Hide Me: Enabling Location Privacy in Heterogeneous Vehicular Networks

Tobias Meuser[1,2], Oluwasegun Taiwo Ojo[2], Daniel Bischoff[1], Antonio Fernández Anta[2(✉)], Ioannis Stavrakakis[3], and Ralf Steinmetz[1]

[1] Multimedia Communications Lab (KOM), Technische Universität Darmstadt, Darmstadt, Germany
{tobias.meuser,daniel.bischoff,ralf.steinmetz}@KOM.tu-darmstadt.de

[2] IMDEA Networks Institute, Madrid, Spain
{oluwasegun.ojo,antonio.fernandez}@imdea.org

[3] National and Kapodistrian University of Athens, Athens, Greece
ioannis@di.uoa.gr

Abstract. In order to support location-based services, vehicles share their location with a server to receive relevant data. Revealing a vehicle's location compromises its privacy. One way to reduce this problem is obfuscating the vehicle's location by adding artificial noise. However, this increases the area where the true location of the vehicle may be. Hence, under limited available bandwidth, the server will provide fewer data relevant to the vehicle's true location, reducing the effectiveness of the location-based service. To compensate for this reduction, we allow that the data relevant to a vehicle is also shared through direct, ad hoc communication between neighboring vehicles. Through such Vehicle-to-Vehicle (V2V) cooperation, the impact of location obfuscation is mitigated. In this set up, and assuming that the data served may have different impact levels, we propose and study a game that determines the data subscription a vehicle should use, without explicit coordination among them. The aim is maximizing the expected impact of the data received, either directly from the server or via V2V. Our analysis and results show that the proposed V2V cooperation and derived strategy lead to significant performance increase compared to other uncoordinated approaches, and largely alleviates the impact of location obfuscation.

Keywords: Floating Car Data · Location-based services · Location privacy · V2V communication

1 Introduction

The vehicles of the future will be required to have increased awareness of their environment, in order to assist the driver or to support autonomous driving.

This work has been supported by the DFG within the CRC 1053 - MAKI (B1), the Spanish grant PID2019-109805RB-I00 (ECID), the Region of Madrid EdgeData-CM program (P2018/TCS-4499), and the NSF of China grant 61520106005.

C. Georgiou and R. Majumdar (Eds.): NETYS 2020, LNCS 12129, pp. 11–27, 2021.
https://doi.org/10.1007/978-3-030-67087-0_2

This awareness has typically been provided by sensors on the vehicles, which measure vital data about the environment of the vehicle. The data provided by these sensors is limited to the vehicle's immediate environment, due to the sensors' inherent physical limitations (e.g., their range). Nevertheless, information from locations away from a given vehicle may also be important to it (e.g., for traffic safety, route planning, or navigation). To make such information available to far away vehicles, passing vehicles may capture it through their own sensors, and communicate it to a server using an appropriate communication infrastructure such as a cellular network. Then, the vehicles desiring to receive such information indicate so to the server, and receive it via a similar infrastructure. By sharing their local perception of the environment via a cellular infrastructure, as described, vehicles can complement their local perception with distant data provided by other vehicles.

In order for vehicles to get this, so-called, Floating Car Data (FCD), they have to share their location with the server, which is usually assumed to be a trusted entity. The server selects the relevant FCD for the vehicles using their location, and distributes it accordingly. This continuous context and location exchange with a server is a risk to the privacy of the vehicles. Consequently, privacy-sensitive users either have to accept this risk, or turn off the option of receiving FCD. Clearly, users that disable the reception of FCD cannot benefit from location-based services and other services enabled by vehicular networks. It is therefore desirable to have a mechanism that allows the reception of FCD while preserving the location privacy of the user.

A technique that is often used to increase the privacy of a vehicle, is adding random noise to its true location (obfuscation). Hence, instead of providing the server with a position, the vehicle provides an area. (We will assume in the rest of the paper that this type of obfuscation to increase privacy is used.) A negative consequence of obfuscation is that the server cannot use the true location to deliver its best FCD to a vehicle, and may hence send it useless data. We assume that different data items may have different value (*impact level*) for a vehicle. A vehicle subscribes to some impact level, and the server provides to the vehicle all available data items with matching impact for the vehicle area. Since the available bandwidth is limited, vehicles using obfuscation end up receiving a smaller portion of data of a given impact that is useful to them. As a result, location-based services would be less effectively provided to privacy-concerned vehicles.

To alleviate this problem, and to increase the amount of location-relevant data provided to a vehicle, we propose that neighboring vehicles can exchange data through direct, ad hoc communication. That is, we assume Vehicle-to-Vehicle (V2V) cooperation for exchanging local relevant data. We assume that vehicles do not use location obfuscation with other neighboring vehicles, only with the server, and hence the messages exchanged via V2V are all relevant and useful. (Trying to hide a vehicle's location to a neighbor seems pointless, since the neighbor can "see" the vehicle with its local sensors.) Hence, through V2V cooperation, the negative impact of location obfuscation could be mitigated to some extent. The use of V2V cooperation has also been considered in [1] combined with vehicle clusters. As it will be discussed later and shown in the results, cluster-based approaches are complex and suffer from connectivity problems, which reduces their performance. For these reasons, in this work the V2V communication is not cluster-based but ad hoc through direct V2V exchanges.

Notice that, without any coordination, neighboring vehicles are expected to subscribe to the same high impact levels, which results in receiving overlapping sets of data. This reduces the potential benefit of V2V cooperation. To prevent that, we develop and study a game among the vehicles. This game drives vehicles to subscribe to certain impact levels, so that the aforementioned overlap is reduced. The design goal is to maximize the expected value of a utilization function as shaped by the participating (neighboring) vehicles as well. Our analysis and results show that the proposed V2V cooperation scheme and derived strategy lead to significant performance increase compared to non-cooperative approaches, while alleviating the impact on privacy of location-based services.

Related Work. Several techniques have been introduced in the literature to protect users' privacy in vehicular networks. Some of the common techniques include the use of pseudonyms [2–4], obfuscation [5,6], and the use of group communications [7–9]. The first technique involves users taking on other identities (pseudonyms) to dissociate their actual identity from their data [10]. The use of a single pseudonym is not very effective, and hence it is often required for users to change pseudonyms periodically, to maintain their level of privacy [11]. Such pseudonym changes are usually done in mix zones where drivers can switch pseudonyms [12]. These mix zones can be fixed [13] or specified dynamically [14]. However, the use of pseudonyms has been shown not to be effective against a global eavesdropper [15], and especially in environments with low car density like highways. Furthermore, the use of pseudonyms usually focuses on eavesdroppers monitoring V2V communications and involves having to deal with a trusted (or semi-trusted) server which coordinates the assignments of pseudonyms [8]. This still involves trusting a central server, which is a risk in the case that an adversary gets hold of such server. Our work focuses on the privacy of users in their communications with the central server.

Likewise, obfuscation has been extensively used in privacy protection in vehicular networks and location-based services. Obfuscation involves users providing (i) an inaccurate location, (ii) an imprecise region including their real location, or (iii) a vague description of their location [16]. To quantify the effectiveness of obfuscation, metrics like k-*anonimity*, which means that a user's shared location data makes it indistinguishable from $k-1$ other users, have been introduced [17,18]. The imprecision added into the location of the user usually leads to users getting less relevant data and, thus, a decrease in efficiency. Our method mitigates against this decrease in performance by implicitly cooperating with other vehicles to get relevant updates through V2V communication.

Game-theory has been applied to modeling aspects of privacy, especially in mobile networks and location-based services [19,20], and in security and privacy assessment of vehicular networks [21]. Distinct from previous studies, our work focuses on privacy of users in their communications with the server considering the impact of the messages to the user. We adopt an obfuscation technique by reporting a region instead of their exact location, and mitigate against the resulting reduction in performance by implicitly coordination the vehicles through a game-theoretic approach, which maximizes the relevant data received by the vehicles.

Contributions. The contributions of this work are the following. First, we introduce privacy considerations in the management of FCD and reveal their impact on location-based services: given a fixed bandwidth availability, some data may not be forwarded to a vehicle due to location obfuscation. Second, in order to alleviate this problem, we propose that vehicles cooperate and forward relevant data to their neighboring vehicles, increasing in principle the data received by a vehicle beyond what is directly received from the server. An ad-hoc, direct V2V cooperation paradigm is employed instead of a cluster-based one, and we show the high performance deterioration of the latter in a real vehicular networking environment. Third, we develop and study of a game determining the strategies (in terms of probabilities that a vehicle is forwarded by the server data of a given impact level) that vehicles should follow, so that the expected utility is maximized. This is shown to lead to a diversification of the data received directly from the server by neighboring vehicles, and increases the effectiveness of V2V cooperation. Finally, the aforementioned contributions are supported through simulation evaluation.

Structure. The rest of the paper is as follows. In Sect. 2, we provide an overview of the system model considered, and describe the influence of location privacy on the network. In Sect. 3, we describe our proposed game theoretic approach for privacy sensitive communication. In Sect. 4, we evaluate the performance of our method. We conclude the paper in Sect. 5 with a discussion about our findings.

2 System Model

Definitions. We provide first an overview of the considered system model. We assume a context-aware vehicular network, in which a central server transmits context-sensitive messages to interested vehicles. In this network, time is assumed to be slotted (a typical slot length is 1 s). Every vehicle has a limited (average) bandwidth A (in bits per time slot) to receive these messages via a cellular network. This assigned bandwidth is generally low compared to the maximum (physically) available bandwidth, such that vehicles may exceed this bandwidth temporarily (as long as the average consumed bandwidth matches the predefined value). A message contains FCD as payload, as well as additional meta-information such as the source location, generation time, and type of FCD. In this work, we assume that FCD carry road-related information (e.g., accidents, traffic jams, traffic flow information) that can be useful for improving the driving behavior of the vehicles in proximity. Let $a(m)$ (in bits) denote the size of a message m, $s(m)$ the source location, $r(m)$ the radius of its dissemination area, and $\mu(m)$ its impact (which depends on the type of FCD: an accident has generally higher impact than traffic flow information). As our bandwidth is limited, the impact per utilized bandwidth is pivotal for our approach. Based on the message impact $\mu(m)$ and the dissemination radius $r(m)$, we divide messages in n_μ impact levels. For simplicity, we assume that every message m of impact level $i \in \{1, \ldots, n_\mu\}$ has the same dissemination radius $r(m) = r_i$ and impact $\mu(m) = \mu_i$. We assume that μ_i is the impact per bit assigned to impact level i. When convenient, we use $\mu_{n_\mu+1} = \infty$.

A vehicle can control the reception of messages from the server by expressing interest in certain *impact* levels and by providing a *representation of its location.* More specifically, a vehicle wants to receive a message m if (i) it has expressed interest in the corresponding impact level i of the message, and (ii) the vehicle's location is at most at distance r_i from the source $s(m)$ of the FCD. Let a_i denote the traffic load of messages of impact level i (in bits per time slot) expected for the vehicle if the provided location is accurate. A vehicle is either interested in an impact level or not, i.e., receives either all or no messages of this impact level. This interest can be changed dynamically at the beginning of every time slot.

Depending on the assumed privacy-sensitivity (referred to as privacy-level) $\phi \in \Phi$ of a vehicle v, the aforementioned *representation of the location* may be accurate or may be imprecise. We implement this imprecision by providing only a (circular) area in which the vehicle is certainly located (uniformly distributed), without actually revealing the exact location to the server. The privacy level ϕ chosen by the respective vehicle determines the radius r_ϕ of this area. That imprecise representation of the location increases the load of received messages due to the less accurate server-side filtering. To capture the additional bandwidth consumption, let $a_{\phi,i} \geq a_i$ denote the expected load (in bits) of messages of impact level i for a vehicle with privacy level ϕ.

The central server uses the announced interest of the vehicles to actively push new messages (i.e., messages containing yet unknown FCD.) via the cellular network to them. Since the available bandwidth is assumed to be limited, a vehicle aims to maximize the total impact of the received messages, which is achieved by dropping low-impact messages if the bandwidth is insufficient. To maximize that total impact of received messages, vehicles may cooperate to share bandwidth for the reception of messages; i.e., vehicles can locally broadcast messages, received via the cellular network, without additional costs to notify vehicles in their proximity. Thus, not every vehicle needs to receive all messages of its interest via the limited cellular bandwidth, as these messages might be provided by its neighbors.

Influence of Location Obfuscation. In the following, we provide an insight on the influence of privacy in our model. Each privacy level $\phi > 1$ adds a certain level of imprecision to the provided location, while $\phi = 1$ refers to no privacy-sensitivity. The privacy-sensitivity and, thus, location imprecision increases with ϕ and reduces the accuracy of the context-based message filtering at the server-side. Thus, a vehicle receives messages not relevant for its current context, while its share of relevant messages is reduced. This influences the number of received messages $n_{\phi,i}$ and their expected impact per bit $\mu_{\phi,i}$ for a privacy state ϕ and an impact level i. The number of messages received typically increases with increasing privacy level, while the expected impact per bit of a message decreases. We reflect this change for every impact level i by the *adaptation factor* $\rho_{\phi,i}$ as follows.

$$a_{\phi,i} = a_i \cdot \rho_{\phi,i} \quad \mu_{\phi,i} = \frac{\mu_i}{\rho_{\phi,i}} \tag{1}$$

$\rho_{\phi,i}$ depends on the context-sensitivity of the distributed messages for a vehicle of privacy level ϕ receiving messages with impact level i. For non-context-sensitive messages, $\rho_{\phi,i} = 1, \forall \phi \in \Phi$. For context-sensitive of messages, i.e., messages with a specific distribution-area with radius r_i, $\rho_{\phi,i} \geq 1, \forall \phi \in \Phi$. These statements are proven in Theorem 1.

Theorem 1. *The adaptation factor for a network with uniformly distributed messages is $\rho_{\phi,i} = \left({}^{r_\phi}/_{r_i} + 1\right)^2$ for a circular geocast-area and a circular location-imprecision, where r_i is the radius of the geocast-area of the message of impact level i and r_ϕ is the radius of the location-imprecision area of privacy-level ϕ.*

Proof. Without location privacy, the vehicle receives all messages with a maximum distance of r_i to its current location. Thus, area of interest for the vehicle is $\pi \cdot r_i^2$. If the vehicle reduces the precision of its location by hiding inside an area of radius r_ϕ, the server will need to transmit all messages within a distance of $r_\phi + r_i$ from the center of the area to ensure that the vehicle receives all relevant messages. The size of this area is $\pi \cdot (r_\phi + r_i)^2$. This leads to $\rho_{\phi,i} = \left({}^{r_\phi}/_{r_i} + 1\right)^2$.

3 Game-Theoretic Model for Privacy-Sensitive Communication

To enhance the performance of our impact-aware vehicular network, we employ a game-theoretic model with the aim to maximize the sum of impact of the received messages. Our innovative approach relies only on the number n_ϕ of vehicles of each privacy-level ϕ in proximity to find a mixed Nash-optimal solution for our developed game-theoretic model, i.e., vehicles receive messages with a certain probability. In our game, each actor (vehicle) aims to find the strategy (receive messages in a certain impact-range via the cellular network) that maximizes its utility (sum of impact values of all received messages, directly via cellular or from the neighbors) while sticking to cellular bandwidth constraints. This game is played periodically in every time slot to adjust the vehicles behavior to environmental changes, i.e. changes in the number of neighbors in proximity and changes in number of messages. Notice that vehicles are assumed to cooperate; thus, a vehicle might additionally receive messages directly by vehicles in proximity. The intuition behind this game model is that high-impact messages are generally prioritized, as their bandwidth usage is more efficient compared to low-impact messages. Thus, vehicles may rely on their neighbors to provide some high-impact messages to them, as a number of neighbors aims to receive these high-impact messages. These vehicles can then use a part of their available cellular bandwidth to receive low-impact messages and share these with their neighbors. The idea is similar to cooperative caching: Instead of storing all high-demand message at every local cache, some nodes fetch low-demand messages instead and satisfy the request of high-demand messages from nearby cooperative caches [22].

The vehicles are the only *actors* in this game; the server is not directly involved, but only determines the set of receivers of messages based on the strategies chosen by the vehicles. For this purpose, the vehicles share their strategy in the form of subscriptions with the server. The *strategy* is represented as a vector $\boldsymbol{p_\phi}$ with n_μ probability entries $p_{\phi,i}$ with $i \in \{1, \ldots, n_\mu\}$, and depends on the chosen privacy level ϕ_e of the vehicle. Each entry $p_{\phi,i}$ refers to the probability of the tagged vehicles to receive messages of the corresponding impact level. Additionally, $0 \leq p_{\phi,i} \leq 1, \forall p_{\phi,i} \in \boldsymbol{p_\phi}$. For the assignment of messages to an impact level, we use the impact μ_i. Note that μ_i does not depend on the privacy level ϕ. The privacy-dependent message impact $\mu_{\phi,i}$ is only used for the calculation of the utility of a vehicle. In the calculation, $\boldsymbol{p_\phi}$ needs to be chosen such that Eq. 2 holds, with $a_{\phi,i}$ being the expected number of bits in the received messages of impact level i and privacy level ϕ according to Eq. 1, and A being the usable bandwidth.

$$\sum_{i=1}^{n_\mu} a_{\phi,i} \cdot p_{\phi,i} \leq A \tag{2}$$

Notice that this differs from previous work, like [1], in which the vehicle is intended to receive all messages in the set $\{m | \mu_i \leq \mu(m)\}$. The advantage of our new model is the additional flexibility provided by removing some of the message redundancy among neighboring vehicles, which improves the total impact of received messages (via cellular and direct neighbor forwarding) by each vehicle.

Each vehicle aims at maximizing its *utility*, which is defined in a way that captures the impact of the messages received. The utility used in this paper is defined in Eq. 3, and is based on the messages sent M_{snt}, the messages received M_{rcv}, and the impact $\mu(m)$ of every message m. $\mathbb{I}_{\{m \in M_{rcv}\}}$ is the indicator function of whether a message m has been received by the vehicle.

$$u = \sum_{m \in M_{snt}} \mu(m) \cdot a(m) \cdot \mathbb{I}_{\{m \in M_{rcv}\}} \tag{3}$$

As the probability of a vehicle receiving a message depends on $\boldsymbol{p_\phi}$, we derive the expectation of the utility based on Eq. 3. For this purpose, we assume that the environment of each vehicle is similar, so that the strategies of two vehicles with the same privacy level are the same. Thus, the strategy of every privacy level can be calculated by every vehicle in proximity, which is the basis of our offloading approach. Thus, we only use the strategies $\boldsymbol{p_\phi}$ along with the number n_ϕ of vehicles for each privacy level ϕ to calculate the probability of receiving a message either via the cellular network or from one of the neighbors. The probability $p(\mu_i)$ to receive a message via any interface (cellular or V2V) with at impact level μ_i can be calculated as shown in Eq. 4. This formula assumes that there is no loss in the network, i.e., every transmitted messages is received by the intended receiver.

$$p(\mu_i) = 1 - \prod_{\phi \in \Phi} (1 - p_{\phi,i})^{n_\phi} \tag{4}$$

We use the probability $p(\mu_i)$ to receive a message to derive the expected utility $\overline{u}(\phi_e, \boldsymbol{p_1}, \ldots, \boldsymbol{p_{|\Phi|}})$. This estimates the set of received messages M_{rcv} using the expected amount of sent messages a_i and the probability $p(\mu_i)$ to receive each message. The resulting expected utility for the tagged vehicle is shown in Eq. 5.

$$\overline{u}(\phi_e, \boldsymbol{p_1}, \ldots, \boldsymbol{p_{|\Phi|}}) = \sum_{i=1}^{n_\mu} \mu_{\phi_e,i} \cdot a_{\phi_e,i} \cdot \left[1 - \prod_{\phi \in \Phi} (1 - p_{\phi,i})^{n_\phi} \right] \tag{5}$$

When clear from context, we refer to $\overline{u}(\phi_e, \boldsymbol{p_0}, \ldots, \boldsymbol{p_{|\Phi|}})$ as $\overline{u}$ to increase readability. In the next section, we describe the process of deriving a utility-maximizing strategy for the described game. The advantage of determining the solution analytically is (i) the possibility to analyze and bound the effects of location privacy to the system, and (ii) the lower computational complexity compared to a non-linear solver.

3.1 Game-Theoretic Solution

We derive now the optimal strategy for a vehicle with privacy level ϕ_e, given that the privacy level and number of vehicles in each privacy level in its environment is known. For this purpose, we calculate the partial derivatives of the expected utility $\overline{u}$ with respect to the probabilities of the tagged vehicle $p_{\phi,i}$. However, it is important to consider the dependency between the probabilities $p_{\phi,i}, \forall \phi \in \Phi$, as Eq. 2 limits the possible values of $p_{\phi,i}$. (This approach would work similarly with any other probability $p_{\phi,i} | i \neq 1$.) We depict this dependency by expressing $p_{\phi,1}$ depending on the other probabilities $\{p_{\phi,i} | i > 1\}$ as shown in Eq. 6. Thus, $p_{\phi,1}$ depends on all other probabilities, i.e., the derivative of $p_{\phi,1}$ with respect to any probability $p_{\phi,i}$ is not always non-zero, which leads to our optimization problem.

$$p_{\phi,1} \leq \frac{A - \sum_{i=2}^{n_\mu} a_{\phi,i} \cdot p_{\phi,i}}{a_{\phi,1}} \tag{6}$$

While the inequality is sufficient to guarantee the bandwidth requirements, we will assume Eq. 6 to be an equation as higher values of $p_{\phi,1}$ cannot decrease the utility. As there is no dependency between any pair of probabilities $p_{\phi,i}$ and $p_{\phi,j}$ if $i \neq j \wedge i \neq 1 \wedge j \neq 1$, the derivative of the utility with respect to $p_{\phi,l}$ depends only on $p_{\phi,1}$ and $p_{\phi,l}$ for every $l > 1$ as shown in Eq. 7. Notice that $\mu_{\phi_e,i} \cdot a_{\phi_e,i} = \mu_i \cdot a_i$ according to Eq. 1. Additionally, we assume that $p_{\phi_e,l} \neq 0$.

We ensure that by considering the cases with $p_{\phi_e,l} = 0, \forall l \in \{1, \ldots, n_\mu\}$ separately as described in Sect. 3.2.

$$\frac{\partial \overline{u}}{\partial p_{\phi_e,l}} = \mu_l a_l n_{\phi_e} \cdot (1 - p_{\phi_e,l})^{n_{\phi_e}-1} \cdot P_l(\varPhi \setminus \{\phi_e\}) \\ + \mu_1 a_1 \left(\frac{\partial p_{\phi_e,1}}{\partial p_{\phi_e,l}}\right) n_{\phi_e} \cdot (1 - p_{\phi_e,1})^{n_{\phi_e}-1} \cdot P_1(\varPhi \setminus \{\phi_e\}) \quad (7)$$

with $P_j(\varPhi) = \prod_{\phi \in \varPhi} (1 - p_{\phi,j})^{n_\phi}$. Equation 1 displays the dependency of $p_{\phi_e,1}$ and $p_{\phi_e,l}$. Thus, the derivative of $p_{\phi_e,1}$ with respect to $p_{\phi_e,l}$ can be calculated according to Eq. 8.

$$\frac{\partial p_{\phi_e,1}}{\partial p_{\phi_e,l}} = -\frac{a_{\phi_e,l}}{a_{\phi_e,1}} \quad (8)$$

By setting the derivative of the utility to 0, we determine all possibly optimal solutions. This leads to Eq. 9 after some minor transformations. Notice that a_l and n_{ϕ_e}, are omitted as they are present on both sides of the equation.

$$\frac{\mu_l}{\mu_1} \cdot \left(\frac{1 - p_{\phi_e,l}}{1 - p_{\phi_e,1}}\right)^{n_{\phi_e}-1} \cdot P_l(\varPhi \setminus \{\phi_e\}) = \frac{\rho_{\phi_e,l}}{\rho_{\phi_e,1}} \cdot P_1(\varPhi \setminus \{\phi_e\}) \quad (9)$$

For a given impact level l, we divide the set of privacy levels $\varPhi$ into $\varPhi^+(l)$, which only contains privacy levels with $p_{\phi,l} > 0$, and $\varPhi^-(l)$, which contains privacy levels with $p_{\phi,l} = 0$. This is necessary, as the derivative of the expected utility with respect to $p_{\phi,l}$ is always 0 if $p_{\phi,l} = 0$, thus, Eq. 9 does not hold. However, Eq. 9 still contains $p_{\phi,l}, \forall \phi \in \varPhi^+(l)$ and $p_{\phi,1}, \forall \phi \in \varPhi(l)$. We need to replace $p_{\phi,l}, \forall \phi \in \varPhi^+(l) \setminus \phi_e$ to calculate $p_{\phi_e,l}$. We can calculate the $p_{\phi_e,l}$ using Eq. 10, according to Theorem 2.

Theorem 2. *For any probability $p_{\phi_e,l}, \forall \phi_e \in \varPhi^+(l)$ with $n_{\phi_e} > 1$, we have that*

$$\frac{\mu_l}{\mu_1} \cdot \prod_{\phi \in \varPhi^+(l) \setminus \{\phi_e\}} \left(\frac{\rho_{\phi_e,l} \cdot \rho_{\phi,1}}{\rho_{\phi,l} \cdot \rho_{\phi_e,1}}\right)^{n_\phi} \cdot \left(\frac{1 - p_{\phi_e,l}}{1 - p_{\phi_e,1}}\right)^{n^+(l)} = \left(\frac{\rho_{\phi_e,l}}{\rho_{\phi_e,1}}\right) \cdot P_1(\varPhi^-(l)) \quad (10)$$

where $n^+(l) = \sum_{\phi \in \varPhi^+(l)} n_\phi - 1$. Hence, $p_{\phi_e,l}$ depends only on $p_{\phi_e,1}$ and previously calculated probabilities.

Proof. We use full induction to prove the correctness of Eq. 10. For the base-case, we consider $\varPhi = \{\phi_e\}$. Based on Eq. 9, we observe that $P_1(\varPhi \setminus \phi_e) = 1$ and $P_l(\varPhi \setminus \phi_e) = 1$, as $\varPhi$ contains only ϕ_e. Additionally, $n^+(l) = n_{\phi_e} - 1$ for the same reason, which immediately leads to Eq. 10. For the induction step, we use $\varPhi^+_+(l) \subseteq \varPhi^+$ and $\varPhi^+_-(l) \subseteq \varPhi^+$ as auxiliary variables with $\phi \in \varPhi^+_+(l) \oplus \varPhi^+_-(l), \forall \phi \in \varPhi^+(l)$, for which the index states if they have already been included in the calculation. Based on Eq. 9 and Eq. 10, we can derive Eq. 11 associated $\phi_e \in \varPhi^+_-(l)$ as intermediate state of the calculation. Notice that $\phi_e \in \varPhi^+_+$ by assumption. Additionally, the privacy levels in $\mathit{\Phi}^-$ are not considered on the left side of the equation, as $p_{\phi,l} = 0, \forall \phi \in \mathit{\Phi}^-$.

$$\begin{aligned}\frac{\mu_l}{\mu_1} \cdot \prod_{\phi \in \mathit{\Phi}^+_+(l) \setminus \{\phi_e\}} \left(\frac{\rho_{\phi_e,l}}{\rho_{\phi,l}}\right)^{n_\phi} \cdot \left(\frac{1-p_{\phi_e,l}}{1-p_{\phi_e,1}}\right)^{n^+_+(l)} \cdot P_l(\mathit{\Phi}^+_- \setminus \{\phi_e\}) \\ = \prod_{\phi \in \mathit{\Phi}^+_+(1) \setminus \{\phi_e\}} \left(\frac{\rho_{\phi_e,1}}{\rho_{\phi,1}}\right)^{n_\phi} \cdot \frac{\rho_{\phi_e,l}}{\rho_{\phi_e,1}} \cdot P_1(\{\mathit{\Phi}^-(l) \cup \mathit{\Phi}^+_-(l)\})\end{aligned} \tag{11}$$

with $n^+_+(l) = \sum_{\phi \in \mathit{\Phi}^+_+(l)} n_\phi - 1$.

We aim to include a privacy level ϕ_n into $\mathit{\Phi}^+_+$. Thus, we solve Eq. 11 associated with ϕ_n for $p_{\phi_n,l}$ and insert it into Eq. 11 associated with all other $\phi_e \in \mathit{\Phi}^+_-(l) \setminus \phi_n$ to obtain Eq. 12.

$$\begin{aligned}\frac{\mu_l}{\mu_1} \cdot \prod_{\phi \in (\mathit{\Phi}^+_+(l) \cup \phi_n) \setminus \{\phi_e\}} \left(\frac{\rho_{\phi_e,l}}{\rho_{\phi,l}}\right)^{n_\phi} \cdot \left(\frac{1-p_{\phi_e,l}}{1-p_{\phi_e,1}}\right)^{n^+_+(l)+n_{\phi_n}} \cdot P_l(\mathit{\Phi}^+_- \setminus \{\phi_n\}) \\ = \prod_{\phi \in (\mathit{\Phi}^+_+(l) \cup \phi_n) \setminus \{\phi_e\}} \left(\frac{\rho_{\phi_e,1}}{\rho_{\phi,1}}\right)^{n_\phi} \cdot \frac{\rho_{\phi_e l}}{\rho_{\phi_e 1}} \cdot P_1(\{\mathit{\Phi}^-(l) \cup \mathit{\Phi}^+_-(l)\} \setminus \phi_n)\end{aligned} \tag{12}$$

This equation is similar to our initial Eq. 11 if we set $\mathit{\Phi}^+_+ = \mathit{\Phi}^+_+ \cup \phi_n$ and $\mathit{\Phi}^+_- = \mathit{\Phi}^+_- \setminus \phi_n$. Additionally, it is evident that Eq. 12 is equal to Eq. 10 if $\mathit{\Phi}^+_+ = \mathit{\Phi}^+$ and $\mathit{\Phi}^+_- = \emptyset$. □

Equation 10 still contains $p_{\phi_e,1}$ as an auxiliary variable. When replacing $p_{\phi_e,1}$ according to its definition in Eq. 6, we can derive the remaining variables $p_{\phi_e,i}, \forall i > 1$ only based on the other variables $p_{\phi_e,i}, \forall i > 1$. For that purpose, we introduce the variable $\mathit{\Lambda}_l$ with $1 < l \leq n_\mu$ as defined in Eq. 14, which encapsulates the constant values and the dependency on other privacy levels ϕ for readability. Thus, we can transform Eq. 10 to Eq. 13 by taking the $n^+(l)$-th root and replacing $p_{\phi_e,1}$.

$$1 - p_{\phi_e,l} = \left[1 - \left(\frac{A}{a_{\phi_e,1}} - \sum_{i=2}^{n_\mu} \frac{a_{\phi_e,i} \cdot p_{\phi_e,i}}{a_{\phi_e,1}}\right)\right] \cdot \mathit{\Lambda}_l \tag{13}$$

with

$$\mathit{\Lambda}_i = \sqrt[n^+(i)]{\left(\frac{\mu_1}{\mu_i}\right) \cdot \left(\frac{\rho_{\phi_e,i}}{\rho_{\phi_e,1}}\right) \cdot \prod_{\phi \in \mathit{\Phi}^+(l) \setminus \{\phi_e\}} \left(\frac{\rho_{\phi,i} \cdot \rho_{\phi_e,1}}{\rho_{\phi,1} \cdot \rho_{\phi_e,i}}\right)^{n_\phi} \cdot \prod_{\phi \in \mathit{\Phi}^-(i)} (1 - p_{\phi,1})^{n_\phi}} \tag{14}$$

The equation system described by Eq. 13 for all $2 \leq l \leq n_\mu$ cannot be solved without considering the dependency on the other privacy levels encapsulated in $\mathit{\Lambda}_l$. However, this dependency is hard to resolve except for some special cases, as it removes the linearity from Eq. 13. Thus, we assume that $\mathit{\Lambda}_l$ is constant for

Algorithm 1: Determining the optimal strategy for all privacy-levels. **recal(...)** recalculates $p_{\phi_e,i}$ based on the current values of $p_{\phi,i}$. ϵ is the infinitesimal.

```
Result: $p_{\phi,i}, \forall \phi \in \Phi, i \in \{1, \ldots, n_\mu\}$
$p_{\phi,i} \leftarrow 0, \forall \phi \in \Phi, i \in \{1, \ldots, n_\mu\}$;
$c \leftarrow \infty$;
for $i \leftarrow 1$; $c > \epsilon$; $i \leftarrow (i \bmod |\Phi|) + 1$ do
    $\mathrm{temp}_j \leftarrow p_{i,j}, \forall j \in \{1, \ldots, n_\mu\}$;
    $\mathrm{recal}(\mathrm{p_{i,j}}), \forall j \in \{1, \ldots, n_\mu\}$;
    $c \leftarrow \sum_{j=1}^{n_\mu} |\mathrm{temp}_j - p_{i,j}|$;
end
return $p_{\phi,i}, \forall \phi \in \Phi, i \in \{1, \ldots, n_\mu\}$;
```

the calculation of $p_{\phi_e,l}, \forall l \in \{2, \ldots, n_\mu\}$. Thus, we can represent $p_{\phi_e,j} \neq 0$ as $p_{\phi_e,i} \neq 0$ by subtracting the representation of $p_{\phi_e,i}$ from the representation of $p_{\phi_e,j}$ according to Eq. 13 and obtain Eq. 15.

$$p_{\phi_e,i} = \Lambda_i \left(\frac{p_{\phi_e,j} - 1}{\Lambda_j} \right) + 1 \tag{15}$$

With this assumption, we can calculate every $p_{\phi_l,l}$ with Eq. 16, which can be derived from Eq. 13 and the representation of any $p_{\phi_e,i}$ as $p_{\phi_e,j}$ from Eq. 15. Notice, that $\Lambda_1 = 1$, as either $p_{\phi,1} = 0$ (then $1 - p_{\phi,1} = 1$ and disappears), or $p_{\phi,1} \neq 0$ (then $\phi \notin \Phi^-(1)$).

$$p_{\phi_e,l} = \frac{\left[A - \sum_{i=1 | i \neq l \wedge \phi_e \notin \Phi^-(i)}^{n_\mu} a_{\phi_e,i} \right] \Lambda_l}{\sum_{i=1 | i \neq l \wedge \phi_e \notin \Phi^-(i)}^{n_\mu} (a_{\phi_e,i} \cdot \Lambda_i)} + 1 \tag{16}$$

Based on Eq. 16, we can determine the strategies for each privacy level using Algorithm 1. This algorithm ensures that the initial error (induced by setting all probabilities to 0) converges, i.e., the initial error constantly reduces for each iteration of Algorithm 1. This algorithm converges immediately if there is no inter-dependency between the privacy levels, i.e., if there is no other privacy level $\phi_o \mid p_{\phi_o,i} = 0$. If there is an inter-dependency, it converges due to three factors: (i) In the calculation of $p_{\phi,1}$, all probabilities $p_{\phi,i}$ with $i > 1$ are utilized, thus, $p_{\phi,1}$ balances the error of the other probabilities. (ii) $p_{\phi,1}$ influences Λ_i of all privacy levels in $\Phi^-(i)$, but we can see that Λ_l in the nominator and Λ_i in the denominator partially cancel out the error of each other in Eq. 16. (iii) $\exists l, \phi \mid n_\phi < n^+(l)$, in which case the error in Λ_l gets reduced based on the errors of the other privacy levels.

3.2 Deriving the Utility-Optimal Strategy

In the previous section, we assumed that every probability under consideration is non-zero. To calculate the overall optimal strategy, we consider every possible combination of zero and non-zero probabilities of every privacy level, i.e., we consider every possible combination of $\Phi^+(l)$ and $\Phi^-(l)$. That is, the computational complexity of our approach is $\mathcal{O}(2^{|\Phi| \cdot n_\mu})$, i.e., is exponential with the number of privacy levels $|\Phi|$ and the number of impact levels n_μ. This exponential growth is justified by the separate consideration of zero probabilities, which leads to 2 tries per probability. While an exponential growth is generally bad, we need to remember the limited size of $|\Phi|$ and n_μ. As every single computation of probabilities is very fast, the total computation time of the probabilities remains comparably small (in our experiments, it stayed around 100 ms). In the calculation, we set the probabilities of all $p_{\phi,l} = 0 \mid \phi \in \Phi^-(l)$ and only calculate the remaining probabilities with our approach proposed in the previous section. The solution found has certain properties.

Optimality. For each possible set of $\Phi^+(j), \forall j \in \{1, \ldots, n_\mu\}$, the partial derivatives of the utility with respect to all probabilities are 0, i.e., are either local optima or saddle points. To prove that the found solutions are global optima, we need to ensure that there is no other optimum with a higher utility than the found solution. For this purpose, we investigate on the second derivative of the utility function.

$$\frac{\partial^2 \overline{u}}{\partial^2 p_{\phi_e,l}} = -\overline{\mu}_l \cdot \Psi_l - \overline{\mu}_1 \cdot \frac{a_l}{a_0} \cdot \left(-\frac{\rho_{\phi_e,l}}{\rho_{\phi_e,1}}\right)^2 \cdot \Psi_1 \tag{17}$$

with $\Psi_j = a_l \cdot n_{\phi_e} \cdot (n_{\phi_e} - 1) \cdot (1 - p_{\phi_e,l})^{n_{\phi_e}-2} \cdot P_j(\phi \in \Phi \setminus \{\phi_e\})$.

As Ψ_i, $\overline{\mu}_i$, and a_i are non negative for all i, the second derivative of the utility with respect to any probability $p_{\phi_e,l}$ is always smaller or equal to 0. Thus, the expected utility presented in Eq. 5 is concave. This guarantees that the found solution maximizes the utility, but is not necessarily unique, i.e., there might be other solutions with similar utility.

Stability. The game solution found is a Nash equilibrium, as shown in the following theorem (the proof is omitted for space limitation).

Theorem 3. *The solution of our non-cooperative game shown in Eq. 16 is a Nash equilibrium, i.e., no vehicle has an incentive to deviate from the found solution.*

Observe that this equilibrium is only reached if every vehicle is aware that its neighbors follow the same strategy.

4 Evaluation

In this section, we evaluate the performance of our approach in a realistic vehicular network under varying environmental conditions. For this purpose, we utilize the vehicular extension of the Simonstrator framework [23] in conjunction with SUMO [24] to simulate a vehicular network in Cologne [25]. We compare our approach with state-of-the-art methods for cooperative communication in large-scale vehicular networks and non-cooperative uncoordinated approaches. In this large-scale vehicular network, messages are provided based on the current location of the vehicle (considering its privacy restrictions).

In our simulation, we generate messages randomly in an area of roughly $220 \times 220\,\mathrm{km}^2$, while the movement of vehicles and their networking is only simulated in an area of $2 \times 2\,\mathrm{km}^2$, to reduce the computational overhead. As all events with a possible influence to the network are simulated, we accurately model the message load in a large-scale vehicular network. Unless otherwise said, the bandwidth A is set to 10% of the total required bandwidth. We use messages of 4 impact levels $(1, 10, 100, 1000)$, with frequencies $(90\%, 9\%, 0.9\%, 0.1\%)$ and ranges (10 km, 1 km, 100 km, 100 km), respectively. The approaches that will be evaluated and compared are the following.

- *Game-Theoretic Privacy-Sensitive Cooperation (GTP).* This is our approach proposed in Sect. 3, which relies on implicit coordination between vehicles.
- *No Cooperation (NC).* The No-Cooperation (NC) approach does not consider cooperation between vehicles. Thus, vehicles using the *NC* approach receive similar messages as their neighbors, i.e., they do not share their messages.
- *Clustering with perfect failure detection (GK).* Clustering is used as follows. A vehicle is chosen as cluster-head, which is the only one communicating directly with the server. The cluster-head distributes the received messages to the vehicles in proximity via V2V communication. In the *GK* approach we assume that the disconnection of the cluster-head (moving out of range) is immediately detected. GK is used as an (unrealistic) upper bound for the performance of our approach.
- *Clustering without perfect failure detection (CL). CL* is similar to *GK*, with the exception that the detection of cluster-head disconnections is now imperfect. Thus, the vehicles need to wait for a timeout until they detect it and reorganize the cluster. This approach is more realistic than *GK*.

We use two metrics to evaluate the performance of our approach: the *achieved relative utility* and the *used bandwidth.* The *achieved relative utility* measures the performance of the network, i.e., how much data is provided to a vehicle in the network. This metric is between 0 and 1, where 1 states that the vehicle has received all the FCD that was sent and 0 states that the vehicle has received nothing. *Used bandwidth* captures whether the approach sticks to the average bandwidth limitation, i.e., if the side condition of the game is fulfilled.

We use box-plots and line-plots to visualize our results. In the box-plots, the boxes show the differences between vehicles inside of one simulation run. Next to each box, there is a line with a dot, visualizing the average value over all vehicles and simulation runs and the standard deviation of the average of all vehicles. In line-plot, the line displays the mean value for the vehicles in one simulation run.

Figure 1 depicts the performance of the approaches under different available bandwidths to each individual vehicle. It is evident that the performance of all approaches increases as the bandwidth increases, as depicted in Fig. 1a. For a full reception of all data available in the network via cellular, a bandwidth of roughly

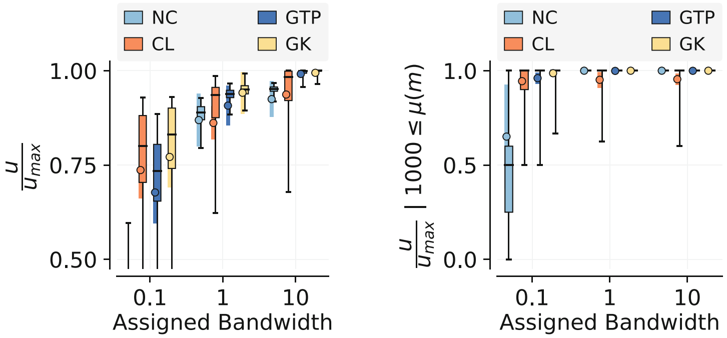

(a) Achieved relative utility.

(b) High-impact.

Fig. 1. Achieved relative bandwidth for different bandwidths (in messages/s).

100 messages per second is required. Even with a much smaller bandwidth of 10 messages per second, all approaches can achieve reasonable utility levels by prioritizing high-impact messages. It can be observed that our *GTP* approach outperforms the *CL* approach as well as the *NC* approach and has much smaller confidence intervals compared to the *CL* approach. Thus, our approach is more resilient and adaptive to different network conditions. Additionally, our approach is very close in performance to the *GK* approach. The same holds for a bandwidth of 1, while our approach decreases in performance for a bandwidth of 0.1. For a bandwidth of 0.1, our approach performs worse than the *CL* approach, as the redundant transmission of high-impact messages and the missing explicit coordination between vehicles decrease the performance of our *GTP* approach. This is also confirmed by Fig. 1b: For the high-impact messages, our approach performs well for both a bandwidth of 1 and 10, but struggles to receives the high-impact messages for a bandwidth of 0.1. That is, a bandwidth of 0.1 is not sufficient to receive the high-impact messages using only the available bandwidth of a single vehicle. Thus, the performance of our approach decreases below the performance of the *CL* approach, as the explicit coordination of vehicles in clustering approaches can handle low bandwidths well. Additionally, all approaches stick to the available bandwidth on average, while the bandwidth is temporarily exceeded by a subset of vehicles. This exceeding of bandwidth is justified by (i) the different number of available messages depending on the event location and (ii) the cooperative reception of messages by vehicles.

Figure 2 displays the influence of the share of privacy (fraction of privacy-sensitive vehicles) on our realistic vehicular network if the privacy-sensitive vehicles use an area of imprecision with radius 10 km. Figure 2a shows the behavior of the relative utility for all of the approaches. The *NC* approach decreases the most, as the privacy-sensitive vehicles have no possibility to compensate for their context imprecision. Additionally, our *GTP* approach constantly outperforms the *CL* approach and the *NC* approach independent of the level of privacy.

Most interestingly, the performance decrease of our *GTP* approach compared to the *GK* approach is not constant, it is lowest around 50% privacy. This can be justified by implicit coordination between privacy levels. This is also visible in Fig. 2b, which displays the relative utility of messages with an impact between 10 and 100. While the *NC* approach is not able to receive these messages at all, the utility of the other approaches decreases constantly. However, for our *GTP* approach, the utility remains constant for a very long duration, which leads to a comparably constant overall utility even for high privacy levels.

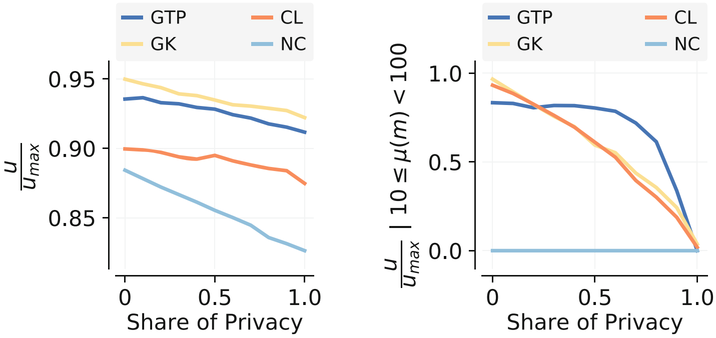

(a) Overall.

(b) Low-impact.

Fig. 2. Achieved relative utility for mixed environments.

5 Conclusion

In this paper we introduce privacy considerations in the management of FCD and have shown its impact on location-based services, since some data are not forwarded to a vehicle due to privacy considerations and the implemented location obfuscation. In order to alleviate this problem, we have introduced cooperation among vehicles so as to forward relevant data to their neighboring vehicles, enhancing in principle the data received by a vehicle only directly from the remote server. In this work, an ad-hoc, direct V2V cooperation paradigm is employed instead of a cluster-based one, also showing the high performance deterioration of the latter in a real vehicular networking environment. A major contribution of this work is the development and study of a game without coordination that determines the strategies (in terms of probabilities that a vehicle is forwarded by the server data of a given impact index) vehicles should follow, so that a properly defined utility is maximized; this is shown to lead to a diversification of the data received directly from the server by neighboring vehicles and increases the effectiveness of V2V cooperation.

In the evaluation, we analyzed the performance of our approach in a realistic vehicular network. Our results show the drastic performance increase compared to non-cooperative uncoordinated approaches, and the improvements over cluster-based approaches. Additionally, our approach performs almost similarly to a perfect clustering approach, which utilizes bandwidth optimally and detects disconnects immediately, but is not realizable in reality. When we analyze the performance of our approach for different privacy levels, we see that the performance remains constant for a long time.

References

1. Meuser, T., Bischoff, D., Richerzhagen, B., Steinmetz, R.: Cooperative offloading in context-aware networks: a game-theoretic approach. In: Proceedings of ACM International Conference on Distributed and Event-Based Systems (DEBS 2019). ACM (2019)
2. Golle, P., Greene, D., Staddon, J.: Detecting and correcting malicious data in VANETs. In: Proceedings of ACM International Workshop on Vehicular Ad Hoc Networks (VANET), VANET 2004, pp. 29–37. ACM, New York (2004)
3. Dötzer, F.: Privacy issues in vehicular ad hoc networks. In: Danezis, G., Martin, D. (eds.) PET 2005. LNCS, vol. 3856, pp. 197–209. Springer, Heidelberg (2006). https://doi.org/10.1007/11767831_13
4. Ying, B., Makrakis, D., Hou, Z.: Motivation for protecting selfish vehicles' location privacy in vehicular networks. IEEE Trans. Veh. Technol. **64**(12), 5631–5641 (2015)
5. Pan, X., Xu, J., Meng, X.: Protecting location privacy against location-dependent attacks in mobile services. IEEE Trans. Knowl. Data Eng. **24**(8), 1506–1519 (2012)
6. Ying, B., Nayak, A.: Social location privacy protection method in vehicular social networks. In: Proceedings of IEEE International Conference on Communications Workshops (ICC Workshops), pp. 1288–1292 (2017)
7. Wasef, A., Shen, X.S.: REP: location privacy for VANETs using random encryption periods. Mobile Netw. Appl. **15**(1), 172–185 (2010)
8. Sampigethaya, K., Huang, L., Li, M., Poovendran, R., Matsuura, K., Sezaki, K.: CARAVAN: providing location privacy for VANET. In: Embedded Security in Cars (ESCAR) (2005)
9. Liu, B., Zhou, W., Zhu, T., Gao, L., Luan, T.H., Zhou, H.: Silence is golden: enhancing privacy of location-based services by content broadcasting and active caching in wireless vehicular networks. IEEE Trans. Veh. Technol. **65**(12), 9942–9953 (2016)
10. Petit, J., Schaub, F., Feiri, M., Kargl, F.: Pseudonym schemes in vehicular networks: a survey. IEEE Commun. Surv. Tutor. **17**(1), 228–255 (2014)
11. Gerlach, M., Guttler, F.: Privacy in VANETs using changing pseudonyms - ideal and real. In: Proceedings of IEEE Vehicular Technology Conference (VTC-Spring), April 2007, pp. 2521–2525 (2007)
12. Palanisamy, B., Liu, L.: MobiMix: protecting location privacy with mix-zones over road networks. In: 2011 IEEE 27th International Conference on Data Engineering, pp. 494–505 (2011)
13. Freudiger, J., Raya, M., Félegyházi, M., Papadimitratos, P., Hubaux, J.-P.: Mix-zones for location privacy in vehicular networks. In: Proceedings of ACM Workshop on Wireless Networking for Intelligent Transportation Systems (WiN-ITS) (2007)

14. Ying, B., Makrakis, D., Mouftah, H.T.: Dynamic mix-zone for location privacy in vehicular networks. IEEE Commun. Lett. **17**(8), 1524–1527 (2013)
15. Wiedersheim, B., Ma, Z., Kargl, F., Papadimitratos, P.: Privacy in inter-vehicular networks: why simple pseudonym change is not enough. In: Proceedings of International Conference on Wireless On-Demand Network Systems and Services (WONS), pp. 176–183 (2010)
16. Duckham, M., Kulik, L.: Location privacy and location-aware computing. In: Dynamic and Mobile GIS, pp. 63–80. CRC Press (2006)
17. Gruteser, M., Grunwald, D.: Anonymous usage of location-based services through spatial and temporal cloaking. In: Proceedings of International Conference on Mobile Systems, Applications and Services (MobiSys), pp. 31–42. ACM, New York (2003)
18. Niu, B., Li, Q., Zhu, X., Cao, G., Li, H.: Achieving k-anonymity in privacy-aware location-based services. In: Proceedings of IEEE International Conference on Computer Communications (INFOCOM), pp. 754–762, April 2014
19. Liu, X., Liu, K., Guo, L., Li, X., Fang, Y.: A game-theoretic approach for achieving k-anonymity in location based services. In: Proceedings of IEEE International Conference on Computer Communications (INFOCOM), pp. 2985–2993 (2013)
20. Freudiger, J., Manshaei, M.H., Hubaux, J.-P., Parkes, D.C.: On non-cooperative location privacy: a game-theoretic analysis. In: Proceedings of the 16th ACM Conference on Computer and Communications Security, CCS 2009, pp. 324–337. ACM, New York (2009)
21. Du, S., Li, X., Du, J., Zhu, H.: An attack-and-defence game for security assessment in vehicular ad hoc networks. Peer-to-Peer Netw. Appl. **7**(3), 215–228 (2012). https://doi.org/10.1007/s12083-012-0127-9
22. Laoutaris, N., Telelis, O., Zissimopoulos, V., Stavrakakis, I.: Distributed selfish replication. IEEE Trans. Parallel Distrib. Syst. **17**(12), 1401–1413 (2006)
23. Meuser, T., Bischoff, D., Steinmetz, R., Richerzhagen, B.: Simulation platform for connected heterogeneous vehicles. In: Proceedings of International Conference on Vehicle Technology and Intelligent Transport Systems (VEHITS). SCITEPRESS, May 2019, pp. 412–419 (2019)
24. Lopez, P.A., et al.: Microscopic traffic simulation using SUMO. In: Proceedings of IEEE ITSC. IEEE (2018)
25. Uppoor, S., Fiore, M.: Large-scale urban vehicular mobility for networking research. In: Proceedings of IEEE Vehicular Networking Conference (VNC), pp. 62–69 (2011)