An Approximately Truthful Mechanism for Electric Vehicle Charging via Joint Differential Privacy

Shuo Han, Ufuk Topcu, George J. Pappas

Abstract— In electric vehicle (EV) charging, the goal is to compute a charging schedule that meets all user specifications while minimizing the influence on the power grid. Usually, an optimal schedule is computed by a central authority (called mediator) according to the specifications reported by the users. A desirable property of this procedure is to ensure that participating users truthfully report their specifications rather than maliciously manipulate the scheduling process by misreporting. In this work, we show that approximate truthfulness can be attained by adopting the popular notion of (joint) differential privacy. Joint differential privacy can limit the power of each user in manipulating the scheduling process by remaining insensitive to changes in user specifications. As a result, a user does not benefit much from misreporting his specifications, which leads to truth-telling behaviors.

I. Introduction

Electric vehicles (EVs) are expected to put significant stress on the power grid in the near future [1], [12]. The stress not only comes from the aggregate demand requested by the vehicles, but also from the fact that most charging activities naturally happen around the same time (in most cases, after rush hours). The key to reducing the influence of EVs on the power grid is to make use of the intrinsic flexibility in vehicle charging. In most cases, users only require that the vehicles should be charged before a certain deadline, which makes it possible to shift the load in time in order to avoid simultaneous charging a large number of vehicles.

This paper considers the scenario of *direct load control*, in which users report their charging specifications to a centralized authority (called mediator), who then computes a coordinated charging schedule over a certain time period for all the participating users. This scenario is different from *indirect load control*, in which a pricing scheme is used to indirectly regulate the charging activities. User specifications include the total charging demand and the maximum charging rates over the given time period. Computation of the charging schedule is cast as an optimization problem, where the objective can be minimizing the peak load, power loss, or load variance [2], [10]. So far, researchers have proposed various efficient and scalable algorithms that are able to handle a large number of vehicles [5], [8].

There are two major concerns in the direct load control scenario. One is the privacy of participating users. It is not difficult to see that user specifications are strongly correlated to daily activities and life patterns, which are normally considered as private information by the users. As a simple example, zero demand from a charging station attached to a single home unit may be a good indication that the home owner is away from home. Trivially, if the mediator is malicious, then reporting the true specifications to the mediator leads to an immediate privacy breach. Perhaps a more surprising fact is that user privacy can still be at risk even if the mediator is trustworthy. This is possible due to the fact that the aggregate load (which is assumed visible to the public) still depends on user specifications. From the aggregate load, an adversary can still potentially decode private information of a single user by collaborating with the remaining participating users. Another concern is that users may untruthfully report their specifications and/or deviate from the charging schedule as computed by the mediator. This may happen especially if the user is able to benefit (e.g., reduce electricity payment) from untruthful behaviors. Such untruthful behaviors may lead to an increase in the social cost and diminish the benefit of coordination.

A connection between these two seemingly unrelated concerns of privacy and truthfulness has been recently discovered in the study of *differential privacy*, which is a rigorous and quantitative framework for database privacy. The original purpose of differential privacy is to protect sensitive user information in the database from potential adversaries [3]. Recently, it has been proposed that differential privacy can also be used as a *mechanism design tool* to promote truthful behaviors in games where a mediator is present [7]. In particular, the notion of differential privacy has been extended to a related one named *joint* differential privacy, which ensures that misreport from any single user cannot significantly influence other users' assignments given by the mediator. In other words, joint differential privacy limits the power of any user manipulating the coordination process.

Contribution: In this paper, we apply the concept of joint differential privacy to develop an EV charging mechanism that ensures approximate truthfulness. There are two major results in the paper. Firstly, we show that joint differential privacy can be applied to ensure η-approximate truthfulness, and we derive bound on η for the EV charging problem. Previous work on the application of joint differential privacy, such as the work by Rogers et al. [9] on routing games, does not directly apply, mainly because the cost function in EV charging also depends on the mediator's assignment (in particular, an additional penalty term to prevent users from deviating from the given assignment). Secondly, we present a

charging mechanism that is able to achieve joint differential privacy, based on our previous work [6] on differentially private distributed EV charging.

Paper organization: The paper is organized as follows. Section II introduces the EV charging problem considered in this paper. In particular, we consider the scenario where a mediator is present to coordinate the charging schedule for all participating users according to the charging specifications reported by the users. We assume that the users are selfish and would like to minimize their cost (i.e., monetary payment for electricity usage) by possibly misreporting their specifications and/or ignoring the mediator's assignment. This naturally defines a game among all users, which we call the *mediator induced EV charging game*. Section III presents one main result of the paper on achieving truthfulness via joint differential privacy. It shows that approximately truthful behavior of the users can be attained if the mediator computes the charging schedule using a joint differentially private mechanism. The result also shows the dependence of truthfulness on the level of joint differential privacy. Section IV presents another main result of the paper on an algorithm that can be used by the mediator to ensure joint differential privacy. An analysis of the algorithm on the tradeoffs between suboptimality and truthfulness is also presented.

II. MEDIATOR INDUCED EV CHARGING GAME

A. Notation

Denote the ℓ_p-norm of any $x \in \mathbb{R}^n$ by $\|x\|_p$. The subscript p is dropped for $p = 2$. 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 $x_i \leq y_i$ for all $1 \leq i \leq n$. For any given set $\mathcal{U}$, positive integer $n \in \mathbb{Z}_{++}$, and $\{u_i\}_{i=1}^n$ such that $u_i \in \mathcal{U}$, consider the n-tuple $u := (u_1, u_2, \ldots, u_n) \in \mathcal{U}^n$. We use the notation u_{-i} to represent the tuple

$$(u_1, \ldots, u_{i-1}, u_{i+1}, \ldots, u_n)$$

generated by removing u_i from u. For any $v \in \mathcal{U}$, we use the notation (v, u_{-i}) to represent the tuple

$$(u_1, \ldots, u_{i-1}, v, u_{i+1}, \ldots, u_n)$$

formed by inserting v into u_{-i} when the position of insertion is clear from the context.

B. Electric vehicle charging with a mediator

We consider the EV charging problem consisting of n participating users over a horizon of T time steps. For simplicity, we assume that each user has only one vehicle to charge, and we will use the terms "user" and "vehicle" interchangeably hereafter. For user i, denote by $r_i \in \mathbb{R}^T$ the temporal *user charging profile* over the horizon. The tuple $r = (r_1, r_2, \ldots, r_i)$ is called a *charging schedule*. By the end of the time horizon, each user i needs to charge his vehicle by a total amount of $E_i \in \mathbb{R}$. In addition, each user i can specify his maximum charging rates over time as a vector $\bar{r}_i \in \mathbb{R}^T$ so that r_i does not exceed $\bar{r}_i$. Both E_i and

$\bar{r}_i$ constitute the *charging specifications* of user i, which can be written as the following constraints:

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

For convenience, we also define

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

In this paper, we consider the scenario of *direct load control*, where there exists a central mediator who is responsible for coordinating the charging activities. The mediator serves two purposes. Firstly, it needs to collect the charging specifications from users and compute a charging schedule that meets the specifications. Secondly, it needs to collect monetary payments from users according to their electricity usage. In this process, we do not assume that each user always reports his true specifications to the mediator, nor do we assume that each user will follow the charging profile computed by the mediator. Suppose that the charging schedule computed by the mediator is $r = (r_1, r_2, \ldots, r_n)$, but user i decides to use r_i' as his actual charging profile instead of r_i. Then the cost of user i is given by

$$c(i, r_i', r) = \left[\mu \left(r_i' + \sum_{j \neq i} r_j \right) + \bar{p} \right]^T r_i' + \lambda \|r_i' - r_i\|^2, \tag{2}$$

for some constants $\mu, \lambda > 0$ and $\bar{p} \in \mathbb{R}_+^T$. The first term in (2) corresponds to the cost of electricity, where $\bar{p}$ is the base electricity price, and μ indicates the increase in electricity price caused by the additional load from electric vehicles; the second term in (2) penalizes user i for deviating from the assigned charging profile r_i computed by the mediator.

C. Mediator induced EV charging game

We assume that the users are selfish; each user is interested in minimizing his cost by possibly manipulating the scheduling process in two ways. Firstly, each user can choose to misreport his specifications. Secondly, upon receiving the assigned charging profile r_i from the mediator, user i can choose to use a different charging profile r_i' (but pay the penalty $\lambda \|r_i' - r_i\|^2$ as described in (2)). For generality, we assume that user i uses a function f_i (called *policy*) that computes r_i' from the charging schedule r given by the mediator, so that $r_i' = f_i(r)$. For example, if user i decides to always follow the charging profile given by the mediator, his policy f_i satisfies $f_i(r) = r_i$ (for all r). We define the *action* of user i as the tuple (E_i, f_i), if user i reports E_i to the mediator and uses the policy f_i for post-mediator decisions.

Throughout the paper, we shall make the following assumption on how users are allowed to misreport their specifications.

Assumption 1 (Limited misreport). *When reporting specifications to the mediator, each user must satisfy the following conditions:*

1) *There exists $E_{\max} > 0$ (independent of i) such that user i can only misreport E_i within the interval $[0, E_{\max}]$;*

2) User i must always report the true $\bar{r}_i$ to the mediator.

For generality, we assume that the mediator computes the charging schedule r using a randomized algorithm M, and each user is interested in minimizing his expected cost (evaluated over the randomness in M). Since we have assumed that users will always report the true $\bar{r}$, we write $r = M(E)$ where $E = (E_1, E_2, \ldots, E_n)$ and leave out the dependence of M on $\bar{r}$. Then, the expected cost of user i is given by

$$c_M(i, E, f) := \mathbb{E}_{r \sim M(E)}[c(i, f_i(r), r)]. \quad (3)$$

Note that the expected cost $c_M(i, \cdot, \cdot)$ of user i not only depends on his own action (E_i, f_i), but also on the the joint action (E, f) of all users. The expected cost function c_M defines a game, which we call the *mediator induced EV charging game*. We say that a joint action (E, f) is an η-approximate equilibrium of the mediator induced game if

$$c_M(i, E, f) \leq c_M(i, (E_i', E_{-i}), (f_i', f_{-i})) + \eta$$

for all $i \in \{1, 2, \ldots, n\}$, $E_i' \in [0, E_{\max}]$, and policy f_i' (i.e., user i changes his action unilaterally). In addition, the game-theoretic setting allows us to define *truthful behavior* in this mediator induced game as follows.

Definition 2 (Truthful behavior)**.** Consider the mediator induced game defined by the cost function (3). Suppose the true user specifications are given by $E = (E_1, E_2, \ldots, E_n)$. A joint action $(\hat{E}, \hat{f})$ is called a *truthful behavior* if for all i we have $\hat{E}_i = E_i$ and $\hat{f}_i(r) = r_i$ for any $r = (r_1, r_2, \ldots, r_n)$.

The definition of truthful behavior implies that any user i will report his true specification E_i and follow the assigned charging profile r_i from the mediator.

The main result of this paper is to present an algorithm that ensures a truthful behavior is an η-approximate equilibrium of the mediator induced game (for some η). Namely, the algorithm guarantees that for any $i \in \{1, 2, \ldots, n\}$ and policy f_i', it holds that

$$\mathbb{E}_{r \sim M(E)}[c(i, r_i, r)] \leq \mathbb{E}_{r' \sim M(E')}[c(i, f_i'(r'), r')] + \eta,$$

where $E_j' = E_j$ for all $j \neq i$. Later, we will show that this can be achieved if the mechanism M preserves joint differential privacy.

III. TRUTHFULNESS VIA JOINT DIFFERENTIAL PRIVACY

In this section, we present one main result of this paper, which shows that truthful behavior can be attained via joint differential privacy. We assume that the mediator is able to implement an allocation algorithm that is ϵ-joint differentially private, and we leave the details of implementation to Section IV. This section begins with several lemmas that provide bounds that are used in the main result. The main result is presented as a theorem in the end of the section.

A. Joint differential privacy

We first introduce the framework of differential privacy before introducing joint differential privacy. Both differential privacy and joint differential privacy consider a mechanism M that acts on a set D (called *database*) consisting of user information. In the EV charging problem considered in this paper, the database is the set of user specifications $E = \{E_i\}_{i=1}^n$. One important concept in differential privacy is *adjacent databases*, which is defined through a binary relation between two databases.

Definition 3 (Adjacent databases)**.** Two databases $D = \{d_i\}_{i=1}^n$ and $D' = \{d_i'\}_{i=1}^n$ are said to be *adjacent with respect to user i*, denoted by $\text{Adj}_i(D, D')$, if $d_j = d_j'$ for all $j \neq i$. Two databases D and D' are said to be *adjacent*, denoted by $\text{Adj}(D, D')$, if there exists $i \in \{1, 2, \ldots, n\}$ such that $\text{Adj}_i(D, D')$.

In the case of EV charging, we define the adjacency relation as follows. Two databases (i.e., user specifications) E and E' are adjacent with respect to user i if and only if

$$E_i, E_i' \in [0, E_{\max}], \quad E_j = E_j' \quad \text{for all } j \neq i. \quad (4)$$

Note the similarity between the adjacency relation described by (4) and condition (1) in Assumption 1. In the context of ensuring truthfulness, the adjacency relation is generally used to define the limitation on how each user is able to misreport their information (specification).

With the definition of adjacent databases, we are ready to give the definition of differential privacy.

Definition 4 (Differential privacy [3])**.** Given $\epsilon > 0$, a randomized mechanism M preserves ϵ-differential privacy if for all $\mathcal{R} \subseteq \text{range}(M)$ and all D and D' such that $\text{Adj}(D, D')$, it holds that

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

The constant ϵ indicates the level of privacy: smaller ϵ implies a higher level of privacy. For the EV charging problem, however, it is impossible for the mediator to implement a differentially private mechanism M that computes the charging schedule as $(r_1, r_2, \ldots, r_n) = M(E)$ for the following reason. In order to satisfy the specifications from user i, the mechanism M must always satisfy $\mathbf{1}^T M_i(E) = E_i$ and $\mathbf{1}^T M_i(E') = E_i'$. When E_i and E_i' are different, it can be verified from Definition 4 that M does not preserve differential privacy. To circumvent this difficulty, the notion of differential privacy has been relaxed to *joint differential privacy* as originally proposed in [7]. Note that the notion of joint differential privacy only applies when the output of the mechanism is an n-tuple, where n is the number of users.

Definition 5 (Joint differential privacy)**.** Given $\epsilon > 0$, a randomized mechanism M whose output is an n-tuple preserves ϵ-joint differential privacy if all $i \in \{1, 2, \ldots, n\}$, all $\mathcal{R} \subseteq \text{range}(M_{-i})$, and all (D, D') such that $\text{Adj}_i(D, D')$, it holds that

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

Informally speaking, when joint differential privacy is preserved, changes in any single user's information does not affect significantly the output of the mechanism M that corresponds to *other users*. In contrast, differential privacy requires that the entire output of the mechanism (i.e., output corresponding to *all users*) should not be affected.

Applying the definition of joint differential privacy to the EV charging problem, we say that the mediator's mechanism M for computing a charging schedule preserves ϵ-joint differential privacy if for all $i \in \{1, 2, \ldots, n\}$, all $\mathcal{R} \subseteq \text{range}(M_{-i})$, and all specifications E and E' that satisfy (4), it holds that

$$\mathbb{P}(M_{-i}(E) \in \mathcal{R}) \leq e^{\epsilon} \mathbb{P}(M_{-i}(E') \in \mathcal{R}).$$

B. Joint differential privacy induces truthfulness

In order to show that a joint differentially private mediator's mechanism induces truthfulness, we begin by bounding the change in the cost function c for user i if the corresponding charging profile given by the mediator is changed from r_i to r'_i. The policy f_i of user i is assumed to be always accepting the mediator's assignment.

Lemma 6. *For any charging schedule r from the mediator and any $r'_i \in \mathbb{R}^T$, we have $c(i, r_i, r) - c(i, r'_i, (r'_i, r_{-i})) \leq \delta$, where*

$$\delta = 2E_{\max} \cdot \|\mu(\sum_{i=1}^{n} r_i + E_{\max}) + \bar{p}\|_{\infty}. \tag{5}$$

Proof: By definition, we have

$$c(i, r_i, r) = \left[\mu\left(r_i + \sum_{j \neq i} r_j\right) + \bar{p}\right]^T r_i,$$

$$c(i, r'_i, (r'_i, r_{-i})) = \left[\mu\left(r'_i + \sum_{j \neq i} r_j\right) + \bar{p}\right]^T r'_i,$$

so that

$$c(i, r_i, r) - c(i, r'_i, (r'_i, r_{-i}))$$
$$= \left(\mu\sum_{j \neq i} r_j + \bar{p}\right)^T (r_i - r'_i) + \mu\|r_i\|^2 - \mu\|r'_i\|^2$$
$$= \left(\mu\sum_{j \neq i} r_j + \bar{p} + \mu(r_i + r'_i)\right)^T (r_i - r'_i)$$
$$= (\mu\sum_{i=1}^{n} r_i + \bar{p} + \mu r'_i)^T (r_i - r'_i)$$
$$\leq \|\mu\sum_{i=1}^{n} r_i + \bar{p} + \mu r'_i\|_{\infty} \cdot \|r_i - r'_i\|_1 \tag{6}$$

Recall that we have $\|r_i\|_1, \|r'_i\|_1 \leq E_{\max}$, which implies $\|r_i - r'_i\|_1 \leq 2E_{\max}$. Applying this to (6) leads to

$$c(i, r_i, r) - c(i, r'_i, (r'_i, r_{-i}))$$
$$\leq 2E_{\max} \cdot \|\mu\sum_{i=1}^{n} r_i + \bar{p} + \mu r'_i\|_{\infty}$$
$$\leq 2E_{\max} \cdot \|\mu\left(\sum_{i=1}^{n} r_i + E_{\max}\right) + \bar{p}\|_{\infty}.$$

∎

The next lemma bounds the maximum individual cost of user i. The policy f_i of user i is still assumed to be always accepting the mediator's assignment.

Lemma 7. *For any charging schedule r from the mediator, we have $c(i, r_i, r) \leq c_{\max}$, where*

$$c_{\max} = \mu n E_{\max}^2 + E_{\max} \|\bar{p}\|_{\infty}. \tag{7}$$

Proof: By definition, for any r we have

$$c(i, r_i, r) = (\mu\sum_{i=1}^{n} r_i + \bar{p})^T r_i$$
$$\leq \|\mu\sum_{i=1}^{n} r_i + \bar{p}\|_{\infty} \|r_i\|_1 \tag{8}$$

From the fact that $\|r_i\|_1 \leq E_{\max}$ for all i, we know that $\|r_i\|_{\infty} \leq E_{\max}$ for all i and

$$\|\mu\sum_{i=1}^{n} r_i + \bar{p}\|_{\infty} \leq \mu\sum_{i=1}^{n} \|r_i\|_{\infty} + \|\bar{p}\|_{\infty}$$
$$\leq \mu n E_{\max} + \|\bar{p}\|_{\infty}.$$

Substitute the above into (8) to complete the proof. ∎

Before presenting the final lemma, we need to define the *optimal policy* of any individual user (after the user receives the mediator's assignment).

Definition 8 (Optimal policy). For any charging schedule r, the *optimal policy* f_i^* of user i is defined as

$$f_i^*(r) = \arg\min_{\hat{r}_i \in \mathcal{C}_i} c(i, \hat{r}_i, r).$$

The final lemma bounds the difference in the individual cost of user i between choosing the optimal policy and choosing to always follow the mediator's assignment.

Lemma 9. *For any charging schedule r, we have $c(i, r_i, r) - c(i, f_i^*(r), r) \leq \gamma$, where*

$$\gamma = \frac{\|\mu r_i + \mu \sum_{i=1}^{n} r_i + \bar{p}\|^2}{4(\mu + \lambda)}. \tag{9}$$

Proof: For notational convenience, define $a = \mu \sum_{j \neq i} r_j + \bar{p}$ so that

$$c(i, x, r) = (a + \mu x)^T x + \lambda \|x - r_i\|^2$$
$$= (\mu + \lambda)\|x\|^2 + (a - 2\lambda r_i)^T x + \lambda \|r_i\|^2.$$

Consider the optimal solution of the unconstrained problem:

$$x^* = \arg\min_{x \in \mathbb{R}^T} c(i, x, r).$$

It can be shown that

$$x^* = \frac{2\lambda r_i - a}{2(\mu + \lambda)}.$$

Substitute the expression of x^* into c to obtain the optimal value of the unconstrained problem as

$$c(i, x^*, r) = (\mu + \lambda)\|x^*\|^2 + (a - 2\lambda r_i)^T x^* + \lambda \|r_i\|^2$$
$$= -\frac{\|2\lambda r_i - a\|^2}{4(\mu + \lambda)} + \lambda \|r_i\|^2.$$

On the other hand, we have $c(i, x^*, r) \leq c(i, f_i^*(r), r)$. Then we have

$$c(i, r_i, r) - c(i, f_i^*(r), r)$$
$$\leq c(i, r_i, r) - c(i, x^*, r)$$
$$= (a + \mu r_i)^T r_i + \frac{\|2\lambda r_i - a\|^2}{4(\mu + \lambda)} - \lambda \|r_i\|^2$$
$$= a^T r_i + \frac{\|2\lambda r_i - a\|^2}{4(\mu + \lambda)} + (\mu - \lambda) \|r_i\|^2$$
$$= \frac{4(\mu + \lambda)a^T r_i + \|2\lambda r_i - a\|^2 + 4(\mu^2 - \lambda^2) \|r_i\|^2}{4(\mu + \lambda)}$$
$$= \frac{\|2\mu r_i + a\|^2}{4(\mu + \lambda)}$$
$$= \frac{\|\mu r_i + \mu \sum_{i=1}^n r_i + \bar{p}\|^2}{4(\mu + \lambda)}.$$

■

With Lemmas 6–9 at hand, we are ready to present the main result of this paper on the truthful behavior of users induced by joint differential privacy.

Theorem 10 (Approximate truthfulness). *Consider the mediator induced EV charging game where $E = (E_1, E_2, \ldots, E_n)$ consists of the true specifications of the users. Suppose the mediator uses a randomized mechanism M to compute the charging schedule r for all users. If M preserves ϵ-joint differentially privacy for some $\epsilon \in (0, 1)$, then the truthful behavior is an η-approximate equilibrium of the mediator induced game. Namely, for any $i \in \{1, 2, \ldots, n\}$ and policy f_i', it holds that*

$$\mathbb{E}_{r \sim M(E)}[c(i, r_i, r)] \leq \mathbb{E}_{r' \sim M(E')}[c(i, f_i'(r'), r')] + \eta,$$

where $E_i' \in [0, E_{\max}]$ and $E_j' = E_j$ for all $j \neq i$. Specifically, we have $\eta = \gamma + 2\epsilon c_{\max} + \delta$, where δ, $c_{\max}$, and γ are given by equations (5)–(9).

Proof: From Lemma 6, we have

$$\mathbb{E}_{r \sim M(E)}[c(i, r_i, r)]$$
$$\leq \mathbb{E}_{r_i' \sim M_i(E'), r_{-i} \sim M_{-i}(E)}[c(i, r_i', (r_i', r_{-i}))] + \delta \quad (10)$$

From the definition of joint differential privacy, we have

$$\mathbb{E}_{r_i' \sim M_i(E'), r_{-i} \sim M_{-i}(E)}[c(i, r_i', (r_i', r_{-i}))]$$
$$\leq e^\epsilon \mathbb{E}_{r_i' \sim M_i(E'), r_{-i}' \sim M_{-i}(E')}[c(i, r_i', (r_i', r_{-i}'))]$$
$$= e^\epsilon \mathbb{E}_{r' \sim M(E')}[c(i, r_i', r')]. \quad (11)$$

Recall that $e^\epsilon \leq 1 + 2\epsilon$ for $\epsilon \in (0, 1)$. Then we have

$$e^\epsilon \mathbb{E}_{r' \sim M(E')}[c(i, r_i', r')]$$
$$\leq (1 + 2\epsilon)\mathbb{E}_{r' \sim M(E')}[c(i, r_i', r')]$$
$$\leq \mathbb{E}_{r' \sim M(E')}[c(i, r_i', r')] + 2\epsilon c_{\max}. \quad (12)$$

The last step uses the result from Lemma 7. Finally, we have from Lemma 9

$$\mathbb{E}_{r' \sim M(E')}[c(i, r_i', r')] \leq \mathbb{E}_{r' \sim M(E')}[c(i, f_i^*(r'), r')] + \gamma$$
$$\leq \mathbb{E}_{r' \sim M(E')}[c(i, f_i'(r'), r')] + \gamma$$
$$\quad (13)$$

for any policy f_i'. Combining equations (10)–(13) completes the proof. ■

The goodness of approximation η depends on three terms. The term γ can be reduced by increasing the level of penalty λ, which essentially prevents users from deviating from the assigned charging profile. The term $2\epsilon c_{\max}$ can be reduced by decreasing ϵ (i.e., increasing the level of privacy). On the other hand, as we will show in the next section, doing so will introduce more randomness in the mechanism M and consequently lead to undesired effects such as increased operating cost (from the perspective of the mediator). The last term is somewhat "intrinsic", since it is related to how user i can potentially benefit from affecting $M_i(E)$, which is not controlled by joint differential privacy.

IV. A JOINT DIFFERENTIALLY PRIVATE CHARGING MECHANISM

In this section, we present a joint differentially private mechanism that computes a charging schedule according to reported user specifications. The mechanism is based on our previous work on differentially private distributed EV charging. After reviewing our previous results, we show that the same algorithm can be used to ensure joint differential privacy.

We will use the following notations throughout this section. For any convex set $\mathcal{C} \subset \mathbb{R}^n$, define the (Euclidean) projection operator $\Pi_{\mathcal{C}}$ such that $\Pi_{\mathcal{C}}(x)$ is the projection of any $x \in \mathbb{R}^n$ onto $\mathcal{C}$ — i.e., $\Pi_{\mathcal{C}}(x) = \arg\min_{z \in \mathcal{C}} \|z - x\|^2$. 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 $p_X(x) = \frac{1}{2\lambda} \exp(-|x|/\lambda)$.

A. A joint differentially private charging mechanism

The charging mechanism used in this paper is presented in Algorithm 1. In the algorithm, the function U can be an arbitrary convex function whose gradient ∇U is L-Lipschitz under the ℓ_2-norm. Namely, there exists $L > 0$ such that

$$\|\nabla U(x) - \nabla U(y)\|_2 \leq L \|x - y\|_2, \quad \forall x, y \in \mathbb{R}^T. \quad (14)$$

The function U is normally selected by the mediator as some objective function to minimize. The choice of U does not, however, influence joint differential privacy, so that we will postpone the discussion on choosing U until Section IV-B.

Algorithm 1 was studied in our previous work in the context of differentially private distributed EV charging. The following proposition shows that the mechanism $M_p(E) := (\hat{p}^{(1)}(E), \hat{p}^{(2)}(E), \ldots, \hat{p}^{(K)}(E))$ is ϵ-differentially private for $\hat{p}^{(k)}$ given by (15) in Algorithm 1. We have explicitly indicated the dependence of $\hat{p}^{(k)}$ on the user specifications $E = \{E_i\}_{i=1}^n$ for clarity.

Proposition 11 (Han et al. [6]). *The mechanism $M_p := (\hat{p}^{(1)}, \hat{p}^{(2)}, \ldots, \hat{p}^{(K)})$ is ϵ-differentially private with respect to the adjacency relation defined by (4). Namely, the mechanism M_p satisfies*

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

Algorithm 1 ϵ-joint differentially private EV charging mechanism.

Input: U, $\{\mathcal{C}_i\}_{i=1}^n$ (i.e., the constants $\{\bar{r}_i, E_i\}_{i=1}^n$), K, $\{\alpha_k\}_{k=1}^K$, L, $E_{\max}$, and ϵ.

Output: $\{r_i^{(K+1)}\}_{i=1}^n$.

Initialize $\{r_i^{(1)}\}_{i=1}^n$ arbitrarily. Let $\tilde{r}_i^{(1)} = r_i^{(1)}$ for all $i \in \{1, 2, \ldots, n\}$ and $\theta_k = 2/(1+k)$ for $k \in \{1, 2, \ldots, K\}$.

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

$$\hat{p}^{(k)} := \nabla U\left(\sum_{i=1}^n r_i^{(k)}\right) + w_k. \tag{15}$$

3) For $i = 1, 2, \ldots, n$, update $r_i^{(k+1)}$ and $\tilde{r}_i^{(k+1)}$ as follows:

$$\tilde{r}_i^{(k+1)} := \Pi_{\mathcal{C}_i}(\tilde{r}_i^{(k)} - \alpha_k \hat{p}^{(k)}) \tag{16}$$

$$r_i^{(k+1)} := (1 - \theta_k)r_i^{(k)} + \theta_k \tilde{r}_i^{(k+1)}. \tag{17}$$

---

for all $\mathcal{R} \subseteq \text{range}(M_p)$ and all E and E' such that equation (4) holds for some $i \in \{1, 2, \ldots, n\}$.

Note that the guarantee of ϵ-differential privacy is for the gradients $\{\hat{p}^{(k)}\}$. In the next, we show that Algorithm 1 also preserves ϵ-joint differential privacy for the output charging schedule $r^{(K+1)}$. The proof makes use of the post-processing theorem from differential privacy.

Proposition 12 (Post-processing [4])**.** *Suppose a mechanism M preserves ϵ-differential privacy. Then, for any function f, the functional composition $f \circ M$ also preserves ϵ-differential privacy.*

The post-processing theorem allows us to construct new differentially private mechanisms from existing ones. Now we are ready to show that the output of Algorithm 1 is a joint differentially private mechanism.

Theorem 13. *Consider the mechanism $M := (r_1^{(K+1)}, r_2^{(K+1)}, \ldots, r_n^{(K+1)})$ acting on the user specifications $E = \{E_i\}_{i=1}^n$, where $\{r_i^{(K+1)}\}_{i=1}^n$ is given by the output of Algorithm 1. Then M is ϵ-joint differentially private under the adjacency relation defined by (4).*

Proof: Observe from Algorithm 1 that for all i and k we can write

$$\tilde{r}_i^{(k+1)} = g_1(E_i, \tilde{r}_i^{(k)}, \hat{p}^{(k)}(E)),$$
$$r_i^{(k+1)} = g_2(r_i^{(k)}, \tilde{r}_i^{(k+1)})$$

for some functions g_1 and g_2. Here we have used $\hat{p}^{(k)}(E)$ to emphasize the dependence of $\hat{p}^{(k)}$ on E. By induction, we can write

$$r_i^{(K+1)} = g(E_i, r_i^{(1)}, \{\hat{p}^{(k)}(E)\}_{k=1}^K)$$

for some function g. Consider E and E' such that $\text{Adj}_i(E, E')$ according to (4). For all $j \neq i$, we have

$$r_j^{(K+1)}(E) = g(E_j, r_j^{(1)}, \{\hat{p}^{(k)}(E)\}_{k=1}^K),$$
$$r_j^{(K+1)}(E') = g(E'_j, r_j^{(1)}, \{\hat{p}^{(k)}(E')\}_{k=1}^K)$$
$$= g(E_j, r_j^{(1)}, \{\hat{p}^{(k)}(E')\}_{k=1}^K).$$

Then, we can view $M_{-i} := r_{-i}^{(K+1)}$ as a post-processing result of the mechanism $M_p := (\hat{p}^{(1)}, \hat{p}^{(2)}, \ldots, \hat{p}^{(K)})$, which is ϵ-differentially private according to Proposition 11. Using the post-processing theorem (Proposition 12), we conclude that M_{-i} is ϵ-differentially private for all i, which is equivalent to M being ϵ-joint differentially private. ∎

Remark 14. In our previous work, Algorithm 1 was used in the context of distributed EV charging, in which the messages $\{\hat{p}^{(k)}\}_{k=1}^K$ are broadcast to all users, and $\{\hat{p}^{(k)}\}_{k=1}^K$ can be potentially eavesdropped. It was shown that Algorithm 1 preserves privacy of the users. Namely, an adversary cannot obtain information on E_i (for any user i) with high confidence, even if the adversary has access to all the messages $\{\hat{p}^{(k)}\}_{k=1}^K$. This implies that Algorithm 1, when used by the mediator, can ensure *both privacy and truthfulness* of the participating users simultaneously.

B. Tradeoffs between truthfulness and suboptimality

Aside from satisfying all user specifications, the mediator is often interested in computing a charging schedule that is optimal with respect to a certain objective such as minimal variance or minimal peak load. Formally, the mediator would like to solve an optimization problem in the following form:

$$\min_{\{r_i\}_{i=1}^n} \quad U\left(\sum_{i=1}^n r_i\right) \tag{18}$$
$$\text{s.t.} \quad r_i \in \mathcal{C}_i, \quad i = 1, 2, \ldots, n.$$

The objective function $U: \mathbb{R}^T \to \mathbb{R}$ in problem (18) is assumed to be convex. This assumption holds for a number of common objectives such as minimal variance and minimal peak load.

As we showed in our previous work [6], Algorithm 1 can be viewed as an implementation of the stochastic gradient descent algorithm. Suppose the step sizes $\{\alpha_k\}_{k=1}^K$ is chosen optimally (see [6] for details). It can be shown that the expected suboptimality of Algorithm 1 is given as follows:

$$\mathbb{E}\left[U\left(\sum_{i=1}^n r_i^{(K+1)}\right) - U^*\right] \leq \mathcal{O}\left(T^{1/8}(E_{\max}/n\epsilon)^{1/4}\right), \tag{19}$$

where U^* is the optimal value of problem (18). Ideally, the mediator wishes to choose ϵ in order to have both small suboptimality gap and small η in the approximately truthful behavior. Unfortunately, there exists an intrinsic trade-off between truthfulness and suboptimality. As ϵ increases, it can be seen from (19) that the suboptimality decreases, whereas the parameter η given by Theorem 10 increases (which implies that it is less likely to obtain truthful behavior).

V. Conclusions and future work

In this paper, we apply the notion of *joint differential privacy*, originally proposed in [7], to the EV charging problem to ensure truthfulness of participating users. In particular, we consider the scenario of *direct load control*, where a mediator is present to collect user specifications and compute a charging schedule for all participating users. Due to their selfish nature, users may misreport their specifications and/or ignore the mediator's assignment in order to minimize their individual cost (i.e., payment for electricity usage).

The paper shows that approximately truthful behavior of the users can be attained if the mediator computes the charging schedule using a joint differentially private mechanism. This is possible since joint differential privacy can limit the power of each user in manipulating the scheduling process by remaining insensitive to changes in user specifications. The paper also presents an algorithm that can be used by the mediator to attain joint differential privacy. The same algorithm has been shown in our previous work [6] to protect the user information (specifications) from potential adversaries. This implies that it is possible to guarantee both privacy and truthfulness of the users simultaneously using the algorithm presented. From the perspective of the mediator, an analysis of the algorithm on the tradeoffs between suboptimality (in terms of operating cost) and truthfulness is also presented. It is found that more truthfulness can be attained at the expense of sacrificing optimality.

One interesting direction for future work is to compare our current results with other mechanisms that promote truthfulness, such as the recent work on faithful distributed optimization by Tanaka et al. [11]. Instead of introducing random perturbations as in our paper, the work by Tanaka et al. achieves truthfulness by designing deterministic tax/subsidy rules, which can lead to a more consistent performance by eliminating randomness as used in our algorithm, although it remains under investigation how this kind of algorithm can be applied to the EV charging problem.

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, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography*, pages 265–284. Springer, 2006.

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

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

[6] S. Han, U. Topcu, and G. J. Pappas. Differentially private distributed protocol for electric vehicle charging. In *Annual Allerton Conference on Communication, Control, and Computing*, 2014.

[7] M. Kearns, M. Pai, A. Roth, and J. Ullman. Mechanism design in large games: Incentives and privacy. In *Conference on Innovations in Theoretical Computer Science*, pages 403–410, 2014.

[8] 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.

[9] R. M. Rogers and A. Roth. Asymptotically truthful equilibrium selection in large congestion games. In *ACM Conference on Economics and Computation*, pages 771–782, 2014.

[10] 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.

[11] T. Tanaka, F. Farokhi, and C. Langbort. Faithful implementations of distributed algorithms and control laws. *arXiv preprint arXiv:1309.4372*, 2013.

[12] 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, 2010.