2016 IEEE 55th Conference on Decision and Control (CDC)
ARIA Resort & Casino
December 12-14, 2016, Las Vegas, USA

Location-dependent Privacy

Fragkiskos Koufogiannis and George J. Pappas

Abstract— **We study the problem of releasing private data where the severity of the privacy concerns depends on the data itself. As a working example, we focus on the problem where a user shares an approximation of her private GPS location with a location-based service under privacy constraints that depend on the population density at user's current location itself; in densely populated areas, less noise is required to preserve privacy. We formalize this notion by extending the definition of differential privacy to *locally Lipschitz* privacy, we establish a connection between differential privacy and the eikonal equation, and we propose a method for computing such privacy-preserving mechanisms. Specifically, this connection allows existing optimized solvers to be used for numerically building private mechanisms and provides a different view of differential privacy. Our approach is illustrated in the scenario where a user within the greater Philadelphia area privately reports her location, where the privacy concerns depend on the population density.**

I. Introduction

Location-based services (LBSs) are daily used by many individuals. In a typical scenario, users retrieve their exact location using a GPS sensor, report it to a provider of an LBS, and receive information regarding this location. For example, a user might request information about nearby places of interest (POI), such as gas stations and restaurants, or subscribe to alert notifications, such as extreme weather and traffic conditions.

From a privacy point of view, reporting the exact GPS location poses a privacy threat to the users and possibly deters them from using LBSs. These privacy issues can be mitigated if users perturb their exact location before using an LBS [1]. By reporting a noisy GPS location, user's exact position cannot be confidently inferred. On the other hand, the utility users receive from using the LBS does not dramatically deteriorate when a perturbation is applied. Indeed, for example, consider a user on a highway inquiring for nearby gas stations. A perturbation of the user's location by a few miles is unlikely to significantly affect the response by such an LBS. Nonetheless, for a user within an urban environment such a perturbation possibly renders the responses from an LBS useless; within city bounds, a perturbation of the user's location by a few blocks is enough to provide privacy without significantly distorting the response of the LBS. Therefore, the amount of the perturbation varies and depends on the private location itself.

Providing privacy guarantees, especially for users' locations, has been studied in the literature. For example, authors in [2] consider *mobile* users and an adversary that, given a training set of traces, attempts to track them. The privacy is then defined by quantifying the effectiveness of the adversary's best inference attack. Another method was proposed in [3] where users *aggregate* their traces using cloaking techniques to provide privacy guarantees. Contrary to the aforementioned approaches which provide privacy guarantees against modeled adversaries, differential privacy [4] —for a literature review of results see [5]— provides strong privacy guarantees without explicitly modeling the adversary. Intuitively, differential privacy injects noise such that an adversary cannot confidently infer the private data. Using differential privacy, authors in [1] considered *stationary* users interacting with LBSs, whereas authors in [6] aggregate the traces of mobile users for traffic estimation purposes.

Our contributions are both theoretical and applied. On the application side, we consider stationary users who report their private locations with a privacy level that depends on the location itself —tighter privacy levels in more sparsely populated areas. In this setting, we extend geo-indistinguishability [1] to the case where different privacy levels are used for different regions. We propose an algorithm to numerically construct such differential private mechanisms by establishing a connection between differential privacy and the eikonal equation and, thus, leveraging existing optimized eikonal equation solvers. On the theoretical side, we formulate the problem of designing a differential private mechanism where the privacy level depends on the private data itself. Our work can be viewed as the "dual" of that in [7]. Specifically, authors in [7] consider differential private mechanisms with a constant *privacy level* that approximate queries whose *sensitivity* varies with the private data. This paper considers mechanisms that approximate identity queries —whose sensitivity is independent of the private data— under a privacy level that varies with the private data.

The paper is structured as follows. Section II motivates and informally introduces the notion of data-dependent privacy level, briefly reviews existing results in differential privacy, presents our notion of privacy as a dual to the notion of "smooth local sensitivity", and formally states the problem of designing mechanisms that satisfy this extended notion of privacy, which we call *locally Lipschitz privacy*. Next, Section III proposes an algorithm to numerically compute locally Lipschitz private mechanisms by reducing the problem to solving multiple eikonal equations and a

Authors are with the Department of Electrical and Systems Engineering, University of Pennsylvania, PA, USA.

This work was supported in part by the TerraSwarm Research Center, one of six centers supported by the STARnet phase of the Focus Center Research Program (FCRP) a Semiconductor Research Corporation program sponsored by MARCO and DARPA and in part by NSF CNS-1505799 and the Intel-NSF Partnership for Cyber-Physical Systems Security and Privacy.

linear system. Section IV illustrate our algorithm in the scenario of reporting a private GPS location to an LBS where the privacy level map is given by Philadelphia's population density. Finally, we conclude this work with a discussion and future directions in Section V.

II. Problem Formulation

In this section, we motivate and formulate the problem of designing differential private mechanisms with "input-dependent" privacy level. Initially, we informally state the problem, whereas in Subsection II-A we review the framework of differential privacy, Subsection II-B provides a formal problem statement, and Subsection II-C explores the connection to the smooth local sensitivity introduced in [7].

Let $(\mathcal{U}, \|\cdot\|)$ be a normed space which includes the set of possible private data and let $q : \mathcal{U} \to \mathcal{Y}$ be a deterministic query of interest, where $\mathcal{Y}$ is the set of possible responses. In our case, we will focus on users reporting their locations to LBSs and, thus, we will mostly focus on Euclidean spaces $\left(\mathbb{R}^2, \|\cdot\|_2\right)$ and identity queries $q(u) = u$. Moreover, we consider a *privacy level map* $\epsilon : \mathcal{U} \to \mathbb{R}_+$ where $\epsilon(u)$ quantifies the need for privacy in a neighborhood of u — smaller values of $\epsilon(u)$ correspond to stronger privacy needs. Then, we wish to design a mechanism Q which outputs a noisy approximation $y = Q(u)$ of the private data u which is *"$\epsilon(u)$-differential private around private data u"*. Note that for constant privacy level maps $\epsilon(u) = \epsilon_0$, our problem reduces to standard ϵ-differential privacy.

There are several practical scenarios where an input-dependent privacy level is meaningful. Specifically, we mention the following two examples:

- *Location-based services:* We consider users interacting with an LBS, as a running example throughout the paper. Whenever the users report their location u with ϵ-privacy, they release an approximation $y = u + V$, where the noise V is proportional to ϵ^{-1}. However, in practice, the desired privacy level ϵ depends on the location u itself. Specifically, as illustrated in Figure 1, densely-populated areas achieve sufficient privacy by using larger values of privacy level. Conversely, a user can more easily be identified in less-crowded areas unless a smaller value of privacy level ϵ is used. By allowing the privacy level to depend on the user's location itself, we can design a single mechanism that satisfies the privacy needs over all regions.
- *Data-dependent incentives:* From the system designer's perspective, having a fixed privacy level for all possible data inputs might not be possible. We depict this idea by sketching the following scenario. Assume that users' private data capture their wealth $u \in [0, 1]$, e.g., quantile of income distribution. Then, when users report their private data, people with $u \to 0$ may require a tight privacy level to protect their privacy, whereas people with $u \to 1$ might benefit by an increased accuracy of the system and, thus, require larger values for privacy levels. Since people may opt out of using such a system, having a flat privacy level is problematic. Instead, a privacy level *map* $\epsilon : [0, 1] \to \mathbb{R}_+$ captures the needs of all users.

Fig. 1: Within densely populated areas (user A), a small perturbation of the exact but private GPS location provides significant privacy. On the contrary, user B requires a larger perturbation in a sparsely-populated area. The figure is adapted from Statistics Canada.

A. Differential Privacy

Differential privacy was introduced in [4] and was, initially, stated in terms of databases. Informally, a randomized algorithm, called mechanism, is differentially private if its outcome does not change significantly for *"adjacent"* private data. Therefore, an adversary that observes the outcome of such a mechanism cannot *confidently* infer the private data. Definition 1 formalizes this concept.

Definition 1 (Differential Privacy)**.** *Consider a set $\mathcal{U}$ of possible private data, a privacy level $\epsilon > 0$, an (symmetric) adjacency relation $\mathcal{A} \subseteq \mathcal{U} \times \mathcal{U}$, and a set $\mathcal{Y}$ of possible responses. Then, the mechanism* $\mathcal{Q}$*

$$\mathcal{Q} : \mathcal{U} \to \Delta\left(\mathcal{Y}\right),$$

is ϵ-differential private if

$$\mathbb{P}\left(\mathcal{Q}(u) \in \mathcal{S}\right) \leq e^{\epsilon}\, \mathbb{P}\left(\mathcal{Q}(u') \in \mathcal{S}\right), \tag{1}$$

for all subsets $\mathcal{S} \subseteq \mathcal{Y}$ and any adjacent inputs u and u', $(u, u') \in \mathcal{A}$.

The privacy level ϵ is a constant that controls the strength of the privacy guarantees; smaller values of ϵ imply stronger privacy and, therefore, more noisy results. Additionally, the constant ϵ in Equation (1) is independent of the inputs u and u'. A frequently used adjacency relation for real-valued private data $u \in \mathbb{R}^n$ is a norm-induced one:

$$(u, u') \in \mathcal{A}_{\|\cdot\|} \;\Leftrightarrow\; \|u - u'\| \leq \alpha,$$

for some parameter $\alpha > 0$. In particular, for such adjacency relations, differential privacy can be almost equivalently

*For a set T and a rich-enough σ-algebra $\mathcal{T}$ on T, we denote the set of all probability measures on $(T, \mathcal{T})$ with $\Delta(T)$.

reduced to a Lipschitz condition [8], [9], [10], as in Definition 2. Specifically, Lipschitz privacy inherits all properties of practical interest from differential privacy with an example of this discrepancy being exploited in [11].

Definition 2 (Lipschitz Privacy)**.** *Consider the normed space $(\mathcal{U}, \|\cdot\|)$ of private data, a privacy level $\epsilon > 0$, and a set $\mathcal{Y}$ of possible responses. Then, the mechanism $\mathcal{Q} : \mathcal{U} \to \Delta(\mathcal{Y})$ is ϵ-Lipschitz private if*

$$\ln \mathbb{P}(\mathcal{Q}(u) \in \mathcal{S}) \text{ is } \epsilon\text{-Lipschitz in } u \text{ for all } \mathcal{S} \subseteq \mathcal{Y}. \quad (2)$$

B. Local Differential Privacy

In the case of users reporting their location, the private data $u \in \mathbb{R}^2$ is their GPS coordinates and we focus on mechanisms $\mathcal{Q}$ that approximate the private data itself

$$\mathcal{Q}(u) \approx u.$$

Nonetheless, the definition of Lipschitz privacy cannot directly capture the problem of input-dependent privacy level, as motivated earlier in this section. Specifically, in Equation (2), the privacy level ϵ is a uniform constant in u. Definition 3 alleviates this by using a *privacy level map* and relaxing the requirement for a uniform Lipschitz constant in Equation 2.

Definition 3 (Local Lipschtiz Privacy)**.** *Consider a normed space $(\mathcal{U}, \|\cdot\|)$ of private data, a privacy level map $\epsilon : \mathcal{U} \to \mathbb{R}_+$, and a set $\mathcal{Y}$ of possible responses. Then, the mechanism $\mathcal{Q} : \mathcal{U} \to \Delta(\mathcal{Y})$ is $\epsilon(\cdot)$-Lipschitz private if, for any $\mathcal{S} \subseteq \mathcal{Y}$, the function*

$$f_{\mathcal{S}}(u) = \ln \mathbb{P}(\mathcal{Q}(u) \in \mathcal{S})$$

is locally Lipschitz continuous with constant $\epsilon(u)$ for any $u \in \mathcal{U}$.

Locally Lipschitz privacy extends Lipschitz privacy, which implies the standard notion of differential privacy (e.g. Proposition 6 in [9]). Specifically, for constant privacy level maps $\epsilon(u) = \epsilon_0$, we retrieve Definition 2. Additionally, Proposition 4 states that, similar to differential privacy, locally Lipschitz privacy is resilient to post-processing; any further processing of the outcome of a locally Lipschitz private mechanism cannot break the privacy guarantees.

Proposition 4. *Let $\mathcal{Q} : \mathcal{U} \to \Delta(\mathcal{Y})$ be a locally Lipschitz mechanism, and $h : \mathcal{Y} \to \mathcal{Z}$ be a (possible randomized) post-processing, where $\mathcal{Y}$ and $\mathcal{Z}$ are two sets of responses. Then, the mechanism $h \circ \mathcal{Q}$ that post-process the outcome of mechanism $\mathcal{Q}$ is ϵ-locally Lipschitz private.*

Proof. The statement follows by re-writing the probability distribution of $h \circ \mathcal{Q}$ in terms of that of $\mathcal{Q}$

$$\mathbb{P}\left((h \circ \mathcal{Q})(u) \in \mathcal{S}\right) = \mathbb{P}\left(\mathcal{Q}(u) \in h^{-1}(\mathcal{S})\right)$$

and noting that the right-hand side is locally Lipschitz at u with constant $\epsilon(u)$. □

Remark 1. Similarly to the privacy level ϵ in differential privacy, the privacy level map $\epsilon : \mathcal{U} \to \mathbb{R}_+$ in Definition 3 is considered public knowledge and is a designer's choice.

In the light of Definition 3, our problem can be naturally formulated as follows.

Problem 1. *Given a set of private data $(\mathcal{U}, \|\cdot\|)$, a privacy level map $\epsilon : \mathcal{U} \to \mathbb{R}_+$, and a query $q : \mathcal{U} \to \mathcal{Y}$, design an ϵ-locally Lipschitz private mechanism $\mathcal{Q}$ that approximates q.*

C. Smooth Local Sensitivity

The notion of locally Lipschitz privacy is related to [7] which introduced the notion of *smooth local sensitivity* as a mean of building differentially private mechanisms. From a theoretical point of view, we consider our present work as the "dual" of [7]. Specifically, let $(\mathbb{R}^n, \|\cdot\|)$ be the space of private data and consider a real-valued deterministic query q which mechanism $\mathcal{Q}$ should approximate:

$$q : \mathbb{R}^n \to \mathbb{R}.$$

The Laplace mechanism [4] allows one to build a private mechanism by adding Laplace-distributed noise as in Proposition 5.

Proposition 5 (Laplace Mechanism)**.** *Consider the Laplace mechanism $\mathcal{Q}$ defined as*

$$\mathcal{Q}(u) = q(u) + V, \text{ with } V \sim Lap\left(\frac{\Delta q^{global}}{\epsilon}\right),$$

where $Lap(b)$ is the Laplace distribution with probability density function $f_V(v) = \frac{1}{2b} e^{-b|v|}$ and Δq^{global} is the global sensitivity defined as

$$\Delta q^{global} = \max_{u,u' :\, (u,u') \in \mathcal{A}} |q(u) - q(u')|.$$

Then, mechanism $\mathcal{Q}$ is ϵ-differentially private.

Proposition 5 shows that the ratio

$$\frac{\text{sensitivity}}{\text{privacy level}} = \frac{\Delta q^{\text{global}}}{\epsilon}$$

is a key quantity, determines the amount of the injected noise, and is *independent* of the input u. Work in [7] considers input-dependent noise by replacing the global sensitivity Δq^{global} by a smooth version of the local sensitivity Δq^{local}, where local sensitivity is defined as

$$\Delta q^{\text{local}}(u) = \max_{u' :\, (u,u') \in \mathcal{A}} |q(u) - q(u')|.$$

In our case, the sensitivity is independent of the input; in fact, we will later focus on identity queries which reduces local sensitivity to a constant. Nonetheless, we allow the privacy level ϵ to depend on the private data u and, thus, add input-dependent noise as well.

Although authors in [7] introduced smooth local sensitivity as a means to create less noisy but still private mechanisms, we introduce the privacy level map to increase the expressitivity of differential privacy. Moreover, authors in [7] use heavy-tailed (polynomially decaying) noise V instead of the exponentially decaying Laplace distribution. In our approach we exploit a link between differential privacy and

the eikonal equation in order to numerically design private mechanisms.

III. Eikonal-based Locally-Lipschitz Private Mechanisms

In the following we focus on Euclidean spaces $(\mathbb{R}^n, \|\cdot\|_2)$ and we focus on building locally Lipschitz private mechanisms that approximate a query $q : \mathbb{R}^n \to \mathcal{Y}$. To this end, we identify the privacy constraint of Definition 3 as an instance of the eikonal equation.

A. The Eikonal Equation

First, we provide a brief overview of the eikonal equation, a PDE that takes the form of Equation (3):

$$\|\nabla u(x)\|_2 = f(x),\ x \in \Omega \text{ and } \ u(x)|_{x\in\partial\Omega} = 0, \quad (3)$$

where $\Omega \subseteq \mathbb{R}^n$. The solution $u(x)$ of Equation (3) can be thought as the shortest path problem in the continuous domain

$$u(x) = \min_{y\in\Omega} d_f(x, y),$$

where d_f is a distance function such that $d_f(x, x + dx) \approx f(x)\,\|dx\|_2$, for small enough dx.

Although the boundary value problem in (3) does not always admit strong solutions, literature provides efficient algorithms for computing weak solutions of it. For example, authors in [12], [13] introduced the *fast-marching methods* for numerically solving such boundary value problems over discretized grids of N points with complexity $\mathcal{O}(N \log N)$. Following work provided improved algorithms for general meshes [14] and approaches with accuracy bounds [15].

By identifying the locally Lipschitz private property in Equation (3) as an eikonal equation, we leverage existing, efficient and accurate numerical solvers in order to build locally private mechanisms.

B. Computing Locally Lipschitz Private Mechanisms

Algorithm 1 proposes a technique to numerically compute locally Lipschitz private mechanisms $\mathcal{Q}$ that approximate a query q.

Theorem 6. *Let $(\mathbb{R}^n, \|\cdot\|_2)$ be the space of possible private data and let $q : \mathbb{R}^n \to \mathcal{Y}$ be a query. Then, in Algorithm 1, if*

$$w(y) \geq 0, \quad \forall y \in \mathcal{Y},$$

then, the mechanism $\mathcal{Q}$ is ϵ-locally Lipschitz private.

Proof. The proof is straightforward. The mechanism $\mathcal{Q}$ such that

$$\mathbb{P}(\mathcal{Q}(u) = y) = g(u, y) = w(y)\, e^{-f_y(u)}$$

has, by assumption, a proper probability density; $g(u, y) \geq 0$ and $\sum_{y\in\mathcal{Y}} g(u, y) = 1$. Moreover, we compute the following derivative in the weak sense

$$\|\nabla_u \ln \mathbb{P}(\mathcal{Q}(u) = y)\|_2 = \|\nabla_u (\ln\, w(y) - f_y(u))\| = \epsilon(u).$$

Therefore, mechanism $\mathcal{Q}$ satisfies Definition 3 and, thus, is ϵ-locally Lipschitz private. □

Algorithm 1 Building a mechanism that satisfies local Lipschitz privacy level map through an eikonal equation solver.

Require: Privacy level map $\epsilon : \mathbb{R}^n \to \mathbb{R}_+$ and query $q : \mathbb{R}^n \to \mathcal{Y}$.

1: **function** PrivacyMapMechanism(Privacy map $\epsilon : \mathcal{U} \to \mathbb{R}_+$, Query $q : \mathcal{U} \to \mathcal{Y}$)
2: **for** each output $y \in \mathcal{Y}$ **do**
3: Compute f_y by solving the eikonal equation problem

$$\|\nabla f_y(u)\| = \epsilon(u) \text{ with } f_y(q^{-1}(y)) = 0.$$

4: **end for**
5: Compute $w(y)$ by solving the linear system

$$\sum_{y\in\mathcal{Y}} e^{-f_y(u)}\, w(y) = 1,\ u \in \mathcal{U}.$$

6: **if** $w(y) \geq 0$, for all $y \in \mathcal{Y}$ **then**
7: Define mechanism $\mathcal{Q}$ as

$$\mathbb{P}(\mathcal{Q}(u) = y) = w(y)\, e^{-f_y(u)}.$$

8: **end if**
9: **end function**

Algorithm 1 works as follows. For each possible response $y \in \mathcal{Y}$, we solve the following boundary value problem stated in line 3, where the exact boundary condition $f_y(q^{-1}(y)) = 0$ is a design choice. This choice stems from the need that the response y should be close to u, although there is no guarantee that the mode of the resulting distribution $\mathbb{P}(\mathcal{Q}(u) = y)$ is at $y = u$. Next, line 5 of the algorithm computes the weights $w(y)$ such that for each input u, the probability $\mathbb{P}(\mathcal{Q}(u) = y) = w(y)\, e^{-f_y(u)}$ is a probability distribution. If there exists a positive solution to this linear system, then, the computed mechanism is locally Lipschitz private. As a guideline, for smooth enough privacy maps ($\|\nabla\epsilon(u)\| \ll 1$) with loose privacy at the edge of map ($\epsilon(u)|_{u\in\partial\mathcal{U}} \gg 1$) Algorithm 1 computes well-defined mechanisms.

In practice, Algorithm 1 fails when the privacy level map is not smooth enough although we do not provide sufficient conditions. Nonetheless, for a constant privacy level map, identity queries, and in the limit, we recover the Laplace mechanism.

IV. Example: GPS Location in Philadelphia

We demonstrate our technique in the scenario of users reporting their private GPS location to a location-based service (LBS). Specifically, we consider an individual in the greater Philadelphia area that observes her private position

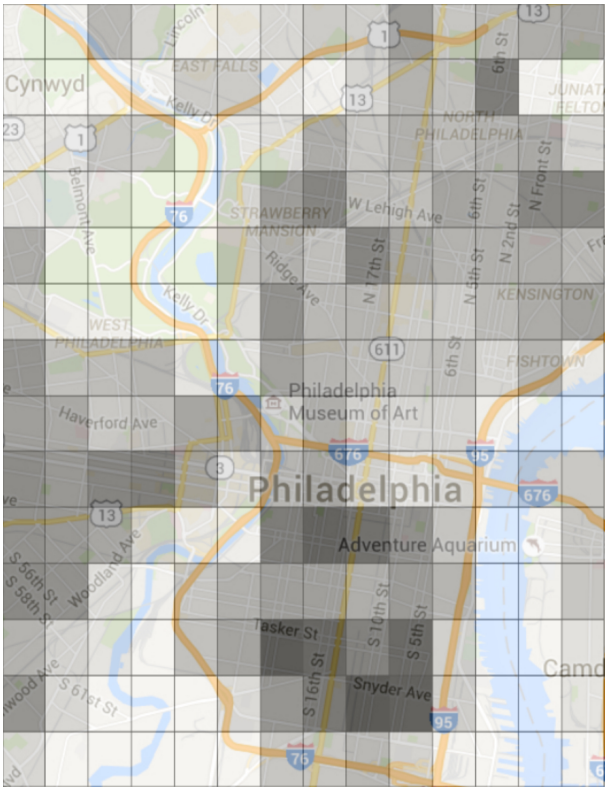

Fig. 2: The population density in Philadelphia's area is shown overlaid with the map is a publicly available knowledge and, thus, has no privacy requirements. In more densely populated areas (darker colored), the privacy level is larger and, thus, less noise is required to mitigate privacy concerns.

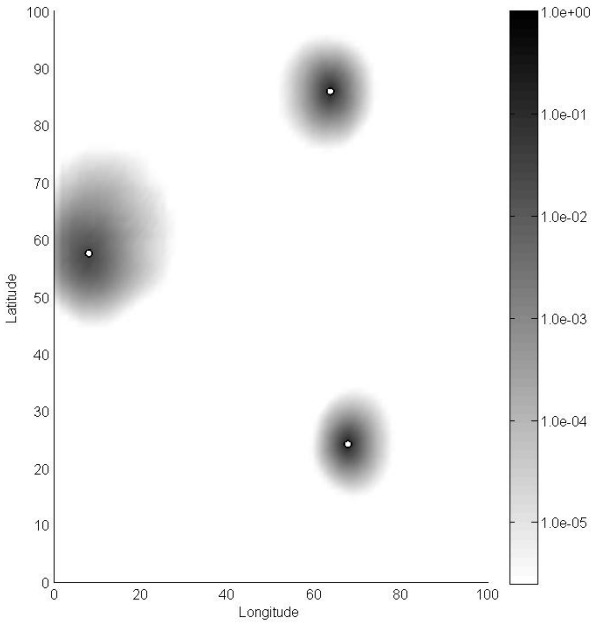

Fig. 3: The figure shows the probability distribution for three points (denoted with white circles) of high, medium, and low population density as shown in Figure 2. Dense areas have higher values of privacy level and, thus, less amount of noise is required to satisfy the privacy constraint.

$u \in \mathbb{R}^2$, reports a proxy location $y \in \mathbb{R}^2$, and receives a response from the LBS. Due to privacy concerns, the proxy location y is a perturbed version of the exact position u with probability density

$$\mathbb{P}(\mathcal{Q}(u) = y) = g(u, y).$$

Under local Lipschitz privacy, we design a privacy level map such that we provide tighter privacy ($\epsilon(u) \to 0$) in sparsely populated areas. To this end, the privacy level map is derived from the population density as

$$\epsilon(u) = 10^{-4}\, d(u) + 0.4,$$

where $d(u)$ is the population density at location u. The constant term provides the tightest possible privacy level and the linear term relaxes the privacy level in densely populated areas. The population density map is originated from the *Global Rural-Urban Mapping Project* (GRUMPv1) [16] and truly is public knowledge. GRUMPv1 provides an estimate of the population of the whole globe up to a grid size of 30 seconds of arc. —in our case, roughly $0.5\,\text{mi} \times 0.7\,\text{mi}$ rectangles. We focus on an area around Philadelphia of size about $9\,\text{mi} \times 6.2\,\text{mi}$ which is shown in Figure 2. Next, we super-sample this patch to a 200×200 grid, and, for simplicity, we re-parametrize it such that $u \in [0, 100]^2$.

We execute Algorithm 1 for the identity query $q(u) = u$ in Matlab using an eikonal equation solver [17]. Algorithm 1 can be run offline and users perturb their private locations by using the stored result. The range of values of the privacy level map is

$$\epsilon(u) \in [0.4, 2.0].$$

Figure 3 shows the probability distributions $\mathbb{P}(\mathcal{Q}(u_i) = y)$ for three different locations. Mechanism $\mathcal{Q}$ adapts to the different values of privacy level for different inputs. Therefore, our approach can provide a single privatizing mechanism without the need to explicitly partition the set of private data.

Finally, we evaluate the performance of the designed mechanism to the following two approaches. To this end, we consider a prior π on the private data u given by the population density itself

$$\pi(u) \propto d(u),$$

where d is the population density and we compute the expected mean-squared error of the mechanism $\mathcal{Q}$

$$\mathtt{mse}_{\text{eikonal}} = \int_{\mathcal{U}\times\mathcal{U}} \pi(u)\, \mathbb{P}(\mathcal{Q}(u) = y)\, \|u - y\|_2^2 \, du\, dy.$$

We compare this to the mean-squared error $\mathtt{mse}_{\text{Laplace}}$ of the Laplace mechanism with constant privacy level $\min_{u\in\mathcal{U}} \epsilon(u)$ and to the mean-squared error $\mathtt{mse}_{\text{optimal}}$ that is computed by

	Mean-squared error
$\texttt{mse}_{\text{Laplace}}$	37.5
$\texttt{mse}_{\text{eikonal}}$	5.78
$\texttt{mse}_{\text{optimal}}$	1.37

TABLE I: We evaluate the performance of the proposed approach to Laplace-based mechanism and the optimal one.

the following optimization problem

$$\begin{aligned} \underset{g:\mathbb{R}^4\to\mathbb{R}_+}{\text{minimize}} \quad & \int_{\mathbb{R}^2\times\mathbb{R}^2} g(u,y)\pi(u)\|y-u\|_2^2\, d^2u\, d^2u \\ \text{s.t.} \quad & \int_{\mathbb{R}^2} g(u,y)\, d^2y = 1,\ \forall u \\ & \|\nabla_u g(u,y)\| \le \epsilon(u) g(u,y), \end{aligned}$$

where $g(u,y) = \mathbb{P}(\mathcal{Q}_{\text{opt}}(u) = y)$. In fact, we solve a coarse (35×35) discritization of Problem IV and report the expected squared-error in Table I. As expected, a Laplace-based approach injects significant amount of noise which depends on the minimum value of the privacy level map; a single area with tight privacy requirements dramatically affects the performance of the mechanism. Moreover, post-processing the responses of our approach can further improve performance.

V. DISCUSSION

In this paper, we extended the notion of differential privacy to that of locally Lipschitz privacy, which allows the private level to depend on the private data used. We established the connection between locally Lipschitz privacy and the notion of smooth local sensitivity. Specifically, our approach allows for input-dependent privacy level, whereas, smooth local sensitivity allows for input-dependent sensitivity; the ration of these quantities defines the scale of the noise required. Next, we observed the connection of locally Lipschitz privacy to the eikonal equation and we proposed an algorithm to numerically build private mechanisms using existing optimized eikonal equation solvers. Finally, we illustrated our approach to a practical scenario where a user interacts with a location-based service.

Future work includes deriving necessary conditions on the privacy level map $\epsilon : \mathbb{R}^2 \to \mathbb{R}_+$ such that Algorithm 1 defines a proper mechanism. Another promising direction is to design the boundary conditions of the eikonal equation solver in Algorithm 1 such that the performance of the mechanism $\mathcal{Q}$ is optimized.

REFERENCES

[1] ME Andrés, NE 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*, 2013.

[2] R Shokri, G Theodorakopoulos, JY Le Boudec, and JP Hubaux. Quantifying location privacy. In *Symposium on Security and Privacy*, pages 247–262. IEEE, 2011.

[3] B Hoh, M Gruteser, R Herring, J Ban, D Work, JC Herrera, AM Bayen, M Annavaram, and Q Jacobson. Virtual trip lines for distributed privacy-preserving traffic monitoring. In *Proceedings of the 6th International Conference on Mobile Systems, Applications, and Services*, pages 15–28, 2008.

[4] 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.

[5] C Dwork and A Roth. The algorithmic foundations of differential privacy. *Theoretical Computer Science*, 9(3-4):211–407, 2013.

[6] J Le Ny, A Touati, and GJ Pappas. Real-time privacy-preserving model-based estimation of traffic flows. In *International Conference on Cyber-Physical Systems (ICCPS)*, pages 92–102. IEEE, 2014.

[7] K Nissim, S Raskhodnikova, and A Smith. Smooth sensitivity and sampling in private data analysis. In *Proceedings of the 39th Symposium on Theory of Computing*, pages 75–84. ACM, 2007.

[8] K Chatzikokolakis, ME Andrés, NE Bordenabe, and C Palamidessi. Broadening the scope of differential privacy using metrics. In *Privacy Enhancing Technologies*, pages 82–102. Springer, 2013.

[9] F Koufogiannis, S Han, and GJ Pappas. Gradual release of sensitive data under differential privacy. *arXiv preprint arXiv:1504.00429*, 2015.

[10] Y Wang, Z Huang, S Mitra, and GE Dullerud. Entropy-minimizing mechanism for differential privacy of discrete-time linear feedback systems. In *IEEE Conference on Decision and Control*, 2014.

[11] Q Geng and P Viswanath. The optimal mechanism in differential privacy. In *Information Theory (ISIT), 2014 IEEE International Symposium on*, pages 2371–2375. IEEE, 2014.

[12] JA Sethian. A fast marching level set method for monotonically advancing fronts. *Proceedings of the National Academy of Sciences*, 93(4):1591–1595, 1996.

[13] JN Tsitsiklis. Efficient algorithms for globally optimal trajectories. *Transactions on Automatic Control*, 40(9):1528–1538, 1995.

[14] JA Sethian and A Vladimirsky. Fast methods for the eikonal and related hamilton–jacobi equations on unstructured meshes. *Proceedings of the National Academy of Sciences*, 97(11):5699–5703, 2000.

[15] SM Hassouna and AA Farag. Multistencils fast marching methods: A highly accurate solution to the eikonal equation on cartesian domains. *Transactions on Pattern Analysis and Machine Intelligence*, 29(9):1563–1574, 2007.

[16] Center for International Earth Science Information Network CIESIN, International Food Policy Research Institute IFPRI, and Centro Internacional de Agricultura Tropical CIAT. Global rural-urban mapping project, version 1 (grumpv1): Urban extents grid (africa), 2011. `http://sedac.ciesin.columbia.edu/data/set/grump-v1-urban-extents/maps`; Accessed on 3/2/2016.

[17] K Dirk-Jan. Accurate fast marching. `http://www.mathworks.com/matlabcentral/fileexchange/24531-accurate-fast-marching`; Accessed on 3/2/2016.