Differentially Private Distributed Protocol for Electric Vehicle Charging

Shuo Han, Ufuk Topcu, George J. Pappas

***Abstract*— In distributed electric vehicle (EV) charging, an optimization problem is solved iteratively between a central server and the charging stations by exchanging coordination signals that are publicly available to all stations. The coordination signals depend on user demand reported by charging stations and may reveal private information of the users at the stations. From the public signals, an adversary can potentially decode private user information and put user privacy at risk. This paper develops a distributed EV charging algorithm that preserves *differential privacy*, which is a notion of privacy recently introduced and studied in theoretical computer science. The algorithm is based on the so-called *Laplace* mechanism, which perturbs the public signal with Laplace noise whose magnitude is determined by the sensitivity of the public signal with respect to changes in user information. The paper derives the sensitivity and analyzes the suboptimality of the differentially private charging algorithm. In particular, we obtain a bound on suboptimality by viewing the algorithm as an implementation of *stochastic gradient descent*. In the end, numerical experiments are performed to investigate various aspects of the algorithm when being used in practice, including the number of iterations and tradeoffs between privacy level and suboptimality.**

I. Introduction

Electric vehicles (EVs), including pure electric and hybrid plug-in vehicles, are believed to be an important component of future power systems [11]. Studies predict that the market share of EVs in the United States can reach approximately 25% by year 2020 [8]. By that time, EVs will become a significant load on the power grid [16], [1], which can lead to undesirable effects such as voltage deviations if charging of the vehicles are uncoordinated.

The key to reducing the impact of EVs on the power grid is to coordinate their charging schedules, which is often cast as an optimization problem with the objective of minimizing the peak load, power loss, or load variance [15], [2]. Due to the large number of vehicles, computing an optimal schedule for all vehicles can be very time consuming if the computation is carried out at a centralized server that collects demand information from charging stations. Instead, it is more desirable that the computation is distributed to individual charging stations. Among others, Ma et al. [12] proposed a distributed charging strategy based on the notion of valley-filling charging profiles, which is guaranteed to be optimal when all vehicles have identical (i.e., homogeneous) demand. Gan et al. [7] proposed a more general algorithm that is optimal for nonhomogeneous demand and allows asynchronous communication.

To aid the coordination among stations, the server is required to publish certain public information that is computed based on the tentative demand information collected from the charging stations. Charging demand often contains private information of the users/consumers. As a simple example, zero demand from a charging station attached to a single home unit is a good indication that the home owner is away from home. Note that the public coordination signal is received by everyone including potential adversaries who can potentially decode a consumer's private information from the public signal and put the consumer's privacy at risk.

Recently, the notion of *differential privacy* proposed by Dwork and her collaborators has received attention due to its mathematically rigorous formulation [4]. The original setting assumes that the sensitive user information is held by a trustworthy party (often called *curator* in related literature), and the curator needs to answer external queries (about the sensitive user information) that potentially come from an adversary who is interested in learning information belonging to some user. For example, in EV charging, the curator is the central server that aggregates user information, and the queries correspond to public coordination signals. 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 are modified or removed. In other words, the adversary should know 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 [3].

In this paper, we propose a differentially private algorithm for distributed EV charging. The setting in this paper is different from the setting used in the differentially private distributed (convex) optimization work by Huang et al. [10]. In the work by Huang et al. [10], it is assumed that the coordination signals *do not* change with changes in the individual cost functions, which may contain private information and needs to be protected. However, in the case of EV charging, the coordination signal is sensitive to changes in private information of the users (i.e., charging demand), so that a different mechanism is needed.

Our first major contribution is computation of the *sensitivity* of the public signal with respect to changes in private information. Computation of sensitivity is often a major challenge in constructing a differentially private mechanism; once the sensitivity is computed, standard mechanisms such as the *Laplace* and the *exponential* mechanisms can be readily implemented. For the EV charging problem, computing sensitivity of the public signal requires analysis on the global sensitivity of the optimization problem solved by individual charging stations. In the paper, we show that the global sensitivity can be computed through local variational analysis on

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

the optimality conditions. The second contribution is analysis on the suboptimality of the differentially private mechanism. In particular, we show that the private EV charging algorithm can be viewed as an implementation of *stochastic gradient descent* [13], whose suboptimality analysis has been studied extensively [17]. Similar techniques have been used in recent work by Hsu et al. [9] on privately solving linear programs using a differentially private variant of the multiplicative weights algorithm (also known as exponentiated gradient descent).

II. Background

A. Notation

Denote the ℓ_p-norm of any $x \in \mathbb{R}^n$ by $\|x\|_p$. The subscript p is dropped in the case of ℓ_2-norm. For any convex set $\mathcal{C} \subset \mathbb{R}^n$ and $x \in \mathbb{R}^n$, denote by $\Pi_{\mathcal{C}}(x)$ the projection operator that projects x onto $\mathcal{C}$ (in ℓ_2-norm). For any function f (not necessarily convex), denote by $\partial f(x)$ the set of subgradients of f at x. For any $\lambda > 0$, denote by $\mathrm{Lap}(\lambda)$ the zero-mean Laplace probability distribution such that the probability density function of $X \sim \mathrm{Lap}(\lambda)$ is $p_X(x) = \frac{1}{2\lambda}\exp(-|x|/\lambda)$. The vector consisting all ones is written as $\mathbf{1}$. The symbol $\preceq$ is used to represent element-wise inequality: for any $x, y \in \mathbb{R}^n$, we have $x \preceq y$ if and only if $x_i \le y_i$ for all $1 \le i \le n$. For any positive integer n, we denote by $[n]$ the set $\{1, 2, \ldots, n\}$.

B. Distributed electric vehicle charging

In the EV charging problem, the goal is to charge n vehicles over a horizon of T time steps with minimal influence on the power grid. For simplicity, we assume that each vehicle is handled by one charging station. For any $i \in [n]$, the vector $r_i \in \mathbb{R}^T$ represents the charging rates of vehicle i over time. In the following, we will denote by $r_i(t)$ the t-th component of r_i. Each vehicle needs to be charged a given amount of electricity $E_i > 0$ by the end of the scheduling horizon; in addition, for any $t \in [T]$, the charging rate $r_i(t)$ cannot exceed the maximum rate $\bar{r}_i(t)$ for some given constant vector $\bar{r}_i \in \mathbb{R}^T$. The constraints on r_i can be expressed as follows:

$$0 \preceq r_i \preceq \bar{r}_i, \qquad \mathbf{1}^T r_i = E_i. \tag{1}$$

For convenience, we also define the set

$$\mathcal{C}_i := \{r_i \colon r_i \text{ satisfies the constraints in (1)}\}.$$

The influence of a charging schedule $\{r_i\}_{i=1}^n$ on the power grid is quantified by a cost function $U \colon \mathbb{R}^T \to \mathbb{R}$. Formally, the EV charging problem can be cast as an optimization problem as described below:

$$\begin{aligned} \min_{\{r_i\}_{i=1}^n}.\quad & U\left(\textstyle\sum_{i=1}^n r_i\right) \\ \text{s.t.}\quad & r_i \in \mathcal{C}_i, \quad i \in [n]. \end{aligned} \tag{2}$$

Throughout the paper, we assume that the objective function U in problem (2) is convex and its gradient ∇U is L-Lipschitz in the ℓ_2-norm. This assumption holds for a number of common objectives such as minimal variance and minimal peak load. The EV charging problem (2) can be solved iteratively using distributed projected gradient descent [7] as described in Algorithm 1, which guarantees that the output converges to the optimal solution as $K \to \infty$ with proper choice of step sizes $\{\alpha_k\}_{k=1}^K$ (see [7] for details).

Algorithm 1 Distributed EV charging algorithm (with a fixed number of iterations).

Input: U, $\{\mathcal{C}_i\}_{i=1}^n$, K, and step sizes $\{\alpha_k\}_{k=1}^K$.
Output: $\{r_i^{(K+1)}\}_{i=1}^n$.
Initialize $\{r_i^{(1)}\}_{i=1}^n$ arbitrarily. For $k = 1, 2, \ldots, K$, repeat:

1) Compute $p^{(k)} := \nabla U\left(\sum_{i=1}^n r_i^{(k)}\right)$.
2) For $i \in [n]$, update $r_i^{(k+1)}$ according to

$$r_i^{(k+1)} := \Pi_{\mathcal{C}_i}(r_i^{(k)} - \alpha_k p^{(k)}). \tag{3}$$

III. A differentially private distributed charging algorithm

A. Privacy in distributed EV charging

In EV charging, both $\bar{r}_i$ and E_i can be associated with personal activities of the owner of vehicle i. For example, $\bar{r}_i(t) = 0$ may indicate that the owner is temporarily away from the charging station (which may be co-located with the owner's residence) so that the vehicle is not ready to be charged. Similarly, $E_i = 0$ may indicate that the owner is not actively using the vehicle so that the vehicle does not need to be charged. In this paper, we only address preserving the privacy of $\{E_i\}_{i=1}^n$ (i.e., assuming that $\{\bar{r}_i\}_{i=1}^n$ is public) and leave that of $\{\bar{r}_i\}_{i=1}^n$ as future work.

In the framework of differential privacy, it is assumed that the output of any algorithms that depend on user information can reveal individual's private information, even for coarse-granularity outputs that correspond to aggregation of user information. Under this assumption, the distributed EV charging algorithm (Algorithm 1) can lead to possible loss of privacy of the users who participate in the scheduling. It can be seen from Algorithm 1 that E_i affects $r_i^{(k)}$ through equation (3) and consequently also $p^{(k)}$. Since $p^{(k)}$ is broadcast publicly to every charging station, with enough side information (such as collaborating with some participating users), an adversary who is interested in learning private information about some vehicle i may be able to infer E_i from the public signals $\{p^{(k)}\}_{k=1}^K$. Therefore, making the exact gradient $p^{(k)}$ public can potentially lead to a loss of privacy.

B. Differential privacy

In this section, we will modify the original distributed charging algorithm (Algorithm 1) to preserve *differential privacy*. Before giving a formal statement of the problem, we first present some preliminaries of differential privacy. Differential privacy considers a set (called *database*) D that contains private user information to be protected. For convenience, we denote by $\mathcal{D}$ the universe of all possible databases of interest. The information that we would like to obtain from a database D is given by $q(D)$ for some mapping q (called *query*) that acts on D. In differential privacy, preserving privacy is equivalent to hiding changes in the database. Formally, changes in database can be defined by a symmetric binary relation between two databases called *adjacency* relation, which is denoted by $\mathrm{Adj}(\cdot,\cdot)$; two databases D and D' that satisfy $\mathrm{Adj}(D, D')$ are called adjacent databases.

Definition 1 (Adjacent databases)**.** Two databases $D = \{d_i\}_{i=1}^n$ and $D' = \{d'_i\}_{i=1}^n$ are said to be *adjacent* if there exists $i \in [n]$ such that $d_j = d'_j$ for all $j \neq i$.

A *mechanism* that acts on a database is said to be differentially private if it is able to ensure that two adjacent databases are nearly indistinguishable from the output of the mechanism. Usually, in order to be useful, the mechanism needs to be an approximation of the query of interest at the same time. 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.

Definition 2 (Differential privacy [4])**.** Given $\epsilon \geq 0$, a mechanism M preserves ϵ-differential privacy if for all $\mathcal{R} \subseteq \mathrm{range}(M)$ and all adjacent databases D and D' in $\mathcal{D}$, it holds that

$$\mathbb{P}(M(D) \in \mathcal{R}) \leq e^{\epsilon}\mathbb{P}(M(D') \in \mathcal{R}). \tag{4}$$

The constant ϵ indicates the level of privacy: smaller ϵ implies higher level of privacy. The notion of differential privacy promises that an adversary cannot tell from the output of M with high probability whether data corresponding to a single user in the database have changed.

C. Differentially private distributed EV charging

Recall that our goal of preserving privacy in distributed EV charging is to protect the user information $\{E_i\}_{i=1}^n$ even if an adversary can collect all public signals $\{p^{(k)}\}_{k=1}^K$. To put this under the framework of differential privacy, we define the database D as the set $\{E_i\}_{i=1}^n$ (so that the universe $\mathcal{D}$ consists of sets of size n whose elements are nonnegative numbers) and the query q as the K-tuple consisting of all the gradients $(p^{(1)}, p^{(2)}, \ldots, p^{(K)})$. Suppose for any user i, his private event can change E_i by at most some constant $E_{\max}$. Then we can define the adjacency relation as follows.

Definition 3 (Adjacency relation for EV charging)**.** For any databases $D = \{E_i\}_{i=1}^n$ and $D' = \{E'_i\}_{i=1}^n$, it holds that $\mathrm{Adj}(D, D')$ if and only if there exists $i \in [n]$ such that

$$|E_i - E'_i| \leq E_{\max}, \qquad E_j = E'_j \quad \text{for all } j \neq i.$$

The problem of designing a differentially private distributed EV charging algorithm is stated as follows.

Problem 4 (Differentially private distributed EV charging)**.** Find a randomized mechanism M_p that approximates $p = (p^{(1)}, p^{(2)}, \ldots, p^{(K)})$ and preserves ϵ-differential privacy, i.e., for any adjacent databases D and D' (as defined by Definition 3), and any $\mathcal{R} \subseteq \mathrm{range}(M_p)$, it holds that

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

Algorithm 2 Differentially private distributed EV charging.

Input: U, $\{\mathcal{C}_i\}_{i=1}^n$, K, $\{\alpha_k\}_{k=1}^K$, $\eta \geq 1$, L, Δ, and ϵ.
Output: $\{\hat{r}_i^{(K+1)}\}_{i=1}^n$.
Initialize $\{r_i^{(1)}\}_{i=1}^n$ arbitrarily. Let $\hat{r}_i^{(1)} = r_i^{(1)}$ for all $i \in [n]$ and $\theta_k = (\eta+1)/(\eta+k)$ for $k \in [K]$.
For $k = 1, 2, \ldots, K$, repeat:

1) If $k = 1$, then set $w_k = 0$; else draw a random vector $w_k \in \mathbb{R}^T$ from the distribution (proportional to) $\exp\left(-\frac{2\epsilon\|w_k\|}{K(K-1)L\Delta}\right)$
2) Compute $\hat{p}^{(k)} := \nabla U\left(\sum_{i=1}^n r_i^{(k)}\right) + w_k$.
3) For $i \in [n]$, compute:
$$r_i^{(k+1)} := \Pi_{\mathcal{C}_i}(r_i^{(k)} - \alpha_k \hat{p}^{(k)}),$$
$$\hat{r}_i^{(k+1)} := (1-\theta_k)\hat{r}_i^{(k)} + \theta_k r_i^{(k+1)}.$$

We present in Algorithm 2 a distributed EV charging protocol that preserves ϵ-differential privacy. The constant Δ is defined as

$$\Delta := \max_{i\in[n]} \max\Big\{\left\|\Pi_{\mathcal{C}_i(E_i)}(z) - \Pi_{\mathcal{C}_i(E'_i)}(z)\right\| : z \in \mathbb{R}^T,\ E_i, E'_i \text{ s.t. } |E_i - E'_i| \leq E_{\max}\Big\}.$$

For clarity, we have used the notation $\mathcal{C}_i(E_i)$ to indicate the dependence of $\mathcal{C}_i$ on E_i. Details on computing Δ will be presented later in Section IV. The purpose of the additional variables $\{\hat{r}_i^{(k)}\}_{k=1}^K$ is to implement the polynomial-decay averaging method in order to improve the convergence rate (a common practice in stochastic gradient descent [14]); introducing $\{\hat{r}_i^{(k)}\}_{k=1}^K$ does not affect privacy.

Algorithm 2 is based on a variant of the widely used Laplace mechanism in differential privacy. The Laplace mechanism requires computing the ℓ_p-sensitivity ($p \geq 1$) of a numerical query $q\colon \mathcal{D} \to \mathbb{R}^m$ (for some dimension m).

Definition 5 (ℓ_p-sensitivity)**.** For any query $q\colon \mathcal{D} \to \mathbb{R}^m$, the *$\ell_p$-sensitivty* of q under the adjacency relation Adj is defined as

$$\Delta_q := \max\{\|q(D) - q(D')\|_p : D, D' \in \mathcal{D} \text{ s.t. } \mathrm{Adj}(D, D')\}.$$

Note that the ℓ_p-sensitivity of q does not depend on a specific database D. In this paper, we will use the Laplace mechanism for bounded ℓ_2-sensitivity. The mechanism operates by introducing additive noise to q according to the ℓ_2-sensitivity of q.

Proposition 6 (Laplace mechanism [4])**.** *Consider a query $q\colon \mathcal{D} \to \mathbb{R}^m$ whose ℓ_2-sensitivity is Δ_q. Define the mechanism M_q as $M_q(D) := q(D) + w$, where w is an m-dimensional random vector whose probability density func-*

tion is given by $p_w(w) \propto \exp(-\epsilon \|w\| / \Delta_q)$. *Then, the mechanism* M_q *preserves* ϵ*-differential privacy.*

New private mechanisms can be constructed from basic mechanisms (such as the Laplace mechanism) by making use of adaptive sequential composition.

Proposition 7 (Adaptive seqential composition [5]). *Consider a sequence of mechanisms* $\{M_k\}_{k=1}^K$, *where* M_k *may depend on* $M_1, M_2, \ldots, M_{k-1}$ *as below:*

$$M_k(D) = M_k(D, M_1(D), M_2(D), \ldots, M_{k-1}(D)).$$

Suppose $M_k(\cdot, a_1, a_2, \ldots, a_{k-1})$ *preserves* ϵ_k*-differential privacy for any* $a_1 \in \mathrm{range}(M_1), \ldots, a_{k-1} \in \mathrm{range}(M_{k-1})$. *Then, the* K*-tuple mechanism* $M := (M_1, M_2, \ldots, M_K)$ *preserves* ϵ*-differential privacy for* $\epsilon = \sum_{k=1}^K \epsilon_k$.

Using the adaptive sequential composition theorem, we can show that Algorithm 2 preserves ϵ-differential privacy. The key is to compute the ℓ_2-sensitivity of $p^{(k)}$, denoted by $\Delta^{(k)}$, when the outputs of $p^{(1)}, p^{(2)}, \ldots, p^{(k-1)}$ are given, so that we can apply the Laplace mechanism on $p^{(k)}$ according to $\Delta^{(k)}$ to ensure differential privacy.

Lemma 8. *When the outputs of* $p^{(1)}, p^{(2)}, \ldots, p^{(k-1)}$ *are given, the* ℓ_2*-sensitivity of* $p^{(k)}$ *satisfies* $\Delta^{(k)} = (k-1)L\Delta$.

Proof: Consider any adjacent D and D' such that $E_j = E'_j$ for all $j \neq i$. We will first show that when the outputs of $p^{(1)}, p^{(2)}, \ldots p^{(k-1)}$ are given, we have

$$\begin{aligned} &\left\| r_i^{(k)}(D') - r_i^{(k)}(D) \right\| \leq (k-1)\Delta, \\ &\left\| r_j^{(k)}(D') - r_j^{(k)}(D) \right\| = 0, \quad \forall j \neq i. \end{aligned}$$

We will prove the above result by induction. For $k = 1$, we have $\left\| r_i^{(1)}(D') - r_i^{(1)}(D) \right\| = 0$ for all $i \in [n]$.

Consider the case when $k > 1$. For notational convenience, we define for $i \in [n]$,

$$v_i^{(k-1)}(D) := r_i^{(k-1)}(D) - \alpha_{k-1} p^{(k-1)}.$$

Here, we have used the fact that the output of $p^{(k-1)}$ is given so that $p^{(k-1)}$ does not depend on D. Then, for all $j \neq i$, we have

$$\begin{aligned} &\left\| r_j^{(k)}(D') - r_j^{(k)}(D) \right\| \\ &\quad = \left\| \Pi_{\mathcal{C}_j(E'_j)}(v_j^{(k-1)}(D')) - \Pi_{\mathcal{C}_j(E_j)}(v_j^{(k-1)}(D)) \right\| \\ &\quad = \left\| \Pi_{\mathcal{C}_j(E_j)}(v_j^{(k-1)}(D')) - \Pi_{\mathcal{C}_j(E_j)}(v_j^{(k-1)}(D)) \right\| \\ &\quad \leq \left\| v_j^{(k-1)}(D') - v_j^{(k-1)}(D) \right\| \\ &\quad = \left\| r_j^{(k-1)}(D') - r_j^{(k-1)}(D) \right\| = 0 \end{aligned}$$

and

$$\begin{aligned} &\left\| r_i^{(k)}(D') - r_i^{(k)}(D) \right\| \\ &\quad = \left\| \Pi_{\mathcal{C}_i(E'_i)}(v_i^{(k-1)}(D')) - \Pi_{\mathcal{C}_i(E_i)}(v_i^{(k-1)}(D)) \right\| \\ &\quad \leq \left\| \Pi_{\mathcal{C}_i(E'_i)}(v_i^{(k-1)}(D')) - \Pi_{\mathcal{C}_i(E_i)}(v_i^{(k-1)}(D')) \right\| \\ &\qquad + \left\| \Pi_{\mathcal{C}_i(E_i)}(v_i^{(k-1)}(D')) - \Pi_{\mathcal{C}_i(E_i)}(v_i^{(k-1)}(D)) \right\| \\ &\quad \leq \Delta + \left\| v_i^{(k-1)}(D') - v_i^{(k-1)}(D) \right\| \\ &\quad = \Delta + \left\| r_i^{(k-1)}(D') - r_i^{(k-1)}(D) \right\| \leq (k-1)\Delta, \end{aligned}$$

where we have used the induction hypothesis

$$\begin{aligned} &\left\| r_i^{(k-1)}(D') - r_i^{(k-1)}(D) \right\| \leq (k-2)\Delta, \\ &\left\| r_j^{(k-1)}(D') - r_j^{(k-1)}(D) \right\| = 0, \quad \forall j \neq i. \end{aligned}$$

Then, the ℓ_2-sensitivity of $p^{(k)}$ can be computed as follows:

$$\begin{aligned} &\left\| p^{(k)}(D') - p^{(k)}(D) \right\| \\ &\quad \leq L \left\| \sum_{i=1}^n \left[r_i^{(k)}(D') - r_i^{(k)}(D) \right] \right\| \leq L(k-1)\Delta. \end{aligned}$$

Since the above results hold for all i such that D and D' satisfy $E_j \neq E'_j$ ($j \neq i$), we have

$$\Delta^{(k)} = \max_{i \in [n]} \max_{E_i, E'_i} \left\| p^{(k)}(D') - p^{(k)}(D) \right\| = L(k-1)\Delta.$$

■

With Lemma 8 at hand, we now show that Algorithm 2 preserves ϵ-differential privacy.

Theorem 9. *Algorithm 2 ensures that* $M_p := (\hat{p}^{(1)}, \hat{p}^{(2)}, \ldots, \hat{p}^{(K)})$ *preserves* ϵ*-differential privacy under the adjacency relation as given by Definition 3.*

Proof: For any $k \in [K]$, we know that $\hat{p}^{(k)}$ satisfies ϵ_k-differential privacy, where ϵ_k satisfies $\epsilon_1 = 0$ and for $k > 1$,

$$\epsilon_k / \Delta^{(k)} = \frac{2\epsilon}{K(K-1)L\Delta}$$

according to Proposition 6. Use the expression for $\Delta^{(k)}$ from Lemma 8 to obtain

$$\epsilon_k = \frac{2(k-1)\epsilon}{K(K-1)}.$$

Using the adaptive sequential composition theorem, we know that the privacy of $M_p := (\hat{p}^{(1)}, \hat{p}^{(2)}, \ldots, \hat{p}^{(K)})$ is given by $\sum_{k=1}^K \epsilon_k = \frac{2\epsilon}{K(K-1)} \sum_{k=1}^K (k-1) = \epsilon$. ■

D. Suboptimality analysis

In Algorithm 3, we present the stochastic gradient descent method with polynomial-decay averaging for solving the following optimization problem:

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

Theorem 10 (due to Shamir and Zhang [14]) gives an upper bound of the expected suboptimality after finitely many steps

Algorithm 3 Stochastic gradient descent with polynomial-decay averaging.

Input: f, $\mathcal{X}$, K, $\{\alpha_k\}_{k=1}^K$, and $\eta \geq 1$.
Output: $\hat{x}^{(K+1)}$.
Initialize $x^{(1)}$ and $k = 1$. Let $\hat{x}^{(1)} = x^{(1)}$ and $\theta_k = (\eta + 1)/(\eta + k)$ for $k \in [K]$.
For $k = 1, 2, \ldots, K$, repeat:

1) Compute an unbiased subgradient $\hat{g}_k$ of f at $x^{(k)}$, i.e., $\mathbb{E}[\hat{g}_k] \in \partial f(x^{(k)})$.
2) Update $x^{(k+1)} := \Pi_{\mathcal{X}}(x^{(k)} - \alpha_k \hat{g}_k)$ and $\hat{x}^{(k+1)} := (1-\theta_k)\hat{x}^{(k)} + \theta_k x^{(k+1)}$.

for the stochastic gradient descent algorithm as presented in Algorithm 3.

Proposition 10 (Shamir and Zhang [14]). *Suppose $\mathcal{X} \subset \mathbb{R}^m$ is a convex set and $f\colon \mathbb{R}^m \to \mathbb{R}$ is a convex function. Assume that there exist some constants ρ and $\widehat{G}$ such that $\sup_{x,x' \in \mathcal{X}} \|x - x'\| \leq \rho$ and $\max_{1 \leq k \leq K} \mathbb{E}\|\hat{g}_k\|^2 \leq \widehat{G}^2$ for $\{\hat{g}^{(k)}\}_{k=1}^K$ given by step 1 of Algorithm 3. If the step sizes are chosen as $\alpha_k = c/\sqrt{k}$ for some constant $c > 0$, then for any $K > 1$, it holds that*

$$\mathbb{E}(f(\hat{x}^{(K+1)}) - f^*) \leq \mathcal{O}\left(\frac{\eta(\rho^2/c + c\widehat{G}^2)}{\sqrt{K}}\right), \tag{5}$$

where $f^* = \inf_{x \in \mathcal{X}} f(x)$.

A tighter upper bound can be obtained from (5) by optimizing the right-hand side of (5) over the constant c.

Corollary 11. *When $c = \rho/\widehat{G}$, the suboptimality bound (5) becomes*

$$\mathbb{E}(f(\hat{x}^{(K+1)}) - f^*) \leq \mathcal{O}\left(\eta \frac{\rho \widehat{G}}{\sqrt{K}}\right). \tag{6}$$

By applying Corollary 11, we are able obtain the bound of suboptimality for Algorithm 2.

Theorem 12. *The expected suboptimality of Algorithm 2 after K iterations is bounded as follows:*

$$\begin{aligned} &\mathbb{E}\left[U\left(\textstyle\sum_{i=1}^n \hat{r}_i^{(K+1)}\right) - U^*\right] \\ &\quad \leq \mathcal{O}\left(\eta\sqrt{n}\rho\left(\frac{G}{\sqrt{K}} + \frac{\sqrt{2T}K^{3/2}L\Delta}{2\epsilon}\right)\right), \end{aligned} \tag{7}$$

where U^* is the optimal value of problem (2), and

$$\begin{aligned} &\rho = \sqrt{\textstyle\sum_{i=1}^n \|\bar{r}_i\|^2}, \\ &G = \max\{\|\nabla U\left(\textstyle\sum_{i=1}^n r_i\right)\| \colon r_i \in \mathcal{C}_i,\ i \in [n]\}, \end{aligned}$$

Proof: In order to apply Corollary 11, we need to compute ρ and $\widehat{G}$ for Algorithm 2. The constant ρ can be obtained as $\rho = \sqrt{\sum_{i=1}^n \|\bar{r}_i\|^2}$. Recall the definition of $\widehat{G}$ as $\widehat{G}^2 := \max_k \mathbb{E}\|\hat{g}_k\|^2$, where $\hat{g}_k = [\hat{p}^{(k)}, \hat{p}^{(k)}, \ldots, \hat{p}^{(k)}]$ is formed by repeating $\hat{p}^{(k)}$ for n times, so that $\widehat{G}^2 = n \cdot \max_k \mathbb{E}\left\|\hat{p}^{(k)}\right\|^2$. Using the expression of $\hat{p}^{(k)}$, we have

$$\begin{aligned} \widehat{G} &= \sqrt{n} \cdot \max_{k \in [K]} \sqrt{\left\|p^{(k)}\right\|^2 + \mathbb{E}\|w_k\|^2} \\ &\leq \sqrt{n} \cdot \max_{k \in [K]} \left\{\|p^{(k)}\| + \sqrt{\mathbb{E}\|w_k\|^2}\right\} \\ &\leq \sqrt{n}(G + \sqrt{2T}K^2 L\Delta/2\epsilon). \end{aligned}$$

Substitute the expression of $\widehat{G}$ into (6) to obtain the result. ■

As K increases, the first term in (7) decreases, whereas the second term in (7) increases. This implies that there exists an optimal choice of K that minimizes the expected suboptimality.

Corollary 13. *The expected suboptimality of Algorithm 2 after K iterations is bounded as follows:*

$$\begin{aligned} &\mathbb{E}\left[U\left(\textstyle\sum_{i=1}^n \hat{r}_i^{(K+1)}\right) - U^*\right] \\ &\quad \leq \mathcal{O}\left(\eta T^{1/8} n^{1/2} \rho (G^3 L\Delta/\epsilon)^{1/4}\right), \end{aligned} \tag{8}$$

where U^*, ρ, and G are defined in Theorem 12.

Proof: The result can be obtained by optimizing the right-hand side of (7) over K. ■

However, since it is generally impossible to obtain a tight bound for ρ and $\widehat{G}$, optimizing K according to (7) usually does not give the best K in practice, and numerical simulation is often needed in order to find the best K for a given problem.

IV. Sensitivity computation

The bound on suboptimality of the private charging algorithm given in the previous section still depends on the unknown sensitivity Δ. The goal of this section is to obtain a tighter bound by analyzing the sensitivity of the projection operation (3) that appears in Algorithm 1.

Recall that the output of the projection operation in (3) is the optimal solution of a least-squares problem in the following form:

$$\begin{aligned} \min_x.\quad & \frac{1}{2}\|x - x_0\|^2 \\ \text{s.t.}\quad & 0 \preceq x \preceq a, \qquad \mathbf{1}^T x = b, \end{aligned} \tag{9}$$

where x_0, a, and b are given constants. For projection onto $\mathcal{C}_i$, the constants a and b are given by $a = \bar{r}_i$ and $b = E_i$. Denote the optimal solution (as a function of b) of problem (9) by $x^*(b)$. Our goal is to bound the (global) solution sensitivity with respect to b in the ℓ_2-norm, i.e.,

$$\|x^*(b_2) - x^*(b_1)\|,$$

where b_1 and b_2 are any constants such that problem (9) is feasible when b is assigned with b_1 and b is assigned with b_2. In fact, we will show a stronger result by bounding the ℓ_1-sensitivity $\|x^*(b_2) - x^*(b_1)\|_1$ and use the inequality

$$\|x^*(b_2) - x^*(b_1)\| \leq \|x^*(b_2) - x^*(b_1)\|_1$$

to obtain a bound on the ℓ_2-sensitivity. In the following, we will obtain a bound on the solution sensitivity by bounding the *local* solution sensitivity, i.e., changes in the optimal solution under infinitesimal changes in b.

A. Local solution sensitivity of optimization problems

We first review existing results on computing local solution sensitivity of nonlinear optimization problems. Consider a generic nonlinear optimization problem parametrized by θ given as follows:

$$\begin{aligned} \min_{x} \quad & f(x;\theta) \\ \text{s.t.} \quad & g_i(x;\theta) \le 0, \quad i \in [p] \\ & h_j(x;\theta) = 0, \quad j \in [q], \end{aligned} \tag{10}$$

whose Lagrangian can be expressed as

$$L(x,\lambda,\nu;\theta) = f(x;\theta) + \sum_{i=1}^{p} \lambda_i g_i(x;\theta) + \sum_{j=1}^{q} \nu_j h_j(x;\theta).$$

If there exists some set Θ such that the optimal solution is unique for all $\theta \in \Theta$, then the optimal solution as a mapping $x^*\colon \Theta \to \mathbb{R}^n$ is well-defined. This uniqueness of solution holds for problem (9) since its objective function is strictly convex. Under certain conditions presented in Theorem 14 below, the derivative of x^* with respect to θ exists. The first-order derivatives of the primal-dual optimal solution (x^*, λ^*, ν^*) are called *local solution sensitivity* and are denoted by $(\dot{x}, \dot{\lambda}, \dot{\nu})$. The local solution sensitivity can be computed using Theorem 14 that is given by Fiacco [6].

Theorem 14 (Fiacco [6])**.** *Let (x^*, λ^*, ν^*) be the primal-dual optimal solution of problem (10). Suppose the following conditions hold.*

1) *x^* is a locally unique optimal primal solution.*
2) *The functions f, $\{g_i\}_{i=1}^p$, and $\{h_j\}_{j=1}^q$ are twice continuously differentiable in x and differentiable in θ.*
3) *The gradients $\{\nabla g_i(x^*)\colon g_i(x^*) = 0,\ i \in [p]\}$ of the active constraints and the gradients $\{\nabla h_j(x^*)\colon j \in [q]\}$ are linearly independent.*
4) *Strict complementary slackness condition holds: $\lambda_i^* > 0$ when $g_i(x^*, \theta) = 0$ for all $i \in [p]$.*

Then the local sensitivity $(\dot{x}, \dot{\lambda}, \dot{\nu})$ of problem (10) exists and is continuously differentiable in a neighborhood of θ . Moreover, $(\dot{x}, \dot{\lambda}, \dot{\nu})$ is uniquely determined by the following:

$$\nabla^2 L \cdot \dot{x} + \sum_{i=1}^{p} \nabla g_i \cdot \dot{\lambda}_i + \sum_{j=1}^{q} \nabla h_j \cdot \dot{\nu}_i + \frac{\partial}{\partial \theta}(\nabla L) = 0$$

and for $i \in [p]$ and $j \in [q]$,

$$\lambda_i^* \nabla g_i \cdot \dot{x} + g_i \dot{\lambda}_i + \lambda_i^* \frac{\partial g_i}{\partial \theta} = 0,$$

$$\nabla h_j \cdot \dot{x} + \frac{\partial h_j}{\partial \theta} = 0.$$

B. Solution sensitivity of the distributed EV charging problem

After the local solution sensitivity is obtained using Theorem 14, the global solution sensitivity of problem (9) (with respect to b) can be obtained by integrating the local sensitivity (over b). When strict complementary slackness holds, the following lemma gives the properties of the local solution sensitivity of problem (9).

Lemma 15. *When strict complementary slackness condition holds, the local solution sensitivity $\dot{x} \in \mathbb{R}^T$ of problem (9) (with respect to b) satisfies the two following properties.*

1) $\mathbf{1}^T \dot{x} = 1$.
2) *There exists $\bar{x} > 0$ such that $\dot{x}_i = 0$ or $\dot{x}_i = \bar{x}$ for all $i \in [T]$.*

Proof: The Lagrangian of problem (9) is

$$L(x,\lambda,\mu,\nu) = \frac{1}{2}\|x - x_0\|^2 - \lambda^T x + \mu^T (x - a) + \nu(b - \mathbf{1}^T x).$$

It can be verified that all conditions in Theorem 14 hold. Apply Theorem 14 to obtain

$$\dot{x} - \dot{\lambda} + \dot{\mu} - \dot{\nu}\mathbf{1} = 0 \tag{11}$$

$$\mathbf{1}^T \dot{x} = 1 \tag{12}$$

$$\lambda_i^* \dot{x}_i + x_i^* \dot{\lambda}_i = 0 \tag{13}$$

$$\mu_i^* \dot{x}_i + (x_i^* - a_i)\dot{\mu}_i = 0. \tag{14}$$

From the complementary slackness condition, i.e., for all $i \in [T]$,

$$\lambda_i^* x_i^* = 0 \qquad \text{and} \qquad \mu_i^*(x_i^* - a_i) = 0,$$

we can rewrite the conditions (13) and (14) as

$$\dot{\lambda}_i \dot{x}_i = 0 \qquad \text{and} \qquad \dot{\mu}_i \dot{x}_i = 0,$$

which imply that one and only one of the following is true for any $i \in [T]$: (1) $\dot{x}_i = 0$; (2) $\dot{\lambda}_i = 0$ and $\dot{\mu}_i = 0$. Define $\mathcal{I} := \{i\colon \dot{x}_i \neq 0\}$, and we have

$$\sum_{i \in \mathcal{I}} \dot{x}_i = 1 \tag{15}$$

from (12). Note that (11) implies that for all $i, j \in [T]$,

$$\dot{x}_i - \dot{\lambda}_i + \dot{\mu}_i = \dot{x}_j - \dot{\lambda}_j + \dot{\mu}_j.$$

Since $\dot{\lambda}_i = 0$ and $\dot{\mu}_i = 0$ for all $i \in \mathcal{I}$, we have $\dot{x}_i = \dot{x}_j$ for all $i, j \in \mathcal{I}$. Therefore, there exists $\bar{x}$ such that $\dot{x}_i = \bar{x}$ for all $i \in \mathcal{I}$ and $\dot{x}_i = 0$ for all $i \notin \mathcal{I}$. The fact that $\bar{x} > 0$ follows from (15). ∎

Unfortunately, the strict complementary slackness condition does not hold for all values of b. However, the following lemma shows that the condition is only violated for a finite number of choices of b.

Lemma 16. *The set of possible values of b in problem (9) for which the strict complementary condition is violated is finite.*

Proof: The optimality conditions for problem (9) imply that

$$x^* - \lambda^* + \mu^* - \nu^* \mathbf{1} = x_0 \tag{16}$$

$$\mathbf{1}^T x^* = b \tag{17}$$

$$\lambda_i^* x_i^* = 0, \quad i \in [T] \tag{18}$$

$$\mu_i^*(x_i^* - a_i) = 0, \quad i \in [T]. \tag{19}$$

Suppose the strict complementary condition violates for a certain value of b. Denote the set of indices of the constraints that violate the strict complementary conditions by $\mathcal{I}_\lambda = \{i\colon \lambda_i^* = 0,\ x_i^* = 0\}$ and $\mathcal{I}_\mu = \{i\colon \mu_i^* = 0,\ x_i^* = a_i\}$. We will only prove for the case where $\mathcal{I}_\lambda$ is nonempty; the case where $\mathcal{I}_\mu$ is nonempty can be proved similarly.

When $\mathcal{I}_\lambda$ is non-empty, from (19), we know that $\mu_i^* = 0$ for all $i \in \mathcal{I}_\lambda$. For any $i \in \mathcal{I}_\lambda$, substitute $x_i^* = 0$, $\lambda_i^* = 0$, and $\mu_i^* = 0$ into (16) to obtain $\nu^* = x_{0,i}$. For any other $j \notin \mathcal{I}_\lambda$, one of the following three cases must hold: (1) $x_j^* = 0$; (2) $x_j^* = a_j$; (3) $0 < x_j^* < a_j$. The last case implies that $\lambda_j^* = \mu_j^* = 0$, so that we have $x_j^* = x_{0,i} + x_{0,j}$, where $i \in \mathcal{I}_\lambda$. Since both a and x_0 are fixed, we know that $b = \mathbf{1}^T x^*$ can take at most finitely many values in order for $\mathcal{I}_\lambda$ to be nonempty. ∎

An important implication of Lemma 16 is that the local solution sensitivity $\dot{x}(b)$ is Riemann integrable so that one can then obtain the global solution sensitivity through integration.

Theorem 17 (Global solution sensivitity)**.** For any b_1 and b_2 such that the problem (9) is feasible both b is assigned with b_1 and b is assigned with b_2, we have

$$\|x^*(b_2) - x^*(b_1)\|_1 = |b_2 - b_1|.$$

Proof: Without loss of generality, assume $b_2 > b_1$. Since $\dot{x}_i$ exists except at finitely many locations according to Lemma 16 and x_i^* is continuous in b, we know that $\dot{x}_i$ is Riemann integrable and

$$x_i^*(b_2) - x_i^*(b_1) = \int_{b_1}^{b_2} \dot{x}_i(b)\, db$$

according to the fundamental theorem of calculus. Hence $x_i^*(b_2) - x_i^*(b_1) \geq 0$ and

$$\begin{aligned}\sum_{i=1}^T (x_i^*(b_2) - x_i^*(b_1)) &= \sum_{i=1}^T \int_{b_1}^{b_2} \dot{x}_i(b)\, db = \int_{b_1}^{b_2} \mathbf{1}^T \dot{x}(b)\, db \\ &= b_2 - b_1 = |b_2 - b_1|.\end{aligned}$$

Therefore we have

$$\begin{aligned}\|x^*(b_2) - x^*(b_1)\|_1 &= \sum_{i=1}^T |x_i^*(b_2) - x_i^*(b_1)| \\ &= \sum_{i=1}^T (x_i^*(b_2) - x_i^*(b_1)) = |b_2 - b_1|.\end{aligned}$$

∎

Note that Theorem 17 holds for any x_0. Recall that in the EV charging problem, the optimal solution $x^*(b)$ corresponds to the projection operation $\Pi_{\mathcal{C}_i(E_i)}(\cdot)$. By applying Theorem 17, we can compute the global solution sensitivity for the EV charging problem as

$$\begin{aligned}\Delta &:= \max_{i\in[n]} \max \Big\{ \big\|\Pi_{\mathcal{C}_i(E_i)}(z) - \Pi_{\mathcal{C}_i(E_i')}(z)\big\| : \\ &\qquad z \in \mathbb{R}^T,\ E_i, E_i' \text{ s.t. } |E_i - E_i'| \leq E_{\max} \Big\} \\ &= \max_{i\in[n]} \max_{E_i, E_i'} |E_i - E_i'| = E_{\max}.\end{aligned}$$

V. Numerical simulations

In this section, we consider the case when the objective function is quadratic:

$$U\left(\textstyle\sum_{i=1}^n r_i\right) = \frac{1}{2}\left\|d + \textstyle\sum_{i=1}^n r_i\right\|^2,$$

where $d \in \mathbb{R}^T$ is a given constant vector and can be viewed as the base load profile in the context of EV charging. The goal is to minimize the variation of the total load, which is the sum of the base load and load incurred by EV charging. The Lipschitz constant L for ∇U can be obtained as $L = 1$. In all simulations, we consider the case of *homogeneous* stations, i.e., there exist $\bar{r}$ and E such that $\bar{r}_i = \bar{r}$ and $E_i = E$ for all $i \in [n]$. We choose $E_{\max} = E$.

A. Suboptimality analysis

In this setting, we have $\rho = \sqrt{n}\,\|\bar{r}\|$, $G = \|d\| + n\,\|\bar{r}\|$, and $\Delta = E$. Substitute ρ, G, and Δ into (8) to obtain

$$\begin{aligned}&\mathbb{E}\left[U\left(\textstyle\sum_{i=1}^n \hat{r}_i^{(K+1)}\right) - U^*\right] \\ &\quad\leq \mathcal{O}\left(\eta T^{1/8} n\,\|\bar{r}\|\,(\|d\| + n\,\|\bar{r}\|)^{3/4} (E/\epsilon)^{1/4}\right) \\ &\quad\leq \mathcal{O}\left(\eta T^{1/8} n\,\|\bar{r}\|\,(n\,\|\bar{r}\|)^{3/4} (E/\epsilon)^{1/4}\right).\end{aligned} \tag{20}$$

Note that the growth of the right-hand side with n is due to the fact that the objective function U also grows with n. In order to eliminate the dependence of U on n, we normalize all load/demand profiles with respect to n. Namely, there exist constants d_u, $\bar{r}_u$ and E_u (which correspond to unnormalized quantities and do not change with n) such that $d = d_u/n$, $\bar{r} = \bar{r}_u/n$ and $E = E_u/n$. Rewrite (20) in d_u, $\bar{r}_u$ and E_u to obtain

$$\begin{aligned}&\mathbb{E}\left[U\left(\textstyle\sum_{i=1}^n \hat{r}_i^{(K+1)}\right) - U^*\right] \\ &\quad\leq \mathcal{O}\left(\eta T^{1/8} n^{-1/4}\,\|\bar{r}_u\|^{7/4}\,(E_u/\epsilon)^{1/4}\right),\end{aligned} \tag{21}$$

i.e., the suboptimality decreases with n and ϵ at the rate $\mathcal{O}((n\epsilon)^{-1/4})$.

B. Results and discussions

We use the data in Gan et al. [7]. The scheduling horizon is divided into 52 time slots of 15 minutes. The maximum charging rate is 3.3 kW for all t, and each vehicle needs 10 kWh of electricity. For convenience, we normalize

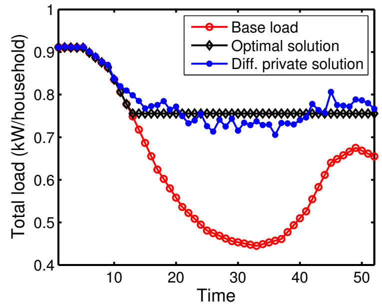

Fig. 1. A typical output of the differentially private distributed EV charging algorithm (Algorithm 2) compared to the optimal solution of problem (2). The number of iterations $K = 4$ and $\epsilon = 0.1$.

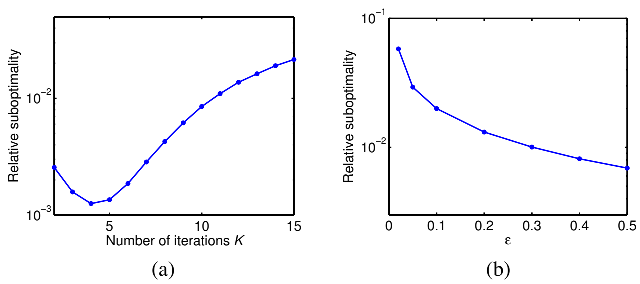

Fig. 2. (a) Relative suboptimality of the distributed EV charging algorithm (Algorithm 2) as a function of the number of iterations K for $\epsilon = 0.1$; (b) Relative suboptimality of Algorithm 2 as a function of ϵ (larger ϵ implies less privacy). The number of iterations K is optimized for every choice of ϵ. All experiments use $c = 0.5/n$.

power/electricity against the total number of households m, so that

$$\bar{r}(t) \text{ [kW/household]} = \frac{3.3 \text{ [kW]}}{m}, \quad t \in [T]$$
$$E \text{ [kW/household]} = \frac{10 \text{ [kWh]}}{m \cdot \Delta T \text{ [h]}}.$$

We consider a large pool of EVs ($n = 100,000$) in a large residential area ($m = 500,000$). A typical output from Algorithm 2 is plotted in Fig. 1 and is compared with the optimal solution. Due to the noise introduced in the gradient, the differentially private solution exhibits some fluctuations compared to the optimal solution. Fig. 2a shows the relative suboptimality as a function of K, where the relative suboptimality is computed by normalizing against the optimal value of problem (2). It can be seen from Fig. 2a that an optimal choice of K exists, which coincides with the the analysis at the end of Section III. In the end, Fig. 2b shows the dependence of relative suboptimality on ϵ. As the privacy requirement becomes less stringent (i.e., as ϵ grows), the suboptimality of Algorithm 2 decreases.

VI. Conclusions

This paper introduces an ϵ-differentially private algorithm for distributed EV charging based on a modification of the original distributed charging algorithm proposed by Gan et al. [7]. The algorithm preserves privacy by publishing the public coordination signal using the Laplace mechanism, which perturbs the public signal with Laplace noise whose magnitude is determined by the global sensitivity of the optimization problem solved by individual charging stations. The paper shows that the global sensitivity can be computed by extending local sensitivity analysis on the optimality conditions. By viewing the private algorithm as an implementation of stochastic gradient descent, we can obtain the expected suboptimality of the differentially private charging algorithm by applying existing results on suboptimality analysis of stochastic gradient descent. In the homogenous case, the suboptimality of the ϵ-differentially private algorithm with n participating users scales as $\mathcal{O}((n\epsilon)^{-1/4})$. Both theoretical analysis and numerical experiments show that there exists an optimal choice of the number iterations: too few iterations affects the convergence behavior, whereas too many iterations lead to too much noise in the coordination signal.

Acknowledgments. 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] K. Clement-Nyns, E. Haesen, and J. Driesen. The impact of charging plug-in hybrid electric vehicles on a residential distribution grid. *IEEE Transactions on Power Systems*, 25(1):371–380, 2010.

[2] S. Deilami, A. S. Masoum, P. S. Moses, and M. A. Masoum. Real-time coordination of plug-in electric vehicle charging in smart grids to minimize power losses and improve voltage profile. *IEEE Transactions on Smart Grid*, 2(3):456–467, 2011.

[3] C. Dwork. Differential privacy: A survey of results. In *Theory and Applications of Models of Computation*, pages 1–19. Springer, 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] A. V. Fiacco. Sensitivity analysis for nonlinear programming using penalty methods. *Mathematical programming*, 10(1):287–311, 1976.

[7] L. Gan, U. Topcu, and S. Low. Optimal decentralized protocol for electric vehicle charging. *IEEE Transactions on Power Systems*, 28(2):940–951, 2013.

[8] S. W. Hadley and A. A. Tsvetkova. Potential impacts of plug-in hybrid electric vehicles on regional power generation. *The Electricity Journal*, 22(10):56–68, 2009.

[9] J. Hsu, A. Roth, T. Roughgarden, and J. Ullman. Privately solving linear programs. *arXiv preprint arXiv:1402.3631*, 2014.

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

[11] A. Ipakchi and F. Albuyeh. Grid of the future. *IEEE Power and Energy Magazine*, 7(2):52–62, 2009.

[12] Z. Ma, D. S. Callaway, and I. A. Hiskens. Decentralized charging control of large populations of plug-in electric vehicles. *IEEE Transactions on Control Systems Technology*, 21(1):67–78, 2013.

[13] H. Robbins and S. Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951.

[14] O. Shamir and T. Zhang. Stochastic gradient descent for non-smooth optimization: Convergence results and optimal averaging schemes. In *Proceedings of the 30th International Conference on Machine Learning (ICML)*, pages 71–79, 2013.

[15] E. Sortomme, M. M. Hindi, S. J. MacPherson, and S. Venkata. Coordinated charging of plug-in hybrid electric vehicles to minimize distribution system losses. *IEEE Transactions on Smart Grid*, 2(1):198–205, 2011.

[16] J. Taylor, A. Maitra, M. Alexander, D. Brooks, and M. Duvall. Evaluations of plug-in electric vehicle distribution system impacts. In *IEEE Power and Energy Society General Meeting*, pages 1–6. IEEE, 2010.

[17] G. G. Yin and H. J. Kushner. *Stochastic approximation and recursive algorithms and applications*. Springer, 2003.