varying number of concurrent vehicles and two different peudonym validity times

Figure 11: Comparison of synchronized vs. non-synchronized pseudonym changing in the blind-spot freeway scenario (beacon rate = 1Hz). Plotted are the averages over all simulation runs. Error bars extend from the 25 % to the 75 % quantiles. Higher values indicate better privacy.

We first study the difference in the freeway blind-spot scenario by altering the distance between the attacker RSUs under different pseudonym validity times (see Figure 11a). As expected, we observe no difference when the combination of attacker gap and pseudonym validity is such that it is guaranteed that each vehicle will at least change its pseudonym once while in the radio blind spot. This is always the case for a validity time of 10s and also for 60s if the gap is large enough. Interestingly, even for the other case the difference between the approaches is not prominent. Of course, only vehicles that changed their pseudonym while in the radio blind spot have a chance of evading the attacker. The chance of avoiding to be tracked then depends on the number of other vehicles who have also changed their pseudonym and on driving manoeuvres that will confuse the tracking mechanism. In the synchronized approach, the periods between the global slot boundaries are periods where no vehicle evades tracking (because nobody changes its pseudonym), while at the global slot boundary, every car changes its pseudonym. In the non-synchronized approach, the number of cars who change their pseudonym depends on their individual slot boundaries and the chance to evade tracking therefore also relies on how many other vehicles have done so while driving in the blind spot together.

The performance of both approaches only show differences for a small number of distance and validity combinations, most visible for pseudonym validity times of 300s and gaps for around 2000m distance. To better understand this effect, we additionally simulated a varying vehicle density and set the RSU distance to a fixed value of 2000m. Our results are shown in Figure 11b. We observe that synchronized changing outperforms non-synchronized changing in all vehicle densities, even to the extent that a 150s validity with synchronized changes provides the same privacy level as 300s with non-synchronized changes. We have to note though, that the level of privacy is very low and only about 4% to 12% of all vehicles were able to avoid being tracked. We conclude that if time-slots are used, synchronized changing should be preferred over non-synchronized changing.

7.4. Results summary

Our results clearly show that confusing a local attacker is not possible at beacon frequencies necessary for the reliable operation of traffic safety applications. This leads to the conclusion that our proposal is in fact not sacrificing local privacy, but merely taking into account that local privacy and traffic safety are not compatible within the parameters of IEEE WAVE and ETSI ITS-G5. Even if they were compatible and an adversary had no means to link two messages based on their address or content, it was shown that physical layer fingerprinting attacks can completely bypass privacy protection mechanisms [43]. Locally announcing pseudonym changes has therefore only marginal impact on privacy protection, yet, it completely overcomes the privacy–safety trade-off problem.

This proposal does also not introduce new attack vectors on privacy: vehicles close enough to receive two or more pseudonym change announcements from the same vehicle are most likely also close enough to visually see this vehicle. To track a vehicle based on pseudonym change announcements either requires global knowledge of all sent messages or to physically follow the vehicle, which can be done regardless of any car-to-car communication.

In terms of overhead, our proposal only requires one additional certificate verification for each nearby vehicle when a new time slot starts. After a vehicle cryptographically proves that it owns both old and new pseudonym, receivers no longer have to check both signatures.

8. Conclusion

In this work, we took a structured approach to deriving a holistic solution for location privacy protection in

VANETs. For this, we carefully selected which aspects of privacy the solution should preserve and which (realistic) adversary model and attack channel it should consider. Our results clearly show that confusing an attacker that receives all messages is not possible at message frequencies necessary for the reliable operation of traffic safety applications. We conclude to defend location privacy against local passive adversaries operating a broad (but not global) network of channel sniffers. We then reviewed which real-world restrictions must be adhered to by a solution for the defense of users' privacy. Most importantly, we identified low overhead (in terms of data size and computational complexity), maintaining accountability of senders, as well as an overruling need to not interfere with traffic safety.

We introduced a model of a state-of-the-art attacker employing a multi-target tracking algorithm and applied this model in extensive computer simulations employing both simulated vehicle movement and real world traces. The study supports our findings that, under reasonable assumptions about an adversary's capability, local privacy is neither required nor can it be achieved without compromising traffic safety.

Consequently, we propose a system consisting of three key components: First, using synchronized time-slotted pseudonym pools, that is, using multiple pseudonyms for communication of which only one is valid at any given time. This simultaneously limits storage overhead and increases adversaries' confusion as well as wards against Sybil attacks (unlike overlapping pseudonym systems). Second, making pseudonym changes visible to direct neighbors, simply by briefly including old pseudonyms after a pseudonym change. This cancels out any negative impact of the proposed system on users' safety without sacrificing privacy; as we have shown a local adversary can easily follow pseudonym changes anyway – either by correlating message contents or by observing physical properties of the transmission. Third, time-slotted pools work well with highly efficient revocation schemes and allow for the preserving of backward privacy. In summary, we overcome the privacy–safety problem while at the same time increasing privacy for all users. Our system is fully compatible with the requirements of envisioned vehicular networks.

References

[1] J. Harding, G. Powell, R. Yoon, J. Fikentscher, C. Doyle, D. Sade, M. Lukuc, J. Simons, J. Wang, Vehicle-to-Vehicle Communications: Readiness of V2V Technology for Application, NHTSA Technical Report DOT HS 812 014, National Highway Traffic Safety Administration (Aug. 2014).

[2] J. Gozalvez, M. Sepulcre, R. Bauza, IEEE 802.11p Vehicle to Infrastructure Communications in Urban Environments, IEEE Communications Magazine 50 (5) (2012) 176–183.

[3] D. Eckhoff, C. Sommer, Driving for Big Data? Privacy Concerns in Vehicular Networking, IEEE Security & Privacy 12 (1) (2014) 77–79.

[4] D. Eckhoff, I. Wagner, Privacy in the Smart City – Applications, Technologies, Challenges and Solutions, IEEE Communications Surveys and Tutorials (2017) .

[5] J.-P. Hubaux, S. Čapkun, J. Luo, The Security and Privacy of Smart Vehicles, IEEE Security and Privacy 2 (3) (2004) 49–55.

[6] J. Petit, F. Schaub, M. Feiri, F. Kargl, Pseudonym Schemes in Vehicular Networks: A Survey, IEEE Communications Surveys and Tutorials 17 (1) (2015) 228–255.

[7] N. Bißmeyer, H. Stübing, E. Schoch, S. Götz, J. P. Stotz, B. Lonc, A generic public key infrastructure for securing Car-to-X communication, in: 18th World Congress on Intelligent Transport Systems, Orlando, FL, 2011.

[8] H. Stübing, M. Bechler, D. Heussner, T. May, I. Radusch, H. Rechner, P. Vogel, simTD: A Car-to-X System Architecture for Field Operational Tests, IEEE Communications Magazine 48 (5) (2010) 148–154.

[9] R. Stahlmann, A. Festag, A. Tomatis, I. Radusch, F. Fischer, Starting European Field Tests for Car-2-X Communication: the DRIVE C2X Framework, in: 18th ITS World Congress and Exhibition, Orlando, FL, 2011.

[10] I. Wagner, D. Eckhoff, Privacy Assessment in Vehicular Networks Using Simulation, in: Winter Simulation Conference (WSC '14), IEEE, Savannah, GA, 2014, pp. 3155–3166.

[11] D. Eckhoff, C. Sommer, Marrying Safety with Privacy: A Holistic Solution for Location Privacy in VANETs, in: 8th IEEE Vehicular Networking Conference (VNC 2016), IEEE, Columbus, OH, 2016, pp. 290–297.

[12] J. Douceur, The Sybil Attack, in: Peer-To-Peer Systems: First International Workshop (IPTPS 2002), Springer, Cambridge, MA, 2002, pp. 251–260.

[13] B. Xiao, B. Yu, C. Gao, Detection and Localization of Sybil Nodes in VANETs, in: 2006 Workshop on Dependability issues in wireless ad hoc networks and sensor networks (DIWANS 2006), ACM, Los Angeles, CA, 2006.

[14] D. Eckhoff, C. Sommer, T. Gansen, R. German, F. Dressler, SlotSwap: Strong and Affordable Location Privacy in Intelligent Transportation Systems, IEEE Communications Magazine 49 (11) (2011) 126–133.

[15] L. Huang, K. Matsuura, H. Yamane, K. Sezaki, Enhancing Wireless Location Privacy Using Silent Period, in: IEEE Wireless Communications and Networking Conference (WCNC 2005), New Orleans, LA, 2005.

[16] K. Sampigethaya, L. Huang, M. Li, R. Poovendran, K. Matsuura, K. Sezaki, CARAVAN: Providing location privacy for VANET, in: Embedded Security in Cars (ESCAR 2005), Tallinn, Estonia, 2005.

[17] M. Gerlach, F. Güttler, Privacy in VANETs Using Changing Pseudonyms - Ideal and Real, in: 65th IEEE Vehicular Technology Conference (VTC2007-Spring), Dublin, Ireland, 2007, pp. 2521–2525.

[18] J. Freudiger, M. Raya, M. Félegyházi, P. Papadimitratos, J.-P. Hubaux, Mix-Zones for Location Privacy in Vehicular Networks, in: First Int. Workshop on Wireless Networking for Intelligent Transportation Systems (WiN-ITS 2007), ACM, Vancouver, Canada, 2007, pp. 1–7.

[19] R. L. Finn, D. Wright, M. Friedewald, Seven Types of Privacy, in: S. Gutwirth, R. Leenes, P. de Hert, Y. Poullet (Eds.), European Data Protection: Coming of Age, Springer, 2013, pp. 3–32.

[20] A. Pfitzmann, M. Hansen, A terminology for talking about privacy by data minimization: Anonymity, Unlinkability, Undetectability, Unobservability, Pseudonymity, and Identity Management, v0.34 (Aug. 2010).

[21] P. Golle, K. Partridge, On the Anonymity of Home/Work Location Pairs, in: 7th International Conference on Pervasive Computing, Vol. LNCS 5538, Springer, Nara, Japan, 2009, pp. 390–397.

[22] I. Wagner, D. Eckhoff, Technical Privacy Metrics: a Systematic Survey, Tech. Rep. 1512.00327, arXiv, arXiv: 1512.00327 (Dec. 2015).

[23] S. Lefèvre, J. Petit, R. Bajcsy, C. Laugier, F. Kargl, Impact of V2X Privacy Strategies on Intersection Collision Avoidance Systems, in: 5th IEEE Vehicular Networking Conference (VNC 2013), IEEE, Boston, MA, 2013, pp. 71–78.

[24] M. Raya, R. Shokri, J.-P. Hubaux, On the Tradeoff Between

Trust and Privacy in Wireless Ad Hoc Networks, in: 3rd ACM Conference on Wireless Network Security (WiSec 2010), ACM, Hoboken, NJ, 2010, pp. 75–80.
[25] National Highway Traffic Safety Administration (NHTSA), Notice of Proposed Rulemaking (NPRM): Federal Motor Vehicle Safety Standards; V2V Communications (NHTSA-2016-0126), Entry 82 FR 3854, Federal Register (Jan. 2017).
[26] M. Feiri, J. Petit, F. Kargl, The Case for Announcing Pseudonym Changes, in: 3rd GI/ITG KuVS Fachgespräch Inter-Vehicle Communication (FG-IVC 2015), Ulm, Germany, 2015.
[27] J. J. Haas, Y.-C. Hu, K. P. Laberteaux, Efficient Certificate Revocation List Organization and Distribution, IEEE Journal on Selected Areas in Communications 29 (3) (2011) 595–604.
[28] D. Eckhoff, F. Dressler, C. Sommer, SmartRevoc: An Efficient and Privacy Preserving Revocation System Using Parked Vehicles, in: 38th IEEE Conference on Local Computer Networks (LCN 2013), IEEE, Sydney, Australia, 2013, pp. 855–862.
[29] W. Whyte, A. Weimerskirch, V. Kumar, T. Hehn, A Security Credential Management System for V2V Communications, in: 5th IEEE Vehicular Networking Conference (VNC 2013), IEEE, Boston, MA, 2013, pp. 1–8.
[30] D. Eckhoff, M. Protsenko, R. German, Towards an Open Source Location Privacy Evaluation Framework for Vehicular Networks, in: 80th IEEE Vehicular Technology Conference (VTC2014-Fall), IEEE, Vancouver, Canada, 2014.
[31] B. Wiedersheim, Z. Ma, F. Kargl, P. Papadimitratos, Privacy in Inter-Vehicular Networks: Why simple pseudonym change is not enough, in: 7th IEEE/IFIP Conference on Wireless On demand Network Systems and Services (WONS 2010), Kranjska Gora, Slovenia, 2010.
[32] C. Sommer, R. German, F. Dressler, Bidirectionally Coupled Network and Road Traffic Simulation for Improved IVC Analysis, IEEE Transactions on Mobile Computing 10 (1) (2011) 3–15.
[33] Federal Highway Administration (FHWA), NGSIM Program US Route 101 data, Version 1, available online: http://www.its-rde.net/ (2016).
[34] D. Eckhoff, N. Sofra, R. German, A Performance Study of Cooperative Awareness in ETSI ITS G5 and IEEE WAVE, in: 10th IEEE/IFIP Conference on Wireless On demand Network Systems and Services (WONS 2013), IEEE, Banff, Canada, 2013, pp. 196–200.
[35] D. Eckhoff, Simulation of Privacy-Enhancing Technologies in Vehicular Ad-Hoc Networks, Ph.D. thesis, University of Erlangen (Mar. 2016).
[36] S. Blackman, R. Popoli, Design and Analysis of Modern Tracking Systems, Artech House Boston, 1999.
[37] R. Kalman, A new approach to linear filtering and prediction problems, Transaction of the ASME Journal of Basic Engineering D (82) (1960) 35–45.
[38] P. C. Mahalanobis, On the Generalized Distance in Statistics, Proceedings of the National Institute of Sciences of India 2 (1) (1936) 49–55.
[39] M. Protsenko, A Framework for Performance Analysis of Tracking Algorithms in Vehicular Networks, Pre-master's thesis (studienarbeit), University of Erlangen (Aug. 2011).
[40] J. Edmonds, Maximum Matching and a Polyhedron with 0, l-vertices, J. Res. Bur. Stand 69B (1-2) (1965) 125–130.
[41] B. Dezső, A. Jüttner, P. Kovács, LEMON - an Open Source C++ Graph Template Library, Electronic Notes in Theoretical Computer Science 264 (5) (2011) 23–45.
[42] N. An, M. Maile, D. Jiang, J. Mittag, H. Hartenstein, Balancing the Requirements for a Zero False Positive/Negative Forward Collision Warnings, in: 10th IEEE/IFIP Conference on Wireless On demand Network Systems and Services (WONS 2013), IEEE, Banff, Canada, 2013, pp. 191–195.
[43] B. Bloessl, C. Sommer, F. Dressler, D. Eckhoff, The Scrambler Attack: A Robust Physical Layer Attack on Location Privacy in Vehicular Networks, in: 4th IEEE International Conference on Computing, Networking and Communications (ICNC 2015), CNC Workshop, IEEE, Anaheim, CA, 2015, pp. 395–400.