53rd IEEE Conference on Decision and Control
December 15-17, 2014. Los Angeles, California, USA

Differentially private convex optimization with piecewise affine objectives

Shuo Han, Ufuk Topcu, George J. Pappas

Abstract— **Differential privacy is a recently proposed notion of privacy that provides strong privacy guarantees without any assumptions on the adversary. The paper studies the problem of computing a differentially private solution to convex optimization problems whose objective function is piecewise affine. Such problems are motivated by applications in which the affine functions that define the objective function contain sensitive user information. We propose several privacy preserving mechanisms and provide an analysis on the trade-offs between optimality and the level of privacy for these mechanisms. Numerical experiments are also presented to evaluate their performance in practice.**

I. Introduction

With the advance in real-time computing and sensor technology, a growing number of user-based cyber-physical systems start to utilize user data for more efficient operation. In power systems, for example, the utility company now has the capability of collecting near real-time power consumption data from individual households through advanced metering infrastructures in order to improve the demand forecast accuracy and facilitate the operation of power plants [1]. At the same time, however, individual customer is exposed to the risk that the utility company or a potential eavesdropper can learn about information that the customer did not intend to share, which may include marketable information such as the type of appliances being used or even sensitive information such as the customer's daily activities. Concerns on such privacy issues have been raised [16] and start to become one major hindrance to effective user participation [10].

Unfortunately, it has been long recognized that *ad-hoc* solutions (e.g., anonymization of user data) are inadequate to guarantee privacy due to the presence of public side information. This fact has been demonstrated through various instances such as identification of Netflix subscribers in the anonymized Netflix prize dataset through linkage with the Internet Movie Database (IMDb) [17]. Providing rigorous solutions to preserving privacy has become an active area of research. In the field of systems engineering, recent work on privacy includes, among others, privacy-preserving filtering of streaming data [14], privacy in smart metering [19], privacy in traffic monitoring [3], privacy in stochastic control [20], etc.

Recently, the notion of *differential privacy* proposed by Dwork and her collaborators has received attention due to its strong privacy guarantees [6]. The original setting assumes that the sensitive database is held by a trustworthy party (often called *curator* in related literature), and the curator needs to answer external queries (about the sensitive database) that potentially come from an adversary who is interested in learning information belonging to some user. Informally, preserving differential privacy requires that the curator must ensure that the results of the queries remain approximately unchanged if data belonging to any single user in the database are modified or removed. In other words, the adversary knows little about any single user's information from the results of queries. Interested readers can refer to recent survey papers on differential privacy for more details on this topic [5].

Aside from privacy, another important aspect to consider is the usefulness of the results of queries. In the context of systems operation, user data are often used for guiding decisions that optimize systems performance. Specifically, the "query" now becomes the *solution* to the optimization problem, whereas "user data" correspond to *parameters* that appear in the objective function and/or constraints of the optimization problem. It is conceivable that preserving user privacy will come at the cost of optimality. Indeed, without any considerations on systems performance, one could protect privacy by choosing to ignore user data, which may lead to solutions that are far from being optimal.

Several researchers have looked into the application of differential privacy to optimization problems. For example, Gupta et al. [8] have studied differential privacy in combinatorial optimization problems and derived information-theoretic bounds on the utility for a given privacy level. Among all related efforts, one that receives increasingly more attention is applying differential privacy to convex optimization problems. Convex optimization problems have traditionally been extensively studied due to the richness in related results in optimization theory and their broad applications. In the case of unconstrained convex optimization, which appears frequently in machine learning (e.g., regression problems), techniques such as output perturbation and objective perturbation have been proposed by, among others, Chaudhuri et al. [4] and Kifer et al. [13]. Huang et al. [12] have studied the setting of private distributed convex optimization, where the cost function of each agent is considered private. Very recently, Hsu et al. [11] have proposed mechanisms for solving linear programs privately using a differentially private variant of the multiplicative weights algorithm.

Rather than focusing on general convex optimization problems or even linear programs, the work in this paper studies the class of convex optimization problems whose objective function is piecewise affine, with the possibility of including linear inequality constraints. This form of optimization problems arises in applications such as ℓ_1/ℓ_∞-norm optimization and resource allocation problems. On one hand, focusing

The authors are with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104. {hanshuo,utopcu,pappasg}@seas.upenn.edu

on this particular class of problems allows us to exploit special structures that may lead to better algorithms. On the other hand, such problems can be viewed as a special form of linear programs, and it is expected that studies on this problem may lead to insights into applying differential privacy to more general linear programs.

Our major result in this paper is the introduction and analysis of several mechanisms that preserve differential privacy for convex optimization problems of this kind. These mechanisms include generic mechanisms such as the *Laplace mechanism* and the *exponential mechanism*, as well as an iterative algorithm named *differentially private subgradient method*, which is inspired by the algorithm in Hsu et al. [11] that is used for solving linear programs. We also provide theoretical analysis on the suboptimality of these mechanisms and show the trade-offs between optimality and privacy.

II. Problem statement

A. Differential privacy

In differential privacy, private user information is modeled as a set D called *database*. Denote by $\mathcal{D}$ the universe of all databases of interest. The information that we would like to obtain from a database $D \in \mathcal{D}$ is represented by a mapping called *query* $q\colon \mathcal{D} \to \mathcal{Q}$ for some target domain $\mathcal{Q}$. Since the database D contains private user information, directly making $q(D)$ available to the public may cause users in the database to lose their privacy. In order to preserve privacy, for any given query q, one needs to develop a *mechanism* $M\colon \mathcal{D} \to \mathcal{Q}$ that approximates q. In the framework of differential privacy, all mechanisms under consideration are *randomized*, i.e., for a given database, the output of such a mechanism obeys a probability distribution.

The fundamental idea of differential privacy is to translate *privacy* of an individual user in the database into *changes* in the database caused by that user (hence the name *differential*). Changes in database can be defined by a symmetric binary relation on $\mathcal{D} \times \mathcal{D}$ called *adjacency* relation, which is denoted by $\mathrm{Adj}(\cdot,\cdot)$, and any two databases that satisfy this adjacency relation are called *adjacent databases*. A differentially private mechanism must ensure that its output distribution does not vary much between two adjacent databases.

Definition 1 (Differential privacy [6])**.** A randomized mechanism $M\colon \mathcal{D} \to \mathcal{Q}$ preserves ϵ-differential privacy if for all $\mathcal{R} \subseteq \mathcal{Q}$ and all pairs of adjacent databases D and D':

$$\mathbb{P}(M(D) \in \mathcal{R}) \leq e^{\epsilon}\mathbb{P}(M(D') \in \mathcal{R}).$$

The constant $\epsilon > 0$ indicates the level of privacy: smaller ϵ implies higher level of privacy. The necessity of randomized mechanisms is evident from the definition, since the output of any non-constant deterministic mechanism will normally change with the input database.

B. Differentially private convex optimization with piecewise affine objectives

We consider minimization problems whose objective function $f\colon \mathbb{R}^d \to \mathbb{R}$ is convex and piecewise affine. Namely, the function f can be written as

$$f(x) = \max_{i=1,2,\ldots,m} \{a_i^T x + b_i\} \tag{1}$$

for some constants $\{a_i \in \mathbb{R}^d, b_i \in \mathbb{R}\}_{i=1}^m$. For generality, we also add additional linear inequality constraints that define a convex polytope $\mathcal{P}$, so that the optimization problem has the following form:

$$\min_x \quad f(x) \qquad \text{s.t.} \quad x \in \mathcal{P}. \tag{2}$$

In this paper, we restrict our attention to the case where user information is in $\{b_i\}_{i=1}^m$, so that the database $D = \{b_i\}_{i=1}^m$. Any other information, including $\{a_i\}_{i=1}^m$ and $\mathcal{P}$, is considered as public and fixed. Define the adjacency relation between two databases $D = \{b_i\}_{i=1}^m$ and $D' = \{b_i'\}_{i=1}^m$ as follows:

$$\mathrm{Adj}(D, D') \quad \text{if and only if} \quad \max_{i\in\{1,2,\ldots,m\}} |b_i - b_i'| \leq b_{\max}. \tag{3}$$

With the definition of adjacency relation, we state the problem of differentially private convex optimization as follows.

Problem 2 (Differentially private convex optimization)**.** For all problems in the form of (2), find a mechanism M that outputs an approximate optimal solution that preserves ϵ-differential privacy under the adjacency relation (3). Namely, for all $\mathcal{R} \subseteq \mathcal{P}$ and all adjacent databases D and D', the mechanism M must satisfy

$$\mathbb{P}(M(D) \in \mathcal{R}) \leq e^{\epsilon}\mathbb{P}(M(D') \in \mathcal{R}).$$

III. Examples of piecewise affine objectives

In this section, we give several examples of convex minimization problems whose objective is piecewise affine. First of all, it can be shown that both the ℓ_∞-norm and ℓ_1-norm are piecewise affine.

Example 3 (ℓ_∞-norm)**.** The ℓ_∞-norm $f(x) = \|x\|_\infty$ can be rewritten in the form of (1) consisting of $2d$ affine functions:

$$f(x) = \max_{i=1,2,\ldots,d} |x_i| = \max\{x_1, -x_1, \ldots, x_d, -x_d\}.$$

Example 4 (ℓ_1-norm)**.** The ℓ_1-norm $f(x) = \|x\|_1$ can be rewritten in the form of (1) consisting of 2^d affine functions:

$$f(x) = \sum_{i=1}^d \max\{x_i, -x_i\} = \max_{\{\alpha_i\in\{0,1\}\}_{i=1}^d} \sum_{i=1}^d (-1)^{\alpha_i} x_i.$$

Piecewise affine objectives can also appear in a particular instance of resource allocation problems.

Example 5 (Resource allocation)**.** Consider the following resource allocation problem, which is one such example where *private optimal solution* may be desired. Suppose we need to purchase a certain kind of resource and allocate it among n agents, and we need to decide the optimal amount of resource to purchase. Agent i, if being allocated z_i amount of resource, can provide utility $c_i z_i$, where c_i is its utility gain. This holds until the maximum desired resource (denoted by $\bar{z}_i$) for agent i is reached.

Suppose the total amount of resource to allocate is given as $x \geq 0$. The maximum utility gain can be determined by the optimal value of the following optimization problem

$$\max_{z} \quad c^T z \qquad \text{s.t.} \quad \mathbf{1}^T z \leq x, \qquad 0 \preceq z \preceq \bar{z}, \quad (4)$$

whose optimal value is denoted as $U(x)$. One can show that $U(x)$ is a concave and piecewise affine function by considering the dual of problem (4):

$$\begin{aligned} \min_{\lambda,\nu} \quad & \nu x + \lambda^T \bar{z} \qquad (5) \\ \text{s.t.} \quad & \nu \geq 0, \qquad \lambda \succeq 0, \qquad \lambda + \nu \mathbf{1} - c \succeq 0. \end{aligned}$$

Strong duality holds since the primal problem (4) is always feasible ($z = 0$ is a feasible solution), which allows us to redefine $U(x)$ as the optimal value of problem (5). In addition, since the optimal value of any linear program can always be attained at a vertex of the constraint polytope, we can rewrite U as the pointwise minimum of affine functions (hence U is concave):

$$U(x) = \min_{i=1,2,\ldots,m} \{\nu_i x + \lambda_i^T \bar{z}\}, \qquad (6)$$

where $\{(\nu_i, \lambda_i)\}_{i=1}^m$ are the vertices of the constraint polytope in problem (5). If we are interested in maximizing the net utility $U(x) - \mu x$ over x, where μ is the price of the resource, the problem becomes equivalent to one in the form of (2).

IV. Review: Useful tools in differential privacy

This section reviews several useful tools in differential privacy that will be used in later sections. Readers who are familiar with common mechanisms used in differential privacy may skip this section.

A. Construction of private mechanisms from existing ones

There are two very useful theorems that enable construction of new differentially private mechanisms from existing ones.

Theorem 6 (Post-processing [7])**.** *Suppose a mechanism $M\colon \mathcal{D} \to \mathcal{Q}$ preserves ϵ-differential privacy. Then for any function f, the (functional) composition $f \circ M$ also preserves ϵ-differential privacy.*

Theorem 7 (Seqential composition [7])**.** *Suppose a mechanism M_1 preserves ϵ_1-differential privacy, and another mechanism M_2 preserves ϵ_2-differential privacy. Define a new mechanism $M(D) := (M_1(D), M_2(D))$. Then the mechanism M preserves $(\epsilon_1 + \epsilon_2)$-differential privacy.*

B. Laplace mechanism

When the range of query $\mathcal{Q}$ is $\mathbb{R}$, one commonly used differentially private mechanism is the *Laplace mechanism* [6]. In this paper, we use a multidimensional generalization of the Laplace mechanism for queries that lie in $\mathbb{R}^d$. Suppose the sensitivity of query q, defined as

$$\Delta := \max_{D,D'} \|q(D) - q(D')\|_\infty ,$$

is bounded. Then one way to achieve ϵ-differential privacy is to add i.i.d. Laplace noise $\mathrm{Lap}(d\Delta/\epsilon)$ to each component of q, which is guaranteed by the sequential composition theorem (Theorem 7). However, a similar mechanism that requires less noise can be adopted in this case by using the fact that the ℓ_2-sensitivity of the query Δ_2 (defined below) is also bounded:

$$\Delta_2 := \max_{D,D'} \|q(D) - q(D')\|_2 \leq \sqrt{d}\Delta.$$

Theorem 8. *For a given query q, let $\Delta_2 = \max_{D,D'} \|q(D) - q(D')\|_2$ be the ℓ_2-sensitivity of q. Then the mechanism $M(D) = q(D) + w$, where w is a random vector whose probability distribution is proportional to $\exp(-\epsilon \|w\|_2 / \Delta_2)$, preserves ϵ-differential privacy.*

We are not aware of the name of the mechanism described in Theorem 8. Although the additive perturbation w in Theorem 8 does not follow the Laplace distribution (in fact, it follows the Gamma distribution), we will still refer to this mechanism as the *vector* Laplace mechanism due to its close resemblance to the (scalar) Laplace mechanism.

C. Exponential mechanism

Another useful and quite general mechanism is the *exponential mechanism*. This mechanism requires a scoring function $u\colon \mathcal{Q} \times \mathcal{D} \to \mathbb{R}$. For minimization problems, one can usually choose the negative objective function as the scoring function. The exponential mechanism $M_E(D; u)$ guarantees ϵ-differential privacy by randomly reporting q according to the probability density function

$$\frac{\exp(\epsilon u(q, D)/2\Delta_u)}{\int_{q' \in \mathcal{Q}} \exp(\epsilon u(q', D)/2\Delta_u)\, dq'},$$

where

$$\Delta_u := \max_x \max_{D,D'\colon \mathrm{Adj}(D,D')} |u(x, D) - u(x, D')|$$

is the (global) sensitivity of the scoring function u.

Theorem 9 (McSherry and Talwar [15])**.** *The exponential mechanism is ϵ-differentially private.*

When the range $\mathcal{Q}$ is finite, i.e., $|\mathcal{Q}| < \infty$, the exponential mechanism has the following probabilistic guarantee on the suboptimality with respect to the scoring function.

Theorem 10 (McSherry and Talwar [15])**.** *Consider the exponential mechanism $M_E(D; u)$ acting on a database D under a scoring function u. If $\mathcal{Q}$ is finite, i.e., $|\mathcal{Q}| < \infty$, then M_E satisfies*

$$\mathbb{P}\left[u_{\mathrm{opt}} - u(M_E(D; u), D) \geq \frac{2\Delta_u}{\epsilon}(\log|\mathcal{Q}| + t)\right] \leq e^{-t},$$

where $u_{\mathrm{opt}} = \max_{q \in \mathcal{Q}} u(q, D)$.

It is also possible to obtain the expected suboptimality using the following lemma.

Lemma 11. *Suppose a random variable X satisfies: (1)$X \geq 0$ and (2) $\mathbb{P}(X \geq t) \leq e^{-\alpha t}$ for some $\alpha > 0$. Then it holds that $\mathbb{E}[X] \leq 1/\alpha$.*

Proof: Use the fact that $X \geq 0$ to write $X = \int_0^\infty I(X \geq t)\, dt$. Then

$$\begin{aligned}\mathbb{E}[X] &= \mathbb{E}\left[\int_0^\infty I(X \geq t)\, dt\right] = \int_0^\infty \mathbb{E}[I(X \geq t)]\, dt \\ &= \int_0^\infty \mathbb{P}(X \geq t)\, dt \leq \int_0^\infty e^{-\alpha t}\, dt = 1/\alpha.\end{aligned}$$

■

Combine Theorem 10 and Lemma 11 to obtain the expected suboptimality.

Theorem 12. *Under the same assumptions in Theorem 10, the exponential mechanism $M_E(D; u)$ satisfies*

$$\mathbb{E}\left[u_{\text{opt}} - u(M_E(D; u), D)\right] \leq 2\Delta_u(1 + \log|\mathcal{Q}|)/\epsilon.$$

V. Privacy-preserving mechanisms

This section presents the main theoretical results of this paper. In particular, we present several mechanisms that are able to obtain a differentially private solution to problem (2). We also give suboptimality analysis for most mechanisms and show the trade-offs between optimality and privacy.

A. The Laplace mechanism acting on the problem data

One straightforward way of preserving differential privacy is to obtain the optimal solution from a privatized version of problem (2) by publishing the *entire* database D privately using the vector Laplace mechanism described in Theorem 8. Privacy is guaranteed by the post-processing rule: once the problem is privatized, obtaining the optimal solution can be viewed as post-processing and does not change the level of privacy due to Theorem 6.

Theorem 13. *The mechanism that outputs $M_P(b) = b + w_P$, where w_P is drawn from the probability density function that is proportional to $\exp(-\epsilon \|w_P\|_2 / \sqrt{m} b_{\max})$, is ϵ-differentially private.*

Proof: In this case, the query is b, whose ℓ_2-sensitivity can be obtained as $\Delta = \max_{b,b'} \|b - b'\|_2 = \sqrt{m} b_{\max}$. Combining with Theorem 8 completes the proof. ■

B. The Laplace mechanism acting on the problem solution

Another way of preserving differential privacy is to apply the vector Laplace mechanism directly on the optimal solution $x_{\text{opt}}(D)$ of the problem: $M_S(D) = x_{\text{opt}}(D) + w_S$. The additive noise w_S is drawn from the distribution proportional to $\exp(-\epsilon \|w_S\|_2 / \sqrt{d}\Delta)$, where Δ is the sensitivity of the optimal solution, i.e.,

$$\Delta = \max_{D, D' :\, \text{Adj}(D, D')} \|x_{\text{opt}}(D) - x_{\text{opt}}(D')\|_2 .$$

This mechanism is ϵ-differentially private also due to Theorem 8.

Unfortunately, it is generally difficult to analyze how the optimal solution $x_{\text{opt}}(D)$ changes with D, and hence the exact value of Δ is often unavailable. However, when the set $\mathcal{P}$ is compact, an upper bound of Δ can be given by the diameter of $\mathcal{P}$, defined as $\text{diam}(\mathcal{P}) := \max_{x,y \in \mathcal{P}} \|x - y\|_2$. Although $\text{diam}(\mathcal{P})$ is still difficult to compute for a generic set $\mathcal{P}$, there are several cases where its exact value or an upper bound can be computed efficiently. One simple case is when $\mathcal{P} = \{x : 0 \leq x_i \leq 1, i = 1, 2, \ldots, d\}$ is a hypercube and hence $\text{diam}(\mathcal{P}) = \sqrt{d}$. In the more general case where $\mathcal{P}$ is described by a set of linear inequalities, an upper bound can be obtained by computing the longest axis of the Löwner-John ellipsoid of $\mathcal{P}$, i.e., the minimum-volume ellipsoid that covers $\mathcal{P}$. The Löwner-John ellipsoid can be approximated from the maximum-volume inscribed ellipsoid, which can be obtained by solving a convex optimization problem (in particular, a semidefinite problem, cf. [2, page 414]).

Suboptimality analysis for this mechanism is given by the following theorem.

Theorem 14. *Define $G = \max_{i \in \{1,2,\ldots,n\}} \|a_i\|_2$. The expected suboptimality for the solution perturbation mechanism M_S is bounded as*

$$\mathbb{E}[f(M_S(D)) - f(x_{\text{opt}})] \leq G d^{3/2} \Delta / \epsilon.$$

Proof: Since $f(x) - f(x_{\text{opt}}) \geq 0$ for all $x \in \mathcal{P}$, we have

$$\mathbb{E}[f(M_S(D)) - f(x_{\text{opt}})] = \mathbb{E}|f(M_S(D)) - f(x_{\text{opt}})|.$$

It is not difficult to show that f is Lipschitz with G as the Lipschitz constant, i.e., $|f(x) - f(y)| \leq G\|x - y\|_2$,which leads to

$$\begin{aligned}\mathbb{E}[f(M_S(D)) - f(x_{\text{opt}})] &\leq G\mathbb{E}\|M_S(D) - x_{\text{opt}}\|_2 \\ &= G\mathbb{E}\|w_S\|_2 = Gd \cdot \sqrt{d}\Delta/\epsilon = Gd^{3/2}\Delta/\epsilon.\end{aligned}$$

■

Theorem 14 shows that the expected suboptimality grows as ϵ decreases (i.e., the level of privacy increases). The suboptimality also grows with d, which is the dimension of the decision variable x. The suboptimality does not show any explicit dependence on the number of users m, except that the constant G may change as the number of users increases.

C. The exponential mechanism

To use the exponential mechanism for privately solving minimization problems, one natural choice of the scoring function is the negative objective function $-f$. However, this choice may not work in all cases, since changes in user data can lead to an infeasible problem, which yields unbounded sensitivity. Even when the problem remains feasible, the sensitivity of the objective function with respect to data can be difficult to compute for a generic optimization problem. Nevertheless, the following shows that the sensitivity for our problem is bounded and can be computed exactly.

Lemma 15. *Suppose the scoring function is given as*

$$u(x, D) = -f(x, D) = -\max_{i=1,2,\ldots,m} \{a_i^T x + b_i\}.$$

Then the sensitivity of u for the adjacency relation defined in (3) is $\Delta_u = b_{\max}$.

The proof is omitted due to space constraints but is available in the long version of this paper [9]. As a result of Theorem 9 and Lemma 15, we know that we can achieve

ϵ-differential privacy by using the exponential mechanism given in the following theorem.

Theorem 16. *The exponential mechanism M_E, which randomly outputs $\tilde{x}_{\rm opt}$ according the probability density function*

$$\frac{\exp(-\epsilon f(\tilde{x}_{\rm opt}, D)/2b_{\max})}{\int_{x\in\mathcal{P}} \exp(-\epsilon f(x, D)/2b_{\max})\, dx}, \tag{7}$$

is ϵ-differentially private.

Remark 17. The denominator in (7) needs to remain bounded in order for (7) to be a valid probability distribution. This trivially holds when $\mathcal{P}$ is bounded. When $\mathcal{P}$ is unbounded, this can be shown by using the fact that $f(x, D)$ is affine in x and the integrand decreases exponentially fast as $\|x\| \to \infty$.

Suboptimality analysis for the exponential mechanism is given by the following theorem.

Theorem 18. *The expected suboptimality for the exponential mechanism M_E is bounded as*

$$\mathbb{E}[f(M_E(D)) - f_{\rm opt}] \le C(0,\epsilon)\cdot 2b_{\max}/\epsilon,$$

where $f_{\rm opt} = \min_{x\in\mathcal{P}} f(x, D)$ and for any $\gamma \ge 0$,

$$C(\gamma,\epsilon) = \frac{\exp(-\epsilon f_{\rm opt}/2b_{\max}) \int_{x\in\mathcal{P}:\ f(x,D)\ge f_{\rm opt}+\gamma} dx}{\int_{x\in\mathcal{P}} \exp(-\epsilon f(x, D)/2b_{\max})\, dx}.$$

Proof: We first prove that for any $\gamma \ge 0$,

$$\mathbb{P}[f(M_E(D)) - f_{\rm opt} \ge \gamma] \le C(\gamma,\epsilon)\exp(-\epsilon\gamma/2b_{\max}). \tag{8}$$

For any given $a \in \mathbb{R}$, the exponential mechanism $M_E(D)$ with scoring function u satisfies

$$\begin{aligned}\mathbb{P}[u(M_E(D)) \le a] &= \frac{\int_{x\in\mathcal{P}:u(x,D)\le a} \exp\left[\frac{\epsilon u(x,D)}{2b_{\max}}\right] dx}{\int_{x\in\mathcal{P}} \exp(\epsilon u(x, D)/2b_{\max})\, dx} \\ &\le \frac{\exp(\epsilon a/2b_{\max}) \int_{x\in\mathcal{P}:u(x,D)\le a} dx}{\int_{x\in\mathcal{P}} \exp(\epsilon u(x, D)/2b_{\max})\, dx}.\end{aligned}$$

Set $u(x, D) = -f(x, D)$ and $a = -(\gamma + f_{\rm opt})$ to obtain (8).

Note that $C(\gamma,\epsilon) \le C(0,\epsilon)$ for all $\gamma \ge 0$. Then

$$\mathbb{P}[f(M_E(D)) - f_{\rm opt} \ge \gamma] \le C(0,\epsilon)\exp(-\epsilon\gamma/2b_{\max}). \tag{9}$$

Apply Lemma 11 on (9) to complete the proof. ∎

It can be shown that $C(0,\epsilon)$ increases as ϵ decreases. Therefore, similar to the solution perturbation mechanism M_S described in Theorem 14, the expected suboptimality of M_E grows as ϵ decreases. The suboptimality does not change explicitly with the number of users m, except that $C(0,\epsilon)$ may change due to changes in $f_{\rm opt}$ and f as a result of adding more affine functions into f.

D. Private subgradient method

It is not difficult to recognize that the optimization problem (2) can be converted to a linear program by introducing additional slack variables, where the algorithm proposed by Hsu et al. [11] can be applied to obtain a differentially private solution. In particular, we note that the algorithm by Hsu et al. [11] can be viewed as a differentially private version of the subgradient method when applied to the optimization problem (2).

The subgradient method iteratively searches for an optimal solution by moving along the direction of a subgradient. Recall that g is a subgradient of f at x_0 if and only if for all x:

$$f(x) \ge f(x_0) + g^T(x - x_0). \tag{10}$$

For a convex and piecewise affine function f, its subgradient at any given x_0 can be obtained as follows. First find $k \in \{1, 2, \ldots, m\}$ such that

$$a_k^T x_0 + b_k = \max_{i=1,2,\cdots,m} \{a_i^T x_0 + b_i\}. \tag{11}$$

Then a subgradient at x_0 is a_k. It can be seen from (11) that computing subgradients requires access to the private data $\{b_i\}_{i=1}^m$. Following from Hsu et al. [11], in order to preserve privacy when applying any iterative method (such as the subgradient method), one must make sure to: (1) privatize the computation during each iteration; (2) limit the total number of iterations.

Similar to Hsu et al. [11], the exponential mechanism can be applied to privatize the computation of subgradients (Algorithm 1). Choose the scoring function $u_{\rm sub}\colon \{1, 2, \cdots, m\} \to \mathbb{R}$ as $u_{\rm sub}(i; x, D) = a_i^T x + b_i$ (in the following, we will sometimes drop D for conciseness). The sensitivity of $u_{\rm sub}$ at any given x_0, which is denoted as $\Delta_{u_{\rm sub}}(x_0)$, can be computed as

$$\max_{i\in\{1,2,\ldots,m\}} \max_{D,D'} |u_{\rm sub}(i; x_0, D) - u_{\rm sub}(i; x_0, D')| = b_{\max}.$$

Algorithm 1 ϵ-differentially private subgradient.

1) Choose the scoring function $u\colon \{1, 2, \cdots, m\} \to \mathbb{R}$ as
$$u_{\rm sub}(i; x_0) = a_i^T x_0 + b_i.$$
2) Select the index i^* using the exponential mechanism:
$$\mathbb{P}(i^* = i) \propto \exp(-\epsilon u_{\rm sub}(i; x_0)/2b_{\max}).$$
3) Output a_{i^*} as the approximate subgradient at x_0.

If the subgradient computation in the regular subgradient method is replaced by Algorithm 1, the modified subgradient method (Algorithm 2) can be shown to preserve ϵ-differential privacy using the sequential composition theorem (Theorem 7), since each iteration preserves (ϵ/k)-differential privacy and the total number of iterations is k. The following theorem bounds the expected suboptimality of Algorithm 2 after k iterations.

Theorem 19. *When the ϵ-differentially private subgradient method is applied, the expected suboptimality after k iterations is bounded as*

$$\mathbb{E}\left[\min_{i=1,2,\ldots,k} f(x^{(i)}) - f_{\rm opt}\right] \le \frac{R^2 + G^2\sum_{i=1}^k \alpha_i^2}{2\sum_{i=1}^k \alpha_i} + \bar{\gamma}(\epsilon/k), \tag{12}$$

where $R = \operatorname{diam}(\mathcal{P})$, $G = \max_{i=1,2,\ldots,m} \|a_i\|_2$, and $\bar{\gamma}(z) = 2b_{\max}(1 + \log m)/z$.

Algorithm 2 ϵ-differentially private subgradient method.

1) Choose the number of iterations k, step sizes $\{\alpha_i\}_{i=1}^k$, and $x^{(1)} \in \mathcal{P}$.
2) For $i = 1, 2, \ldots, k$, repeat:
   a) Obtain an (ϵ/k)-private subgradient $g^{(i)}$ using Algorithm 1;
   b) Update $x^{(i+1)} := x^{(i)} - \alpha_i g^{(i)}$.
3) Output $x^{(k+1)}$ as the solution.

The proof follows a similar procedure as the convergence proof for the stochastic subgradient descent method, except for the presence of additional terms due to violation of the subgradient condition (10). Due to space constraints, the complete proof is omitted and can be found in the long version of this paper [9].

Theorem 19 shows a tradeoff between privacy and suboptimality. The first term, which also appears in the convergence analysis for the regular subgradient method, implies that the optimal gap vanishes as $k \to \infty$. However, if k becomes too large, inaccuracy in the private subgradients will start to act as a dominant factor in suboptimality as the second term indicates. In particular, Theorem 19 implies that there exists an optimal number of iterations: as the number of iterations grows, the first term in (12) decreases, whereas $\bar{\gamma}(\epsilon/k)$ increases due to increased level of privacy for each iteration. Similar to previous results given by Theorem 14 and 18, the second term (due to privacy) grows as ϵ decreases. When the optimal k and $\{\alpha_i\}_{i=1}^k$ that minimize the right-hand side of (12) are used, it can be shown that the suboptimality depends on m and ϵ as $\mathcal{O}((\log m/\epsilon)^{1/3})$, which increases more slowly as ϵ decreases than the ones in Theorems 14 and 18 but has additional explicit dependence on m.

VI. Numerical experiments

A. Implementation details

In all simulations, the problem data $\{(a_i, b_i)\}_{i=1}^m$ are generated from i.i.d. Gaussian distributions. The constraint set is chosen to be a d-dimensional hypercube centered at the origin: $\mathcal{P} = \{x\colon -c \preceq x \preceq c\}$, whose diameter $\operatorname{diam}(\mathcal{P}) = 2\sqrt{d}c$. The level of privacy ϵ is set at 0.1. The expected objective value for different privacy-preserving mechanisms is approximated by the sample average from 1,000 runs.

a) Implementation of the exponential mechanism: The exponential mechanism requires drawing samples from a distribution proportional to a non-negative function. Such sampling is usually performed using Markov chain Monte Carlo (MCMC) methods [18]. In this paper, we use the Metropolis algorithm with a multivariate Gaussian proposal distribution. Due the shape of the constraint set, the covariance matrix Σ of the Gaussian distribution is chosen to be isotropic, and its magnitude is set to be proportional to the size of the constraint set: $\Sigma = \eta c I_{d\times d}$, where $I_{d\times d}$ is the $d\times d$ identity matrix, and $\eta = 0.1$. Each sample is generated by running 5,000 MCMC steps, after which the Markov chain is assumed to have reached its stationary distribution.

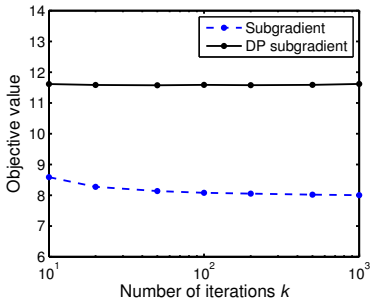

Fig. 1. Objective values as a function of the number of iterations. Blue: Objective values obtained from the regular subgradient method. Black: Objective values obtained from the differentially private subgradient method.

b) Number of iterations for the subgradient method: Although Theorem 19 implies that an optimal number of iterations exists for a given choice of ϵ, the bound is often loose for a given problem so that optimizing the bound does not provide direct guidance for choosing k. In practice, we observe that the objective value is quite robust to k, as shown in Fig. 1. The plot also includes the objective values obtained from the regular subgradient method, which decrease slightly as k grows. Due to this robustness, we choose $k = 100$ in all subsequent simulations.

B. Results and discussions

The simulations investigate the effects of changing c (the size of the constraint set $\mathcal{P}$) and m (the number of member affine functions) on all the mechanisms presented in Section V. Fig. 2 shows the expected objective value as a function of c as well as the true optimal value obtained by solving the original problem. For all privacy preserving mechanisms, the expected optimal value eventually grows as c increases, except that it shows some initial decrease for the exponential mechanism and the differentially private subgradient method. This non-monotonic behavior can be explained by noticing two factors that contribute to the objective value. One factor is the effect of c on the (original) optimization problem itself. As c increases, it leads to a more relaxed optimization problem and consequently decreases the true optimal value (magenta dashed line). Another factor of c is on the amount of perturbation introduced by the mechanisms. For example, for the mechanism that perturbs the solution directly, the magnitude of the vector Laplace noise grows with c. For the exponential mechanism, the distribution from which the solution is drawn will become less concentrated around the optimal solution as c grows. We are unable to provide a definitive explanation for the other two mechanisms, but it is expected that changes in c have a similar effect.

The effect of m on the objective value is illustrated in Fig. 3. As m increases, more affine functions will be added (i.e., the affine functions used for a smaller m is a subset of those used for a larger m). Unlike changing c, increasing m causes the objective value to monotonically increase. First of all, adding more affine functions causes the optimal value (magenta dashed line) to increase even

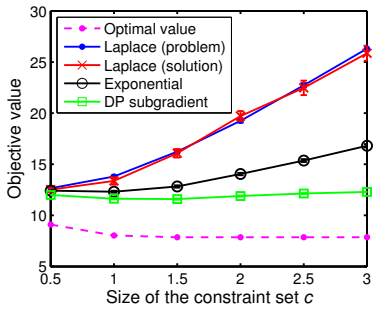

Fig. 2. Objective values (with error bars corresponding to 2σ error) as a function of c (the size of the constraint set $\mathcal{P}$) for different mechanisms. The true optimal value is shown for comparison.

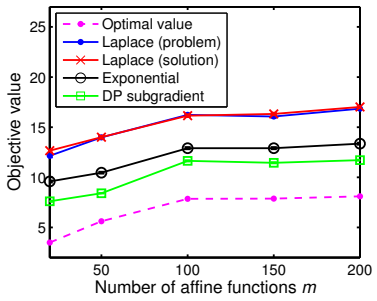

Fig. 3. Objective values (with error bars corresponding to 2σ error) as a function of m (the number of member affine functions) for different mechanisms. The true optimal value is shown for comparison.

in the absence of privacy constraints. In addition, at least for the case when the problem is perturbed, the magnitude of the vector Laplace noise grows with m. For the differentially subgradient method, increase in m also causes the suboptimality gap $\bar{\gamma}$ (that has $\log m$ dependence) to increase.

As an interesting observation from all simulations, the differentially subgradient method is superior to other mechanisms. It achieves the lowest expected suboptimality among all mechanisms. Also, when c increases, it has the slowest growth rate of suboptimality. The reason that why subgradient method works best is not evident from the suboptimality analysis presented in Section V. Although it is known that the subgradient method is quite robust to *unbiased* noise in subgradients, the noise introduced by the exponential mechanism in Algorithm 1 is biased so that the robustness analysis does not directly apply. This remains an interesting question for future investigations.

VII. Conclusions

In this paper, we study the problem of preserving differential privacy for the solution of convex optimization problems with a piecewise affine objective. Several privacy-preserving mechanisms are presented, including the Laplace mechanism applied on either the problem data or the problem solution, the exponential mechanism, and the differentially private subgradient method. Theoretical analysis on the suboptimality of these mechanisms shows the trade-offs between optimality and privacy: more privacy can be provided at the expense of sacrificing optimality. Empirical numerical experiments show that the differentially private subgradient method has the least adverse effect on optimality for a given level of privacy. In addition, it is likely that the scheme of providing privacy by iteratively solving an optimization problem privately (as used by the private subgradient method) can be applied to more general convex optimization problems beyond the specific form studied in this paper. This appears to be an interesting direction for future investigations.

Acknowledgments. The authors would like to thank Aaron Roth for providing early access to the draft on differentially private linear programming [11] and helpful discussions on differential privacy. This work was supported in part by the NSF (CNS-1239224) and TerraSwarm, one of six centers of STARnet, a Semiconductor Research Corporation program sponsored by MARCO and DARPA.

References

[1] The Benefits of Smart Meters. http://www.cpuc.ca.gov/PUC/energy/Demand+Response/benefits.htm (retrieved: March 19, 2014).

[2] S. P. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[3] E. S. Canepa and C. G. Claudel. A framework for privacy and security analysis of probe-based traffic information systems. In *ACM International Conference on High Confidence Networked Systems*, pages 25–32, 2013.

[4] K. Chaudhuri, C. Monteleoni, and A. D. Sarwate. Differentially private empirical risk minimization. *The Journal of Machine Learning Research*, 12:1069–1109, 2011.

[5] C. Dwork. Differential privacy: A survey of results. In *Theory and Applications of Models of Computation*, pages 1–19. Springer, 2008.

[6] 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.

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

[8] A. Gupta, K. Ligett, F. McSherry, A. Roth, and K. Talwar. Differentially private combinatorial optimization. In *ACM-SIAM Symposium on Discrete Algorithms*, pages 1106–1125, 2010.

[9] S. Han, U. Topcu, and G. J. Pappas. Differentially private convex optimization with piecewise affine objectives. *arXiv preprint arXiv:1403.6135*, 2014.

[10] R. Hoenkamp, G. B. Huitema, and A. J. C. de Moor-van Vugt. The neglected consumer: The case of the smart meter rollout in the Netherlands. *Renewable Energy Law & Policy Review*, page 269, 2011.

[11] J. Hsu, A. Roth, T. Roughgarden, and J. Ullman. Privately solving linear programs. *Automata, Languages, and Programming*, 8572:612–624, 2014.

[12] Z. Huang, S. Mitra, and N. Vaidya. Differentially private distributed optimization. *arXiv preprint arXiv:1401.2596*, 2014.

[13] D. Kifer, A. Smith, and A. Thakurta. Private convex empirical risk minimization and high-dimensional regression. *Journal of Machine Learning Research*, 1:41, 2012.

[14] J. Le Ny and G. Pappas. Differentially private filtering. *IEEE Transactions on Automatic Control*, 59(2):341–354, 2014.

[15] F. McSherry and K. Talwar. Mechanism design via differential privacy. In *IEEE Symposium on Foundations of Computer Science*, pages 94–103, 2007.

[16] A. Molina-Markham, P. Shenoy, K. Fu, E. Cecchet, and D. Irwin. Private memoirs of a smart meter. In *ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building*, 2010.

[17] A. Narayanan and V. Shmatikov. How to break anonymity of the Netflix prize data set. *The University of Texas at Austin*, 2007.

[18] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes: The Art of Scientific Computing*. Cambridge University Press, 3rd edition, 2007.

[19] L. Sankar, S. Rajagopalan, S. Mohajer, and H. Poor. Smart meter privacy: A theoretical framework. *IEEE Transactions on Smart Grid*, 4(2):837–846, 2013.

[20] P. Venkitasubramaniam. Privacy in stochastic control: A markov decision process perspective. In *Annual Allerton Conference on Communication, Control, and Computing*, pages 381–388, 2013.