

Prolonging the Hide-and-Seek Game: Optimal Trajectory Privacy for Location-Based Services

George Theodorakopoulos¹, Reza Shokri², Carmela Troncoso³,
Jean-Pierre Hubaux⁴ and Jean-Yves Le Boudec⁴

¹Cardiff University, UK, ²ETH Zurich, Switzerland, ³Gradiant, Spain, ⁴EPFL, Lausanne, Switzerland,
¹g.theodorakopoulos@cs.cardiff.ac.uk, ²reza.shokri@inf.ethz.ch, ³ctroncoso@gradient.org,
⁴firstname.lastname@epfl.ch

ABSTRACT

Human mobility is highly predictable. Individuals tend to only visit a few locations with high frequency, and to move among them in a certain sequence reflecting their habits and daily routine. This *predictability* has to be taken into account in the design of location privacy preserving mechanisms (LPPMs) in order to effectively protect users when they expose their whereabouts to location-based services (LBSs) continuously. In this paper, we describe a method for creating LPPMs tailored to a user's mobility profile taking into her account privacy and quality of service requirements. By construction, our LPPMs take into account the sequential correlation across the user's exposed locations, providing the maximum possible *trajectory* privacy, i.e., privacy for the user's past, present location, and expected future locations. Moreover, our LPPMs are optimal against a *strategic adversary*, i.e., an attacker that implements the strongest inference attack knowing both the LPPM operation and the user's mobility profile. The optimality of the LPPMs in the context of trajectory privacy is a novel contribution, and it is achieved by formulating the LPPM design problem as a Bayesian Stackelberg game between the user and the adversary. An additional benefit of our formal approach is that the design parameters of the LPPM are chosen by the optimization algorithm.

Categories and Subject Descriptors

C.2.0 [Computer-Communication Networks]: General—Security and protection; K.4.1 [Computers and Society]: Public Policy Issues—Privacy

Keywords

Location Privacy; Trajectory Privacy; Location Transition Privacy; Optimal Location Obfuscation; Privacy-Utility Tradeoff; Bayesian Stackelberg Game

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
WPES'14, November 3, 2014, Scottsdale, Arizona, USA.
Copyright 2014 ACM 978-1-4503-2948-7/14/11 ...\$15.00.
<http://dx.doi.org/10.1145/2665943.2665946>

1. INTRODUCTION

Location-Based Services (LBSs) provide users with valuable information about their surroundings such as traffic status (e.g., Beat the Traffic, or INRIX Traffic Maps, Routes & Alerts), nearby points of interest (e.g., Google Maps), or friends' activities (e.g., Foursquare or Google Latitude). Despite the potential benefits of location sharing, information about our current, future, or frequently-visited locations is highly sensitive, as it can be used to infer our habits, preferences, political and religious affiliations, and even to endanger our physical security if it falls in the wrong hands.

Hence, the need arises to protect users' location privacy, while maintaining the usability and quality of LBSs. This task must account for the three following considerations:

First and foremost, as Shannon's maxim states, "One ought to design systems under the assumption that the enemy will immediately gain full familiarity with them." In other words, the adversary will adapt his attack to the protection mechanism. This in turn shall lead to an updated mechanism, then to a novel attack, and so on ad infinitum. This is commonly known as the *arms-race problem*. To cut the arms race short, in our approach the defender anticipates the adversary's reaction, and so the initial design is already robust against an informed adversary.

Second, in order to ease the deployment of a Location Privacy Preserving Mechanism (LPPM) it must be user-centric, i.e., users must be able to use it independently of other users' behavior, and without the permission or collaboration of a third party. Hence, decisions taken to protect privacy (e.g., hiding, perturbing, or faking locations) need to be made locally to the users. Our approach only requires users to perform a local look-up in a pre-computed table, and hence it can be easily integrated in mobile devices frequently used to access LBSs.

Third, the protection of users current location is intricately bound with protecting also past and future locations. The privacy protection offered by an LPPM is severely impacted by the frequency with which locations are revealed to the LBS, since locations exposed in quick succession are highly correlated. When correlated locations are exposed, inferring the user's current location provides the adversary with tools to reduce the uncertainty on the user's immediate past or future whereabouts. Our LPPMs take location correlation into account to effectively protect a user's location privacy along her trajectory.

We propose a framework to design user-centric LPPMs that – given a user's quality requirements, privacy requirements, and mobility profile – can (1) protect the privacy of

past locations (i.e., the current obfuscation is chosen to be compatible with past ones), (2) protect the privacy of future locations (i.e., the current obfuscation is chosen to be compatible with likely next locations), (3) protect the privacy of transitions between locations (i.e., the current obfuscation is chosen to hide sensitive movements between successive locations, even if individually they are not sensitive), (4) protect locations that the user visits between two LBS accesses (i.e., the current obfuscation is chosen to protect locations that the user visits without issuing an LBS queries). To the best of our knowledge, this is the first work that addresses the two latter objectives as separate targets in need of protection. The two former objectives have already been addressed in the literature, but proposals do not consider informed adversaries, or they do not operate in a user-centric manner.

Our LPPMs can find the optimal privacy protection mechanism in any given scenario, i.e., they achieve the best protection among all possible mechanisms against a strategic adversary with knowledge of mobility profiles and LPPM algorithm. In other words, LPPMs designed under our framework provide privacy level that constitutes an upper bound on the privacy that is achievable by any other defense. The proposed framework can handle a wide range of correlation levels without making any prior assumption on the user mobility and LBS access patterns: From cases where there is high correlation between exposed locations, e.g., LBSs that require continuous requests to the provider (e.g., navigating with Google maps), to cases in which there is complete independence among exposed locations, e.g., LBS that require sporadic location updates (e.g., checking-in in Foursquare).

The key technique used in our solution is a Bayesian Stackelberg game between the privacy defender and adversary, launched by the user every time she wishes to share her location with the LBS. The two main benefits of this game theory technique are the following. First, it can naturally express the objective of optimizing privacy under the double constraint of anticipating the adversary’s attack and respecting the user’s quality requirements. Second, it allows us to efficiently search in an infinite space of potential solutions, guaranteeing that the computed solution is optimal without testing each and every single one of them (computationally impossible) and without limiting ourselves to the heuristic (but undeniably creative) solutions that human ingenuity can concoct. We provide a general design method that can be instantiated for particular privacy objectives, periods of observation, and patterns of exposure. For the sake of illustration, we provide examples to specifically protect two plausible privacy objectives: protecting the most recent locations (including the current one) and protecting the current and future locations.

We run our solution on real users’ trajectories to obtain optimal location obfuscation mechanisms that maximize privacy, defined as the adversary’s error in estimating the user’s true location. Our results show that our method is more effective at protecting location privacy than mechanisms that only consider the current exposed locations. We also show that quality of service can be traded off for privacy, but the maximum privacy achievable strongly depends on users’ behavior (i.e., on the predictability of their movements).

2. RELATED WORK

In this section we qualitatively compare our approach to previous work on location privacy. The comparison focuses

on schemes that provide trajectory privacy [8] rather than sporadic privacy [22, 9, 7], i.e., privacy of locations exposed to the LBS independently of each other. A quantitative comparison with the latter is provided in Section 5, where we show that our trajectory-aware approach outperforms sporadic privacy-preserving mechanisms.

A first class of trajectory-aware mechanisms in the literature aims at protecting user privacy when trajectories are published in bulk. Protection is achieved by grouping different users’ trajectories in a wide area such that the aggregate trajectory can be ascribed to at least k users [1]; mixing the trajectories of k users [16]; eliminating some events from the published dataset [14, 23]; or replacing locations with larger regions defined by a pre-defined grid [13]. Other algorithms need access to the complete trajectory before protection can be applied [26], or they delay the location exposure so as to gather additional information about subsequent user locations [12, 3]. In contrast, our approach decides in real time how to protect the location that the user is about to expose.

Other trajectory-aware mechanisms assume either the existence of a trusted third party (e.g., the cellular service provider) [17, 11], or of nearby users that can be leveraged to achieve joint privacy protection [4, 10, 15]. Both scenarios violate the user-centricity design requirement in this paper.

In addition to addressing trajectory privacy in a user-centric and real-time manner, our major qualitative difference from prior work is that we address the problem formally. We formalize the privacy-preserving mechanism design problem as a Bayesian Stackelberg game, similarly to Shokri *et al.* [22]. This has two main advantages. First, it allows to provide optimal privacy against any adversary describable by our model. Second, it allows to define fine-grained knobs for expressing users’ privacy and quality requirements. We would like to emphasize that our approach is not an extension to Shokri *et al.*, but uses the same methodology to achieve a different goal, i.e., to protect trajectory privacy instead of sporadically exposed locations.

Other formal approaches to the location privacy problem extend the concept of differential privacy to location privacy, defining a new privacy metric: geo-indistinguishability [2, 5, 6, 20]. The solutions in [2, 5, 20] focus on sporadic locations and are not applied on trajectory privacy. Andrés *et al.* [2] state that in the case of successive location disclosures, the geo-indistinguishability that their mechanism provides decreases linearly in the number of disclosed locations. However, this solution is not optimal and can be outperformed by [22] when the adversary has prior knowledge. In Section 5, we show that our algorithm compares favorably against that of Shokri *et al.* when protecting trajectory privacy, hence we can conclude that our algorithm also offers better protection than that of Andrés *et al.* The solution in [6] deals with mobility traces, but is non optimal.

3. PROBLEM STATEMENT

User Mobility and LBS Access Pattern. Consider a user moving within M discrete locations $\mathcal{R} = \{r_1, \dots, r_M\}$. The user’s movements are represented as a discrete-time trajectory of locations at times $\mathcal{T} = \{1, 2, \dots\}$. An *event* $\langle r, t \rangle$ denotes that the user is at location $r \in \mathcal{R}$ at time $t \in \mathcal{T}$. Slightly abusing notation, the time-subscripted variable r_t will denote the user’s location at time t . Typical values from




Figure 1: A user moves from location $(3, 4)$ at time $t - 2$, to $(4, 3)$ at $t - 1$, to $(4, 2)$ at current time t . The user wants to protect locations at times $t - 1$ and t , and these are denoted by a_{trg} (target events). At past times $t - 2$ and $t - 1$, the LPPM exposed pseudolocations $(3, 3)$ and $(3, 2)$ (denoted by o_{pre}), and to protect a_{trg} the LPPM currently exposes location $(3, 1)$ (denoted by o_{post}).

\mathcal{R} are r, r_i, r_j , whereas typical time-subscripted locations are r_t, r_{t-1}, r_{t+1} .

User mobility is modeled probabilistically. Any model is possible, as long as it allows to compute probabilities of the user visiting various sequences of locations. In this paper, we choose to model mobility as a first-order Markov chain on \mathcal{R} , but we must stress that our contribution is not tied to this choice since our method can compute the optimal protection for any other mobility model.

We assume an LBS access pattern in which the user accesses the LBS at each time instant $t \in \mathcal{T}$, i.e., from each location that she visits. We assume this for simplicity, and at the end of this section we show that we can accommodate sparser LBS-access patterns. After all, the most interesting case is when LBS-access times are close enough to each other so that successively exposed locations *are correlated*. If they are not, the problem becomes equivalent to sporadic location disclosure, studied in [22].

LPPM Functionality. The user wants to protect her privacy from an adversary who observes the locations exposed to the LBS (so the adversary could be the LBS provider itself, an eavesdropper, other LBS users, etc.). Hence, she uses an LPPM that obfuscates her real locations before they are sent to the LBS. We model obfuscation as a replacement operation in which a fake location from a set \mathcal{R}' is sent to the LBS instead of the real location. We take \mathcal{R}' to be the same as \mathcal{R} . We call these fake locations *pseudolocations* or *obfuscated locations* and denote them by r' . The corresponding events $\langle r', t \rangle$ are termed *pseudoevents*.

The problem we tackle is the design of an LPPM algorithm $f(\cdot)$ that takes as input the real location (or locations) to be protected and the pseudolocations previously sent to the adversary, and then probabilistically selects the pseudoloca-




Figure 2: Information available to the LPPM and the adversary: The LPPM wants to protect location(s) a_{trg} by producing appropriate pseudolocation(s) o_{post} . The adversary observes the output o_{post} of the LPPM and, using his knowledge of the LPPM function f , estimates a_{trg} ; the adversary's estimate is \hat{a}_{trg} . The prior knowledge of the adversary and of the LPPM consists of the transition matrix P and the pseudolocations o_{pre} that have been produced in the past.

tion (or pseudolocations) to expose to the adversary. The pseudolocation set is fixed, and it coincides with the location set. The adversary knows the probabilistic model that describes user mobility, and also knows the LPPM algorithm $f(\cdot)$. Notice here the self-reference: We design $f(\cdot)$ assuming an adversary who knows the $f(\cdot)$ that will be designed and can respond optimally to it.

We introduce the notation used in the paper with the example shown in Figure 1. Assuming that the current time is t , the elements of the framework are defined as follows:

- a_{trg} denotes the target events that the user wants to protect, or equivalently, the events that the adversary wants to infer. In the example, the user wants to protect her location at times $t - 1$ and t and thus $a_{\text{trg}} = (r_{t-1}, r_t) = \{\langle (4, 3), t - 1 \rangle, \langle (4, 2), t \rangle\}$.
- o_{pre} is a subset of the pseudoevents that the LPPM created and sent to the LBS *up to but before* the current time. These are the pseudoevents that matter for the estimation of a_{trg} : Typically, o_{pre} would be a sequence of consecutive pseudoevents starting with a recent time instant (as old ones do not matter for estimating a_{trg}) and leading up to the current time. These are known both to the adversary and to the LPPM. In the example, the relevant pseudolocations are those exposed at times $t - 1$ and $t - 2$ and thus $o_{\text{pre}} = (r'_{t-2}, r'_{t-1}) = \{\langle (3, 3), t - 2 \rangle, \langle (3, 2), t - 1 \rangle\}$.
- o_{post} is the pseudolocation (or set of pseudolocations) that the LPPM produces to protect a_{trg} and that will be sent to the LBS at current time. In the example the user exposes pseudolocation $(3, 1)$ at time t , thus $o_{\text{post}} = (r'_t) = \{\langle (3, 1), t \rangle\}$.
- $f(o_{\text{post}}|a_{\text{trg}}, o_{\text{pre}})$ is the probability that the LPPM produces o_{post} , given its knowledge o_{pre} and the locations a_{trg} it is trying to protect. This function encodes the defensive mechanism. It can be viewed as a codebook that prescribes, for each value of a_{trg} and o_{pre} , a randomization over the possible values of o_{post} .

Notice that a_{trg} need not be the same length as o_{post} : In the example, the LPPM exposes the current pseudoevent (only time t), while the goal is to protect events at the current and the previous time instants (t and $t - 1$).

Attacker Model and Privacy Metric. In short, privacy is quantified as the adversary's error in estimating the user's

true location(s) a_{trg} . Figure 2 illustrates the information flow of events and pseudoevents to the LPPM and to the adversary. The detailed notation is as follows:

- $\psi(a_{\text{trg}}|o_{\text{pre}})$ is the adversary’s prior probability distribution on the inference target a_{trg} , given his prior knowledge o_{pre} . It encodes what the adversary can deduce about a_{trg} *before* observing the LPPM’s current output o_{post} .
- \hat{a}_{trg} denotes the adversary’s estimate of a_{trg} . Similarly to a_{trg} , it can be seen as a time-indexed vector whose elements belong to the set \mathcal{R} of locations.
- $h(\hat{a}_{\text{trg}}|o_{\text{pre}}, o_{\text{post}})$ is the probability that the adversary estimates \hat{a}_{trg} to be the true value of a_{trg} , given his knowledge of prior pseudolocations o_{pre} and given the pseudolocation(s) o_{post} exposed at current time t . Note that, by definition, $h(\cdot)$ can contain multiple observed exposed pseudolocations that happen across any time period (e.g., multiple days).
- $d_p(\hat{a}_{\text{trg}}, a_{\text{trg}}) \geq 0$ is the privacy gain when the adversary’s estimate is \hat{a}_{trg} and the true value of the inference target is a_{trg} . It is zero only if $\hat{a}_{\text{trg}} = a_{\text{trg}}$. The value of d_p for each pair of locations reflects the sensitivity of the user with respect to different location. The user must assign low values if she is sensitive towards a particular location a_{trg} . We treat d_p as an input to our framework.

The privacy that an LPPM $f(\cdot)$ achieves against an adversary implementing attack $h(\cdot)$ is then the expected value of $d_p(\hat{a}_{\text{trg}}, a_{\text{trg}})$, given prior observations o_{pre} :

$$\begin{aligned} \text{Privacy}(\psi, f, h, d_p; o_{\text{pre}}) &= \mathbb{E}\{d_p(\hat{a}_{\text{trg}}, a_{\text{trg}})|o_{\text{pre}}\} \quad (1) \\ &= \sum_{a_{\text{trg}}, \hat{a}_{\text{trg}}} \mathbb{P}\{a_{\text{trg}}, \hat{a}_{\text{trg}}\} d_p(\hat{a}_{\text{trg}}, a_{\text{trg}}) \\ &= \sum_{\substack{a_{\text{trg}} \\ o_{\text{post}} \\ \hat{a}_{\text{trg}}}} \psi(a_{\text{trg}}|o_{\text{pre}}) f(o_{\text{post}}|a_{\text{trg}}, o_{\text{pre}}) h(\hat{a}_{\text{trg}}|o_{\text{pre}}, o_{\text{post}}) \\ &\quad d_p(\hat{a}_{\text{trg}}, a_{\text{trg}}). \end{aligned}$$

This formula represents the adversary’s expected estimation error. An alternative pessimistic quantification of privacy, we could take the minimum estimation error over all possible values of a_{trg} , which would be a worst-case scenario.

Regarding $d_p(\hat{a}_{\text{trg}}, a_{\text{trg}})$, intuitively it can be seen as a distance between \hat{a}_{trg} and a_{trg} that measures the similarity between the inferred and real locations with respect to the user’s privacy concerns. For instance, it can be the sum, the minimum, or the maximum of the Euclidean distances between the corresponding locations of each vector, i.e., the total, minimum, or maximum error of the adversary over each pair of corresponding locations in the two vectors. Contributions to $d_p(\hat{a}_{\text{trg}}, a_{\text{trg}})$ can be *weighted* to encode the *privacy sensitivity of individual locations* $r \in a_{\text{trg}}$, e.g., if r is a very sensitive location the contribution of estimating $\hat{r} \in \hat{a}_{\text{trg}}$ to $d_p(\hat{a}_{\text{trg}}, a_{\text{trg}})$ could be large even if r and \hat{r} differ by very little. Alternatively, $d_p(\hat{a}_{\text{trg}}, a_{\text{trg}})$ could be the Hamming distance between the two vectors, i.e., the number of locations at which the estimate differs from the true value.

Furthermore, $d_p(\hat{a}_{\text{trg}}, a_{\text{trg}})$ can encode the privacy sensitivity of *transitions between locations*, rather than individual locations. For instance, visiting the bank and visiting a government official may not be very sensitive if considered separately, but visiting the official immediately after the bank may be much more sensitive, especially if the user just made a large withdrawal from the bank and the official is in charge of land development licensing. Moreover, transitions between regions reveal the direction of travel. For instance, the adversary may learn whether the user enters or exits a building, e.g., a hospital.

Appropriate $d_p(\cdot)$ needs to be chosen depending on the user’s privacy concerns, and whether she wants to protect sensitive locations or sensitive transitions. We note that some $d_p(\cdot)$ definitions are more general than others: a $d_p(\cdot)$ that protects the transition between two successive locations automatically protects the locations themselves, so it could be used to protect both objectives simultaneously.

Quality Metric. Sending pseudoevents instead of true events to the LBS may help with privacy, but it can degrade the quality of the LBS’s response. We model the quality loss stemming from exposing pseudolocations as follows:

- q_{trg} denotes the relevant events with respect to quality. Similarly to a_{trg} , q_{trg} is a time-indexed vector. However, its time indices are not necessarily the same as those of a_{trg} : The events that matter for quality may be different from the ones that matter for privacy.
- $d_q(q_{\text{trg}}, o_{\text{post}}, o_{\text{pre}})$ represents the quality loss when q_{trg} is the true value of the quality-relevant events, the LPPM currently reports o_{post} and it has reported o_{pre} in the past. The function d_q is an input to our framework and it reflects the user’s requirements with respect to quality of service, and is related to need for accurate location information at the service provider to return useful service to the user.

The expected quality loss caused by an LPPM $f(\cdot)$ is the expected value of $d_q(q_{\text{trg}}, o_{\text{post}}, o_{\text{pre}})$ over all q_{trg} and o_{post} , for a given history o_{pre} :

$$\begin{aligned} Q_{\text{loss}}(f, d_q, o_{\text{pre}}) &= \mathbb{E}\{d_q(q_{\text{trg}}, o_{\text{post}}, o_{\text{pre}})|o_{\text{pre}}\} \quad (2) \\ &= \sum_{q_{\text{trg}}, o_{\text{post}}} \mathbb{P}\{q_{\text{trg}}, o_{\text{post}}|o_{\text{pre}}\} d_q(q_{\text{trg}}, o_{\text{post}}, o_{\text{pre}}) \\ &= \sum_{q_{\text{trg}}, o_{\text{post}}} \mathbb{P}\{q_{\text{trg}}|o_{\text{pre}}\} \mathbb{P}\{o_{\text{post}}|q_{\text{trg}}, o_{\text{pre}}\} \\ &\quad d_q(q_{\text{trg}}, o_{\text{post}}, o_{\text{pre}}) \end{aligned}$$

In the equation above, $f(\cdot)$ is hidden in $\mathbb{P}\{o_{\text{post}}|q_{\text{trg}}, o_{\text{pre}}\}$, which can be unwrapped as

$$\begin{aligned} \mathbb{P}\{o_{\text{post}}|q_{\text{trg}}, o_{\text{pre}}\} &= \sum_{a_{\text{trg}}} \mathbb{P}\{o_{\text{post}}, a_{\text{trg}}|q_{\text{trg}}, o_{\text{pre}}\} \quad (3) \\ &= \sum_{a_{\text{trg}}} \mathbb{P}\{o_{\text{post}}|a_{\text{trg}}, q_{\text{trg}}, o_{\text{pre}}\} \mathbb{P}\{a_{\text{trg}}|q_{\text{trg}}, o_{\text{pre}}\} \\ &= \sum_{a_{\text{trg}}} f(o_{\text{post}}|a_{\text{trg}}, o_{\text{pre}}) \mathbb{P}\{a_{\text{trg}}|q_{\text{trg}}, o_{\text{pre}}\}. \end{aligned}$$

We assume that there is a maximum expected quality loss $Q_{\text{loss}}^{\text{max}}$ that users are willing to tolerate. Formally,

$$Q_{\text{loss}}(f, d_q, o_{\text{pre}}) \leq Q_{\text{loss}}^{\text{max}}. \quad (4)$$

The quality loss $d_q(q_{\text{trg}}, o_{\text{post}}, o_{\text{pre}})$ can be seen as a distance between two vectors: q_{trg} and the combination of $o_{\text{post}}, o_{\text{pre}}$. It will be zero if an accurate, noiseless, trajectory is reported by the LPPM (i.e., if these two vectors are identical), and positive otherwise. If the application needs high location precision to function well, then $d_q(\cdot)$ will be large even for small differences between q_{trg} and $o_{\text{post}}, o_{\text{pre}}$.

Moreover, $d_q(\cdot)$ can encode variable sensitivity across locations. For example, in locations with many nearby restaurants, an application that finds the nearest restaurant can tolerate a lot of noise, while isolated areas may require more precision. The quality loss function could be instantiated, among other possibilities, as a Euclidean or Hamming distance between real (q_{trg}) and reported locations ($o_{\text{post}}, o_{\text{pre}}$). The versatility of $d_q(\cdot)$ extends to encoding quality loss for applications that depend on the whole *trajectory of locations*, rather than just on a single location, e.g., consider a car insurance company that monitors the driving behavior of a customer: quantities such as speed and sudden acceleration or deceleration cannot be evaluated on single locations.

3.1 Sparse LBS Access Pattern

So far we have assumed that users access the LBS at every single time instant although in reality one cannot expect that LBS accesses are uninterrupted (e.g., users may access an online navigation system to travel around some parts of a city, but not in others). In this scenario, an adversary may be concerned about the user’s whereabouts *between* two LBS accesses, i.e., at times when there are no corresponding exposed pseudolocations. Therefore, these intermediate locations also need to be protected (recall objective (4) in the introduction).

Our dense-LBS-access assumption can also express such a privacy objective. We can accommodate these inference targets by extending the definition of $d_p(\hat{a}_{\text{trg}}, a_{\text{trg}})$ to incorporate the privacy sensitivity of any intermediate locations that the adversary can infer from \hat{a}_{trg} . For example, assume that, from the estimates \hat{r}_{t-1} and \hat{r}_t of locations r_{t-1} and r_t , the adversary can estimate two intermediate locations visited at $t-1+\epsilon$ and $t-\epsilon$ (e.g., the attacker can use Viterbi decoding [19] to infer the most likely trajectory, or the forward-backward algorithm [19] to compute the probability distribution of locations at some times between $t-1$ and t). In this case, $d_p((\hat{r}_{t-1}, \hat{r}_t), (r_{t-1}, r_t))$ would be expressed as $d_p((\hat{r}_{t-1}, \hat{r}_{t-1+\epsilon}, \dots, \hat{r}_{t-\epsilon}, \hat{r}_t), (r_{t-1}, r_{t-1+\epsilon}, \dots, r_{t-\epsilon}, r_t))$.

We note that this approach to handle the estimation of locations where the LBS is not accessed is not particular to LPPMs designed under our framework, hence it can be used to complement other privacy-preserving solutions in the literature.

4. TRAJECTORY PRIVACY AS A STACKELBERG GAME

4.1 General Privacy Scenario

As we have discussed in Section 3, our goal is to design an LPPM that protects user privacy (by maximizing d_p), while preserving quality of service (by respecting the maximum quality loss threshold $Q_{\text{loss}}^{\text{max}}$ for d_q). Designing an LPPM reduces to choosing appropriate values for the probabilities $f(o_{\text{post}}|a_{\text{trg}}, o_{\text{pre}})$. In addition, the design process must anticipate that the adversary will know the values chosen for f , which means that she will choose the attack h accordingly.




Figure 3: The LPPM can choose among an infinite selection of probability functions f . For each one of the f_i , the adversary chooses one of his infinite selection of attacks h_j to minimize P_{ij} ; for instance, h_2 is the best response to f_1 , resulting in Privacy P_{12} , and h_1 is the best response to f_2 , resulting in Privacy P_{21} . Call $h_{\min}(f)$ the minimizing h for a given f . Anticipating the adversary’s choice for each f , the LPPM chooses the f that maximizes $\text{Privacy}(f, h_{\min}(f))$; for instance, if $P_{12} > P_{21}$, the LPPM would choose f_1 over f_2 . Call f_{\max} the maximizing f . The resulting pair $f_{\max}, h_{\min}(f_{\max})$ is the Stackelberg equilibrium, and the privacy achieved is $\text{Privacy}(f_{\max}, h_{\min}(f_{\max}))$.

Figure 3 details the reasoning involved in solving our task, which is equivalent to solving a Stackelberg game. The distinguishing feature of this game is that there is a *leader*, who commits to a choice, and a *follower*, who observes the leader’s choice and then makes a choice of his own. In our task, the leader is the user and her choice is the LPPM, the follower is the adversary who chooses an attack given the user’s choice. The Stackelberg equilibrium is a pair of choices (f^* for the LPPM and h^* for the attack) such that neither the user nor the adversary would gain anything by changing their respective choices. In other words, h^* is the choice that minimizes privacy against f^* , and f^* is the choice that maximizes privacy against an adversary who will make her choice after observing f^* while respecting the quality constraint. Note that f^* does not necessarily maximize privacy against h^* , i.e., if the user could be certain that the adversary would choose h^* , then a better choice than f^* could exist. The LPPM design tries to limit the worst possible privacy loss, knowing only that the adversary will choose the most effective attack against whatever f the LPPM implements.

Formally, the general LPPM design task is to choose f and h that solve the Stackelberg game max-minimization

$$\max_f \min_h \text{Privacy}(\psi, f, h, d_p; o_{\text{pre}}) \quad (5)$$

subject to

$$Q_{\text{loss}}(f, d_q, o_{\text{pre}}) \leq Q_{\text{loss}}^{\text{max}}. \quad (6)$$

Other than f and h , all functions and parameters $\psi, d_p, o_{\text{pre}}, d_q, Q_{\text{loss}}^{\text{max}}$ are inputs to the problem: ψ – the user’s mobility, and d_p – the user’s location sensitivity, are specific to the user we aim to protect; the last two $d_q, Q_{\text{loss}}^{\text{max}}$ are specific to the LBS application and perhaps also depend on the user’s tolerance to quality deterioration; the prior observations o_{pre} depend on the particular time when the user wants to protect her privacy.

We now give two specific examples of the general LPPM design task for two plausible privacy objectives. The first

objective is to protect the $k + 1$ most recent locations (including the current one at time t), having already exposed pseudolocations for the k past time instants, by choosing an appropriate pseudolocation to expose at time t . The second objective is to protect the current *and* future locations, assuming nothing has been exposed so far, by choosing a pseudolocation for time t *and* for future time instants. We emphasize that one can instantiate any number of objectives by selecting the time period one wishes to protect (a_{trg}), the events that have been exposed already (o_{pre}), and the time period for which the LPPM can expose pseudolocations (o_{post}).

4.2 Joint Protection of Past-Present Locations

Consider a user who, at time t , wants to publish her location r_t . She has already published her locations r_τ at the k previous time instants $\tau = \{t - 1, t - 2, \dots, t - k\}$. These locations have been sent by the LPPM to the LBS as pseudolocations o_τ . As explained in Section 3, locations exposed prior to time $t - k$, i.e., from 1 to $t - k - 1$, are considered to have no influence on the choice of the user at time t .

The simplest privacy objective (a_{trg}) that the user could have is to protect her current location only: $a_{\text{trg}} = r_t$. We call this *single location privacy*. But as we have argued in [24], the *transition* from the previous location r_{t-1} to the current position may be sensitive, or, in general, the transition from the k -tuple r_τ to r_t may be sensitive. A different objective, therefore, is to protect the whole vector r_τ in addition to r_t : $a_{\text{trg}} = (r_t, r_\tau)$. Observe that the latter objective (transition privacy) is more general than the previous one (single location privacy): Choosing a function $d_p((\hat{r}_t, \hat{r}_\tau), (r_t, r_\tau))$ that just ignores r_τ and \hat{r}_τ makes the two cases equivalent.

The prior observations o_{pre} are the set of pseudolocations o_τ , and the pseudolocation o_{post} that the LPPM produces is just the one corresponding to the current time t : o_t . The quality loss d_q is a function of past and present exposed pseudolocations, (o_t, o_τ) , and of q_{trg} . As stated in Section 3 q_{trg} does not need to coincide with a_{trg} , and can consist of any subset of events from time 1 up to and including t . In fact, it does not even need to overlap with (r_t, r_τ) .

Making the appropriate substitutions in (1), we derive the privacy definition for the specific case of protecting past and present locations as follows:

$$\begin{aligned} \text{Privacy}(\psi, f, h, d_p; o_\tau) = & \sum_{r_t, r_\tau, o_t, \hat{r}_t, \hat{r}_\tau} \psi(r_t, r_\tau | o_\tau) \\ & f(o_t | r_t, r_\tau, o_\tau) \\ & h(\hat{r}_t, \hat{r}_\tau | o_t, o_\tau) \\ & d_p((\hat{r}_t, \hat{r}_\tau), (r_t, r_\tau)) \end{aligned} \quad (7)$$

The quality loss is also straightforward to define using (2):

$$\begin{aligned} Q_{\text{loss}}(f, d_q, o_{\text{pre}}) &= \mathbb{E}\{d_q(q_{\text{trg}}, o_t, o_\tau) | o_\tau\} \\ &= \sum_{q_{\text{trg}}, o_t} \mathbb{P}\{q_{\text{trg}}, o_t | o_\tau\} d_q(q_{\text{trg}}, o_t, o_\tau) \\ &= \sum_{q_{\text{trg}}, o_t} \mathbb{P}\{q_{\text{trg}} | o_\tau\} \mathbb{P}\{o_t | q_{\text{trg}}, o_\tau\} \\ & \quad d_q(q_{\text{trg}}, o_t, o_\tau) \end{aligned}$$

Notice that setting $k = 0$ eliminates o_{pre} (which would mean that none of the previous pseudolocations exposed to the adversary are assumed to correlate with the current lo-

cation), and the target events a_{trg} reduce to the current time t only. In total, only the current location matters for privacy and for quality, and the design task reduces to the sporadic case handled by Shokri et al.'s framework [21].

4.3 Joint Protection of Present-Future Locations

We now consider a user who, as before, wants to publish her location r_t at time t . However, in contrast with the previous case, she is not concerned about the past locations she has visited, but rather about future ones. This concern can be motivated as follows:

Disclosing the current location might not be important in and of itself, but it might make it much easier for the adversary to infer the next location, which happens to be very sensitive. For instance, the user might currently be on a street that only leads to an abortion clinic. Hence, disclosing her current location is almost equivalent to disclosing that she will go to the clinic. Symmetrically, her current location might be very sensitive, and her next (expected) location can be linked easily to her current one. For instance, she might about to leave the abortion clinic and enter a street that is only used as the clinic's exit. Furthermore, as argued in [24], neither the current nor the next location might be particularly sensitive separately, but the transition from one to the other might be.

The conclusion in all these cases is that the current location must be protected jointly with the (possible) next one(s), where the user will be at time $t + 1$ and later. For this reason, a_{trg} includes time subscripts larger than t , and so does o_{post} , i.e., the LPPM should take into account at the present time t what it is likely to output in future times, so that the current choice of o_t does not limit future choices. The intuition is that the LPPM should choose the current pseudolocation o_t so that future paths that the user will likely take can be protected with pseudolocations that are compatible with o_t .

For simplicity, we consider an example where the LPPM anticipates only the next location, rather than several successive future locations, therefore a_{trg} is (r_{t+1}, r_t) and o_{post} is (o_{t+1}, o_t) , and no prior locations matter for privacy, so o_{pre} is omitted.

Substituting in (1), privacy in this example is defined as

$$\begin{aligned} \text{Privacy}(\psi, f, h, d_p) = & \sum_{\substack{r_{t+1}, r_t \\ o_{t+1}, o_t \\ \hat{r}_{t+1}, \hat{r}_t}} \psi(r_{t+1}, r_t) f(o_{t+1}, o_t | r_{t+1}, r_t) \\ & h(\hat{r}_{t+1}, \hat{r}_t | o_{t+1}, o_t) d_p((\hat{r}_{t+1}, \hat{r}_t), (r_{t+1}, r_t)), \end{aligned} \quad (8)$$

and the quality loss, from (2), is defined for a general q_{trg} as

$$\begin{aligned} Q_{\text{loss}}(f, d_q) &= \mathbb{E}\{d_q(q_{\text{trg}}, o_t, o_{t+1})\} \\ &= \sum_{q_{\text{trg}}, o_t, o_{t+1}} \mathbb{P}\{q_{\text{trg}}, o_t, o_{t+1}\} d_q(q_{\text{trg}}, o_t, o_{t+1}). \end{aligned}$$

4.4 Optimal Attacks and Defenses via Linear Programming

Having reduced the LPPM design to a Stackelberg game max-minimization, we now compute the equilibrium of the game, which is equivalent to computing the optimal defense f and attack h .

Note that there is an infinity of candidate f s and h s (all possible probability distributions), so enumeration (as shown in Figure 3) cannot be used directly as an algorithm to find the equilibrium. To overcome this problem, we use a standard technique for transforming the computation of a game theoretic equilibrium to a linear program. This technique introduces auxiliary variables $x_{o_{\text{post}}}$, for each possible value of o_{post} , which roughly correspond to the amount of privacy gained when the LPPM reports each particular value of o_{post} . As previous research [22] has expanded on this transformation technique, we merely present the key features of the resulting linear program:

We want to maximize $\sum_{o_{\text{post}}} x_{o_{\text{post}}}$ under the constraint

$$x_{o_{\text{post}}} \leq \sum_{a_{\text{trg}}} \psi(a_{\text{trg}}|o_{\text{pre}}) f(o_{\text{post}}|a_{\text{trg}}, o_{\text{pre}}) d_p(\hat{a}_{\text{trg}}, a_{\text{trg}}),$$

$$\forall \hat{a}_{\text{trg}}, o_{\text{post}}, \quad (9)$$

and under the constraint

$$\sum_{\substack{q_{\text{trg}} \\ o_{\text{post}}}} \Pr\{q_{\text{trg}}|o_{\text{pre}}\} \Pr\{o_{\text{post}}|q_{\text{trg}}, o_{\text{pre}}\} d_q(q_{\text{trg}}, o_{\text{post}}, o_{\text{pre}})$$

$$\leq Q_{\text{loss}}^{\max}. \quad (10)$$

Equation (9) is equivalent to the min-maximization (5), and equation (10) is just the quality constraint (6).

In this way, we can compute the optimal LPPM f^* . The optimal attack h^* can either be computed via (5), or by solving another linear program called the dual. We do not go into detail about the dual here and point the interested reader to prior research [22] and to standard textbooks on linear programming [25].

Having computed the optimal LPPM f^* and the optimal attack h^* , the resulting $\text{Privacy}(f^*, h^*)$ is the level of privacy achieved by the user, and the resulting $Q_{\text{loss}}(f^*)$ is the application's incurred quality loss.

Alternative quantifications of privacy. As we mentioned at the beginning of this section, the optimal LPPM f^* depends on all inputs $\psi, d_p, o_{\text{pre}}, d_q, Q_{\text{loss}}^{\max}$, so if any of these functions or variables changes, a new LPPM must be computed to maximize privacy. Alternatively, it is possible to compute an LPPM that maximizes the *average* privacy across a range of different values, e.g., o_{pre} , or across a range of different applications represented by different functions d_q . One could also be conservative and compute instead an LPPM that maximizes the *minimum* privacy across a range of values, e.g., o_{pre} . In this paper, we take $\psi, d_p, o_{\text{pre}}, d_q, Q_{\text{loss}}^{\max}$ to be given inputs, but we note that it is not difficult to adapt the framework to accommodate alternatives.

5. EVALUATION

By formulating the LPPM design as an optimization problem, the LPPM algorithm $f(\cdot)$ is guaranteed to be optimal among all possible algorithms that respect the same constraints. Hence, there is no point in evaluating our design with simulations or any other heuristic evaluation method. We nevertheless compare to a sporadic LPPM to stress the importance of using a trajectory-aware LPPM. The sporadic LPPM that we use is the optimal one, as presented in prior work [22]. Due to incompatible assumptions, we cannot compare to the trajectory-aware LPPMs in the literature (see discussion of related work in Section 2). We also

show, for two illustrative scenarios, how our LPPM design allows to trade off privacy and quality. Finally, we discuss the run-time complexity of our design in Section 5.3.

For the comparison to the sporadic LPPM and for the illustration of the privacy-quality tradeoff, we use a real data set of location traces. These traces, which are one day long, belong to 10 randomly chosen mobile users (vehicles) in the San Francisco Bay area from the epfl/mobility dataset at CRAWDDAD [18]. These 10 examples serve to illustrate the optimality of the LPPMs designed by our method, since the technique is user-centric and does not need information about other users.

We discretize both time and location: we divide the Bay Area into 10×25 equal-size locations, and consider a day to be composed by 288 time units, one per each 5 minutes. We emphasize that the granularity of both time and locations can be arbitrarily selected depending on the required accuracy in quantifying privacy and service quality.¹ We consider all the locations that are visited by each user, which on average is 23.4 locations per user. We also consider all the transitions that each user has made between these locations in our dataset.

For both the comparison to the sporadic LPPM and for the privacy-quality tradeoff, we need to specify all the input parameters/functions $d_p, d_q, Q_{\text{loss}}^{\max}, o_{\text{pre}}, \psi$.

Without loss of generality, we select the privacy gain d_p and the quality loss d_q functions to be the Hamming distance: $d_p(\hat{a}_{\text{trg}}, a_{\text{trg}}) = 1_{\hat{a}_{\text{trg}} \neq a_{\text{trg}}}$ and $d_q(q_{\text{trg}}, o_{\text{post}}, o_{\text{pre}}) = 1_{q_{\text{trg}} \neq (o_{\text{post}}, o_{\text{pre}})}$. Using the Hamming distance means, taking the privacy gain as an example, that the only bad case for privacy is when the attacker correctly estimates the *exact* value of the target locations (i.e., when \hat{a}_{trg} is exactly equal to a_{trg}). All other estimates are equally good for privacy, regardless, e.g., of the physical distance between the attacker's estimate and the true value of a_{trg} . As our quantification of privacy is the *expected* value of $d_p(\hat{a}_{\text{trg}}, a_{\text{trg}})$ – and the expected value of $1_{\hat{a}_{\text{trg}} \neq a_{\text{trg}}}$ is just the probability of $\hat{a}_{\text{trg}} \neq a_{\text{trg}}$ – in effect we quantify privacy as the probability that the adversary will make an erroneous estimate.

For the maximum tolerable quality loss Q_{loss}^{\max} , we do not specify a single value, but rather compute the achievable privacy for multiple values, so as to observe the privacy-quality tradeoff.

For the previously reported events o_{pre} , we do not specify a single value. Instead, the privacy values that we compute and present in the following figures are averaged over all possible values of o_{pre} , because such an average is more representative of the privacy that a user can expect to achieve:

$$\sum_{o_{\text{pre}}} \Pr\{o_{\text{pre}}\} \text{Privacy}(\psi, f, h, d_p; o_{\text{pre}}).$$

We must stress that the actual values obtained in the evaluation depend on context (e.g., user mobility). In this sense, they amount to the “total” privacy of the user, not just the

¹Note that locations need not necessarily form a grid. In general, a higher number of locations corresponds to higher granularity, and, as a result, to more precise and more accurate quantification and protection of location privacy. The only requirement is for time and locations to be discrete. Of course, more locations and more time instants make the computation of the LPPM more demanding in resources. The run-time computation discussion in Section 5.3 elaborates on the effect of the number of time instants and locations on the complexity of the problem.




Figure 4: Users’ single-location privacy, using a sporadic LPPM against two attacks: sporadic attack vs. correlation-aware attack. For 10 different users (lines), and for various values of the service quality threshold Q_{loss}^{max} (dots), we see that the privacy against a correlation-aware attack (vertical axis) is always less than the privacy against a sporadic attack (horizontal axis).

privacy that is due to the LPPM alone. However, the evaluation shows that, other context being the same, our LPPMs achieve the highest “total” privacy value among all mechanisms.

To compute the prior probability $\psi(a_{trg})$ on the target events, we use the aforementioned traces to build a first-order Markov chain on the discretized set of locations. As mentioned in Section 3, choosing a Markov chain over a different mobility model is arbitrary, but our method can handle other mobility models as well.

In relation to ψ , notice that, in general, we need to specify the *conditional* prior on the target events $\psi(a_{trg}|o_{pre})$, i.e., the prior *given* the previously reported events o_{pre} (see Equation (1)). But from the traces we can only compute the *unconditional* prior $\psi(a_{trg})$. The connection between o_{pre} and a_{trg} will typically be given by whatever LPPM was in use when o_{pre} was reported, in conjunction with the unconditional prior. In [24], we show how this connection can be established, and how $\psi(a_{trg}|o_{pre})$ can be computed, for the joint protection of past and present privacy that we illustrate in this section.

5.1 Comparison to Optimal Sporadic LPPM

A trajectory-oblivious (sporadic) LPPM is typically evaluated against an attack that is also sporadic, i.e., an attack in which location correlation is not taken into account. To provide quantitative justification for the inadequacy of such LPPMs and their evaluation when the exposed locations are correlated, we show in Figure 4 that a correlation-aware attack can achieve much lower privacy than a sporadic attack.

Of course, a sporadic LPPM protects single locations only, so to compare meaningfully, we pick as objective of the correlation-aware attack the single-location privacy objective, i.e., $a_{trg} = r_t$ (see Section 4.2). The difference between the correlation-aware attack and the sporadic attack is that the former uses the conditional prior probability on the tar-




Figure 5: Privacy-quality tradeoff in the first scenario (single-location privacy): Protecting the current location ($a_{trg} = r_t$), when o_{pre} is the pseudolocation reported in the previous time instant o_{t-1} . Each curve corresponds to one user.

get location $\psi(r_t|o_{pre})$ (for $o_{pre} = o_{t-1}$), whereas the latter uses the unconditional prior $\psi(r_t)$.

Each attack is paired against the same sporadic LPPM (the optimal one [22], as mentioned earlier), and the results are plotted across the 10 mobile users and for various values of the service quality threshold Q_{loss}^{max} . As all data points are below the $x = y$ diagonal, we conclude that privacy in the correlation-aware attack (x -axis) is lower than privacy in the sporadic attack (y -axis). The only cases where the two attacks are equally (un-)successful are when the quality loss threshold is so high that the sporadic LPPM can inject enough noise to blur even the inference of a correlation-aware attack.

5.2 Privacy-Quality Tradeoff

In this section, we illustrate the privacy-quality tradeoff of our LPPMs for two particular scenarios: Protecting single-location privacy for the current location, taking into account the immediately previous pseudolocation ($a_{trg} = r_t$ and $o_{pre} = o_{t-1}$), shown in Figure 5; and protecting transition privacy for the current and future locations ($a_{trg} = (r_t, r_{t+1})$), as described in Section 4.3, shown in Figure 6.

Under each of these two scenarios, we construct the optimal protection mechanism for each of the 10 users in our traces (i.e., the mechanism that provides the maximum privacy for her). We plot this maximum privacy as a function of the service quality threshold Q_{loss}^{max} . We see in both figures that the achievable privacy increases as Q_{loss}^{max} increases. This is not surprising, as higher values of Q_{loss}^{max} let the LPPM inject more and more noise.

However, in both scenarios we observe two effects: First, a saturation effect takes place for most users as Q_{loss}^{max} increases. Their privacy reaches a plateau beyond which any further increase in Q_{loss}^{max} does not contribute to a corresponding increase in privacy. Second, the privacy plateau, as well as the privacy level for any value of Q_{loss}^{max} , is not the same for all users. This suggests that not all users can be protected




Figure 6: Privacy-quality tradeoff in the second scenario (transition privacy): Protecting the current and the next location ($a_{\text{trg}} = (r_t, r_{t+1})$). Each curve corresponds to one user.

equally well, but rather there is some inherent privacy limit for each user that is connected to the user’s mobility. Users with more predictable mobility cannot be protected as efficiently as less predictable ones, regardless of the amount of noise that the LPPM injects. Looking at the figures, more predictable users correspond to the lowest curves in Figures 5 and 6, whereas users with more random mobility correspond to the highest curves.

It is worth noting at this point that both these effects are not artefacts of our LPPM. Our LPPMs provide the best possible protection, so these effects are inherent in the mobility patterns of the users.

5.3 Computational Considerations

Our mechanism is intended to be computed offline and used online: The LPPM function $f(o_{\text{post}}|a_{\text{trg}}, o_{\text{pre}})$ is pre-computed offline and then downloaded to the device. Then, whenever the user attempts to expose a location, the LPPM looks up and performs the appropriate randomization on pseudolocations o_{post} , based on the actual values of the target events to be protected a_{trg} and the previously exposed vector of pseudolocations o_{pre} . In this way, the only computational burden of the resource-constrained mobile device is a look-up and a randomized selection of o_{post} .

The offline computation of the LPPM function f requires solving a separate linear program for each value of o_{pre} that may arise in practice. But most of the theoretically possible values of the vector o_{pre} are nonsensical sequences of locations, e.g., sequences where successive locations are too far away from each other, so these need not be taken into account, which saves considerable time. Similarly, the number of variables in each linear program is theoretically equal to the total number of pairs of a_{trg} and o_{post} vectors, since a value for f must be computed for each such combination. This number is $M^{\text{length}(a_{\text{trg}})+\text{length}(o_{\text{post}})}$ (recall that M is the total number of locations – see Section 3), but in practice it is much smaller. The actual number of linear programs and of variables is closer to the number of likely trajectories

of the corresponding length (the number of linear programs is equal to the number of trajectories of length $\text{length}(o_{\text{pre}})$, whereas the number of variables is equal to the number of trajectories of length $\text{length}(a_{\text{trg}}) + \text{length}(o_{\text{post}})$).

It is very important to notice also that the computation of f needs to be done only once, so the associated cost only needs to be incurred once. A recomputation of f is only necessary if, for example, the user parameters or application parameters $d_p, d_q, Q_{\text{loss}}^{\text{max}}$ change, or if the user wants to protect a different aspect of her privacy (e.g., previous, present, and next location, instead of just present and next location), which would translate to a change in a_{trg} , or if one wishes to take into account different prior knowledge of previously reported pseudolocations o_{pre} (e.g., take into account the 3 previously reported pseudolocations instead of just one).

6. CONCLUSIONS

Existing location privacy-preserving mechanisms either ignore the information leaked by the exposure of correlated locations, or ignore that the adversary will adapt his attack to the protection mechanism. Hence, in practice, these schemes do not provide the promised level of privacy. In this paper, we have proposed a framework that simultaneously considers correlation and the background knowledge of the adversary, namely the mobility profile of the user, the previously exposed locations, and the internal algorithm implemented by the protected mechanism; while at the same time respecting the user’s service quality requirements.

Our framework allows users to design LPPMs that protect not only her current location, but also her past and future whereabouts. Furthermore, our solution is the first to deal with protecting the privacy of transitions between locations, and with preserving the privacy of locations from which the user does not access the location based service. Two key advantages of the framework are that it is not limited to a particular scenario, but can be used to compute optimal defenses for different privacy and quality user preferences; and that it finds an optimal defense among a wide variety of conceivable mechanisms, effectively any mechanism that can be modeled as a probability distribution describing how obfuscated locations are produced from real locations.

Using real mobility traces, we show that users can relax their quality of service requirements in exchange for privacy, but the predictability of their movements determines the maximum protection they can obtain. The privacy level achieved by the LPPMs computed using our framework can be considered an upper bound on the privacy achievable by any defense in presence of a strategic adversary who knows the users’ mobility patterns. Hence, our solution is ideal to be used as benchmark to measure the effectiveness of future defenses.

Acknowledgments. Carmela Troncoso’s research is partially supported by EU FP7 projects PRIPARE (GA No 610613) and LIFTGATE (GA No 285901). George Theodorakopoulos’s research was partially supported by EU FP7 EINS (GA No 288201).

7. REFERENCES

- [1] O. Abul, F. Bonchi, and M. Nanni. Never walk alone: Uncertainty for anonymity in moving objects

- databases. In *24th International Conference on Data Engineering (ICDE 2008)*, pages 376–385. IEEE, 2008.
- [2] M. E. Andrés, N. E. Bordenabe, K. Chatzikokolakis, and C. Palamidessi. Geo-indistinguishability: Differential privacy for location-based systems. In *ACM Conference on Computer and Communications Security (CCS'13)*, pages 901–914. ACM, 2013.
- [3] C. A. Ardagna, G. Livraga, and P. Samarati. Protecting privacy of user information in continuous location-based services. In *15th International Conference on Computational Science and Engineering (CSE)*, pages 162–169. IEEE, 2012.
- [4] A. R. Beresford and F. Stajano. Location privacy in pervasive computing. *IEEE Pervasive Computing*, 2(1):46–55, 2003.
- [5] N. E. Bordenabe, K. Chatzikokolakis, and C. Palamidessi. Optimal geo-indistinguishable mechanisms for location privacy. *arXiv preprint arXiv:1402.5029*, 2014.
- [6] K. Chatzikokolakis, C. Palamidessi, and M. Stronati. A predictive differentially-private mechanism for mobility traces. In *Privacy Enhancing Technologies*, pages 21–41. Springer International Publishing, 2014.
- [7] R. Cheng, Y. Zhang, E. Bertino, and S. Prabhakar. Preserving user location privacy in mobile data management infrastructures. In *6th International Workshop on Privacy Enhancing Technologies (PET 2006)*, pages 393–412, 2006.
- [8] C.-Y. Chow and M. F. Mokbel. Trajectory privacy in location-based services and data publication. *SIGKDD Explorations Newsletter*, 13(1):19–29, June 2011.
- [9] M. L. Damiani, E. Bertino, and C. Silvestri. The PROBE framework for the personalized cloaking of private locations. *Transactions on Data Privacy*, 3(2):123–148, 2010.
- [10] J. Freudiger, R. Shokri, and J.-P. Hubaux. On the optimal placement of mix zones. In *PETS '09: Proceedings of the 9th International Symposium on Privacy Enhancing Technologies*, pages 216–234, 2009.
- [11] S. Gao, J. Ma, W. Shi, G. Zhan, and C. Sun. TrPF: A trajectory privacy-preserving framework for participatory sensing. *IEEE Transactions on Information Forensics and Security*, 8(6):874–887, June 2013.
- [12] G. Ghinita, M. L. Damiani, C. Silvestri, and E. Bertino. Preventing velocity-based linkage attacks in location-aware applications. In *17th ACM SIGSPATIAL International Symposium on Advances in Geographic Information Systems (ACM-GIS 2009)*, pages 246–255. ACM, 2009.
- [13] G. Gidófalvi, X. Huang, and T. B. Pedersen. Privacy-preserving data mining on moving object trajectories. In *8th International Conference on Mobile Data Management (MDM 2007)*, pages 60–68. IEEE, 2007.
- [14] B. Hoh, M. Gruteser, H. Xiong, and A. Alrabady. Achieving guaranteed anonymity in gps traces via uncertainty-aware path cloaking. *IEEE Transactions in Mobile Computing*, 9(8):1089–1107, 2010.
- [15] L. Huang, H. Yamane, K. Matsuura, and K. Sezaki. Silent cascade: Enhancing location privacy without communication QoS degradation. In *Security of Pervasive Computing (SPC)*, pages 165–180, 2006.
- [16] M. E. Nergiz, M. Atzori, Y. Saygin, and B. Güç. Towards trajectory anonymization: a generalization-based approach. *Transactions on Data Privacy*, 2(1):47–75, 2009.
- [17] X. Pan, X. Meng, and J. Xu. Distortion-based anonymity for continuous queries in location-based mobile services. In *17th ACM SIGSPATIAL International Symposium on Advances in Geographic Information Systems (ACM-GIS 2009)*, pages 256–265, 2009.
- [18] M. Piorkowski, N. Sarafijanovic-Djukic, and M. Grossglauer. CRAWDAD data set epl/mobility (v. 2009-02-24).
- [19] L. R. Rabiner. A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, 1989.
- [20] R. Shokri. Optimal user-centric data obfuscation. *arXiv preprint arXiv:1402.3426*, 2014.
- [21] R. Shokri, G. Theodorakopoulos, G. Danezis, J.-P. Hubaux, and J.-Y. Le Boudec. Quantifying location privacy: the case of sporadic location exposure. In *Proceedings of the 11th international conference on Privacy enhancing technologies (PETS'11)*, pages 57–76, 2011.
- [22] R. Shokri, G. Theodorakopoulos, C. Troncoso, J.-P. Hubaux, and J.-Y. Le Boudec. Protecting location privacy: optimal strategy against localization attacks. In *ACM Conference on Computer and Communications Security (CCS'12)*, pages 617–627, 2012.
- [23] M. Terrovitis and N. Mamoulis. Privacy preservation in the publication of trajectories. In *9th International Conference on Mobile Data Management (MDM '08)*, pages 65–72. IEEE, 2008.
- [24] G. Theodorakopoulos, R. Shokri, C. Troncoso, J.-P. Hubaux, and J.-Y. Le Boudec. Prolonging the hide-and-seek game: Optimal trajectory privacy for location-based services. *arXiv preprint*, 2014.
- [25] R. J. Vanderbei. *Linear programming*, volume 114. Springer, 2008.
- [26] T.-H. You, W.-C. Peng, and W.-C. Lee. Protecting moving trajectories with dummies. In *8th International Conference on Mobile Data Management (MDM '07)*, pages 278–282, May 2007.