

Reza Shokri*

Privacy Games: Optimal User-Centric Data Obfuscation

Abstract: Consider users who share their data (e.g., location) with an untrusted service provider to obtain a personalized (e.g., location-based) service. Data obfuscation is a prevalent user-centric approach to protecting users' privacy in such systems: the untrusted entity only receives a noisy version of user's data. Perturbing data before sharing it, however, comes at the price of the users' utility (service quality) experience which is an inseparable design factor of obfuscation mechanisms. The entanglement of the utility loss and the privacy guarantee, in addition to the lack of a comprehensive notion of privacy, have led to the design of obfuscation mechanisms that are either suboptimal in terms of their utility loss, or ignore the user's information leakage in the past, or are limited to very specific notions of privacy which e.g., do not protect against adaptive inference attacks or the adversary with arbitrary background knowledge.

In this paper, we design user-centric obfuscation mechanisms that impose the minimum utility loss for guaranteeing user's privacy. We optimize utility subject to a joint guarantee of differential privacy (indistinguishability) and distortion privacy (inference error). This double shield of protection limits the information leakage through obfuscation mechanism as well as the posterior inference. We show that the privacy achieved through joint differential-distortion mechanisms against optimal attacks is as large as the maximum privacy that can be achieved by either of these mechanisms separately. Their utility cost is also not larger than what either of the differential or distortion mechanisms imposes. We model the optimization problem as a leader-follower game between the designer of obfuscation mechanism and the potential adversary, and design adaptive mechanisms that anticipate and protect against optimal inference algorithms. Thus, the obfuscation mechanism is optimal against any inference algorithm.

Keywords: Data Privacy, Obfuscation, Utility, Differential Privacy, Distortion Privacy, Inference Attack, Prior Knowledge, Optimization, Game Theory

DOI 10.1515/popets-2015-0024

Received 2015-02-15; revised 2015-05-10; accepted 2015-05-15.

*Corresponding Author: Reza Shokri: University of Texas at Austin, E-mail: shokri@cs.utexas.edu

1 Introduction

Data obfuscation is a mechanism for hiding private data by using misleading, false, or ambiguous information with the intention of confusing an adversary [14]. A data obfuscation mechanism acts as a noisy information channel between a user's private data (secret) and an untrusted observer [16]. The noisier this channel is, the higher the privacy of the user will be. We focus on *user-centric* mechanisms, in which each user independently perturbs her secret before releasing it. Note that we are not concerned with database privacy, but with the privacy issues of releasing a single sensitive data sample (which however could be continuously shared over time). For example, consider a mobile user who is concerned about the information leakage through her location-based queries. In this case, obfuscation is the process of randomizing true locations so that the location-based server only receives the user's perturbed locations.

By using obfuscation mechanisms, the *privacy* of a user and her *utility* experience are at odds with each other, as the service that the user receives is a function of what she shares with the service provider. There are problems to be addressed here. One is how to design an obfuscation mechanism that protects privacy of the user and imposes a *minimum* utility cost. Another problem is how to *guarantee* the user's privacy, despite the lack of a single best metric for privacy.

Regarding utility optimization, we define utility loss of obfuscation as the degradation of the user's service-quality expectation due to sharing the noisy data instead of its true value. Regarding privacy protection, there are two major metrics proposed in the literature. *Differential* privacy limits the information leakage through observation. But, it does not reflect the absolute privacy level of the user, i.e., what actually is learned about the user's secret. So, user would not know how close the adversary's estimate will get to her secret if she releases the noisy data, despite being sure that the relative gain of observation for adversary is bounded. *Distortion* privacy (inference error) metric overcomes this issue and measures the error of inferring user's secret from the observation. This requires assumption of a prior knowledge which enables us to quantify absolute privacy, but is not robust to adversaries with arbitrary

knowledge. Thus, either of these metrics alone is incapable of capturing privacy as a whole.

The problem of optimizing the tradeoff between privacy and utility has already been discussed in the literature, but notably for differential privacy in the context of statistical databases [11, 25, 26, 28, 36]. Regarding user-centric obfuscation mechanisms, [52] solves the problem of maximizing distortion privacy under a constraint on utility loss. The authors construct the optimal adaptive obfuscation mechanism as the user’s best response to the adversary’s optimal inference in a Bayesian zero-sum game. In the same context, [9] solves the opposite problem, i.e., optimizing utility but for differential privacy. In both papers, the authors construct the optimal solutions using linear programming.

Differential and distortion metrics for privacy complement each other. The former is sensitive to the likelihood of observation given data. The latter is sensitive to the joint probability of observation and data. Thus, by guaranteeing both, we encompass all the defense that is theoretically possible. In this paper, we model and solve the optimal obfuscation mechanism that: (i) minimizes utility loss, (ii) satisfies differential privacy, and (iii) guarantees distortion privacy, given a public knowledge on prior leakage about the secrets. We measure the involved metrics based on separate distance functions defined on the set of secrets. We model prior leakage as a probability distribution over secrets, that can be estimated from the user’s previously released data. Ignoring such information leads to overestimating the user’s privacy and thus designing a weak obfuscation mechanism (against adversaries who include such exposed information in their inference attack).¹

A protection mechanism for distortion privacy metric can be designed such that it is optimal against a particular inference algorithm (e.g., Bayesian inference [7, 38] as privacy attacks [51, 54]). But, by doing so, it is not guaranteed that the promised privacy level can be achieved in practice: an adversarial observer can run inference attacks that are optimally tailored against the very obfuscation mechanism used by the user (regardless

of the algorithm that the user assumes a priori). In fact, the adversary has the upper hand as he infers the user’s secret (private information) *after* observing the output of the obfuscation mechanism. Thus, the obfuscation mechanisms must *anticipate* the adaptive inference attack that will follow the observation. This enables us to design an obfuscation mechanism that is independent of the adversary’s inference algorithm.

To address this concern, we adapt a game-theoretic notion of privacy for designing optimal obfuscation mechanisms against adaptive inference. We formulate this game as a Stackelberg game and solve it using linear programming.² We then add the differential privacy guarantee as a constraint in the linear program and solve it to construct the optimal mechanism. The result of using such obfuscation mechanism is that, not only the perturbed data samples are indistinguishable from the true secret (due to differential privacy bound), but also they cannot be used to accurately infer the secret using the prior leakage (due to distortion privacy measure). To the best of our knowledge, this work is the first to construct utility maximizing obfuscation mechanisms with such formal privacy guarantees.

We illustrate the application of optimal protection mechanisms on a real data set of users’ locations, where users want to protect their location privacy against location-based services. We evaluate the effects of privacy guarantees on utility cost. We also analyze the robustness of our optimal obfuscation mechanism against inference attacks with different algorithms and background knowledge. We show that our joint differential-distortion mechanisms are robust against adversaries with optimal attack and background knowledge. Moreover, the utility loss is at most equal to the utility loss of differential or distortion privacy, separately.

The novelty of this paper in the context of user-centric obfuscation is twofold:

- We construct optimal obfuscation mechanisms that provably limit the user’s privacy risk (i.e., by guaranteeing the user’s distortion privacy) against *any* inference attack, with minimum utility cost.
- We design obfuscation mechanisms that optimally balance the tradeoff between utility and joint distortion-differential privacy. The solution is robust against adversary with arbitrary knowledge, yet it guarantees a required privacy given the user’s estimation of the prior information leakage.

¹ Note that the prior leakage is *not* equivalent to the adversary’s knowledge. An adversary might have access to some information about the user’s data through channels where the user is unaware of and has no control over. No protection mechanism can guarantee the distortion privacy against adversaries with arbitrary knowledge. Imagine the worst case where adversary knows the exact secret but through channels other than observation of the obfuscation mechanism. Therefore, our focus is on the user, and we incorporate what the user thinks has been leaked so far.

² As opposed to [52], the game is not zero-sum anymore given that here user maximizes utility and observer minimizes privacy.

2 Related Work

This paper contributes to the broad area of research that concerns designing obfuscation mechanisms, e.g., in the context of quantitative information flow [34, 40, 41], quantitative privacy in data sharing systems [4, 51, 52], as well as differential privacy [21, 26, 28, 36]. The conflict between privacy and utility has been discussed in the literature [12, 31]. We build upon prevalent notions of privacy and protect it with respect to information leakage through both observation (differential privacy) and posterior inference (distortion privacy) while optimizing the tradeoff between utility and privacy. We also formalize this problem and solve it for user-centric obfuscation mechanisms, where it's each individual user who perturbs her secret data before sharing it with external observers (e.g., service providers).

The problem of perturbing data for differential and distortion privacy, separately, and optimizing their effect on utility has already been discussed in the literature. Original metric for differential privacy measures privacy of output perturbation methods in statistical databases [21]. Assuming two statistical databases to be neighbor if they differ only in one entry, [25] and [26] design utility maximizing perturbation mechanisms for the case of counting queries. In [28, 36], authors propose different approaches to designing perturbation mechanisms for counting queries under differential privacy. However, [11] presents some impossibility results of extending these approaches to other types of database queries. Under some assumptions about the utility metric, [24] shows that the optimal perturbation probability distribution has a symmetric staircase-shaped probability density function. [6, 15, 49] extend differential privacy metric using generic distance functions on the set of secrets. Some extensions of differential privacy also consider the problem of incorporating the prior knowledge into its privacy definition [29, 32].

The most related paper to our framework, in this domain, is [9] where the authors construct utility-maximizing differentially private obfuscation mechanisms using linear programming. The authors prove an interesting relation between utility-maximizing differential privacy and distortion-privacy-maximizing mechanisms that bound utility, when distance functions used in utility and privacy metrics are the same. This, however, cannot guarantee distortion privacy for general metrics. The optimal differentially private mechanisms, in general, do not incorporate the available knowledge about the secret while achieving differential privacy.

Distortion privacy, which evaluates privacy as the inference error [51], is a follow-up of information-theoretic metrics for anonymity and information leakage [16, 20, 34, 50]. This class of metrics is concerned with what can be inferred about the true secret of the user by combining the observation (of obfuscated information) and prior knowledge. The problem of maximizing privacy under utility constraint, assuming a prior, is proven to be equivalent to the user's best strategy in a zero-sum game against adaptive adversaries [52]. With this approach, one can find the optimal strategies using linear programming. In fact, linear programming is the most efficient solution for this problem [18]. However, if we want to guarantee a certain level of privacy for the user and maximize her utility, the problem cannot be modeled as a zero-sum game anymore and there has been no solution for it so far. We formalize this game, and construct a linear programming solution for these privacy games too.

Regarding the utility metric, we consider the expected distance between the observation and the secret as the utility metric [11, 15, 25, 52]. The distance function can depend on the user and also the application.

In the case of applying obfuscation over time, we need to update the user's estimation of the prior leakage according to what has been shared by the user [19, 53]. We might also need to update the differential privacy budget over time [17]. In this paper, we model one time sharing of a secret, assuming that the prior leakage and the differential privacy budget are properly computed and adjusted based on the previous observations.

Our problem is also related to the problem of adversarial machine learning [5, 30] and the design of security mechanisms, such as intelligent spam detection algorithms [13, 35, 37], against adaptive attackers. It is also similar to the problem of placing security patrols in an area to minimize the threat of attackers [48], and faking location-based queries to protect against localization attack [52]. The survey [39] explores more examples of the relation between security and game theory.

3 Definitions

In this section, we define different parts of our model. We assume a user shares her data through an information sharing system in order to obtain some service (utility). We also assume that users want to protect their sensitive information, while they share their data with untrusted entities. For example, in the case of shar-

ing location-tagged data with a service provider, a user might want to hide the exact visited locations, their semantics, or her activities that can be inferred from the visited locations. We refer to the user’s sensitive information as her *secret*. To protect her privacy, we assume that user obfuscates her data before sharing or publishing it. Figure 1 illustrates the information flow that we assume in this paper.

The input to the protection mechanism is a secret $s \in S$, where S is the set of all possible values that s can take (for example, the locations that the user can visit, or the individuals that she is acquainted with). Let prior leakage π be the probability distribution over values of s to reflect the data model and the a priori exposed information about the secret.

$$\pi(s) = \Pr\{S = s\} \quad (1)$$

The probability distribution π is estimated by the user to be the predictability of the user’s secret given her exposed information in the past. Thus, anytime that user shares some (obfuscated) information, she needs to update this probability distribution [19, 53]. This is how we incorporate the correlation between users’ data shared over time.

3.1 Obfuscation Mechanism

We assume that a user wants to preserve her privacy with respect to s . To protect her privacy, a user obfuscates her secret s and shares an inaccurate version of it through the system. We assume that this obfuscated data $o \in O$ is observable through the system. We consider a generic class of obfuscation mechanisms, in which the observable o is sampled according to the following probability distribution.

$$p(o|s) = \Pr\{O = o|S = s\} \quad (2)$$

Thus, we model the privacy preserving mechanism as a noisy channel between the user and the untrusted observer. This is similar to the model used in quantitative information flow and quantitative side-channel analysis [34, 41]. The output, i.e., the set of observables O , can in general be a member of the powerset of S . As an example, in the most basic case, $O = S$, i.e., the protection mechanism can only perturb the secret by replacing it with another possible secret’s value. This can happen through adding noise to s . In a more generic case, the members of O can contain a subset of secrets. For example, the protection mechanism can generalize a location coordinate, by reducing its granularity.

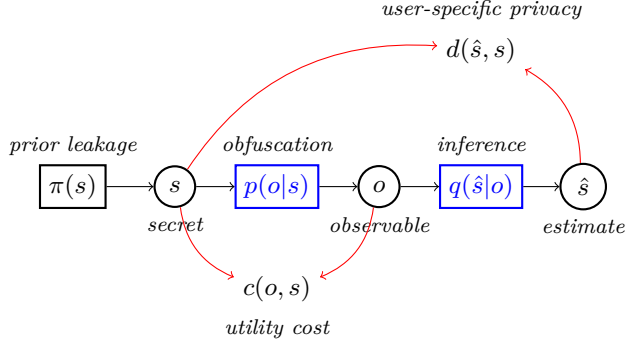


Fig. 1. The Information Sharing Framework. Probability distribution π encodes the user’s estimation of a priori leaked information about secret s . The secret is obfuscated by the protection mechanism p whose output is an observable o . The adaptive adversary (anticipated by the user) runs inference attack q on o and draws a probability distribution over estimates \hat{s} . Distance function c denotes the utility cost of the protection mechanism due to obfuscation. Distance function d denotes the privacy of user (for distortion privacy metric) or the required indistinguishability between secrets (for differential privacy metric). User defines the distance function d to reflect her privacy sensitivities.

3.2 Utility Cost

Users incur a utility loss due to obfuscation. Let the distance function $c(o, s)$ determine the utility cost (information usefulness degradation) due to replacing a secret s with an observable o . The cost function is dependent on the application of the shared information, on the specific service that is provided to the user, and also on the user’s expectations. We compute the expected utility cost of a protection mechanism p as

$$\sum_s \pi(s) \sum_o p(o|s) \cdot c(o, s). \quad (3)$$

We can also compute the worst (maximum) utility cost over all possible secrets as

$$\max_s \sum_o p(o|s) \cdot c(o, s). \quad (4)$$

In this work, we do not plan to determine which metrics are the best representative utility loss metrics for different types of services or users. We only assume that the designer of optimal obfuscation mechanism is provided with such a utility function, for example, by constructing it according to the application [43], or by learning it automatically [8] from the users’ preferences and application profile.

3.3 Inference Attack

We stated that the user wants to protect her privacy with respect to secret s against untrusted observers. To be consistent with this, we define the adversary as an entity who aims at finding the user's secret by observing the outcome of the protection mechanism and minimizing the user's privacy with respect to her privacy sensitivities. For any observation o , then we determine the probability distribution over the possible secrets $\hat{s} \in S$ as to be the true secret of the user.

$$q(\hat{s}|o) = \Pr\{S = \hat{s} | O = o\} \quad (5)$$

The goal of the inference algorithm q is to invert a given protection mechanism p to estimate \hat{s} . The error of adversary, in this estimation process, determines the effectiveness of the inference algorithm, which is captured by the distortion privacy metric.

3.4 Distortion Privacy Metric

As stated above, the user's privacy and the adversary's inference error are two sides of the same coin. We define the privacy gain of the user with secret s as a distance between the two data points: $d(\hat{s}, s)$, where \hat{s} is the a posteriori estimation of the secret [51]. The distance function d is determined by the sensitivity of the user towards each secret s when estimated as \hat{s} . A user would be less worried about revealing $o \sim p(o|s)$, if the portrait of her secret s in the eyes of adversary is an estimate \hat{s} with a large distance $d(\hat{s}, s)$.

This distance function is defined by the user. It could be a semantic distance between different values of secrets to reflect the privacy risk of \hat{s} on user when her secret is s . Usually, the highest risk is associated with the case where the estimate \hat{s} is equal to the secret s . However, sometimes even wrong estimates can impose a high risk on the user, for example when they leak information about the semantic of the secret.

We compute the user privacy obtained through a protection mechanism p , with respect to a given inference algorithm q , for a specific secret s as

$$\sum_o p(o|s) \sum_{\hat{s}} q(\hat{s}|o) \cdot d(\hat{s}, s). \quad (6)$$

By averaging this value over all possible secrets, we compute the expected distortion privacy of the user as

$$\sum_s \pi(s) \sum_o p(o|s) \sum_{\hat{s}} q(\hat{s}|o) \cdot d(\hat{s}, s). \quad (7)$$

This metric shows the average estimation error, or how distorted the reconstructed user's secret is. Thus, we refer to it as the *distortion* privacy metric.

What associates a semantic meaning to this metric is the distance function d . Many distance functions can be defined to reflect distortion privacy. This depends on the type of the secret and to the sensitivity of the user. For example, if the user's secret is her age, function d could be the absolute distance between two numbers. If the secret is the user's location, function d could be a Euclidean distance between locations, or their semantic dissimilarity. If the secret is the movies that she has watched, function d could be the Jaccard distance between two sets of movies.

3.5 Differential Privacy Metric

The privacy that is achieved by an obfuscation mechanism can be computed with respect to the information leakage through the mechanism, regardless of the secret's inference. For example, the differential privacy metric, originally proposed for protecting privacy in statistical databases [21], is sensitive only to the difference between the probabilities of obfuscating multiple secrets to the same observation (which is input to the attack).

According to the original definition of differential privacy, a randomized function \mathcal{K} (that acts as the privacy protection mechanism) provides ϵ -differential privacy if for all data sets D and D' , that differ on at most one element, and all $Y \subseteq \text{Range}(\mathcal{K})$, the following inequality holds.

$$\Pr\{\mathcal{K}(D) \in Y\} \leq \exp(\epsilon) \cdot \Pr\{\mathcal{K}(D') \in Y\} \quad (8)$$

Differential privacy is not limited to statistical databases. It has been used in many different contexts where various types of adjacency relations capture the context dependent privacy. A typical example is edge privacy in graphs [46]. It has also been proposed for arbitrary distance function between secrets [15].

This notion can simply be used for measuring information leakage [3]. It has been shown that differential privacy imposes a bound on information leakage [1, 2]. And, this is exactly why we are interested in this metric. Let $d^\epsilon(s, s')$ be a distinguishability metric between $s, s' \in S$. A protection mechanism is defined to be differentially private if for all secrets $s, s' \in S$, where $d^\epsilon(s, s') \leq d_m^\epsilon$, and all observables $o \in O$, the following inequality holds.

$$p(o|s) \leq \exp(\epsilon) \cdot p(o|s') \quad (9)$$

In this paper, we use a generic definition of differential privacy, assuming arbitrary distance function $d^\epsilon()$ on the secrets [6, 15, 22, 49]. In this form, a protection mechanism is differentially private if for all secrets $s, s' \in S$, with distinguishability $d^\epsilon(s, s')$, and for all observables $o \in O$, the following holds.

$$p(o|s) \leq \exp(\epsilon \cdot d^\epsilon(s, s')) \cdot p(o|s') \quad (10)$$

In fact, the differential privacy metric guarantees that, given the observation, there is not enough convincing evidence to prefer one secret to other similar ones (given d^ϵ). In other words, it makes multiple secret values indistinguishable from each other.

4 Problem Statement

The problem that we address in this paper is to find an optimal balance between privacy and utility, and to construct the protection mechanisms that achieve such optimal points. More precisely, we want to construct utility-maximizing obfuscation mechanisms with joint differential-distortion privacy guarantees.

The problem is to find a probability distribution function p^* such that it minimizes utility cost of the user, on average,

$$p^* = \underset{p}{\mathbf{argmin}} \sum_s \pi(s) \sum_o p(o|s) \cdot c(o, s) \quad (11)$$

or, alternatively, over all the secrets

$$p^* = \underset{p}{\mathbf{argmin}} \max_s \sum_o p(o|s) \cdot c(o, s) \quad (12)$$

under the user's privacy constraints.

4.1 Distortion Privacy Constraint

Let d_m be the minimum desired distortion privacy level. The user's average distortion privacy is guaranteed if the obfuscation mechanism p^* satisfies the following inequality.

$$\sum_s \pi(s) \sum_o p^*(o|s) \sum_{\hat{s}} q^*(\hat{s}|o) \cdot d(\hat{s}, s) \geq d_m \quad (13)$$

where q^* is the optimal inference attack against p^* .

4.2 Differential Privacy Constraint

Let ϵ_m be the differential privacy budget associated with the minimum desired privacy of the user, and d_m^ϵ be

the distinguishability threshold. The user's privacy is guaranteed if p^* satisfies the following inequality.

$$p^*(o|s) \leq \exp(\epsilon_m) \cdot p^*(o|s'), \forall o, s, s' : d^\epsilon(s, s') \leq d_m^\epsilon \quad (14)$$

Or, alternatively (following [15]'s definition of differential privacy):

$$p^*(o|s) \leq \exp(\epsilon_m \cdot d^\epsilon(s, s')) \cdot p^*(o|s'), \forall s, s', o \quad (15)$$

In this paper, we mainly use the latter definition, but make use of the former one as the basis to reduce the computation cost of optimizing differential privacy (see Appendix C).

5 Solution: Privacy Games

The flow of information starts from the user where the secret is generated. The user then selects a protection mechanism, and obfuscates her secret according to its probabilistic function. After the adversary observes the output, he can design an optimal inference attack against the obfuscation mechanism to invert it and estimate the secret. We assume the obfuscation mechanism is not oblivious and is known to the adversary. This gives the adversary the upper hand against the user in their conflict. So, designing an obfuscation mechanism against a fixed attack is always suboptimal.

The best obfuscation mechanism is the one that *anticipates* the adversary's attack. Thus, the obfuscation mechanism should be primarily designed against an *adaptive* attack which is tailored to each specific obfuscation mechanism. So, by assuming that the adversary designs the best inference attack against each protection mechanism, the user's goal (as the defender) must be to design the obfuscation mechanism that maximizes her (privacy or utility) objective against an adversary that optimizes the conflicting objective of guessing the user's secret. The adversary is an entity assumed by the user as the entity whose objective is exactly the opposite of the user's. So, we do not model any particular attacker but the one that minimizes user's privacy according to distance functions d and d^ϵ .

For each obfuscation mechanism there is an inference attack that optimizes the adversary's objective and leads to a certain privacy and utility payoff for the user. The optimal obfuscation mechanism for the user is the one that brings the maximum payoff for her, against the mechanism's corresponding optimal inference attack.

Enumerating all pairs of user-attacker mechanisms to find the optimal obfuscation function is infeasible. We

model the joint user-adversary optimization problem as a leader-follower (Stackelberg) game between the user and the adversary. The user leads the game by choosing the protection mechanism p , and the adversary follows by designing the inference attack q . The solution to this game is the pair of user-adversary best response strategies p^* and q^* which are mutually optimal against each other. If the user implements p^* , we have already considered the strongest attack q^* against it. Thus, p^* is robust against *any* algorithm used as inference attack.

For any secret $s \in S$, the strategy space of the user is the set of observables O . For any observable $o \in O$, the strategy space of the adversary is the set of secrets S (all possible adversary's estimates $\hat{s} \in S$). For a given secret $s \in S$, we represent a mixed strategy for the user by a vector $p(\cdot|s) = (p(o_1|s), p(o_2|s), \dots, p(o_m|s))$, where $\{o_1, o_2, \dots, o_m\} = O$. Similarly, a mixed strategy for the adversary, for a given observable $o \in O$ is a vector $q(\cdot|o) = (q(\hat{s}_1|o), q(\hat{s}_2|o), \dots, q(\hat{s}_n|o))$, where $\{\hat{s}_1, \hat{s}_2, \dots, \hat{s}_n\} = S$. Note that the vectors $p(\cdot|s)$ and $q(\cdot|o)$ are respectively the conditional distribution functions associated with an obfuscated function for a secret s and an inference algorithm for an observable o . Let P and Q be the sets of all mixed strategies of the user and the adversary, respectively.

$$P = \{p(\cdot|s) = (p(o_1|s), p(o_2|s), \dots, p(o_m|s)), \forall s \in S : \\ p(o_i|s) \geq 0, \forall o_i \in O, \sum_i p(o_i|s) = 1\} \quad (16)$$

$$Q = \{q(\cdot|o) = (q(\hat{s}_1|o), q(\hat{s}_2|o), \dots, q(\hat{s}_n|o)), \forall o \in O : \\ q(\hat{s}_j|o) \geq 0, \forall \hat{s}_j \in S, \sum_j q(\hat{s}_j|o) = 1\} \quad (17)$$

A member vector of sets P or Q with a 1 for the k th component and zeros elsewhere is the pure strategy of choosing action k . For example, an obfuscation function $p(\cdot|s)$ for which $p(o_i|s) = 0, \forall i \neq k$ and $p(o_k|s) = 1$ is the pure strategy of exclusively and deterministically outputting observable o_k for secret s . Thus, the set of pure strategies of a player is a subset of mixed strategies of the player.

In the case of the distortion privacy metric, the game needs to be formulated as a *Bayesian Stackelberg game*. In this game, we assume the probability distribution π on the secrets and we find $p^* \in P$ and $q^* \in Q$ that create the equilibrium point. If user deviates from this strategy and chooses $p' \neq p^*$, there would be an inference attack q'^* against it such that (p', q'^*) leads to a lower privacy for the user, i.e., p^* is optimal.

In the case of a differential privacy metric, as the metric is not dependent to the adversary's inference at-

tack, the dependency loop between finding optimal p^* and q^* is broken. Nevertheless, it is still the user who plays first by choosing the optimal protection mechanism. In the following sections, we solve these games and provide solutions on how to design the optimal user-adversary strategies.

6 Stackelberg Privacy Games

Assume that the nature draws secret s according to the probability distribution $\pi(s)$. Given s , the user draws o according to her obfuscation mechanism $p(o|s)$, and makes it observable to the adversary. Given observation o , the adversary draws \hat{s} according to his inference attack $q(\hat{s}|o)$. We assume that $\pi(s)$ is known to both players. We want to find the mutually optimal (p^*, q^*) : The solution of the Bayesian Stackelberg privacy game.

To this end, we first design the optimal inference attack against any given protection mechanism p . This will be the *best response* of the adversary to the user's strategy. Then, we design the optimal protection mechanism for the user according to her objective and constraints, as stated in Section 4. This will be the user's best utility-maximizing strategy that anticipates the adversary's best response.

6.1 Optimal Inference Attack

The adversary's objective is to minimize (the user's privacy and thus) the inference error in estimating the user's secret. Given a secret s , the distance function $d(\hat{s}, s)$ determines the error of an adversary in estimating the secret as \hat{s} . In fact, this distance is exactly what a user wants to maximize (or put a lower bound on) according to the distortion privacy metric. We compute the expected error of the adversary as

$$\sum_s \pi(s) \sum_{\hat{s}} \Pr\{\hat{s}|s\} \cdot d(\hat{s}, s) = \\ = \sum_{s, o, \hat{s}} \pi(s) \cdot p(o|s) \cdot q(\hat{s}|o) \cdot d(\hat{s}, s) \quad (18)$$

Therefore, we design the following linear program, through which we can compute the adversary's inference strategy that, given the probability distribution π and obfuscation p , minimizes his expected error with respect to a distance function d .

$$q^* = \underset{q}{\operatorname{argmin}} \sum_{s, o, \hat{s}} \pi(s) \cdot p(o|s) \cdot q(\hat{s}|o) \cdot d(\hat{s}, s) \quad (19a)$$

under the constraint that the solution is a proper conditional probability distribution function.

In the next subsection, we will show that the optimal deterministic inference (that associates one single estimate with probability one to each observation) results in the same privacy for the user (33). Alternative ways to formulate this problem is given in Appendix A.

6.2 Optimal Protection Mechanism

In this case, we assume the user would like to minimize her utility cost (11) under a (lower bound) constraint on her privacy (13). Therefore, we can formulate the problem as

$$p^* = \underset{p}{\operatorname{argmin}} \sum_{s,o} \pi(s) \cdot p(o|s) \cdot c(o, s) \quad (20a)$$

$$\text{s. t. } \sum_{s,o,\hat{s}} \pi(s) \cdot p(o|s) \cdot q^*(\hat{s}|o) \cdot d(\hat{s}, s) \geq d_m \quad (20b)$$

However, solving this optimization problem requires us to know the optimal q^* against p^* , for which we need to know p^* as formulated in (19). So, we have two linear programs (one for the user and one for the adversary) to solve. But, the solution of each one is required in solving the other. This optimization dependency loop reflects the game-theoretic concept of *mutual best response* of the two players. This game is a *nonzero-sum Stackelberg game* as the user (leader player) and adversary (follower player) have different optimization objectives (one maximizes utility, and the other minimizes privacy). We break the dependency loop between the optimization problems using the game-theoretic modeling, and we prove that the user's best strategy can be constructed using linear programming.

Theorem 1. *Given a probability distribution π , the distance functions d and c , and the threshold d_m , the solution to the following linear program is the optimal protection strategy p^* for the user, which is the solution to (20) with respect to adversary's best response (19).*

$$p^* = \underset{p}{\operatorname{argmin}} \sum_{s,o} \pi(s) \cdot p(o|s) \cdot c(o, s) \quad (21a)$$

$$\text{s. t. } \sum_s \pi(s) \cdot p(o|s) \cdot d(\hat{s}, s) \geq x(o), \forall o, \hat{s} \quad (21b)$$

$$\sum_o x(o) \geq d_m \quad (21c)$$

Proof. See Appendix B. \square

7 Optimal Differential Privacy

In this section, we design optimal differentially private protection mechanisms. We solve the optimization problems for maximizing utility under privacy constraint.

We design the following linear program to find the user strategy p^* that guarantees user differential privacy (15), for a maximum privacy budget ϵ_m , and minimizes the utility cost (11) of the obfuscation mechanism.

$$\underset{p}{\min} \sum_{s,o} \pi(s) \cdot p(o|s) \cdot c(o, s) \quad (22a)$$

$$\text{s. t. } \frac{p(o|s)}{p(o|s')} \leq \exp(\epsilon_m \cdot d^\epsilon(s, s')), \forall s, s', o \quad (22b)$$

Or, alternatively, for a distinguishability bound d_m^ϵ , we can solve the following.

$$\underset{p}{\min} \sum_{s,o} \pi(s) \cdot p(o|s) \cdot c(o, s) \quad (23a)$$

$$\text{s. t. } \frac{p(o|s)}{p(o|s')} \leq \exp(\epsilon_m), \forall o, s, s' : d^\epsilon(s, s') \leq d_m^\epsilon \quad (23b)$$

8 Optimal Joint Differential and Distortion Privacy Mechanism

Obfuscation mechanisms designed based on distortion and differential privacy protect the user's privacy from two different angles. In general, for arbitrary d and d^ϵ , there is no guarantee that a mechanism with a bound on one metric holds a bound on the other.

Distortion privacy metric reflects the *absolute* privacy of the user, based on the posterior estimation on the obfuscated information. Differential privacy metric reflects the *relative* information leakage of each observation about the secret. However, it is not a measure on the extent to which the observer, who already has some knowledge about the secret from the previously shared data, can guess the secret correctly. So, the inference might be very accurate (because of the background knowledge) despite the fact that the obfuscation in place is a differentially-private mechanism.

As distortion and differential metrics guarantee different dimensions of the user's privacy requirements, we respect both in a protection mechanism. This assures that not only the information leakage is limited, but also the absolute privacy level is at the minimum required level. Thanks to our unified formulation of privacy optimization problems as linear programs, the problem of

jointly optimizing and guaranteeing privacy with both metrics can also be formulated as a linear program.

The solution to the following linear program is a protection mechanism p^* that maximizes the user's utility and guarantees a minimum distortion privacy d_m and a minimum differential privacy ϵ_m , given probability distribution π and distance functions c and d and distinguishability metric d^ϵ . The value of the optimal solution is the utility cost of the optimal mechanism.

$$\min_p \sum_{s,o} \pi(s) \cdot p(o|s) \cdot c(o, s) \quad (24a)$$

$$\text{s. t. } \sum_s \pi(s) \cdot p(o|s) \cdot d(\hat{s}, s) \geq x(o), \forall o, \hat{s} \quad (24b)$$

$$\sum_o x(o) \geq d_m \quad (24c)$$

$$\frac{p(o|s)}{p(o|s')} \leq \exp(\epsilon_m \cdot d^\epsilon(s, s')), \forall s, s', o \quad (24d)$$

9 Analysis

We have implemented all our linear program solutions in a software tool that can be used to process data for different applications, in different settings. In this section, we use our tool to design privacy protection mechanisms, and also to make a comparison between different optimal mechanisms, i.e., distortion, differential, and joint distortion-differential privacy preserving mechanisms. We study the properties of these mechanisms and we show how robust they are with respect to inference attack algorithms as well as to the adversary's knowledge on secrets. We also investigate their utility cost for protecting privacy. Furthermore, we show that the optimal joint distortion-differential mechanisms are more robust than the two mechanisms separately. In Appendix C, we discuss and evaluate approximations of the optimal solution for large number of constraints.

We run experiments on location data, as today they are included in most of data sharing applications. We use a real data-set of location traces collected through the Nokia Lausanne Data Collection Campaign [33]. The location information belong to a 15×8 km area. We split the area into 20×15 cells. We consider location of a mobile user in a cell as her secret. Hence, the set of secrets is equivalent to the set of location cells. We assume the set of observables to be the set of cells, so the users obfuscate their location by perturbation (i.e., replacing their true location with any location in the map). We run our experiments on 10 randomly selected users, to see the difference in the results due to difference

in user's location distribution π based on users' different location access profiles. We build π for each user separately given their individual location traces, using maximum likelihood estimation (normalizing the user's number of visits to each cell in the trace).

We assume a Euclidean distance function for d and d^ϵ . This reflects the sensitivity of user towards her location. By using this distance function for distortion privacy, we guarantee that the adversary cannot guess the user's true location with error lower than the required privacy threshold (d_m). Choosing Euclidean distance function as the metric for distinguishability ensures that the indistinguishability between locations is larger for locations that are located closer to each other.

We assume a Hamming distortion function for c (i.e., the utility cost is 0 only if the user's location and the observed location are the same, otherwise the cost is 1). The utility metric can vary depending on the location-based sharing application and also the purpose for which the user shares her location [8]. Choosing the Hamming function reflects the utility requirement of users who want to inform others about their current location in location check-in applications.

We evaluate utility-maximizing optimal protection mechanisms with three different privacy constraints:

- *Distortion Privacy Protection*, (21).
- *Differential Privacy Protection*, (22).
- *Joint Distortion-Differential Privacy Protection*, (24).

We compare the effectiveness of these protection mechanisms against inference attacks by using the distortion privacy metric (7). We consider two inference attacks:

- *Optimal Attack*, (19).
- *Bayesian Inference Attack*, using the Bayes rule:

$$q(\hat{s}|o) = \frac{\pi(\hat{s}) \cdot p(o|\hat{s})}{\Pr\{o\}} = \frac{\pi(\hat{s}) \cdot p(o|\hat{s})}{\sum_s \pi(s) \cdot p(o|s)} \quad (25)$$

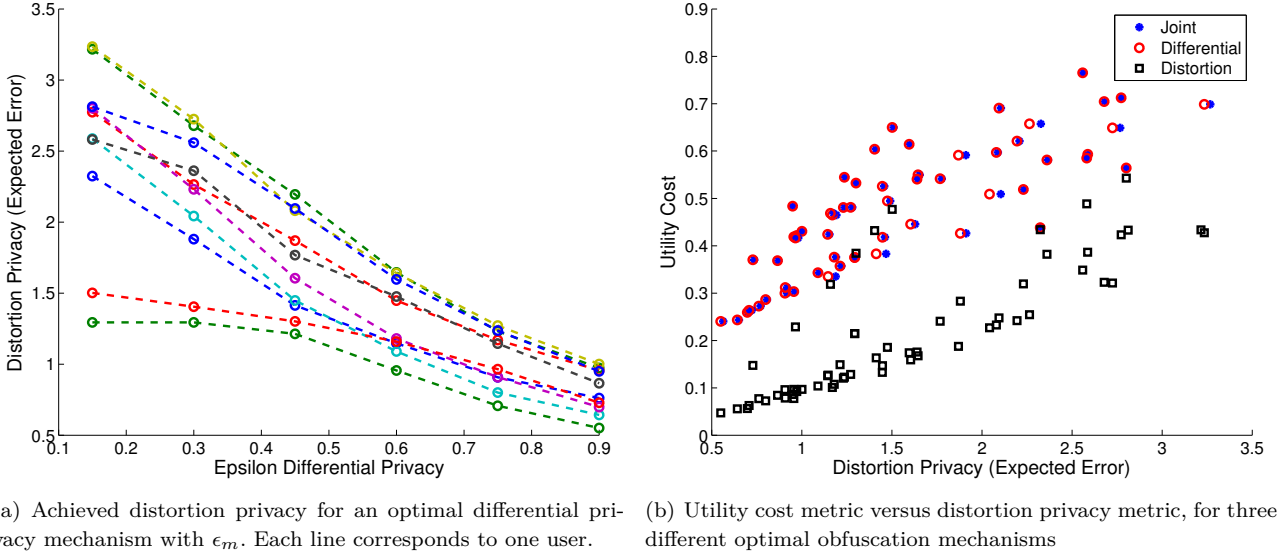
9.1 Comparing Obfuscation Mechanisms

Scenario 1.

Our first goal is to have a fair comparison between optimal distortion privacy mechanism and optimal differential mechanism. To this end, we set the privacy parameter ϵ_m to $\{0.15, 0.3, \dots, 0.9\}$.

For each user and each value of ϵ_m ,

1. We compute the optimal differential privacy mechanism using (22). Let $p_{\epsilon_m}^*$ be the optimal mechanism.
2. We run optimal attack (19) on $p_{\epsilon_m}^*$, and compute the user's absolute distortion privacy as $AP(p_{\epsilon_m}^*)$.



(a) Achieved distortion privacy for an optimal differential privacy mechanism with ϵ_m . Each line corresponds to one user.

(b) Utility cost metric versus distortion privacy metric, for three different optimal obfuscation mechanisms

Fig. 2. Privacy and Utility of optimal protection mechanisms, computed based on the evaluation scenario number 1 in Section 9.1. Each dot in the plots corresponds to privacy of one user for one value of ϵ_m .

3. We compute the optimal distortion privacy mechanism $p_{d_m}^*$ using (21). For this, we set the privacy lower-bound d_m to $AP(p_{\epsilon_m}^*)$. This enforces the distortion privacy mechanism to guarantee what the differential privacy mechanism provides.
4. We compute the optimal joint distortion-differential privacy mechanism p_{ϵ_m, d_m}^* using (24). We set the privacy lower-bounds to ϵ_m and d_m for the differential and distortion constraints, respectively.
5. We run optimal attack (19) on both $p_{d_m}^*$ and p_{ϵ_m, d_m}^* , and compute the user's absolute distortion privacy as $AP(p_{d_m}^*)$ and $AP(p_{\epsilon_m, d_m}^*)$, respectively.
6. As a baseline for comparison, we run Bayesian inference attack (25) on the three optimal mechanisms $p_{\epsilon_m}^*$, $p_{d_m}^*$, and p_{ϵ_m, d_m}^* .

Figure 2 shows the results of our analysis, explained above. Distortion privacy is measured in km and is equivalent to the expected error of adversary in correctly estimating location of users. Figure 2(a) shows how expected privacy of users $AP(p_{\epsilon_m}^*)$ decreases as we increase the value of the lower-bound on differential privacy ϵ_m . Users have different secret probability distribution, with different randomness. However, as ϵ_m increases, expected error of adversary (the location privacy of users) converges down to below 1km. Figure 2(b) plots the utility cost versus distortion privacy of each optimal protection mechanism. As we have set the privacy bound of the optimal distortion mechanism (and of course the optimal joint mechanism) to the privacy

achieved by the optimal differential mechanism, we can make a fair comparison between their utility costs. We observe that the utility cost for achieving some level of distortion privacy is much higher for optimal differential and joint mechanisms compared with the optimal distortion mechanism. Note that the utility cost of differential and joint mechanisms are the same. So, distortion privacy bound does not impose more cost than what is already imposed by the differential privacy mechanism.

As we set d_m to $AP(p_{\epsilon_m}^*)$, the user's distortion privacy in using optimal distortion and optimal differential mechanism is the same, when we confront them with the optimal attack (19). In Figure 3, however, we compare the effectiveness of these two mechanisms against Bayesian inference attack (25). It is interesting to observe that the optimal differential mechanism is more robust to such attacks compared to the optimal distortion mechanisms. This explains the extra utility cost due to optimal differential mechanisms.

In Figure 4, we compare the effectiveness of Bayesian inference attack (25) and optimal attack (19). We show the results for all three optimal protection mechanisms. It is clear that optimal attack outperforms the Bayesian attack, as users have a relatively higher privacy level under the Bayesian inference. However, the difference is more obvious for the case of differential protection and joint protection mechanisms. The Bayesian attack overestimates users' privacy, as it ignores the distance function d , whereas the optimal attack minimizes the expected value of d over all secrets and estimates.

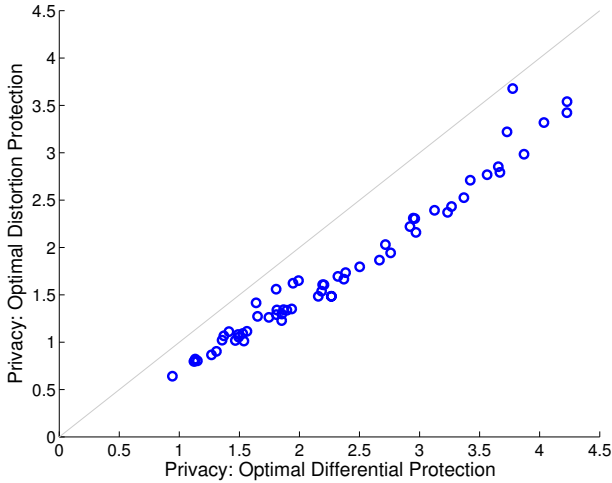


Fig. 3. Distortion privacy of users against the Bayesian inference attack (25) when using optimal differential privacy obfuscation versus using optimal distortion privacy obfuscation. Each dot represents privacy of one user for one value of ϵ_m .

Scenario 2.

In this paper, we introduce the optimal joint distortion-differential protection mechanisms to provide us with the benefits of both mechanisms. Figure 2(b) shows that the optimal joint mechanism is not more costly than the two optimal distortion and differential mechanisms. It also shows that it guarantees the highest privacy for a certain utility cost. To further study the effectiveness of optimal joint mechanisms, we run the following evaluation scenario.

We design optimal differential mechanisms for some values of ϵ_m . And, we design optimal distortion mechanisms for some values of d_m that are higher than the distortion privacy resulted from those differential privacy mechanisms. We also construct their joint mechanisms given the ϵ_m and d_m parameters. Figure 5 shows how the optimal joint mechanism adapts itself to guarantee the maximum of the privacy levels guaranteed by optimal Bayesian and optimal differential mechanisms individually. This is clear from the fact that users' privacy for the optimal joint mechanism is equal to their privacy for distortion mechanism (that as we set in our scenario, they are higher than that of differential mechanisms).

Thus, by adding the distortion privacy constraints in the design of optimal mechanisms, we can further increase the privacy of users (with the same utility cost) that cannot be otherwise achieved by only using differential mechanisms.




Fig. 4. Distortion privacy of users using any of the three optimal mechanisms against the basic Bayesian inference attack (25) versus their privacy against the optimal attack (19). Each dot represents privacy of one user for one value of ϵ_m .

Scenario 3.

In order to further investigate the relation between the privacy (and utility) outcome of the optimal joint mechanism and that of individual differential or distortion privacy mechanisms, we run the following set of experiments on all the available user profiles.

1. For any value of ϵ_m in $\{0.2, 0.4, \dots, 1\}$, we compute the utility of optimal differential privacy mechanism as well as its privacy against optimal attack.
2. For any value of d_m in $\{0.5, 1, \dots, d_m^{max}\}$, we compute the utility of optimal distortion privacy mechanism as well as its privacy against optimal attack. d_m^{max} is dependent on π and is the maximum value that the threshold can take (beyond which there is no solution to the optimization problem).
3. For any value of ϵ_m in $\{0.2, 0.4, \dots, 1\}$, and for any value of d_m in $\{0.5, 1, \dots, d_m^{max}\}$, we compute the utility and privacy of the optimal joint mechanism.

Figure 7 shows the results. By an experiment we refer to the comparison of privacy (or utility) of a joint mechanism (with bounds ϵ_m, d_m) with the corresponding differential privacy mechanism (with bound ϵ_m) and the corresponding distortion privacy mechanism (with bound d_m). Note that here the thresholds ϵ_m and d_m are chosen independently as opposed to scenarios 1 (and also 2). We put the results of all the experiments next to each other in the x-axis. Therefore, any vertical cut on the Figure 7's plots contain three points for privacy/utility of p_{ϵ_m, d_m}^* , $p_{\epsilon_m}^*$, and $p_{d_m}^*$. To better visualize

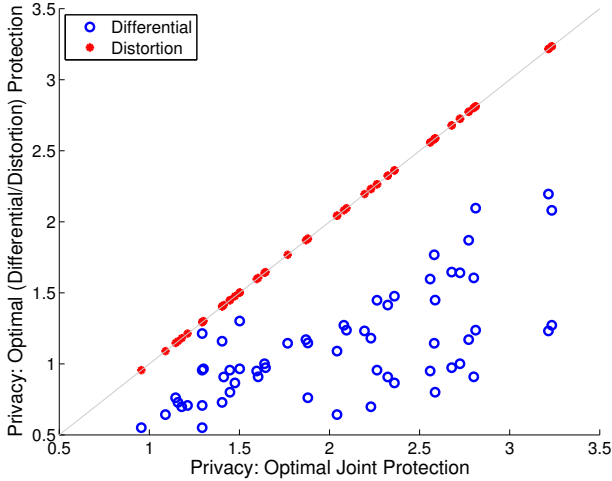


Fig. 5. Users' privacy using optimal differential or distortion protection mechanism versus using optimal joint protection mechanism. Distortion privacy is computed using optimal attack (19).

the results, we have sorted all the experiments based on the privacy/utility of the joint mechanism.

As the results show, the privacy achieved by the optimal joint mechanism is equal to the maximum privacy that each of the individual differential/distortion mechanisms provides separately. This means that the user would indeed benefit from including a distortion privacy constraint based on her prior leakage into the design criteria of the optimal obfuscation mechanism. This comes at no extra utility cost for the user, as the utility graph shows. In fact, the utility cost of an optimal joint mechanism is not additive and instead is the maximum of the two components, which is the differential privacy mechanism in all tested experiments. The reason behind this is that the differential privacy component makes the joint obfuscation mechanism robust to the case where the background knowledge of the adversary includes not only the prior leakage but also other auxiliary information available to him.

9.2 Evaluating the Effects of Prior

When using distortion metric in protecting privacy, we achieve optimal privacy given the user's estimated prior leakage modeled by probability distribution π over the secrets. In the optimal attack against various protection mechanisms, a real adversary makes use of a prior distribution over the secrets. In this subsection, we evaluate to what extent a more informed adversary can harm privacy of users further than what is promised by the optimal protection mechanisms. Note that no matter




Fig. 6. Users' privacy against the optimal attack using optimal differential protection versus using optimal distortion protection. Each circle represents privacy of a user for a different ϵ_m and for a different prior assumed in the attack. The red dots correspond to the cases where the probability π assumed in designing the protection mechanism is the same as the attacker's knowledge.

what protection mechanism is used by the user, a more knowledgeable adversary will learn more about the secret. In this section, our goal is not to show this obvious fact, but to evaluate how robust our mechanisms are with respect to adversaries with different knowledge accuracy levels.

To perform this analysis, we consider a scenario in which the adversary's assumption on π , for each user, has a lower level of uncertainty compared to π . This can happen in the real world when an adversary obtains new evidence about a user's secret that is not used by user for computing π . Let $\hat{\pi}$ be the other version of π assumed by adversary, for a given user. For the sake of our analysis, we generate $\hat{\pi}$ by providing the adversary with more evidence about most frequently visited locations, e.g., home and work. This is equivalent to the scenario in which the adversary knows the user's significant locations, e.g., where the user lives and works. The entropy of $\hat{\pi}$ is less than that of π , hence it contains more information about the user's mobility.

We construct the protection mechanisms assuming π , and we attack them by optimal inference attacks, but assuming the lower entropy $\hat{\pi}$ priors. Figure 6 illustrates privacy of users for different assumptions of $\hat{\pi}$, using optimal differential protection versus optimal distortion protection (assuming π). We observe that a more informed adversary has a lower expected error. However, it further shows that an optimal differential protection mechanism compared to an optimal distortion mecha-

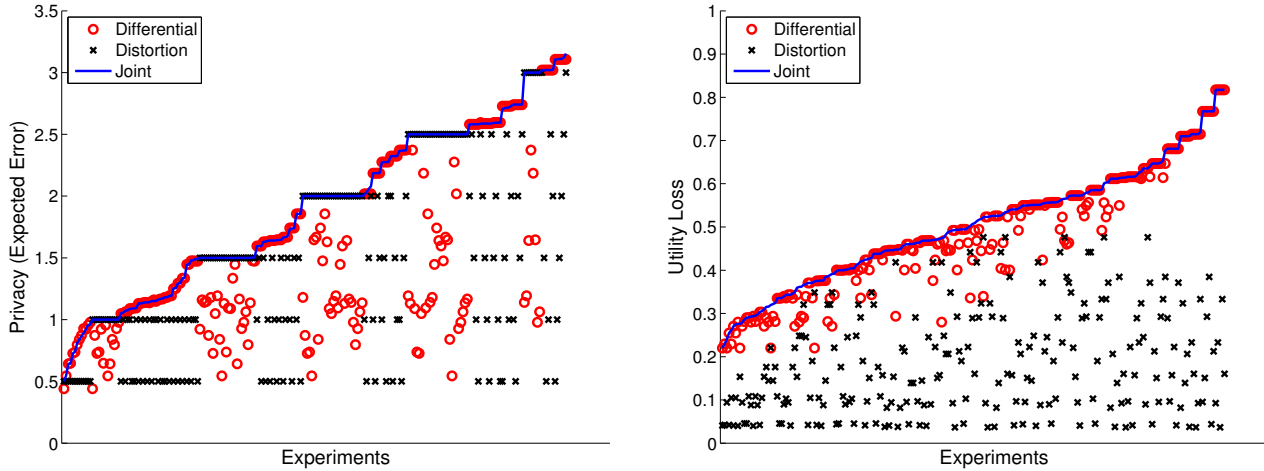


Fig. 7. The comparison between privacy and utility of optimal joint mechanism p_{ϵ_m, d_m}^* with the individual protection mechanisms $p_{\epsilon_m}^*$ and $p_{d_m}^*$, i.e., the mechanisms whose bounds are jointly respected in the optimal joint mechanism. The three points on each vertical line represent the results of one such comparison experiment for different values of pairs of privacy thresholds (ϵ_m, d_m) in $\{0.2, 0.4, \dots, 1\} \times \{0.5, 1, \dots, d_m^{max}\}$.

nism is more robust to knowledgeable adversaries. Note that we set d_m to $AP(p_{\epsilon_m}^*)$, according to scenario 1 in Section 9.1. So, when $\hat{\pi} = \pi$, both optimal protection mechanisms guarantee the same level of privacy. However, as there is more information in $\hat{\pi}$ than in π , more information can be inferred from the optimal distortion mechanism compared to the differential mechanism.

10 Conclusions

We have solved the problem of designing *optimal* user-centric obfuscation mechanisms for data sharing systems. We have proposed a novel methodology for designing such mechanisms against any *adaptive* inference attack, while maximizing users' utility. We have proposed a generic framework for quantitative privacy and utility, using which we formalize the problems of maximizing users' utility under a lower-bound constraint on their privacy. The major novelty of the paper is to solve these optimization problems for both state-of-the-art distortion and differential privacy metrics, for the generic case of any distance function between the secrets. Being generic with respect to the distance functions, enables us to formalize any sensitivity function on any type of secrets. We have also proposed a new privacy notion, joint distortion-differential privacy, and constructed its optimal mechanism that has the strengths of both metrics. We have provided linear program solutions for our optimization problems that provably achieve minimum utility loss under those privacy bounds.

Acknowledgements

We would like to thank the PC reviewers for their constructive feedback, and Kostas Chatzikokolakis for very useful discussions on this work.

References

- [1] M. S. Alvim, M. E. Andrés, K. Chatzikokolakis, P. Degano, and C. Palamidessi. Differential privacy: on the trade-off between utility and information leakage. In *Formal Aspects of Security and Trust*, pages 39–54. Springer, 2012.
- [2] M. S. Alvim, M. E. Andrés, K. Chatzikokolakis, and C. Palamidessi. On the relation between differential privacy and quantitative information flow. In *Automata, Languages and Programming*, pages 60–76. Springer, 2011.
- [3] M. S. Alvim, M. E. Andrés, K. Chatzikokolakis, and C. Palamidessi. Quantitative information flow and applications to differential privacy. In *Foundations of security analysis and design VI*. 2011.
- [4] M. E. Andrés, N. E. Bordenabe, K. Chatzikokolakis, and C. Palamidessi. Geo-indistinguishability: Differential privacy for location-based systems. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, pages 901–914. ACM, 2013.
- [5] M. Barreno, B. Nelson, R. Sears, A. D. Joseph, and J. Tygar. Can machine learning be secure? In *Proceedings of the ACM Symposium on Information, computer and communications security*, 2006.
- [6] G. Barthe, B. Köpf, F. Olmedo, and S. Zanella Béguelin. Probabilistic relational reasoning for differential privacy. *ACM SIGPLAN Notices*, 2012.

- [7] J. O. Berger. *Statistical decision theory and Bayesian analysis*. Springer, 1985.
- [8] I. Bilogrevic, K. Huguenin, S. Mihaila, R. Shokri, and J.-P. Hubaux. Predicting users' motivations behind location check-ins and utility implications of privacy protection mechanisms. In *In Network and Distributed System Security (NDSS) Symposium*, 2015.
- [9] N. E. Bordenabe, K. Chatzikokolakis, and C. Palamidessi. Optimal geo-indistinguishable mechanisms for location privacy. In *Proceedings of the 16th ACM conference on Computer and communications security*, 2014.
- [10] S. P. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [11] H. Brenner and K. Nissim. Impossibility of differentially private universally optimal mechanisms. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*, pages 71–80. IEEE, 2010.
- [12] J. Brickell and V. Shmatikov. The cost of privacy: Destruction of data-mining utility in anonymized data publishing. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '08, pages 70–78, New York, NY, USA, 2008. ACM.
- [13] M. Brückner and T. Scheffer. Stackelberg games for adversarial prediction problems. In *17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2011)*, 2011.
- [14] F. Brunton and H. Nissenbaum. Vernacular resistance to data collection and analysis: A political theory of obfuscation. *First Monday*, 16(5), 2011.
- [15] K. Chatzikokolakis, M. E. Andrés, N. E. Bordenabe, and C. Palamidessi. Broadening the scope of differential privacy using metrics. In *Privacy Enhancing Technologies*, pages 82–102. Springer, 2013.
- [16] K. Chatzikokolakis, C. Palamidessi, and P. Panangaden. Anonymity protocols as noisy channels. *Information and Computation*, 206(2-4):378–401, 2008.
- [17] 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.
- [18] V. Conitzer and T. Sandholm. Computing the optimal strategy to commit to. In *Proceedings of the 7th ACM conference on Electronic commerce*, 2006.
- [19] G. Danezis and C. Troncoso. You cannot hide for long: de-anonymization of real-world dynamic behaviour. In *Proceedings of the 12th ACM workshop on Workshop on privacy in the electronic society*, pages 49–60. ACM, 2013.
- [20] C. Diaz, S. Seys, J. Claessens, and B. Preneel. Towards measuring anonymity. In *Privacy Enhancing Technologies*, pages 54–68. Springer Berlin Heidelberg, 2003.
- [21] C. Dwork. Differential privacy. In *Automata, languages and programming*, pages 1–12. Springer, 2006.
- [22] C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography*, pages 265–284. Springer, 2006.
- [23] V. F. Farias and B. Van Roy. Tetris: A study of randomized constraint sampling. In *Probabilistic and Randomized Methods for Design Under Uncertainty*. 2006.
- [24] Q. Geng and P. Viswanath. The optimal mechanism in differential privacy. *arXiv preprint arXiv:1212.1186*, 2012.
- [25] A. Ghosh, T. Roughgarden, and M. Sundararajan. Universally utility-maximizing privacy mechanisms. In *Proceedings of the 41st annual ACM symposium on Theory of computing*, pages 351–360. ACM, 2009.
- [26] A. Ghosh, T. Roughgarden, and M. Sundararajan. Universally utility-maximizing privacy mechanisms. *SIAM Journal on Computing*, 41(6):1673–1693, 2012.
- [27] M. Grötschel, L. Lovász, and A. Schrijver. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica*, 1981.
- [28] M. Gupte and M. Sundararajan. Universally optimal privacy mechanisms for minimax agents. In *Proceedings of the twenty-ninth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, 2010.
- [29] X. He, A. Machanavajjhala, and B. Ding. Blowfish privacy: Tuning privacy-utility trade-offs using policies. In *Proceedings of the 2014 ACM SIGMOD international conference on Management of data*, pages 1447–1458. ACM, 2014.
- [30] L. Huang, A. D. Joseph, B. Nelson, B. I. Rubinfeld, and J. Tygar. Adversarial machine learning. In *Proceedings of the 4th ACM workshop on Security and artificial intelligence*, 2011.
- [31] S. Ioannidis, A. Montanari, U. Weinsberg, S. Bhagat, N. Fawaz, and N. Taft. Privacy tradeoffs in predictive analytics. *arXiv preprint arXiv:1403.8084*, 2014.
- [32] D. Kifer and A. Machanavajjhala. No free lunch in data privacy. In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of data*, pages 193–204. ACM, 2011.
- [33] N. Kiukkonen, J. Blom, O. Dousse, D. Gatica-Perez, and J. Laurila. Towards rich mobile phone datasets: Lausanne data collection campaign. *Proc. ICPS, Berlin*, 2010.
- [34] B. Köpf and D. Basin. An information-theoretic model for adaptive side-channel attacks. In *Proceedings of the 14th ACM conference on Computer and communications security*, 2007.
- [35] 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.
- [36] C. Li, M. Hay, V. Rastogi, G. Miklau, and A. McGregor. Optimizing linear counting queries under differential privacy. In *Proceedings of the twenty-ninth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pages 123–134. ACM, 2010.
- [37] W. Liu and S. Chawla. A game theoretical model for adversarial learning. In *IEEE International Conference on Data Mining Workshops (ICDM 2009)*, 2009.
- [38] D. J. MacKay. *Information theory, inference and learning algorithms*. Cambridge university press, 2003.
- [39] M. Manshaei, Q. Zhu, T. Alpcan, T. Basar, and J.-P. Hubaux. Game theory meets network security and privacy. *ACM Computing Surveys*, 45(3), 2012.
- [40] P. Mardziel, M. S. Alvim, M. Hicks, and M. R. Clarkson. Quantifying information flow for dynamic secrets. In *IEEE Symposium on Security and Privacy*, 2014.
- [41] S. A. Mario, K. Chatzikokolakis, C. Palamidessi, and G. Smith. Measuring information leakage using generalized gain functions. *2012 IEEE 25th Computer Security*

- Foundations Symposium*, 2012.
- [42] R. T. Marler and J. S. Arora. Survey of multi-objective optimization methods for engineering. *Structural and multidisciplinary optimization*, 26(6):369–395, 2004.
- [43] K. Micinski, P. Phelps, and J. S. Foster. An empirical study of location truncation on android. *Weather*, 2:21, 2013.
- [44] K. Miettinen. *Nonlinear multiobjective optimization*, volume 12. Springer, 1999.
- [45] Y. E. Nesterov and A. Nemirovskii. *Interior point polynomial methods in convex programming: Theory and algorithms*. SIAM Publications. SIAM, Philadelphia, UsA, 1993.
- [46] K. Nissim, S. Raskhodnikova, and A. Smith. Smooth sensitivity and sampling in private data analysis. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 75–84. ACM, 2007.
- [47] V. Pareto. *Manuale di economia politica*, volume 13. Societa Editrice, 1906.
- [48] 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 *Conference on Artificial Intelligence*, 2008.
- [49] J. Reed and B. C. Pierce. Distance makes the types grow stronger: a calculus for differential privacy. *ACM Sigplan Notices*, 2010.
- [50] A. Serjantov and G. Danezis. Towards an information theoretic metric for anonymity. In *Privacy Enhancing Technologies*, pages 41–53. Springer Berlin Heidelberg, 2003.
- [51] R. Shokri, G. Theodorakopoulos, J.-Y. Le Boudec, and J.-P. Hubaux. Quantifying location privacy. In *Proceedings of the IEEE Symposium on Security and Privacy*, 2011.
- [52] R. Shokri, G. Theodorakopoulos, C. Troncoso, J.-P. Hubaux, and J.-Y. Le Boudec. Protecting location privacy: optimal strategy against localization attacks. In *Proceedings of the ACM conference on Computer and communications security*, 2012.
- [53] G. Theodorakopoulos, R. Shokri, C. Troncoso, J.-P. Hubaux, and J.-Y. L. Boudec. Prolonging the hide-and-seek game: Optimal trajectory privacy for location-based services. In *ACM Workshop on Privacy in the Electronic Society (WPES 2014)*, 2014.
- [54] C. Troncoso and G. Danezis. The bayesian traffic analysis of mix networks. In *Proceedings of the 16th ACM conference on Computer and communications security*, 2009.
- [55] L. Zadeh. Optimality and non-scalar-valued performance criteria. *Automatic Control, IEEE Transactions on*, 8, 1963.

A Optimal Inference Attacks

Given the user’s protection mechanism p^* , the inference attack (19) is a valid strategy for the adversary, as there is no dependency between the defender and attacker strategies in the case of differential privacy metric.

However, as the differential privacy metric (used in the protection mechanism) does not include any probability distribution on secrets, we can design an inference

attack whose objective is to minimize the conditional expected error E_s :

$$E_s = \sum_{o, \hat{s}} p^*(o|s) \cdot q(\hat{s}|o) \cdot d(\hat{s}, s) \quad (26)$$

for all secrets s . This is a multi-objective optimization problem [42] that does not prefer any of the E_s (for any secret) to another. Under no such preferences, the objective is to minimize $\sum_s E_s$, using weighted sum method with equal weight for each secret.

Thus, the following linear program constitutes the optimal inference attack, under the mentioned assumptions.

$$\min_q \sum_{s, o, \hat{s}} p^*(o|s) \cdot q(\hat{s}|o) \cdot d(\hat{s}, s) \quad (27)$$

As all the weights of E_s are positive ($= 1$), the minimum of (27) is Pareto optimal [55]. Thus, minimizing (27) is sufficient for Pareto optimality. The optimal point in a multi-objective optimization (as in our case) is Pareto optimal “if there is no other point that improves at least one objective function without detriment to another function” [42, 47].

An alternative approach is to use the min-max formulation, and minimize the maximum conditional expected error E_s over all secrets s . For this, we introduce a new unknown parameter y (that will be the maximum E_s). The following linear program solves the optimal inference attack using the min-max formulation. This also provides a necessary condition for the Pareto optimality [44].

$$\min_q y \quad (28a)$$

$$\text{s. t. } \sum_{o, \hat{s}} p^*(o|s) \cdot q(\hat{s}|o) \cdot d(\hat{s}, s) \leq y, \forall s \quad (28b)$$

We can also consider the expected error conditioned on both secret s and estimate \hat{s} as the adversary’s objective to minimize. So, we can use $E_{\hat{s}, s} = \Pr\{\hat{s}|s\} \cdot d(\hat{s}, s)$ instead of $\sum_{\hat{s}} \Pr\{\hat{s}|s\} \cdot d(\hat{s}, s)$ in (26), and use the same approach as in (28). The following linear program finds the optimal inference attack that minimizes the conditional expected estimation error over all s and \hat{s} , using the min-max formulation.

$$\min_q y \quad (29a)$$

$$\text{s. t. } \sum_o p^*(o|s) \cdot q(\hat{s}|o) \cdot d(\hat{s}, s) \leq y, \forall s, \hat{s} \quad (29b)$$

Overall, we prefer the linear program (27) as it has the least number of constraints among the above three. We can also use (19) for comparison of optimal protection mechanisms based on distortion and differential metrics.

B Proof of Theorem 1

Proof. We construct (21) from (20). In (20), we condition the optimal obfuscation p^* on its corresponding optimal inference (best response) attack q^* . So, for any observable o , the inference strategy $q^*(\cdot|o)$ is the one that, by definition of the best response, minimizes the expected error

$$\sum_{\hat{s}} q(\hat{s}|o) \sum_s \pi(s) \cdot p(o|s) \cdot d(\hat{s}, s) \quad (30)$$

Thus, the privacy value (20b) to be guaranteed is

$$\begin{aligned} & \sum_{s,o,\hat{s}} \pi(s) \cdot p(o|s) \cdot q^*(\hat{s}|o) \cdot d(\hat{s}, s) = \\ & = \sum_o \min_{q(\cdot|o)} \sum_{\hat{s}} q(\hat{s}|o) \sum_s \pi(s) \cdot p(o|s) \cdot d(\hat{s}, s) \quad (31) \end{aligned}$$

Note that (30) is an average of $\sum_s \pi(s) \cdot p(o|s) \cdot d(\hat{s}, s)$ over \hat{s} , and thus it must be larger or equal to the smallest value of it for a particular \hat{s} .

$$\begin{aligned} & \min_{q(\cdot|o)} \sum_{\hat{s}} q(\hat{s}|o) \sum_s \pi(s) \cdot p(o|s) \cdot d(\hat{s}, s) \\ & \geq \min_{\hat{s}} \sum_s \pi(s) \cdot p(o|s) \cdot d(\hat{s}, s) \quad (32) \end{aligned}$$

Let $q'(\cdot|o)$ be a conditional probability distribution function such that for any given observable o ,

$$q'(s'|o) = \begin{cases} 1 & \text{if } s' = \mathbf{argmin}_{\hat{s}} \sum_s \pi(s) \cdot p(o|s) \cdot d(\hat{s}, s) \\ 0 & \text{otherwise} \end{cases} \quad (33)$$

Note that $q' \in Q$ is a pure strategy that represents one particular inference attack. Moreover, (31) constructs q^* such that it optimizes (30) over the set of all mixed strategies Q that include all the pure strategies. The minimum value for the optimization over the set of all mixed strategies is clearly less than or equal to the minimum value for the optimization over its subset (the pure strategies). Thus, the following inequality holds.

$$\begin{aligned} & \sum_{\hat{s}} q^*(\hat{s}|o) \sum_s \pi(s) \cdot p(o|s) \cdot d(\hat{s}, s) \\ & = \min_{q(\cdot|o)} \sum_{\hat{s}} q(\hat{s}|o) \sum_s \pi(s) \cdot p(o|s) \cdot d(\hat{s}, s) \\ & \leq \sum_{\hat{s}} q'(\hat{s}|o) \sum_s \pi(s) \cdot p(o|s) \cdot d(\hat{s}, s) \\ & = \min_{\hat{s}} \sum_s \pi(s) \cdot p(o|s) \cdot d(\hat{s}, s) \quad (34) \end{aligned}$$

Therefore, from inequalities (32) and (34) we have

$$\begin{aligned} & \sum_{s,o,\hat{s}} \pi(s) \cdot p(o|s) \cdot q^*(\hat{s}|o) \cdot d(\hat{s}, s) \\ & = \sum_o \min_{\hat{s}} \sum_s \pi(s) \cdot p(o|s) \cdot d(\hat{s}, s) = \sum_o x(o) \quad (35) \end{aligned}$$

where $x(o) = \min_{\hat{s}} \sum_s \pi(s) \cdot p(o|s) \cdot d(\hat{s}, s)$, or equivalently $x(o) \leq \sum_s \pi(s) \cdot p(o|s) \cdot d(\hat{s}, s), \forall \hat{s}$.

Thus, the constraint (20b) in the linear program (20) is equivalent to (and can be replaced by) the constraints (21b) and (21c) in the linear program (21). \square

C Approximating The Optimal Mechanisms

Here, we briefly discuss the computational aspects of the design of optimal protection mechanisms. Although the solution to linear programs provides us with the optimal protection mechanism, their computation cost is quadratic (for distortion mechanisms) and cubic (for differential mechanisms) in the cardinality of the set of secrets and observables. Providing privacy for a large set of secrets needs a high computation budget. To establish a balance between the computation budget and privacy requirements, we can make use of approximation techniques to design optimal protection mechanisms. We explore some possible approaches.

Linear programming [10] is one of the fundamental areas of mathematics and computer science, and there is a variety of algorithms to solve a linear program. Surveying those algorithms and evaluating their efficiencies is out of the scope of this paper. These algorithms search the set of feasible solutions of a problem for finding the optimal solution that meets the constraints. Many of these algorithms are iterative and they converge to the optimal solution as the number of iterations increases [27, 45]. Thus, a simple approximation method is to stop the iterative algorithm when our computation budget is over. Other approximation methods exist. For example, [23] suggests a sampling algorithm to select a subset of constraints in an optimization problem to speed up the computation. Moreover, we can rely on the particular structure of secrets to reduce the set of constraints [9].

We can implement those approximation techniques to solve approximately optimal protection mechanisms in an affordable time. Furthermore, we can rely on the definition of privacy to find the constraints that have a minor contribution to the design of the protection mechanism. In this section, we study one approximation

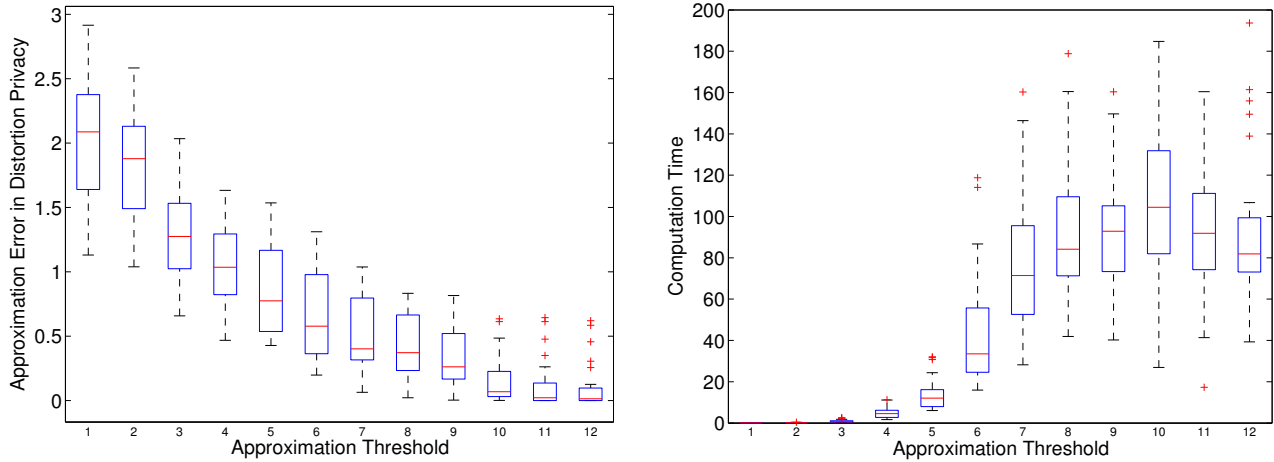


Fig. 8. Error (in km) and computation time (in sec) of approximate optimal joint distortion-differential privacy protection mechanisms, by considering a subset of optimization constraints. We consider only the constraints for which the distance $d(\cdot)$ (in km) between observation and secret and the distance between two secrets is less than the x -axis. In the left-hand side figure, the y -axis shows the distribution of the difference between privacy of users with and without approximation. In the right-hand side figure, the y -axis represents the total computation time of solving the linear program of the approximate optimal joint mechanism in Matlab on a machine with 4 core CPU model Intel(R) Xeon(R) 2.40GHz. The central mark in each box shows the median value, computed over all users. The boxplot also shows the 25th, and 75th percentiles as well as the outliers.

method, following the intuition behind the differential privacy bound: we remove the constraints for which the distance $d(s, s')$ is larger than a threshold. We can justify this by observing that, in the definition of differential privacy metric (10), the privacy is more protected when for secrets s, s' , the distance $d(s, s')$ is small. To put this in perspective, note that if we use the original definition of differential privacy, there would not be any constraint if $d(s, s') > 1$. We also apply this approximation to the distance between observables and secrets.

In Figure 8, we show the privacy loss of users as well as the speed-up of their computation due to approximation. We performed the computation on a machine with 4 core CPU model Intel(R) Xeon(R) 2.40GHz. As we increase the approximation threshold (which is the distance beyond it we ignore the constraints), the approximation error goes to zero. This suggests that, for a large set of secrets, if we choose a relatively small threshold the approximated protection mechanism provides almost the same privacy level as in the optimal solution. The computation time, however, increases as the approximation error decreases (due to increasing the approximation threshold). Figure 8 captures such a tradeoff of our approximation method.