Protecting Location Privacy: Optimal Strategy against Localization Attacks

Reza Shokri[†], George Theodorakopoulos[‡], Carmela Troncoso[*],
Jean-Pierre Hubaux[†], and Jean-Yves Le Boudec[†]

[†]LCA, EPFL, Lausanne, Switzerland,
[*]ESAT/COSIC, K.U.Leuven, Leuven-Heverlee, Belgium,
[‡]School of Computer Science and Informatics, Cardiff University, Cardiff, UK
[†]firstname.lastname@epfl.ch, [‡]g.theodorakopoulos@cs.cardiff.ac.uk,
[*]carmela.troncoso@esat.kuleuven.be

ABSTRACT

The mainstream approach to protecting the location-privacy of mobile users in location-based services (LBSs) is to alter the users' actual locations in order to reduce the location information exposed to the service provider. The location obfuscation algorithm behind an effective location-privacy preserving mechanism (LPPM) must consider three fundamental elements: the privacy requirements of the users, the adversary's knowledge and capabilities, and the maximal tolerated *service quality* degradation stemming from the obfuscation of true locations. We propose the first methodology, to the best of our knowledge, that enables a designer to find the *optimal* LPPM for a LBS given each user's service quality constraints against an adversary implementing the *optimal* inference algorithm. Such LPPM is the one that maximizes the expected distortion (error) that the optimal adversary incurs in reconstructing the actual location of a user, while fulfilling the user's service-quality requirement. We formalize the mutual optimization of user-adversary objectives (location privacy vs. correctness of localization) by using the framework of Stackelberg Bayesian games. In such setting, we develop two linear programs that output the best LPPM strategy and its corresponding optimal inference attack. Our optimal user-centric LPPM can be easily integrated in the users' mobile devices they use to access LBSs. We validate the efficacy of our game theoretic method against real location traces. Our evaluation confirms that the optimal LPPM strategy is superior to a straightforward obfuscation method, and that the optimal localization attack performs better compared to a Bayesian inference attack.

Categories and Subject Descriptors

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

CCS'12, October 16–18, 2012, Raleigh, North Carolina, USA.
Copyright 2012 ACM 978-1-4503-1651-4/12/10 ...$10.00.

1. INTRODUCTION

The widespread use of smart mobile devices with continuous connection to the Internet has fostered the development of a variety of successful location-based services (LBSs). Even though LBSs can be very useful, these benefits come at a cost of users' privacy. The whereabouts users' disclose to the service provider expose aspects of their private life that is not apparent at first, but can be inferred from the revealed location data [8, 11, 18].

A large body of research has focused on developing location-privacy protection mechanisms (LPPMs) that allow users to make use of LBSs while limiting the amount of disclosed sensitive information [1, 3, 9, 10, 14, 16, 22]. These protection mechanisms are based on hiding or perturbing the real locations of a user, or even sending fake locations to the LBS, in order to increase the uncertainty of the adversary about a user's true whereabouts. However, the evaluation of these designs usually disregards that the adversary might have some knowledge about the users' access pattern to the LBS and also about the algorithm implemented by the LPPM. Such information allows the attacker to reduce his uncertainty on the user's true location [25]. Hence, prior evaluations overestimate the location privacy offered by a given protection system.

In this paper, we focus on a broad range of LBSs and location sharing services in which users reveal their location in a *sporadic* manner, e.g., location check-in, location-tagging, or applications for finding nearby points-of-interests, local events, or nearby friends. We consider an adversary interested in uncovering the location of a user at the time when she sends the LBS query (i.e., an adversary performing localization attacks [25, 26]). We focus on *user-centric* LPPMs in which decision taken to protect privacy (e.g., hiding, perturbing, or faking the location) is made locally by the user. Hence, these LPPMs can be easily integrated in the mobile device that she uses to access LBS. We note, however, that the principles behind our protection mechanism design are applicable to LBSs where users reveal their location continuously (rather than sporadically), and where the adversary's aim is to track users continuously over space and time [26].

We propose an analytical framework that allows system designers to find the optimal LPPM against a strategic adversary who, knowing each user's LBS access pattern and the underlying obfuscation algorithm, employs the optimal attack to localize them. The challenge is to design such optimal protection mechanism when the inference attack, dependent on the mechanism being designed, is unknown to the designer. As opposed to making any assumption about the adversary's inference algorithm (i.e., limiting his power), we co-infer the optimal attack while finding the defense mechanism. Additionally, our methodology constrains the search space to LPPMs that obfuscate locations in such a way that the quality of the LBS response is not degraded below a threshold imposed by the user, hence guaranteeing required service quality for the user. We assume that the adversary is also aware of this user-specified service quality constraint.

We formalize the problem of finding the optimal LPPM anticipating the optimal inference attack as an instance of a zero-sum Bayesian Stackelberg game. In this game, a leader and a follower interact strategically with each one's gain being the loss of the other. The leader decides on her strategy knowing that it will be observed by the follower, who will optimize his choice based on this observation. In our scenario the user is the leader and the adversary is the follower. Then, this game precisely models that the adversary knows the user's choice of protection mechanism and will use that knowledge to improve his attack's effectiveness. We extend the classic formulation of a Stackelberg game with an extra constraint to ensure that the service quality is satisfactory for the user. This enables us to find the optimal point in the tradeoff curve between privacy and service quality that satisfies both user privacy and service quality requirements.

We build on the probabilistic model proposed by Shokri *et al.* [24, 25, 26] to find the best localization attack against a given LPPM and to measure the users' location privacy. This privacy measure is in turn used in the Stackelberg game to find the optimal LPPM for each user, i.e., the one that offers the best location privacy subject to the user's service quality requirements. Ours is, to the best of our knowledge, the first analytical framework that allows engineers to methodologically integrate adversarial knowledge in the design of optimal user-centric privacy protection mechanisms.

We evaluate the LPPMs generated by our method using real location traces. We show how, for a given user's LBS access pattern and service-quality threshold, our game-theoretic approach enables us to simultaneously find the optimal LPPM and the optimal attack against it. We confirm that there is a trade-off between the maximum achievable privacy and the service quality but once a certain privacy level is reached, loosening the quality requirements does not necessarily result in a privacy gain. We also find that the location-privacy gain of using the optimal LPPM, with respect to a suboptimal one, is larger when the quality constraint is tighter (compared to the case where users' quality requirements allow the LPPM to significantly perturb locations before sending them to the LBS).

The rest of the paper is organized as follows. We present the elements of our framework and describe the objectives of the user and adversary in the next section. We formalize the problem of finding an LPPM that offers optimal location privacy in terms of a Bayesian Stackelberg game in Section 3, and develop the best solution for both user and adversary in Section 4. We evaluate our method in Section 5 against real location traces. Section 6 revisits previous work on location privacy protection mechanisms, as well as on game theory applied to security-related scenarios. Finally, we conclude the paper in Section 7.

2. THE PROBLEM STATEMENT

In this section, we explain our system model based on the probabilistic framework proposed in [25, 26], as well as our assumptions and adversarial model. We conclude by sketching the problem this work aims at solving. In Table 1, we summarize the notations introduced throughout the section.

2.1 User and Adversary

We consider a scenario in which users move in an area partitioned into M discrete regions $\mathcal{R} = \{r_1, r_2, \cdots, r_M\}$. We also assume that time is discrete and it is partitioned into different time periods (e.g., morning, afternoon). We denote the spatiotemporal position of a user u at time t as the *actual event* $a_u(t) = \langle u, t, r\rangle$. We do not make any specific assumption about the users' mobility patterns. Users connect *sporadically* to an LBS provider to which they need to share their current location in order to obtain a service, i.e., there is a non-negligible time gap between two successive accesses a user to the LBS. The *access profile* $\psi_u^\tau(r)$ of user u is the probability distribution of the location r from which user u accesses the LBS in time period τ. For a given user u in a given time period τ, we have $\sum_{r\in\mathcal{R}} \psi_u^\tau(r) = 1$. We note that this profile is time-dependent (i.e., users may have different access patterns in the morning than in the afternoon). This dependency also affects users' location privacy requirements, and service quality requirements. For the sake of simplicity, in this paper, we omit the time-period τ and provide a solution for each user in a given time period. But, we note that the method is easily adaptable to more complex access patterns and privacy/quality requirements elicitation that account for such changes in time (e.g., by applying the method to each time period separately).

We assume that the LBS to which the user connects, or any entity that can eavesdrop on the user-LBS communications, is a passive and curious adversary whose aim is to discover the location of the user at the query time. As the LBS use is sporadic, the knowledge that the adversary can accumulate with repeated observations/eavesdropping is the frequency with which the user issues queries from regions in $\mathcal{R}$, i.e., $\psi_u(r)$. We assume that the adversary learns the user's profile $\psi_u(.)$ for example by using the algorithm explained in [26]. As we focus on user-centric mechanisms, which give protection to each user separately, in the remainder of the paper we omit the user identity u and present the model for this user with profile $\psi(.)$.

2.2 Location-Privacy Protection Mechanism

We consider that users want to preserve their location privacy when they use the LBS. Users implement a local and user-centric LPPM that transforms each true location r into a pseudolocation $r' \in \mathcal{R}'$, which is then sent to the LBS instead of the actual location. We set $R' = R$ (however, in the most general case R' is the powerset of R). The spatiotemporal position of a user as perceived by the LBS, denoted $o(t) = \langle t, r'\rangle$, is called an *observed event*. For each actual event $a(t) = \langle t, r\rangle$ the LPPM chooses a pseudolocation r' by sampling from the following probability distribution:

$$f(r'|r) = \Pr\left\{o(t) = \langle t, r'\rangle | a(t) = \langle t, r\rangle\right\} \qquad (1)$$

The adversary's knowledge is modeled as the user's access profile $\psi(.)$. As accesses to the LBS are sporadic, two successive query locations of the user are *conditionally* independent given $\psi(.)$. The larger the inter-query time is, the

Table 1: Summary of notations

Symbol	Meaning
u	Identity of the user
r, $\mathcal{R}$	Actual location of the user, set of possible locations for the user
$\psi(r)$	Location access profile of the user (probability of being at location r when accessing the LBS)
$a(t) = \langle t, r \rangle$	Actual location r of the user at time t
r', $\mathcal{R}'$	Pseudolocation output by the LPPM, set of possible pseudolocations output by the LPPM
$o(t) = \langle t, r' \rangle$	Observed pseudolocation r' of the user at time t
$f(r'\|r)$	Location obfuscation function implemented by the LPPM: Probability of replacing r with r'.
$d_q(r', r)$	Incurred service-quality loss by the user if LPPM replaces location r with pseudolocation r'
$Q_{loss}(\psi, f, d_q)$	Expected quality loss of an LPPM with location obfuscation function f
$Q_{loss}^{\max}$	Maximum tolerable service quality loss
$\hat{r}$	Adversary's estimate of the user's actual location
$h(\hat{r}\|r')$	Adversary's attack function: Probability of estimating $\hat{r}$ as user's actual location, if r' is observed
$d_p(\hat{r}, r)$	Distance between locations $\hat{r}$ and r: Privacy of the user at location r if adversary's estimate is $\hat{r}$
$Privacy(\psi, f, h, d_p)$	Expected location privacy of the user with profile $\psi(.)$ using protection f against attack h

more independent the two locations of the user in her successive LBS accesses are. This is also reflected in the LPPM's obfuscation algorithm that outputs pseudolocations that depend only on the user's current location.

2.3 Service Quality Metric

In the aforementioned setting, the LBS response quality depends on the pseudolocation output by the LPPM and not on the user's actual location. The distortion introduced in the observed pseudolocations determines the quality of service that the user experiences. The more similar the actual and the observed location are, the higher the service quality is. The expected *quality loss* due to an LPPM $f(.)$ is computed as an average of $d_q(r', r)$ over all r and r':

$$Q_{loss}(\psi, f, d_q) = \sum_{r,r'} \psi(r) f(r'|r) d_q(r', r). \tag{2}$$

Function $d_q(.)$ determines the dissimilarity between location r and pseudolocation r'. The semantics of this dissimilarity depend on the LBS under consideration, and also on the user's specific service-quality expectations. In many applications, the service quality can be considered inversely proportional to the physical distance between r and r'. For example, applications that find nearby points of interest could give very different responses to r and to r' even if they are only a couple of kilometers apart. In contrast, there exist LBSs in which the service quality depends on other criteria, such as on whether r' is within a region of interest. For a weather forecast application, for instance, any pseudolocation r' in the same city as the actual location r would result in a high quality LBS response.

We assume that users impose a maximum tolerable service quality loss, $Q_{loss}^{\max}$, caused by sharing pseudolocations instead of their actual locations. Formally,

$$Q_{loss}(\psi, f, d_q) \leq Q_{loss}^{\max}. \tag{3}$$

This constraints the LPPM obfuscation function $f(r'|r)$, that must not output pseudolocations that, on average, result in lower quality. We note that the influence of threshold $Q_{loss}^{\max}$ on the LPPM depends on the function $d_q(.)$, hence it is also dependent on the type of the LBS the user is querying. In the case of an LBS that finds nearby points of interest, where $d_q(.)$ is proportional to the physical distance between r and r', enforcing the quality threshold could result in ensuring a maximum allowed distance between these two locations. For the weather application, enforcing the quality threshold could result in setting region boundaries within which locations lead to the same forecast. For other location-based applications, the function $d_q(.)$ and the threshold $Q_{loss}^{\max}$ can be defined in the same vein.

2.4 Location Privacy Metric

The adversary's goal is to infer the user's actual events $a(t) = \langle t, r \rangle$ given the observed events $o(t) = \langle t, r' \rangle$. Recall that the adversary knows the user's profile, $\psi(.)$. He uses this background knowledge to run an inference attack on the observed events in order to output estimations $\hat{r}$ of the user's actual locations. Formally, the attack result can be described as a probability density function $h(.)$ such that

$$h(\hat{r}|r') = \Pr\left\{a(t) = \langle t, \hat{r} \rangle | o(t) = \langle t, r' \rangle\right\} . \tag{4}$$

As the adversary's prior information is the probability that the user is at a given location when she accesses the LBS, the current (query) location of the user is conditionally independent of her observed past and future locations. This is reflected in that the computation of the estimated location $\hat{r}$ at time t only depends on the pseudolocation r' observed at the same time t.

We note that the attack formulation is independent of whether the considered LPPM anonymizes the events or not. In this work, we assume that the adversary knows the identity of the users behind the events, but the framework can be adapted to anonymous LPPMs as well. Note that even when users are anonymous, our optimal solution provides a guarantee for their location privacy (even after a potential re-identification attack).

We follow the definition in [26] and quantify the user's location privacy as the adversary's expected error in his inference attack, i.e., the expected distortion in the reconstructed event. We compute the expectation over all r, r', and $\hat{r}$:

$$Privacy(\psi, f, h, d_p) = \sum_{\hat{r},r',r} \psi(r) f(r'|r) h(\hat{r}|r') d_p(\hat{r}, r) \tag{5}$$

The distortion function quantifies the loss of privacy stemming from the inference attack. The privacy loss depends on the locations' semantics and also on the privacy requirements of the user (i.e., users might consider locations inside a hospital more sensitive than other places), and $d_p(.)$ must be defined accordingly. For instance, if the user wants to hide just her exact current location (as opposed to hiding her location area), the appropriate distortion function could be the Hamming distance (probability of error) between the estimated location $\hat{r}$ and the actual location r:

$$d_p(\hat{r}, r) = \begin{cases} 0, & \text{if } \hat{r} = r \\ 1, & \text{otherwise} \end{cases} \tag{6}$$

In this case, any location different from the user's actual location results in a high level of location privacy. Alternatively, the user's privacy might depend on the physical distance between the estimated and actual locations, hence the distortion function can be modeled as the Euclidean distance between these locations, i.e., the squared-error distortion:

$$d_p(\hat{r}, r) = (\hat{r} - r)^2 \tag{7}$$

2.5 Problem Statement

Given

1. a maximum tolerable service-quality loss $Q_{loss}^{\max}$ imposed by the user as a bound for $Q_{loss}(.)$, computed using the quality function $d_q(.)$, and
2. a prior adversarial knowledge of the user's profile $\psi(.)$,

the problem is finding the LPPM obfuscation function $f(.)$ that maximizes the user's location privacy as defined in (5). The solution must consider that the adversary

1. observes the LPPM's output r', and
2. is aware of the LPPM's internal algorithm $f(.)$.

Hence, the adversary implements the *optimal* attack $h(.)$ that estimates the true location of the user with the least distortion as measured by $d_p(.)$.

3. GAME FORMULATION

The problem of finding an LPPM that offers optimal location privacy given the knowledge of the adversary is an instance of a zero-sum Bayesian Stackelberg game. In a Stackelberg game the *leader*, in our case the user, plays first by choosing an LPPM and committing to it by running it on her actual location. The *follower*, in our case the adversary, plays next by estimating the user's location, knowing the LPPM that the user has committed to. It is a Bayesian game because the adversary has incomplete information about the user's true location, and plays according to his hypothesis about this location. It is also an instance of a zero-sum game, as the adversary's gain (or loss) of utility is exactly balanced by the losses (or gains) of the utility of the user: the information gained (lost) by the adversary is the location privacy lost (gained) by the user. We now proceed to define the game adapted to our problem:

Step 0 Nature selects a location $r \in \mathcal{R}$ for the user to access the LBS, according to a probability distribution $\psi(.)$. That is, location r is selected with probability $\psi(r)$.

Step 1 Given r, the user runs the LPPM $f(r'|r)$ to select a pseudolocation $r' \in \mathcal{R}'$, subject to $f(.)$ complying with the service quality constraint (3).

Step 2 Having observed r', the adversary selects an estimated location $\hat{r} \sim h(\hat{r}|r'), \hat{r} \in \mathcal{R}$. The adversary knows the probability distribution $f(r'|r)$ used by the LPPM; he also knows the user's profile $\psi(.)$, but not the true location r.

Final Step The adversary pays an amount $d_p(\hat{r}, r)$ to the user. This amount represents the adversary's error (equivalently, the location privacy gained by the user).

The above description is common knowledge to both the adversary and the user. They both aim to maximize their payoff, i.e. the adversary tries to minimize the expected amount that he will pay, while the user tries to maximize it.

4. SOLUTION

In this section, we describe a precise optimization problem that formalizes the objectives of the user and of the adversary. We construct two linear programs that, given $\psi(.)$, $d_p(.)$ and $d_q(.)$, we can compute the user's optimal choice of protection mechanism $f(.)$, and the adversary's optimal choice of inference attack $h(.)$.

4.1 Optimal Strategy for the User

The adversary observes the pseudolocation r' output by the LPPM, he knows the function $f(r'|r)$ implemented by the LPPM, and he also knows the user's profile $\psi(.)$. Thus, he can form the posterior distribution

$$\Pr(r|r') = \frac{\Pr(r, r')}{\Pr(r')} = \frac{f(r'|r)\psi(r)}{\sum_r f(r'|r)\psi(r)} \tag{8}$$

on the true location r of the user, conditional on the observation r'. The adversary's objective is then to choose $\hat{r}$ to minimize the user's conditional expected privacy, where the expectation is taken under $\Pr(r|r')$. The user's conditional expected privacy for an arbitrary $\hat{r}$ is

$$\sum_r \Pr(r|r') d_p(\hat{r}, r), \tag{9}$$

and for the minimizing $\hat{r}$ it is

$$\min_{\hat{r}} \sum_r \Pr(r|r') d_p(\hat{r}, r). \tag{10}$$

If there are multiple values of $\hat{r}$ that satisfy (10), then the adversary randomizes arbitrarily among them. The probability with which $\hat{r}$ is chosen in this randomization is $h(\hat{r}|r')$. Of course, $h(\hat{r}|r')$ will be positive only for minimizing values of $\hat{r}$; for all other values $h(\hat{r}|r')$ will be zero. When randomizing, (10) is rewritten as

$$\sum_{r,\hat{r}} \Pr(r|r') h(\hat{r}|r') d_p(\hat{r}, r). \tag{11}$$

Note that if there is only one value of $\hat{r}$ satisfying (10), then this value is selected with probability 1 in the randomization, whereas all other values are selected with probability 0, so (11) reduces to (10). In this sense, (11) is a generalization of (10), but it should be noted that both expressions compute the same conditional expected privacy.

We see that for a given r', the user's conditional privacy is given by (10). The probability that r' is output by the LPPM is $\Pr(r') = \sum_r f(r'|r)\psi(r)$. Hence, the user's *unconditional* expected privacy (averaged over all r') is

$$\begin{aligned} &\sum_{r'} \Pr(r') \min_{\hat{r}} \sum_r \Pr(r|r') d_p(\hat{r}, r) \\ &= \sum_{r'} \min_{\hat{r}} \sum_r \psi(r) f(r'|r) d_p(\hat{r}, r). \end{aligned} \tag{12}$$

To facilitate the computations, we define

$$x_{r'} \triangleq \min_{\hat{r}} \sum_{r} \psi(r) f(r'|r) d_p(\hat{r}, r). \tag{13}$$

Incorporating $x_{r'}$ into (12), we rewrite the unconditional expected privacy of the user as

$$\sum_{r'} x_{r'}, \tag{14}$$

which the user aims to maximize by choosing the optimal $f(r'|r)$. The minimum operator makes the problem nonlinear, which is undesirable, but (13) can be transformed to a series of linear constraints:

$$x_{r'} \leq \sum_{r} \psi(r) f(r'|r) d_p(\hat{r}, r), \ \forall \hat{r}. \tag{15}$$

It turns out that maximizing (14) under (13) is equivalent to maximizing (14) under (15) [4, Ch. 7, p. 224].

We construct the linear program for the user from (14) and (15). Note that variable $x_{r'}$ is a *decision* variable in the linear program, i.e. it is among the quantities chosen by the solver. This might appear counterintuitive, as $x_{r'}$ is defined in (13) as a function of $f(.)$, rather than as an independent variable that can be freely selected. But, because of the transformation, it is always guaranteed that (13) will hold.

The linear program for the user is the following: Choose $f(r'|r), x_{r'}, \forall r, r'$ in order to

$$\textbf{Maximize} \ \sum_{r'} x_{r'} \tag{16}$$

subject to

$$x_{r'} \leq \sum_{r} \psi(r) f(r'|r) d_p(\hat{r}, r), \ \forall \hat{r}, r' \tag{17}$$

$$\sum_{r} \psi(r) \sum_{r'} f(r'|r) d_q(r', r) \leq Q_{loss}^{\max} \tag{18}$$

$$\sum_{r'} f(r'|r) = 1, \ \forall r \tag{19}$$

$$f(r'|r) \geq 0, \ \forall r, r' \tag{20}$$

Inequalities (17) are the series of linear constraints (15), one series for each value of r'; inequality (18) reflects the service quality constraint; constraints (19) and (20) reflect that $f(r'|r)$ is a probability distribution function.

4.2 Optimal Strategy for the Adversary

The reasoning is similar for the formalization of the adversary's optimization problem. When the LPPM's output is pseudolocation r', the adversary will solve (10) to find an estimate $\hat{r}$. More generally, the adversary will find many minimizing values of $\hat{r}$, and each of them will be selected with some probability $h(\hat{r}|r')$. Given that the true location is r and that the observed pseudolocation is r', the conditional expected user privacy is

$$\sum_{\hat{r}} h(\hat{r}|r') d_p(\hat{r}, r). \tag{21}$$

The user chooses r' to maximize (21). So, given that the true location is r, the conditional expected user privacy for the maximizing r' is

$$y_r \triangleq \max_{r'} \sum_{\hat{r}} h(\hat{r}|r') d_p(\hat{r}, r). \tag{22}$$

Similarly as before, the maximization can be generalized to a randomization among maximizing values of r'. The probability with which r' is chosen is $f(r'|r)$.

The prior distribution $\psi(r)$ contains the adversary's knowledge of r. Thus, the unconditional expected user privacy is

$$\sum_{r} \psi(r) \, y_r, \tag{23}$$

that the adversary aims to minimize by choosing $h(\hat{r}|r')$. Similarly as before, (22) can be transformed to an equivalent series of linear constraints:

$$y_r \geq \sum_{\hat{r}} h(\hat{r}|r') d_p(\hat{r}, r), \forall r'. \tag{24}$$

We construct the linear program for the adversary (which is the dual of the user's linear program) from (23) and (24): Choose $h(\hat{r}|r'), y_r, \forall r, r', \hat{r}$, and $z \in [0, \infty)$ in order to

$$\textbf{Minimize} \ \sum_{r} \psi(r) \, y_r + z Q_{loss}^{\max} \tag{25}$$

subject to

$$y_r \geq \sum_{\hat{r}} h(\hat{r}|r') d_p(\hat{r}, r) + z d_q(r', r), \forall r, r' \tag{26}$$

$$\sum_{\hat{r}} h(\hat{r}|r') = 1, \forall r' \tag{27}$$

$$h(\hat{r}|r') \geq 0, \forall r', \hat{r} \tag{28}$$

$$z \geq 0 \tag{29}$$

Note the role of variable z: In linear programming parlance, it is the *shadow price* of the service quality constraint. Intuitively, z is the "exchange rate" between service quality and privacy. Its value in the optimal solution indicates the amount of privacy (in privacy units) that is lost (gained) if the service quality threshold $Q_{loss}^{\max}$ increases (decreases) by one unit of quality.

For example, if $z > 0$ in the optimal solution, then any change $\Delta Q_{loss}^{\max}$ in $Q_{loss}^{\max}$ will affect the privacy achieved by $z \Delta Q_{loss}^{\max}$. In this case, constraint (18) is satisfied as a strict equality. In contrast, if constraint (18) is satisfied as a strict inequality, then, intuitively, the selection of $f(r'|r)$ has not been constrained by $Q_{loss}^{\max}$. In this case, any (small) changes in $Q_{loss}^{\max}$ will have no effect on $f(r'|r)$, nor on the privacy achieved. So, z would be zero.

Note that both linear programs compute the unconditional expected privacy of the user (5), which we repeat here for convenience.

$$Privacy(\psi, f, h, d_p) = \sum_{\hat{r}, r', r} \psi(r) f(r'|r) h(\hat{r}|r') d_p(\hat{r}, r). \tag{30}$$

Previous expressions can be derived from this one. For instance, if there is a single best choice of a pseudolocation r' for each given location r, then $f(r'|r)$ is always either 0 or 1, so (10) is obtained.

The optimal solution of each linear program results in the same value for the privacy of the user. Hence, in principle, we only need to compute one of the two to quantify maximum level of privacy of the user. We choose to present both, because the user's linear program incorporates the service quality constraint in a more straightforward manner, whereas the adversary's linear program explicitly computes the "exchange rate" between service quality and privacy.

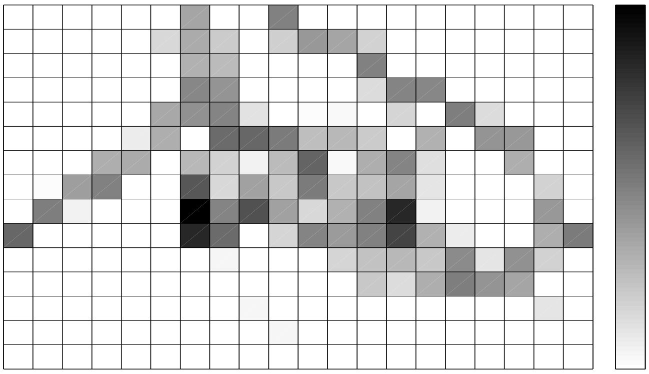

Figure 1: Spatial histogram showing the density of users per region (in log scale) in Lausanne. The area size is 15.32km $\times$ 7.58km, divided into 20×15 regions.

5. EVALUATION

The proposed optimization framework enables us to determine the most effective location-privacy protection mechanism (LPPM) against optimal inference attacks. The optimal LPPM is designed under the constraint of guaranteeing a minimum service quality such that the location-based service remains useful for the user. In this section, we evaluate the relation between location privacy and service quality for a few example location-based services (Recall that the service-quality sensitivity of a LBS to location obfuscation is encoded through the dissimilarity function $d_q(.)$). Moreover, we evaluate the performance of non-optimal LPPMs and non-optimal inference attacks against the optimal strategies.

We use real location traces of people (in Lausanne, Switzerland) who use various means of transportation.[1] We select 11 users at random, and we focus on their location traces during the day (8am to 8pm), when it is more probable that user use location-based services. The length of the considered traces is one month. The location area, within which they move, is divided into 300 regions. Figure 1 shows the density of users across all the regions. The grayness of the cells shows the density of its corresponding region in log scale. As many of the regions are abandoned (or very rarely visited) by many individual users, we compute each user's profile $\psi(.)$ by considering only the 30 most popular regions across the whole population. This prevents sparse user profiles. A user's profile is the normalized number of her visits to each region.

Given distance functions $d_p(.)$ and $d_q(.)$ and service-quality loss threshold $Q_{loss}^{\max}$, we compute the optimal LPPM and its corresponding optimal attack by solving (16) and (25) using Matlab's linear programming solver. We then compare the obtained optimal protection mechanism and the optimal inference attack against obfuscation LPPMs and Bayesian inference attacks, respectively.

Basic Obfuscation LPPM.

The basic obfuscation LPPM, with an obfuscation level $k = 1, 2, 3, \ldots$, is constructed in the following way: For each location r, we find its $k-1$ closest locations (using the Euclidean distance between the centers of the regions). The probability distribution function $f(.|r)$ will be the uniform probability distribution on the set of the $k-1$ selected locations together with the location r. That is, location r is replaced by each of the k locations, as a pseudolocation, with the same probability $\frac{1}{k}$, and all the rest of locations have probability 0. Thus, in practice, an actual location r is hidden among its $k-1$ nearest locations. We choose this mechanism, as it has been very popular in the literature.

Given the user profile $\psi(.)$ and quality distance function $d_q(.)$, we use (2) to compute the expected service-quality loss $Q_{loss}(\psi, f, d_q)$ for any LPPM obfuscation $f(.)$, whether it be optimal or not.

Bayesian Inference Attack on an LPPM.

We compare the effectiveness of our optimal attack with the Bayesian inference attack, which has been shown effective before in [26]. In the Bayesian approach, for each pseudolocation r', the posterior probability distribution over the locations is used to invert the noise added by the LPPM and, thus, to estimate the actual location:

$$h(\hat{r}|r') = \frac{\Pr(\hat{r}, r')}{\Pr(r')} = \frac{f(r'|\hat{r})\psi(\hat{r})}{\sum_r f(r'|r)\psi(r)} \tag{31}$$

We use (5) to compute the expected location privacy of a user who adopts a given (obfuscation or optimal) LPPM $f(.)$ against a (Bayesian or optimal) inference attack $h(.)$. The expected location privacy also depends on the distortion function $d_p(.)$ that we choose to use.

Briefly, if $d_p(.)$ is the Hamming distance, then the Bayesian attack chooses the location with the highest posterior probability $\Pr(\hat{r}|r')$. If $d_p(.)$ is the Euclidean distance, the Bayesian attack chooses the conditional expected value $\mathrm{E}[\hat{r}|r']$.

Optimal Inference Attack on an Arbitrary LPPM.

In order to make a fair comparison between the effectiveness of the optimal and obfuscation LPPM, we need to run the same attack on both of them. The Bayesian inference attack described by (31) can be performed against both. However, we still need to design an optimal attack against arbitrary LPPMs that have not been constructed in our game-theoretic framework.

The optimal inference attack is the one that minimizes the expected user privacy:

$$h(\hat{r}|r') = \arg\min_h Privacy(\psi, f, h, d_p). \tag{32}$$

Given the user profile $\psi(.)$, an LPPM $f(.)$ and distortion function $d_p(.)$, the following linear program finds the optimal attack $h(.)$. Note that, compared to (25), there is no service quality constraint here, as the LPPM has been assumed to be arbitrary.

$$\textbf{Minimize} \sum_{\hat{r}, r', r} \psi(r) f(r'|r) h(\hat{r}|r') d_p(\hat{r}, r) \tag{33}$$

$$\textbf{subject to} \sum_{\hat{r}} h(\hat{r}|r') = 1, \forall r', \text{and } h(\hat{r}|r') \geq 0, \forall \hat{r}, r' \tag{34}$$

Location-Privacy Protection Mechanism Output.

Consider a LBS user making use of our optimal LPPM on her mobile device. The way her location appears in the eyes of the adversary is shown in Figure 2. For the sake of comparison, Figure 2 also shows how a basic obfuscation LPPM distributes the pseudolocations over space. In order to make

[1] The traces are obtained from the Lausanne Data Collection Campaign dataset, http://research.nokia.com/page/11367

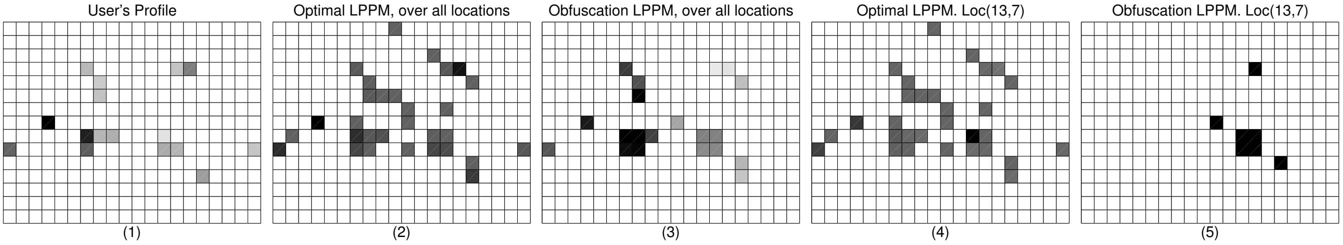

Figure 2: Input/Output of LPPM. Profile of a user for whom the subsequent calculations are made (sub-figure 1). Distribution $\Pr(r')$ of observed pseudolocations when using the optimal LPPM with $Q_{loss}^{\max} = 0.8690$ (sub-figure 2). Distribution $\Pr(r')$ of observed pseudolocations when using obfuscation LPPM with $Q_{loss}(\psi, f, d_q) = 0.8690$ (sub-figure 3). Conditional distribution $\Pr(r'|r)$ when using the optimal LPPM on location $r = (13, 7)$ (sub-figure 4). Conditional distribution $\Pr(r'|r)$ when using obfuscation LPPM on location $r = (13, 7)$ (sub-figure 5). Column 1 is the leftmost column, and row 1 is the bottom row. (Euclidean d_p, Hamming d_q)

a fair comparison, we need to make sure that the cost of the two LPPMs, in terms of service quality, is the same. To do so, we compute the quality loss Q_{loss} of the obfuscation LPPM and assign this loss as the quality threshold $Q_{loss}^{\max}$ of the optimal LPPM. Hence, the optimal LPPM cannot sacrifice the service quality more than the obfuscation LPPM to gain higher location privacy.

Figures 2(2) and 2(3) show $\Pr(r')$, the distribution of pseudolocations averaged over all locations for optimal and obfuscation LPPMs, respectively. Given arbitrary LPPM location obfuscation function $f(.)$ and user profile $\psi(.)$, the probability distribution of pseudolocations is

$$\Pr(r') = \sum_{r} \psi(r) f(r'|r). \tag{35}$$

As it is shown, the distribution corresponding to the optimal LPPM is more uniform, making it more difficult for the adversary to invert it effectively.

In Figures 2(4) and 2(5), we show the distribution of pseudolocations for specific location $r = loc(13, 7)$. By observing how uniform their outputs are, we can easily make the comparison between the two LPPMs. The obfuscation LPPM is obviously more concentrated around the actual location, whereas the optimal LPPM (with the same service-quality loss as the obfuscation method) broadens the set of pseudolocations to most of possible regions including highly probable regions (i.e. regions r with a large $\psi(r)$). This higher diversity brings higher privacy as we will see later in this section.

Tradeoff between Privacy and Service Quality.

We now study the tradeoff between the level of privacy that the optimal LPPM provides, against the optimal attack, and the service-quality loss that it causes. We plot in Figure 3(a) the evolution of the service quality loss, as the optimal LPPM is configured to guarantee different levels of service quality (for users with diverse profiles and for various service quality thresholds). Each line in the figure represents one user and each $\circ$ represents one $Q_{loss}^{\max}$. We plot $Privacy(\psi, f, h, d_p)$ versus $Q_{loss}(\psi, f, d_q)$.

Unsurprisingly, increasing the level of location-privacy protection significantly degrades the service quality. Also, as expected, we can observe that the maximum achievable location privacy is strongly dependent on the user profile. This is reflected by the separation between the different lines. Each user can have up to a certain level of privacy regardless of the quality threshold (represented by $\circ$ in the figure). Hence, the service-quality loss remains constant once this level has been reached. This is due to the presence of the optimal attack that squeezes the location-privacy gain.

This effect is further illustrated in Figure 3(b), where the service-quality loss of optimal LPPM is plotted against the service-quality threshold. Once the optimal LPPM offers the maximal location privacy for a given user profile, loosening the service-quality constraint does not significantly change the LPPM's underlying function f, and thus there is no reduction in service quality. In other words, there is no need to sacrifice the service quality, because doing so does not increase the user's location privacy.

Effectiveness of the Optimal Strategies.

Given Euclidean distance functions $d_p(.)$ and $d_q(.)$, we compute the optimal LPPM and attack methods for a set of service quality thresholds $Q_{loss}^{\max}$. For each user, we run the Bayesian inference attack on her optimal LPPM. We also evaluate the location privacy offered by the basic obfuscation LPPM with respect to the optimal attack. We vary the obfuscation level from 1 (minimum) to 30 (maximum), and for each case we compute the corresponding quality loss. Then, this value is set as the threshold $Q_{loss}^{\max}$ for finding the optimal attack mechanism.

Figure 4(a) shows the superiority of the optimal attack to the Bayesian attack, when location privacy of users is protected using the optimal LPPM: For any given user and service-quality threshold, the location privacy that the user obtains is smaller when the adversary implements the optimal strategy rather than the Bayesian inference attack.

Figure 4(b) shows the superiority of the optimal LPPM to the obfuscation LPPM, against the optimal attack: For any given user and service-quality threshold, a user has a higher privacy level when the LPPM implements the optimal strategy. As expected, obtained privacy by both mechanisms become equal when no service quality is guaranteed for the user (i.e., $Q_{loss}^{\max}$ is set to its maximum value).

Consider a single user. To further investigate the effectiveness of optimal strategies, we evaluate her privacy under four different combinations of optimal and non-optimal protection/attack methods, that have been explained before.

Similar to Figure 2, we consider the basic obfuscation LPPM as the basis for generating the service-quality thresh-

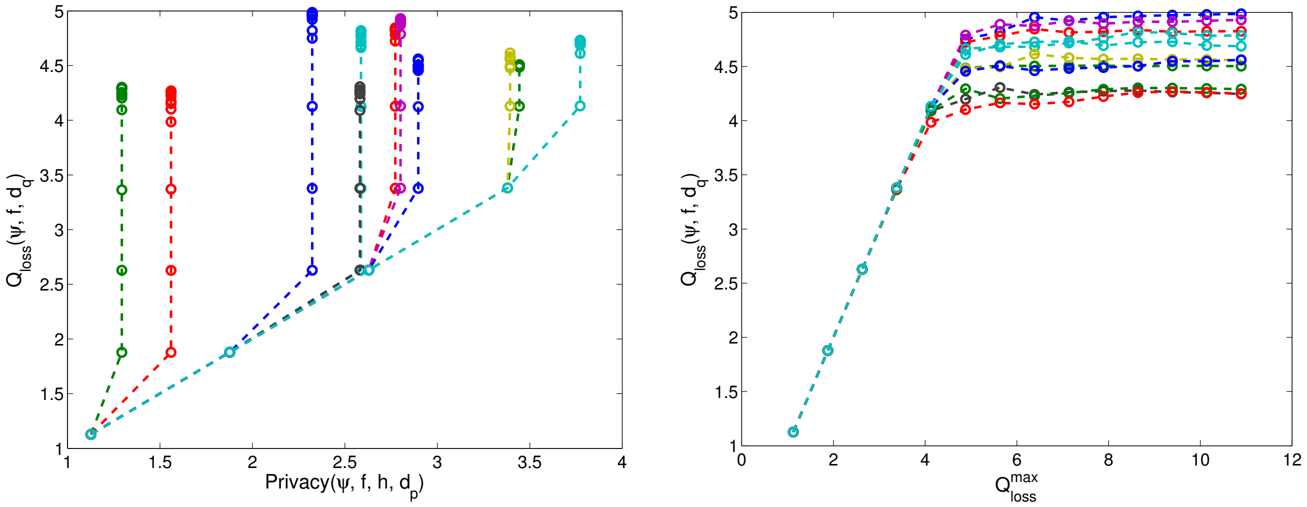

(a) Location privacy $Privacy(\psi, f, h, d_p)$ vs. Service-quality loss $Q_{loss}(\psi, f, d_q)$ for a given service-quality threshold Q_{loss}^{max}. The circles ○ represent different values of Q_{loss}^{max}.

(b) Service-quality threshold Q_{loss}^{max} vs. Service-quality loss $Q_{loss}(\psi, f, d_q)$, for a given level of location privacy $Privacy(\psi, f, h, d_p)$. The circles ○ represent different values of $Privacy(\psi, f, h, d_p)$.

Figure 3: Tradeoff between Privacy and Service Quality: Optimal LPPM against the optimal attack. The different lines represent users with diverse profiles $\psi(.)$. (Euclidean $d_q(.)$ and Euclidean $d_p(.)$.)

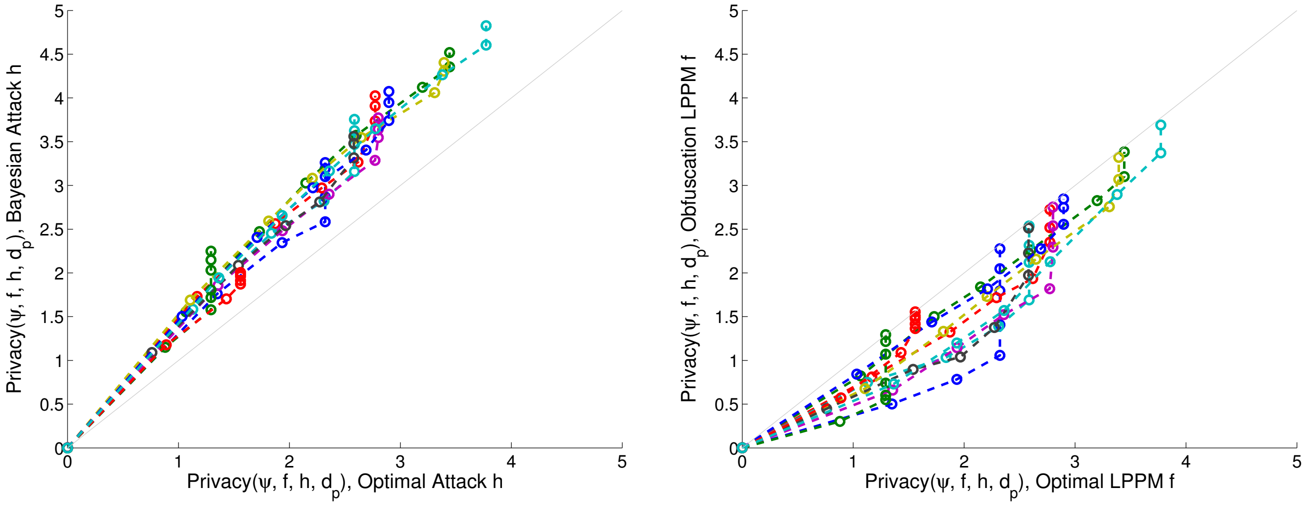

(a) Location privacy $Privacy(\psi, f, h, d_p)$ offered by the optimal LPPM against the optimal attack derived using the game theoretic approach vs. against the Bayesian-inference attack.

(b) Location privacy $Privacy(\psi, f, h, d_p)$ offered by the optimal LPPM vs. location privacy offered by the basic obfuscation LPPM, both evaluated against the optimal attack.

Figure 4: Effectiveness of the optimal attack and optimal LPPM strategies. Different lines represent users with diverse profiles $\psi(.)$, and the circles ○ represent different values of Q_{loss}^{max}. (Euclidean $d_q(.)$ and $d_p(.)$.)

old Q_{loss}^{max}. In all graphs of Figure 5 each dot represents one obfuscation level used in the basic obfuscation LPPM. The corresponding service-quality loss for each obfuscation level is shown on the x-axis of all four plots. As it can be easily observed from the figures, the optimal attack, compared with the Bayesian attack, always results in a higher degradation of the user's location privacy. Moreover, the optimal LPPM always provides a higher level of privacy for the user (regardless of the service-quality threshold) compared with the basic obfuscation LPPM.

The figures well illustrate how both user and adversary converge to use optimal strategies against each other. The user's favorite setting, i.e. the one that brings her a high level of privacy, is (Optimal, Bayesian). Inversely, the (Obfuscation, Optimal) combination is the favorite setting for the adversary, in which he pays the minimum cost of estimation-error. However, neither of these two settings is a stable state. In the (Optimal, Bayesian) combination, the adversary would gain more by choosing the Optimal attack. In the (Obfuscation, Optimal) combination, the user would gain

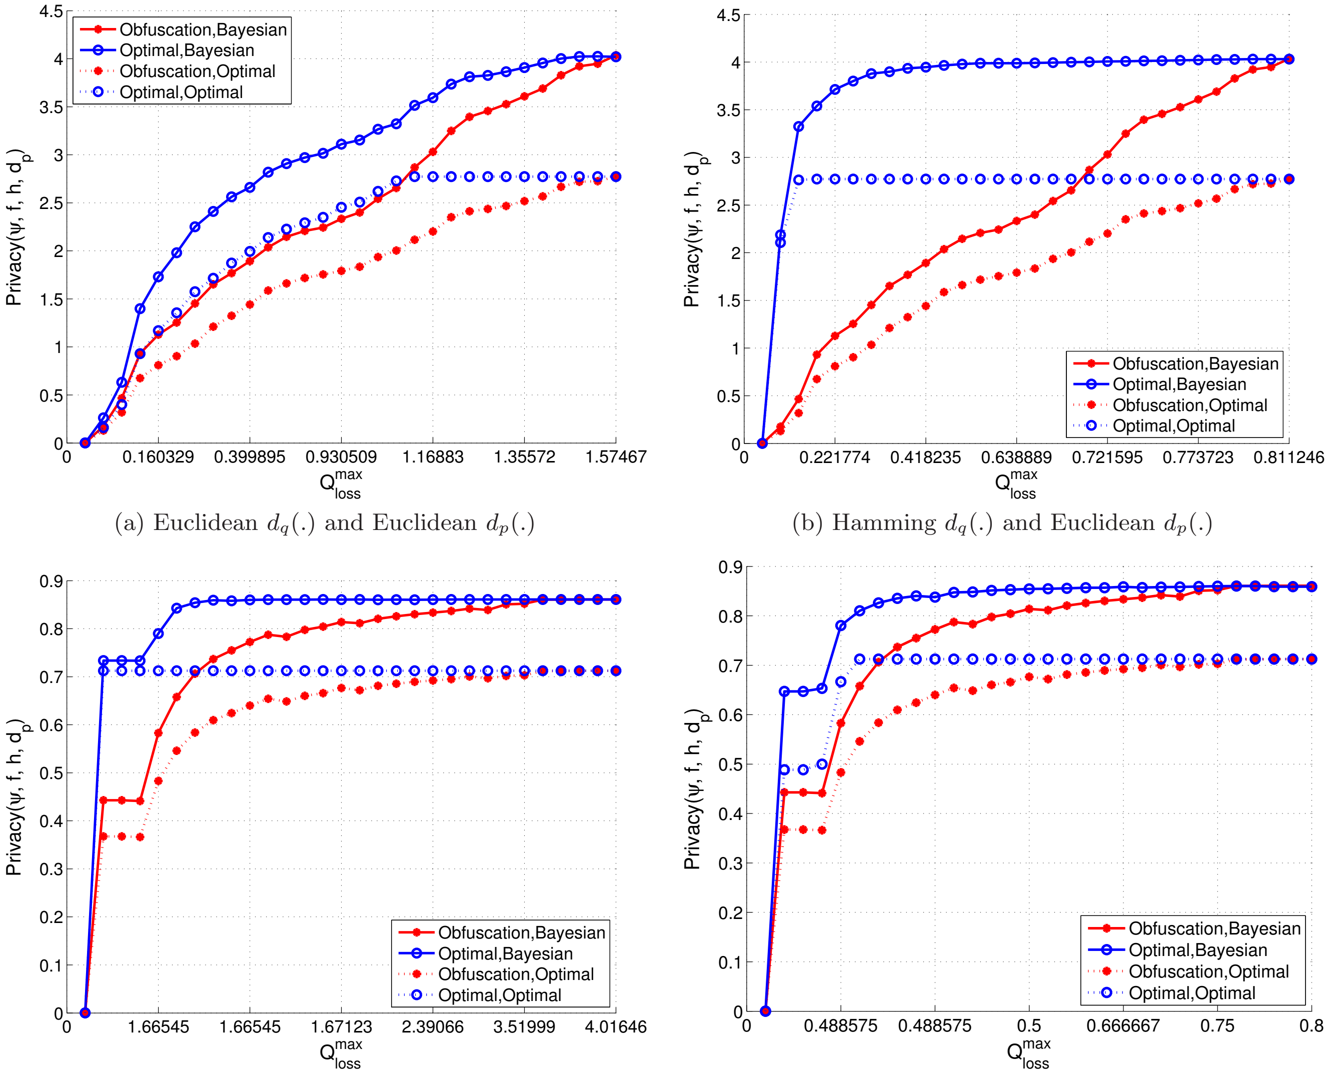

(a) Euclidean $d_q(.)$ and Euclidean $d_p(.)$

(b) Hamming $d_q(.)$ and Euclidean $d_p(.)$

(c) Euclidean $d_q(.)$ and Hamming $d_p(.)$

(d) Hamming $d_q(.)$ and Hamming $d_p(.)$

Figure 5: Service-quality threshold $Q_{loss}^{\max}$ vs. Location privacy $Privacy(\psi, f, h, d_p)$, for one single user. The different lines represent combinations of optimal ($\circ$) and basic obfuscation ($\bullet$) LPPMs tested against optimal ($\cdots$) and Bayesian-inference ($-$) attacks. The service-quality threshold $Q_{loss}^{\max}$ is equal to the service quality obtained by the basic obfuscation LPPM when the number of obfuscation levels used to perturb the location varies from 1 to 30 (its maximum value).

more by choosing the Optimal LPPM. Hence, the (Optimal, Optimal) combination is a stable equilibrium for both.

The fourth combination (Obfuscation, Bayesian) illustrates an interesting behavior. For small quality thresholds $Q_{loss}^{\max}$ (or, equivalently, smaller obfuscation levels) the user's privacy is lower compared with the (Optimal, Optimal) case. However, at some middle point its provided privacy increases and surpasses the privacy obtained from the optimal methods. Indeed, for small $Q_{loss}^{\max}$, the optimal LPPM uses all its available capacity to increase privacy by distributing the user's pseudolocations over a higher number of locations. So, it performs better than the basic obfuscation LPPM, which is limited to distributing pseudolocations only in a small set of regions. But when the obfuscation level (or, similarly, the service-quality threshold) increases, the basic obfuscation LPPM does better. First, because it is no longer severely limited, and, second, because it is paired against the Bayesian inference attack, which is weaker than the optimal inference attack.

6. RELATED WORK

The field of location privacy has been a very active area of research in recent years. Work on this topic can be roughly classified in three categories: mainly focused on the design of LPPMs; mainly focused on recovering actual user trajectories from anonymized or perturbed traces; or mainly focused on the formal analysis and the search for an appropriate location privacy metric that allows for the fair comparison between LPPMs.

Existing LPPMs are built according to different design principles. The most popular approach to obtaining location privacy is to send a space- or time-obfuscated version of the users' actual locations to the service provider [10, 12, 14, 16]. A different approach consists in hiding some of the users' locations by using mix zones [1, 9], or silent periods [15]. These are regions where users do not communicate with the provider while changing their pseudonym. Provided that several users traverse the zone simultaneously, this mechanism prevents an adversary from tracking them, as he cannot link those who enter with those who exit the region. A third line of work protects location privacy by adding dummy requests, indistinguishable from real requests, issued from fake locations to the service provider [3]. The purpose of these fake locations is to increase the uncertainty of the adversary about the users' real movements.

A number of papers show that the predictability of users' location traces, and the particular constraints of users' movements, are sufficient to reconstruct and/or identify anonymous or perturbed locations. For instance, an adversary can, to name but a few possibilities, infer users' activities from the frequency of their visits to certain locations [19]; re-identify anonymous low-granularity location traces given the users' mobility profiles [5]; or derive [13], and re-identify [11, 18] the home address of individuals from location traces.

Several authors have made efforts towards formalizing the desirable location privacy requirements that LPPMs should fulfill, as well as towards finding suitable metrics to evaluate the degree to which these requirements are fulfilled. Examples of these lines of work are Krumm [18], Decker [6], and Duckham [7]. Shokri *et al.* [24] revisit existing LPPMs and the location-privacy metrics used in their evaluation. They classify these metrics in three categories: uncertainty-based (entropy), error-based and k-anonymity. The authors conclude, by means of a qualitative evaluation, that metrics such as entropy and k-anonymity are not suitable for measuring location privacy. In a follow-up of this work, Shokri *et al.* provide a framework [25, 26] to quantify location privacy. The framework allows us to specify an LPPM and then to evaluate various questions about the location information leaked. Our design methodology uses this analytical framework as an evaluation tool to quantifying the LPPMs' offered privacy against the localization attack.

Despite the extent to which location privacy has been studied, there is a patent disconnection between these different lines of work. Most of the aforementioned papers use different models to state the problem and evaluate location privacy. This hinders the comparison of systems and slows down the design of robust LPPMs. Further, in some of these papers there is a detachment between the proposed design and the adversarial model against which it is evaluated. Often the considered adversary is static in its knowledge and disregards the information leaked by the LPPM algorithm; or adversarial knowledge is not even considered in the evaluation. The works by Freudiger *et al.* [9] and Shokri *et al.* [24, 25, 26] do consider an strategic adversary that exploits the information leaked by the LPPM in order to compute location privacy. Nevertheless, their work, which we build on in this paper, does not address how this privacy computation can be integrated in the design of location-privacy preserving mechanisms.

In this work, we bridge the gap between design and evaluation of LPPMs. We provide a systematic method for developing LPPMs; it maximize users' location privacy while guaranteeing a minimum service quality. We formalize the optimal design problem as a Bayesian Stackelberg game similar to previous work on security in which, as in the location-privacy scenario, the defender can be modeled as a Stackelberg game leader, and the adversary as the follower. The common theme is that the defender must commit to a defense strategy/protocol, which is then disclosed to the adversary, who can then choose an optimal course of action *after* observing the defender's strategy. Paruchuri *et al.* [23] propose an efficient algorithm for finding the leader's optimal strategy considering as a main case study a patrolling agent who searches for a robber in a limited area. In their case, the defender is unsure about the type of the adversary (i.e. where the adversary will attack). In contrast, in our work it is the adversary who is unsure about the type (i.e. the true location) of the user/defender. A similar approach is used by Liu and Chawla [20] in the design of an optimal e-mail spam filter, taking into account that spammers adapt their e-mails to get past the spam detectors. The same problem is tackled by Brückner and Scheffer [2], who further compare the Stackelberg-based approach with previous spam filters based on support vector machines, logistic regression, and Nash-logistic regression. Korzhyk *et al.* [17] contrast the Stackelberg framework with the more traditional Nash framework, within a class of security games. A recent survey [21] explores the connections between security and game theory more generally. To the best of our knowledge, our work is the first that uses Bayesian Stackelberg games to design optimal privacy-protection mechanisms.

7. CONCLUSION

Accessing location-based services from mobile devices entails a privacy risk for users whose sensitive information can be inferred from the locations they visit. This information leakage raises the need for robust location-privacy protecting mechanisms (LPPMs). In this paper, we have proposed a game-theoretic framework that enables a designer to find the optimal LPPM for a given location-based service, ensuring a satisfactory service quality for the user. This LPPM is designed to provide user-centric location privacy, hence it is ideal to be implemented in the users' mobile devices.

Our method accounts for the fact that the strongest adversary not only observes the perturbed location sent by the user but also knows the algorithm implemented by the protection mechanism. Hence, he can exploit the information leaked by the LPPM's algorithm to reduce his uncertainty about the user's true location. However, the user is only aware of the adversary's knowledge and does not make any assumption about his inference attack. Hence, she prepares the protection mechanism against the most optimal attack. By modeling the problem as a Bayesian Stackelberg competition, we ensure that the optimal LPPM is designed anticipating such strong inference attack.

We have validated our method using real location traces. We have demonstrated that our approach finds the optimal attack for a given LPPM and service-quality constraint, and we have shown that it is superior to other LPPMs such as basic location obfuscation. We have also shown that the superiority of the optimal LPPM over alternatives is more significant when the service-quality constraint imposed by the user is tightened. Hence, our solution is effective exactly where it will be used. Finally, our results confirm that loosening the service-quality constraint allows for increased privacy protection, but the magnitude of this increase strongly depends on the user profile, i.e., on the degree to which a user's location is predictable from her LBS access profile.

To the best of our knowledge, this is the first framework that explicitly includes the adversarial knowledge into a privacy-preserving design process, and considers the *common knowledge* between the privacy protector and the attacker. Our obtained result is a promising step forward in the quest for robust and effective privacy preserving systems.

8. REFERENCES

[1] A. R. Beresford and F. Stajano. Location privacy in pervasive computing. *IEEE Pervasive Computing*, 2(1):46–55, 2003.

[2] M. Brückner and T. Scheffer. Stackelberg games for adversarial prediction problems. In C. Apté, J. Ghosh, and P. Smyth, editors, *17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 2011.

[3] R. Chow and P. Golle. Faking contextual data for fun, profit, and privacy. In *WPES '09: Proceedings of the 8th ACM workshop on Privacy in the electronic society*, New York, NY, USA, 2009.

[4] S. Dasgupta, C. Papadimitriou, and U. Vazirani. *Algorithms*. McGraw-Hill, New York, NY, 2008.

[5] Y. De Mulder, G. Danezis, L. Batina, and B. Preneel. Identification via location-profiling in gsm networks. In *WPES '08: Proceedings of the 7th ACM workshop on Privacy in the electronic society*, New York, NY, USA, 2008.

[6] M. Decker. Location privacy - an overview. In *International Conference on Mobile Business*, 2009.

[7] M. Duckham. Moving forward: location privacy and location awareness. In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on Security and Privacy in GIS and LBS*, New York, NY, USA, 2010.

[8] J. Freudiger, R. Shokri, , and J.-P. Hubaux. Evaluating the privacy risk of location-based services. In *Financial Cryptography and Data Security (FC)*, 2011.

[9] 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*, Berlin, Heidelberg, 2009.

[10] B. Gedik and L. Liu. Location privacy in mobile systems: A personalized anonymization model. In *ICDCS '05: Proceedings of the 25th IEEE International Conference on Distributed Computing Systems*, Washington, DC, USA, 2005.

[11] P. Golle and K. Partridge. On the anonymity of home/work location pairs. In *Pervasive '09: Proceedings of the 7th International Conference on Pervasive Computing*, Berlin, Heidelberg, 2009.

[12] M. Gruteser and D. Grunwald. Anonymous usage of location-based services through spatial and temporal cloaking. In *MobiSys '03: Proceedings of the 1st international conference on Mobile systems, applications and services*, New York, NY, USA, 2003.

[13] B. Hoh, M. Gruteser, H. Xiong, and A. Alrabady. Enhancing security and privacy in traffic-monitoring systems. *IEEE Pervasive Computing*, 5(4):38–46, 2006.

[14] B. Hoh, M. Gruteser, H. Xiong, and A. Alrabady. Preserving privacy in gps traces via uncertainty-aware path cloaking. In *CCS '07: Proceedings of the 14th ACM conference on Computer and communications security*, New York, NY, USA, 2007.

[15] T. Jiang, H. J. Wang, and Y.-C. Hu. Preserving location privacy in wireless lans. In *MobiSys '07: Proceedings of the 5th international conference on Mobile systems, applications and services*, New York, NY, USA, 2007.

[16] P. Kalnis, G. Ghinita, K. Mouratidis, and D. Papadias. Preventing location-based identity inference in anonymous spatial queries. *Knowledge and Data Engineering, IEEE Transactions on*, 19(12):1719–1733, Dec. 2007.

[17] D. Korzhyk, Z. Yin, C. Kiekintveld, V. Conitzer, and M. Tambe. Stackelberg vs. Nash in security games: An extended investigation of interchangeability, equivalence, and uniqueness. *Journal of Artificial Intelligence Research*, 41:297–327, May–August 2011.

[18] J. Krumm. Inference attacks on location tracks. In *In Proceedings of the Fifth International Conference on Pervasive Computing (Pervasive)*, 2007.

[19] L. Liao, D. J. Patterson, D. Fox, and H. A. Kautz. Learning and inferring transportation routines. *Artif. Intell.*, 171(5-6):311–331, 2007.

[20] W. Liu and S. Chawla. A game theoretical model for adversarial learning. In Y. Saygin, J. X. Yu, H. Kargupta, W. Wang, S. Ranka, P. S. Yu, and X. Wu, editors, *IEEE International Conference on Data Mining Workshops (ICDM 2009)*, 2009.

[21] M. Manshaei, Q. Zhu, T. Alpcan, T. Basar, and J.-P. Hubaux. Game theory meets network security and privacy. *ACM Computing Surveys*, 2011.

[22] J. Meyerowitz and R. Roy Choudhury. Hiding stars with fireworks: location privacy through camouflage. In *MobiCom '09: Proceedings of the 15th annual international conference on Mobile computing and networking*, New York, NY, USA, 2009.

[23] P. Paruchuri, J. P. Pearce, J. Marecki, M. Tambe, F. Ordóñez, and S. Kraus. Efficient algorithms to solve Bayesian Stackelberg games for security applications. In D. Fox and C. P. Gomes, editors, *23rd AAAI Conference on Artificial Intelligence (AAAI 2008)*, 2008.

[24] R. Shokri, J. Freudiger, M. Jadliwala, and J.-P. Hubaux. A distortion-based metric for location privacy. In *WPES '09: Proceedings of the 8th ACM workshop on Privacy in the electronic society*, New York, NY, USA, 2009.

[25] 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)*, Berlin, Heidelberg, 2011.

[26] R. Shokri, G. Theodorakopoulos, J.-Y. Le Boudec, and J.-P. Hubaux. Quantifying location privacy. In *IEEE Symposium on Security and Privacy*, Oakland, CA, USA, 2011.