Marrying Safety with Privacy: A Holistic Solution for Location Privacy in VANETs

David Eckhoff*† and Christoph Sommer‡
* Computer Networks and Communication Systems, Dept. of Computer Science, University of Erlangen, Germany
† Computing and Information Systems Department, The University of Melbourne, Australia
‡ Heinz Nixdorf Institute and Dept. of Computer Science, Paderborn University, Germany
david.eckhoff@unimelb.edu.au, sommer@ccs-labs.org

***Abstract*—Existing privacy measures often conflict with the requirements of future vehicular networks. First and foremost, in attempting to achieve local privacy, they interfere with the primary purpose of vehicular networks: to improve road safety. Other solutions undermine accountability or introduce too high overhead to be of practical use. For this reason, proper privacy protection is disregarded in many field experiments, proposals, and standardization documents. In this paper, we take a structured approach to deriving a holistic solution for location privacy protection in Vehicular Ad-Hoc Networks (VANETs): Under reasonable assumptions about an adversary's capability, local privacy is neither required nor can it be achieved without compromising traffic safety. Our approach is therefore 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.**

I. Introduction

North America, Europe, and Japan are progressing toward the age of wirelessly communicating vehicles as an important part of future intelligent transportation systems. In an effort to both improve traffic safety and also to bring new applications to modern vehicles, the American IEEE and the European ETSI are finalizing standardization documents for the operation of future vehicular networks. In recent events, the US DOT has announced its intent to make ad-hoc communication systems mandatory for new cars [1].

One of the key features in all of these systems is that vehicles establish a virtual view of their surroundings. This is achieved by creating so-called cooperative awareness through periodic (with a frequency of 1 Hz to 10 Hz) broadcast transmissions, or beacons: In IEEE WAVE and ETSI ITS-G5, each vehicle informs all other cars in its vicinity (up to 300 m to 800 m [2]) of its current state including its position, heading, and velocity.

These unencrypted transmissions can be received not just by other vehicles but anyone with a receiver physically close enough to the sender. 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].

This privacy problem in vehicular networks has been understood from the very beginning [4]. 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 [5] has been proposed since, some of them also considered in various field trials [6], [7], 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 [8]: 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 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 paper, 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 maximize 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.

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.

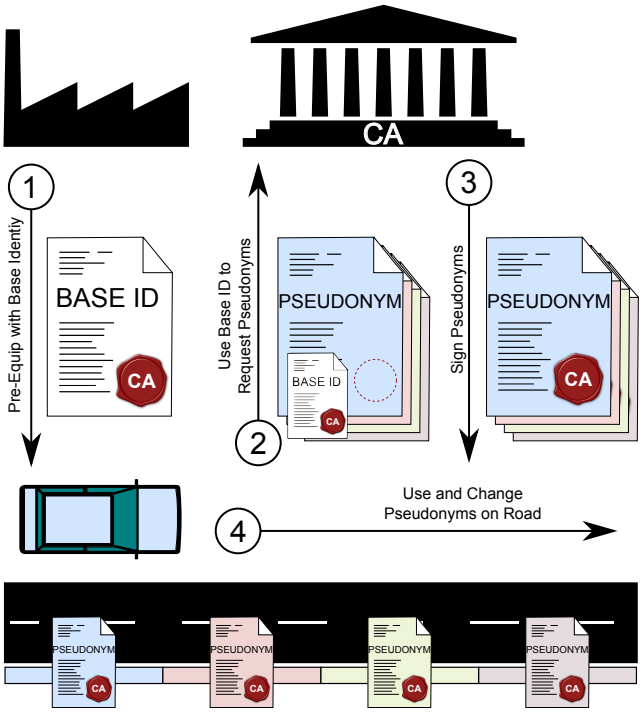

Figure 1. A simplified vehicular Public Key Infrastructure (PKI).

The remainder of the paper is structured as follows: In Section II 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 paper, we will also refer to and discuss related work. Section III discusses the privacy threats of vehicular networks; Section IV explains the constraints that privacy protection mechanisms must work within. In Section V 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.

II. 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 [9], [10], 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 [11].

Other proposals for privacy protection include the use of silent periods, that is, not transmitting beacons after a pseudonym change [12] or the use of group cryptography [13] 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 and Güttler [14] have proposed to consider the context of a vehicle to determine when a pseudonym change can be effective [14], Freudiger et al. [15] presented their concept of mix-zones, that is, geographic areas for pseudonym changes [15]. The results presented in this paper 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 V-C, 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.

This paper contributes to finding these recommendations by first identifying the exact requirements and constraints and then presenting a proposal which we believe satisfies these requirements.

III. 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. [16] 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 [17]. 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 [8]. 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 [8], 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, 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 [18], 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 [19]. 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.

IV. Requirements for Privacy Protection in VANETs

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).

A. 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.

B. 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 V-B.

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 [12], it was shown that the effectiveness of traffic safety application is significantly reduced during these intervals [20].

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.

C. 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. 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.

D. 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 [10]. 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, if not fully secured, exchanging pseudonyms between cars [11] would empower one malicious user to collect pseudonyms and thereby compromise the entire pseudonym infrastructure. This user could then not be identified, potentially breaking system security.

V. 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 provides a basis for higher layer privacy protection as it secures the basic function of future VANETs, that is, cooperative awareness.

A. 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 [11], [21]. 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 fact that this is true for *all* vehicles outside of the adversary's reach is an important property of the privacy protection mechanism, as it maximizes the adversary's confusion [22].

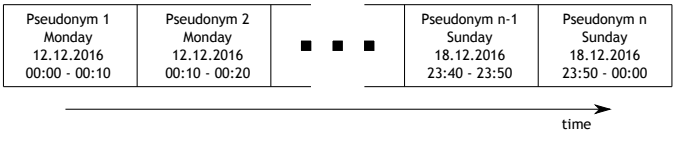

Figure 2. Time slotted pseudonym pool of 1 week length and 10 minute pseudonym validity.

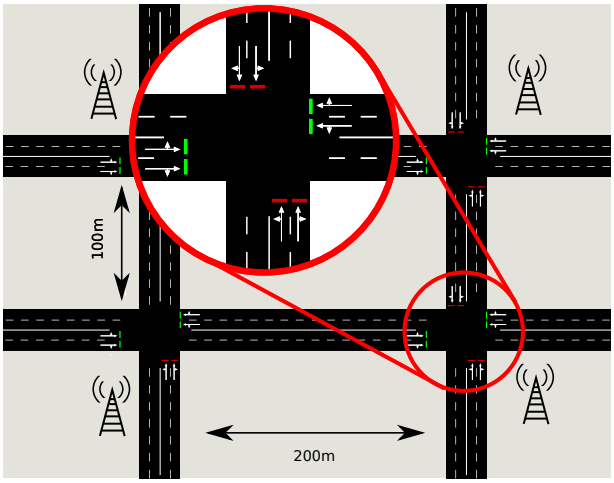

Figure 3. A grid scenario consisting of four intersection, road segments leading to the intersections are 400 meters long. An adversary has placed 4 antennas and is able to fully monitor the scenario.

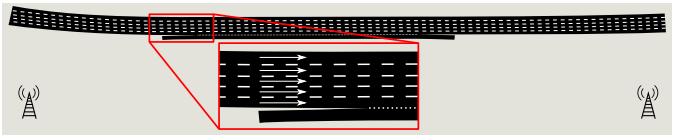

Figure 4. Highway 101 Scenario. Node movement is based on 900 s of real traffic on a 640 m stretch of Hollywood Fwy near Universal City Plaza, Los Angeles, CA. Five lanes are running in the same direction, temporarily joined by a sixth lane in the middle. Traffic was recorded by eight video cameras and post-processed to derive trace files [26]. An adversary is able to fully monitor the scenario.

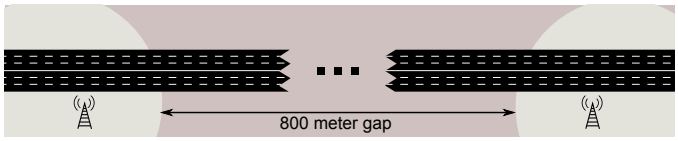

Figure 5. Freeway blind spot scenario: an adversary set up two access points on a highway, partially monitoring the scenario, unable to receive messages from vehicles in the 800 m blind spot.

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 V-C.

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.

B. 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 [23]). We further claim that this has almost no negative impact on privacy, as it is almost impossible to confuse eavesdropping attackers. In the following, we want to back up this strong claim by an extensive simulation study.

First indications of the difficulty to confuse local adversaries were given in [22] and [24]. The underlying scenarios, however, were simplistic and purely synthetic. In our simulation study, we make use of both real-world traces and synthetic scenarios. We implemented a modern multi-target tracking algorithm within our Veins simulation framework [25]. A more detailed description of this tracking framework can be found in [22].

In total, we investigated three different scenarios. A simulated combination of four intersections (Figure 3), real-world mobility traces recorded in the NGSIM project [26] 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. The realistic mobility trace 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. Throughout all scenarios, we measured the tracking fail rate, that is, the chance of a vehicle evading tracking. A higher fail rate means a higher level of privacy.

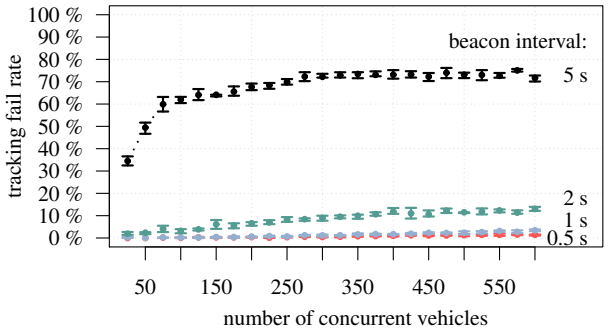

Figure 6. Chance of evading tracking in the synthetic grid scenario. Plotted are the averages over all simulation runs. Error bars extend from the 25 % to the 75 % quantiles.

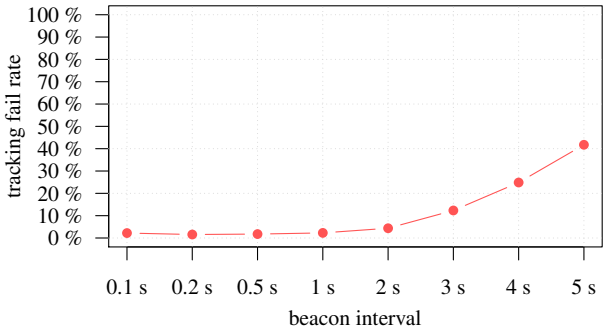

Figure 7. Chance of evading tracking in the (fully deterministic) real-world highway scenario.

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 6 shows our results. 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 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 [27] 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 [26]. 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.

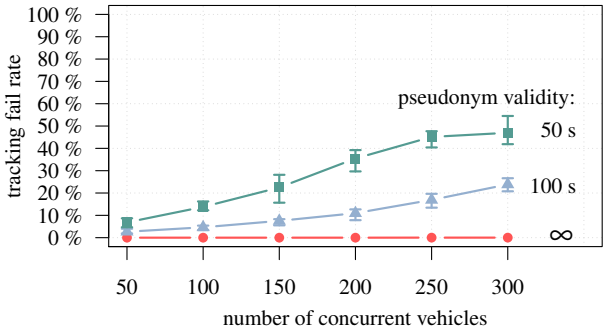

Figure 8. Chance of evading tracking in the blind-spot freeway scenario. Plotted are the averages over all simulation runs. Error bars extend from the 25 % to the 75 % quantiles.

Results are shown in Figure 7. 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.

This confirms 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 maximize 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 8 shows our results, comparing different pseudonym validities. As a ground truth, we show that without pseudonym changes, i.e., $t_p = \infty$, vehicles enjoy no privacy as tracking becomes trivial. With shorter times of pseudonym validity, the tracking fail rate considerably increased. 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 maximize attacker confusion.

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 [28]. 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. As soon as a vehicle could cryptographically prove that it owns both old and new pseudonym, receivers do not have to check both signatures anymore.

C. 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 to revoke a vehicle's pseudonyms include the intentional or unintentional transmission of false messages or a change in ownership [29]. This is a challenging process with regard to both efficiency and privacy. 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. Therefore, the published CRL must not include past pseudonyms that are no longer valid to preserve backward privacy of revoked vehicles.

We propose the use of linkage values as first introduced in [29], refined by [30], and later by [31]. 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 $e(\cdot)$ (see Figure 9).

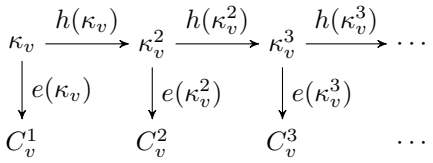

Figure 9. Graphical representation of approach based on linkage values.

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 (see Figure 9, top half). 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 $< j$ 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.

VI. Conclusion

In this paper, 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. 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 showed that, under reasonable assumptions about an adversary's capability, local privacy is neither required nor can it be achieved without compromising traffic safety. We base these conclusions also on the results of extensive simulations utilizing a realistic attacker model, applied to both artificial vehicle movement and real-world movement traces.

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 maximizes 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, and J. Wang, "Vehicle-to-Vehicle Communications: Readiness of V2V Technology for Application," National Highway Traffic Safety Administration, DOT HS 812 014, Aug. 2014.

[2] J. Gozalvez, M. Sepulcre, and R. Bauza, "IEEE 802.11p vehicle to infrastructure communications in urban environments," *IEEE Communications Magazine*, vol. 50, no. 5, pp. 176–183, May 2012.

[3] D. Eckhoff and C. Sommer, "Driving for Big Data? Privacy Concerns in Vehicular Networking," *IEEE Security & Privacy*, vol. 12, no. 1, pp. 77–79, Feb. 2014.

[4] J.-P. Hubaux, S. Čapkun, and J. Luo, "The Security and Privacy of Smart Vehicles," *IEEE Security & Privacy*, vol. 2, no. 3, pp. 49–55, May 2004.

[5] J. Petit, F. Schaub, M. Feiri, and F. Kargl, "Pseudonym Schemes in Vehicular Networks: A Survey," *IEEE Communications Surveys & Tutorials*, vol. 17, no. 1, pp. 228–255, Mar. 2015.

[6] H. Stübing, M. Bechler, D. Heussner, T. May, I. Radusch, H. Rechner, and P. Vogel, "simTD: A Car-to-X System Architecture for Field Operational Tests," *IEEE Communications Magazine*, vol. 48, no. 5, pp. 148–154, May 2010.

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

[8] I. Wagner and D. Eckhoff, "Privacy Assessment in Vehicular Networks Using Simulation," in *Winter Simulation Conference (WSC '14)*, Savannah, GA, Dec. 2014, pp. 3155 –3166.

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

[10] B. Xiao, B. Yu, and 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'06)*, ACM, Los Angeles, CA, USA, Sep. 2006, pp. 1–8.

[11] D. Eckhoff, C. Sommer, T. Gansen, R. German, and F. Dressler, "SlotSwap: Strong and Affordable Location Privacy in Intelligent Transportation Systems," *IEEE Communications Magazine*, vol. 49, no. 11, pp. 126–133, Nov. 2011.

[12] L. Huang, K. Matsuura, H. Yamane, and K. Sezaki, "Enhancing Wireless Location Privacy Using Silent Period," in *IEEE Wireless Communications and Networking Conference (WCNC 2005)*, New Orleans, LA, USA: IEEE, Mar. 2005, pp. 1187–1192.

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

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

[15] J. Freudiger, M. Raya, M. Félegyházi, P. Papadimitratos, and 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)*, Vancouver, Canada: ACM, Aug. 2007.

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

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

[18] P. Golle and K. Partridge, "On the Anonymity of Home/Work Location Pairs," in *7th International Conference on Pervasive Computing*, Nara, Japan: Springer, May 2009, pp. 390–397.

[19] I. Wagner and D. Eckhoff, "Technical Privacy Metrics: a Systematic Survey," arXiv, cs.CR 1512.00327, Dec. 2015.

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

[21] M. Raya, R. Shokri, and 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)*, Hoboken, NJ, USA: ACM, Mar. 2010, pp. 75–80.

[22] D. Eckhoff, M. Protsenko, and R. German, "Towards an Open Source Location Privacy Evaluation Framework for Vehicular Networks," in *80th IEEE Vehicular Technology Conference Fall (VTC 2014-Fall)*, Vancouver, Canada: IEEE, Sep. 2014.

[23] M. Feiri, J. Petit, and F. Kargl, "The Case for Announcing Pseudonym Changes," in *3rd GI/ITG KuVS Fachgespräch Inter-Vehicle Communication (FG-IVC 2015)*, Ulm, Germany, Mar. 2015.

[24] B. Wiedersheim, Z. Ma, F. Kargl, and 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: IEEE, Feb. 2010, pp. 176–183.

[25] C. Sommer, R. German, and F. Dressler, "Bidirectionally Coupled Network and Road Traffic Simulation for Improved IVC Analysis," *IEEE Transactions on Mobile Computing*, vol. 10, no. 1, pp. 3–15, Jan. 2011.

[26] Federal Highway Administration (FHWA), *NGSIM Program US Route 101 data*, From http://www.its-rde.net/, Version 1, 2016.

[27] N. An, M. Maile, D. Jiang, J. Mittag, and 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)*, Banff, Canada: IEEE, Mar. 2013, pp. 191–195.

[28] B. Bloessl, C. Sommer, F. Dressler, and 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*, Anaheim, CA: IEEE, Feb. 2015, pp. 395–400.

[29] J. J. Haas, Y.-C. Hu, and K. P. Laberteaux, "Efficient Certificate Revocation List Organization and Distribution," *IEEE Journal on Selected Areas in Communications*, vol. 29, no. 3, pp. 595–604, Mar. 2011.

[30] D. Eckhoff, F. Dressler, and C. Sommer, "SmartRevoc: An Efficient and Privacy Preserving Revocation System Using Parked Vehicles," in *38th IEEE Conference on Local Computer Networks (LCN 2013)*, Sydney, Australia: IEEE, Oct. 2013, pp. 855–862.

[31] W. Whyte, A. Weimerskirch, V. Kumar, and T. Hehn, "A Security Credential Management System for V2V Communications," in *5th IEEE Vehicular Networking Conference (VNC 2013)*, Boston, MA: IEEE, Dec. 2013, pp. 1–8.