

Computation of privacy-preserving prices in smart grids

Fragkiskos Koufogiannis, Shuo Han, George J. Pappas
Electrical and Systems Engineering, University of Pennsylvania

Abstract—¹ Demand management through pricing is a modern approach that can improve the efficiency of modern power networks. However, computing optimal prices requires access to data that individuals consider private. We present a novel approach for computing prices while providing privacy guarantees under the differential privacy framework. Differentially private prices are computed through a distributed utility maximization problem with each individual perturbing their own utility function. Privacy concerning temporal localization and monitoring of an individual’s activity is enforced in the process. The proposed scheme provides formal privacy guarantees and its performance-privacy trade-off is evaluated quantitatively.

I. INTRODUCTION

Modern power networks are equipped with smart meters that allow sensing in a very fine spatial and temporal scale, which enables new features and allows for more efficient operation. Demand response is an emerging feature of power networks that publishes real-time prices as a means of providing incentives to individuals to modify their power consumption towards a more efficient network operation [13], [7]. The increasing amount of information captured and communicated by new techniques has raised concerns about the privacy of individuals [8]. In particular, habits and activities can be inferred from an individual’s power consumption trace [4]. However, the computation of electricity prices, which is formulated as an optimization problem, requires each individual to report their preference and can lead to severe loss of privacy.

In this work, we employ the formal notion of differential privacy [2], [3] in order to provide privacy guarantees for individuals while computing electricity prices. From a theoretical point of view, providing privacy in distributed optimization problems can severely deteriorate the performance of the system when incorporated in a naive way [5]. Intuitively, successive queries to private data can quickly deplete the privacy budget, while making only small steps towards optimality. However, there have been applications of differential privacy in aggregation schemes that sacrifice almost no performance while providing privacy [10]. This is in accordance with the intuition that massive aggregation schemes should provide privacy with minimal noise injection.

¹This work was supported in part by the TerraSwarm Research Center, one of six centers supported by the STARnet phase of the Focus Center Research Program (FCRP) a Semiconductor Research Corporation program sponsored by MARCO and DARPA.

Existing work on privacy in smart grids spans a variety of directions. A signal processing treatment of providing privacy of an individual’s activities given her power consumption trace was proposed in [14], which provides information-theoretic privacy guarantees. Besides being difficult in handling physical constraints, such approaches enforce privacy in a probabilistic sense, i.e., there may be severe privacy breaches for rare realizations of the system that cannot be captured by these approaches. A different approach employs differential privacy [1] to hide an individual’s activity by using a rechargeable battery that adds noise to the individual’s power consumption trace. Other approaches include computing an individual’s billing information under privacy constraints [12]. However, our work focuses on a different problem. Individuals are interested in coordinating their power consumption in order to maximize the total utility they enjoy, while having formal privacy guarantees on private information they communicate. Note that the scheme presented here can be used in tandem with aforementioned techniques for a more complete privacy-aware smart grid paradigm.

This paper is structured as follows. In Section II, we present the system model which is based on maximizing the total utility that society enjoys. The desired privacy guarantees are presented in Section III and the proposed architecture in Section IV. Simulations depicting how efficiency degrades with increasing privacy and discussion on future directions follow in Sections V and VI, respectively.

II. SMART GRID PRICING

In this section, we present the model of a smart grid which we will work with. In Subsection II-A a utility-oriented model is described and in Subsection II-B the problem of operating the model in order to maximize utility is revisited.

A. Utility Modelling

We use a utility maximization scheme similar to the ones displayed in [13] and [7]. Consider a set $V = \{1, \dots, n\}$ of nodes, where each node represents an individual which, according to the micro-grids paradigm, either consumes or provides power to the network. In this context, there is no underlying network on the set of nodes. Each node $i \in V$ has a type $\sigma(i) \in \{1, \dots, K\}$ that can represent a residential unit, a power plant, or an HVAC system. A node i with type $\sigma(i)$ possesses a parameter vector $x_i \in \mathbb{R}^{T \times n_{\sigma(i)}}$ which captures the activity of the node, where T is the finite time

horizon and \mathbb{R}^{n_σ} is the parameter space of type σ for each time instant. Finally, each node i with parameter vector x_i provides utility $U_{\sigma(i)}(y; x_{i,t})$ when consuming y amount of electricity.

Consider the following examples:

- *Example 1 (PHEV)*: Efficiently scheduling the charging period of a *Plug-in Hybrid Electric Vehicle (PHEV)* is one promising application of the smart grid paradigm. Assume in this case that $n_{\sigma(i)} = 1$ and $x_{i,t}$ captures the availability of the vehicle for charging. The corresponding utility function can then be expressed in the following form:

$$U_{\text{PHEV}}(y_{i,t}; x_{i,t}) = x_{i,t} y_{i,t},$$

where:

$$x_{i,t} = \begin{cases} 1, & \text{when PHEV is available for charging} \\ 0, & \text{otherwise} \end{cases}$$

It should be highlighted that an individual's activity can be inferred from the vector of parameters x_i . In particular, the time instances t that individual i plans a trip are designated with the parameter $x_{i,t}$ being zero. Note that parameters x capture only the usage of a PHEV and not its existence.

- *Example 2 (HVAC)*: A similar example is the operation of an HVAC unit. In this case, a natural parametrization $x_i \in \mathbb{R}^{T \times 2}$, with $n_{\sigma(i)} = 2$ dimensions, is the number of occupants in a building $x_{i,t}^{(1)}$ and the temperature set point $x_{i,t}^{(2)}$ at each time instance t . Note that such information discloses the activity at the node, which raises privacy concerns. The system, when consuming y amount of electricity, provides $U_{\text{HVAC}}(y; x_{i,t})$ level of comfort as utility at time t .
- *Example 3*: The previous examples were cases where there was a natural parametrization of the utility functions. A utility function that models the activity of a household with numerous low-power electrical appliances may possess a high-dimensional parametrization. Such parametrization requires modelling every single device separately. To avoid enumerating all devices, we propose the following way of handling the case of an arbitrary utility function $U(y)$. Specifically, we approximate the original utility function U_{σ_i} with a piece-wise linear version $U_{\text{PWL},M} : [y_{\min}, y_{\max}] \times \mathbb{R}^{n_{\sigma(i)}} \rightarrow \mathbb{R}$ as the solution to the following optimization problem:

$$U_{\text{PWL},M}(y; x_{i,t}) = U(y_{\min}) + \max_{q_1, \dots, q_M} \sum_{j=1}^M x_{i,t}^{(j)} q_j \quad (1)$$

$$\text{s.t. } \sum_{j=1}^M q_j = y, \quad q_j \in [0, \Delta y], \quad \forall j$$

where $\Delta y = \frac{y_{\max} - y_{\min}}{n_{\sigma(i)}}$. In this formulation each node is modelled as operating M devices and each device j yields q_j units of utility per unit of power consumed. The utility function takes the form of a minimum of

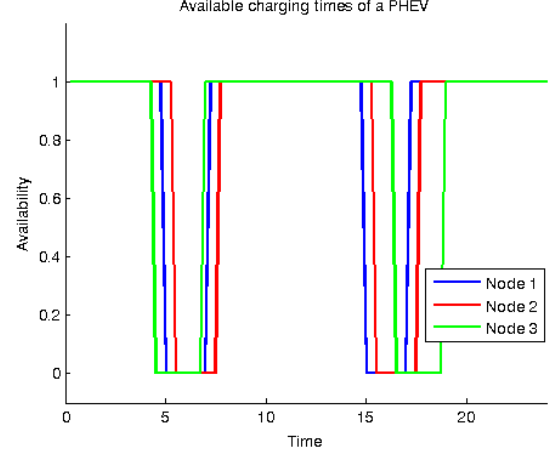


Fig. 1. Sample parameters for a Plug-in Hybrid Electric Vehicle (PHEV) charging schedule.




Fig. 2. Overview of total utility maximizing demand response system.

linear functions and the parameters $x_{i,t}$ are defined as $x_{i,t}^{(j)} = U_{\sigma_i}(j\Delta y; x_{i,t}) - x_{i,t}^{(j-1)}$, with $x_{i,t}^{(0)} = 0$.

B. Pricing Algorithm

We are interested in computing the loads $y_{i,t}$ for each node i and time t that maximize the total utility which can be formulated as the following optimization problem:

$$\begin{aligned} & \text{maximize}_{\{y_{i,t}\}_{i,t}} \sum_{t=1}^T \left\{ \sum_{i=1}^n U_{\sigma(i)}(y_{i,t}; x_{i,t}) \right\} \\ & \text{s.t. } \sum_{i=1}^n y_{i,t} = 0, \quad \forall t \in \{1, \dots, T\}. \end{aligned} \quad (2)$$

The objective is to maximize the total utility that nodes enjoy, whereas the constraint dictates that power consumption should be balanced for each time instance. As shown independently in [7] and [13], it is possible to solve this problem in a decentralized way. This is accomplished by considering the dual problem:

$$\text{maximize}_{\{p_t\}_t} \sum_{i=1}^n d_i(p), \quad (3)$$

where:

$$d_i(p) = \min_{\{y_{i,t}\}_t} \sum_{t=1}^T p_t y_{i,t} - U_{\sigma(i)}(y_{i,t}, x_{i,t}) \quad (4)$$

The functions $d_i(p)$ and the required directions $\nabla d_i(p)$ depend on information local to each node i and are computed by solving the discounted utility maximization problem (4)

for each node i .

Specifically, in successive iterations, a server publishes the price signals $\{p_t(k)\}_{t \in \{1, \dots, T\}}$, each node i computes its descent direction $\nabla d_i(p(k))$. The descent directions are then aggregated to produce a descent direction of the prices:

$$p(k+1) = p(k) - \beta \sum_{i=1}^n \nabla d_i(p(k)) \quad (5)$$

As shown in [7] and [13], under concavity assumptions of utility functions $U_\sigma(y; x)$ with respect to y for any parameter x , process (5) converges to the optimal electrical prices that provide incentives to nodes to operate in a globally efficient manner. In the next section, we raise privacy concerns and formally state a framework to handle them.

III. DEALING WITH PRIVACY

Computing electrical prices requires access to the utility function of every node. Part of this information may be sensitive data that a node is unwilling to share with the community. In this work we interested in computing prices that provide incentives for users to adapt their consumption to the global optimum while having control over the amount of private information that they disclose in the process. In order to motivate the privacy concerns related with the process of computing prices, let's think of a concrete example. Consider that data from Figure (1) are accessed while computing electrical prices. Data such as the scheduling time of charging of a PHEV, the number of occupants in a building and the temperature set points of an HVAC system contain sensitive information that is disclosed in the process of computing prices. In order to mitigate the privacy concerns raised, we employ the formal notion of *differential privacy*.

A. Differential Privacy

In the seminal work [3], the notion of differential privacy was introduced. This notion features a concrete formulation and makes few assumptions on the form of side information that an adversary has access to. Moreover, it is tailored with practical algorithms and composition theorems that allow existing primitives to be extend to more complex systems. Within the context of differential privacy, a *mechanism*, such as the process of computing electrical prices, is allowed to respond stochastically in order to smooth out the dependency of the outputs on the inputs. Roughly speaking, differential privacy is related to the sensitivity of a system and is built upon the dogma that a differentially private mechanism should respond *almost* identically to nearby inputs. The aspect of privacy that one is interested in is captured by an *adjacency relation* over the inputs. Formally stated:

Definition 1 (Differential Privacy [3]): Given an adjacency relation $\mathcal{A} \subseteq \mathcal{D} \times \mathcal{D}$, a mechanism $Q : \mathcal{D} \times \Omega \rightarrow \mathcal{Y}$ is ϵ -differential private iff:

$$\mathbb{P}(Qu \in S) \leq e^\epsilon \mathbb{P}(Qu' \in S), \forall (u, u') \in \mathcal{A}, S \subseteq \mathcal{Y} \quad (6)$$

where \mathcal{D} is the set of inputs, \mathcal{Y} is the set of outputs, Ω is the sample space of the coin flips of the mechanism and the probability \mathbb{P} is taken over the measure of Ω .² The level of privacy is controlled by the parameter ϵ , with smaller values leading to stronger privacy guarantees. Increased levels of privacy ($\epsilon \rightarrow 0$) result in the injection of higher levels of noise, which in general degrades the performance of the mechanism.

In the setting of utility maximization, the set of inputs \mathcal{D} is the set of the parameters x and the mechanism Q is the process of computing prices. The adjacency relation \mathcal{A} is a design decision that captures what features of the inputs are considered private.

The notion of differential privacy is accompanied by numerous results [3]. One way to enforce privacy is through the exponential mechanism:

Theorem 1 (Exponential Mechanism [9]): Given input $u \in \mathcal{U}$, consider the mechanism Q that responds with $y \in \mathcal{Y}$ with probability:

$$\mathbb{P}(Qu = y) \propto e^{\frac{\epsilon}{\Delta q} q(u, y)} \quad (7)$$

where $q(u, y)$ is a scoring function, with higher values indicating more favourable responses. Δq is defined as:

$$\Delta q = \max_{(u, u') \in \mathcal{A}, y} |q(u, y) - q(u', y)|. \quad (8)$$

Then Q is ϵ -differentially private.

We can also compose ϵ -differentially private mechanisms that act on a partitioning of the input space in a single ϵ -differentially private mechanism.

Theorem 2 (Parallel Composition [3]): Consider a set of ϵ -differentially private mechanisms Q_1, \dots, Q_n , where each $Q_i : \mathcal{U}_i \times \Omega_i \rightarrow \mathcal{Y}_i$ is defined with respect to an adjacency relation \mathcal{A}_i . Then, the mechanism

$$Q : (\times_{i=1}^n \mathcal{U}_i) \times (\times_{i=1}^n \Omega_i) \rightarrow (\times_{i=1}^n \mathcal{Y}_i) \quad (9)$$

that returns:

$$Q(u_1 \dots u_n) = (Q_1 u_1, \dots, Q_n u_n) \quad (10)$$

is also ϵ -differentially private with respect to the adjacency relation:

$$(u, u') \in \mathcal{A} \Leftrightarrow \{\exists j \text{ s.t. } u_i = u'_i, \forall i \neq j \text{ and } (u_j, u'_j) \in \mathcal{A}_j\} \quad (11)$$

Another comforting fact that we will use is that any post-processing on the output of a differentially private mechanism does not weaken the privacy guarantees.

Theorem 3: Let $Q : \mathcal{U} \times \Omega \rightarrow \mathcal{Y}$ be an ϵ -differentially private mechanism and $f : \mathcal{Y} \rightarrow \mathcal{Z}$ be any (measurable) function. Then the mechanism $f \circ Q$ is also ϵ -differentially private.

²Formally, we focus only on a rich enough σ -algebra of the set of outputs \mathcal{Y} . For the sake of clarity, this technical detail is omitted.

B. Pricing with Privacy

We proceed by describing an application of the differential privacy framework in the case of publishing electricity price signals. We are interested in protecting the parameters $\{x_i\}_{i \in \{1, \dots, N\}}$ and consequently the utility functions U_i . Choosing a suitable adjacency relation is crucial and can drastically affect the performance of the resulting mechanism. As a general rule, it is better to employ sparsely knitted adjacency relations with strict privacy budget. Specifically, we are interested in the following aspects of privacy:

- *Temporal localization*: An adversary should not be able to monitor when an individual wishes to perform an action. For example, consider the case of available slots for charging a PHEV, with parameters as shown in (1). The times that the individual travels between her home and work are designated by the unavailability of the PHEV for charging. The existence of the PHEV itself is not considered sensitive information that needs to be protected. Instead, the exact time that transportation occurs is regarded as private information. As such, we wish to provide ϵ -differential privacy with respect to the adjacency relation $\mathcal{A}_{i,1}$ that allows time-shifting activity by a single unit of time. Strictly speaking, we define two parameters x and \tilde{x} to be adjacent with respect to $\mathcal{A}_{i,1}$ if the one can be generated from the other by shifting the time indexing by one unit:

$$(x, \tilde{x}) \in \mathcal{A}_{i,1} \Leftrightarrow x_t = \tilde{x}_{t+1}, \forall t \quad (12)$$

where we assume periodic boundary conditions:

$$x_t = \tilde{x}_{(t+1) \bmod T+1}, \text{ for } t \in \{1, \dots, T\} \quad (13)$$

Note that the ϵ -differential privacy guarantee can be extended to a $T\epsilon$ privacy guarantee with respect to time-shifting by an arbitrary amount of time units. On the downside, it is worth noting that if an adversary collects information about a single event, such as the time an individual leaves from work, there will be no more privacy guarantees regarding other events within the day.

- *Activity monitoring*: Another aspect of privacy is that an adversary should not be able to infer the exact activity taking place in a node. This notion can be easily captured by requiring ϵ -differential privacy with respect to an adjacency relation that allows small variations in the parameters:

$$(x, \tilde{x}) \in \mathcal{A}_{i,2} \Leftrightarrow \|x - \tilde{x}\|_2 \leq \alpha \quad (14)$$

Multiple adjacency relations is not typical in differential privacy³. One way to merge the two adjacency relations is by considering the following adjacency relation:

$$(x, \tilde{x}) \in \mathcal{A}_i \Leftrightarrow \min_{\tau \in \mathbb{Z}} \frac{1}{\alpha} \|D^\tau x - \tilde{x}\|_2 + |\tau| \leq 1, \quad (15)$$

³It has been noted that adjacency relations do not differ much from queries [11]. Thus, this should not be surprising.

where D^τ is the operator of time-shifting the signal x by τ units of time, assuming periodic boundary conditions:

$$\{D^\tau x\}_t = x_{(t+\tau) \bmod T+1}, \text{ for } t \in \{1, \dots, T\}$$

In words, two parameters x and \tilde{x} are considered adjacent if a time shifted version of x is close to \tilde{x} . We include a penalty term for the magnitude $|\tau|$ of the time shift. The parameter α controls the importance of time shifts and parameter perturbation.

IV. DIFFERENTIALLY PRIVATE PRICING

Each node uses the exponential mechanism to generate a proxy version of its utility function, which it uses in collaboratively computing electrical prices. This technique is referred to as *input perturbation* which consists of generating a private version of the data before computing the quantities of interest. Another advantage of input perturbation is that it requires minimal modifications to the original scheme of computing prices. This dictates that privacy guarantees can be incorporated into an existing pricing scheme. Beginning with the adjacency relation \mathcal{A}_i in (15), we consider the following scoring function:

$$q(x, \tilde{x}) = -\min_{\tau \in \mathbb{Z}} \left[\frac{1}{\alpha} \|D^\tau x - \tilde{x}\|_2 + |\tau| \right] \quad (16)$$

Unfortunately, it is not trivial to sample the proxy parameters according to the distribution

$$\mathbb{P}(\tilde{x}|x) \propto e^{\frac{\epsilon}{\alpha} q(x, \tilde{x})}$$

Instead, we approximate the scoring function using a soft-maximum function:

$$\hat{q}(x, \tilde{x}) = \ln \left[\sum_{\tau \in \mathbb{Z}} e^{-\frac{1}{\alpha} \|D^\tau x - \tilde{x}\|_2 - |\tau|} \right]. \quad (17)$$

Scoring function \hat{q} is particularly appealing. It captures the aspects of privacy that we wish to enforce. Also, sampling from the exponential mechanism is simple since it can be viewed as sampling a parameter perturbation δx , conditioned on a discrete exponentially distributed delay τ :

$$\mathbb{P}(\tilde{x}|x) \propto e^{\epsilon \hat{q}(x, \tilde{x})} = \quad (18)$$

$$\sum_{\tau \in \mathbb{Z}} e^{-\epsilon \tau} \cdot e^{-\frac{\epsilon}{\alpha} \|D^\tau x - \hat{x}\|_2} \quad (19)$$

Specifically, we draw the following noise samples:

$$\tau \sim \text{Lap}\left(\frac{1}{\epsilon}\right), \quad \{\delta x_t\}_t \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1), \quad r \sim \Gamma\left(T, \frac{\alpha}{\epsilon}\right) \quad (20)$$

where $\text{Lap}(b)$ is the Laplace distribution with variance $2b^2$, $\mathcal{N}(0, 1)$ is the standard normal distribution, and $\Gamma(N, \theta)$ is the gamma distribution with shape N and scale θ . The proxy parameter \hat{x} is then computed as follows:

$$\hat{x} = D^\tau x + r \cdot \frac{\delta x}{\|\delta x\|_2}, \quad (21)$$

where τ is the time shift applied, with periodic boundary conditions assumed once again, and δx_t is the perturbation




Fig. 3. Generating a proxy utility function from the private one.

of the utility function at each time step. The proxy version \tilde{U}_i of the utility function U_i is computed as depicted in Figure 3. Node i first shifts the utility function by τ time units. Afterwards, the parameter x is perturbed by δx .

Remark 1: Due to the infinite-dimensionality of the space of utility functions we cannot simply bound the difference $\|U_i - \tilde{U}_i\|_2$. This is problematic for differential privacy, as noted later. Instead, we project utility functions onto a finite-dimensional subspace and enforce differential privacy there.

Algorithm 1 Privatization of utility function.

Sample $\tau \sim \text{Lap}(\frac{1}{\epsilon})$
 Sample $\{\delta x_t\}_t \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$
 Normalize $[\delta x_t]_t \leftarrow \frac{[\delta x_t]_t}{\|[\delta x_t]_t\|_2}$
 Sample $r \sim \Gamma(T, \frac{\alpha}{\epsilon})$
 Shift $\hat{x}_{i,t} \leftarrow x_{(t-\tau) \bmod T}$
for all $t \in \{1, \dots, T\}$ **do**
 $\tilde{x}_t = \hat{x}_t + r\delta x_t$
end for
 Use $\tilde{U}(y) = U_{\sigma(i)}(y; \tilde{x})$

Theorem 4: The utility functions $\tilde{U}_i : \mathbb{R} \times \{1, \dots, T\} \rightarrow \mathbb{R}$, $i \in \{1, \dots, n\}$, defined as $\tilde{U}_i(y) = U_{\sigma(i)}(y; \tilde{x})$ that result from algorithm (1) provide ϵ -differential privacy guarantees against time localization and activity monitoring up to level α . In particular, it guarantees ϵ dp with respect to the adjacency relation $\mathcal{A}' \ni (x, \tilde{x})$:

$$\exists j \text{ s.t. } x_i = \tilde{x}_i, \forall i \neq j \text{ and}$$

$$\text{SOFTMIN}_{\tau} \left(\frac{1}{\alpha} \|D^{\tau} x_j - \tilde{x}\|_2 + |\tau| \right) \leq 1$$

where the soft minimum function used is defined as $\text{SOFTMIN}_{\tau}(z_{\tau}) = -\ln \sum_{\tau} e^{-z_{\tau}}$.

Proof: The proof of this theorem is a direct application of the exponential mechanism and the parallel composition theorem. ■

These privacy guarantees can be extended to the electrical prices, according to the resilience to post-processing theorem [2]:

Corollary 1: The electrical prices that result from the optimization problem bear ϵ -differential privacy guarantees. The performance of the system is affected when providing formal privacy guarantees. This introduces a trade-off between performance and privacy, where it is expected that performance degrades with increasing privacy level ($\epsilon \rightarrow 0$). In particular, the performance is measured by the total utility

that nodes enjoy:

$$\text{PERFORMANCE}(p) = \sum_{i=1}^n \sum_{t=1}^T U_{\sigma(i)}(y_{i,t}^*; x_{i,t}), \quad (22)$$

where:

$$\{y_{i,t}\}_t^* = \arg \max_{\{y_{i,t}\}_t} \sum_{t=1}^T U_{\sigma(i)}(y_{i,t}; x_{i,t}) - p_t \cdot y_{i,t} \quad (23)$$

is the optimal power consumption of node i , given electrical prices $\{p_t\}_{t \in \{1, \dots, T\}}$. Note that nodes use their private parameters x while consuming electricity. Privacy concerns that arise from the way billing is performed can be efficiently handled by the work in [12], which provides a secure and private technique computing bills.

V. PERFORMANCE - PRIVACY TRADE-OFF

Note that the prices will no longer be optimal. We evaluate the performance of this scheme numerically in order to depict the performance loss as a function of increasing privacy guarantees. A concrete example of a smart grid system with the associated utility functions and pricing schemes is presented in this section. We consider the case where $N = 41$ nodes compute prices that result in maximizing the total utilization. Node 1 acts as a power plant with no privacy concerns and reports the exact utility function. The setting is similar to the one used in [13]. We employ truncated quadratic utility functions from the family:

$$U_{\text{load}}(u) = \begin{cases} \omega u - \frac{1}{4}u^2, & u \in [0, 2\omega] \\ 2\omega, & u \in (2\omega, \infty) \end{cases}$$

where the family parameter $\omega \in [0, 1]$ controls the level of activity, with higher values denoting increased activity and lower values indicating that the node is less interested in computing large amounts of electricity. We compute the prices hourly for a single day $t \in \{1, \dots, 24\}$. Residential nodes stay relatively inactive during the morning hours and become more active in the afternoon. Specifically, each node becomes active at a random time $t_{\text{on}} \sim \mathcal{N}(10, 1.5)$, i.e., i.i.d. normally distributed around $\mu = 10$ with variance $\sigma^2 = 1.5$, and become inactive at $t_{\text{off}} \sim \mathcal{N}(17, 1.5)$. Parameter ω is chosen i.i.d. across different time slots t and nodes i with $\omega_{\text{off}} \sim \mathcal{U}[0, .4]$, i.e., uniformly distributed between the values $l = 0$ and $u = .4$, and $\omega_{\text{on}} \sim \mathcal{U} [.7, 1]$ for inactive and active time slots, respectively. Node 1 is assumed to be the power plant which lacks any privacy concerns and, thus, reports its utility function exactly:

$$U_{\text{generator}}(u) = -.1u^2$$

The nodes are interested in providing privacy guarantees against temporal localization and activity monitoring of the parameter vector $\{\omega_{i,t}\}_{t \in \{1, \dots, T\}}$. Figure 4 depicts how price signals deviate from the optimal prices for a variety of privacy parameters, whereas Figure 5 depicts how the total utility that nodes enjoy degrades with increasing privacy parameters.




Fig. 4. Electricity prices for different levels of privacy ($\alpha = \frac{\epsilon}{5}$).

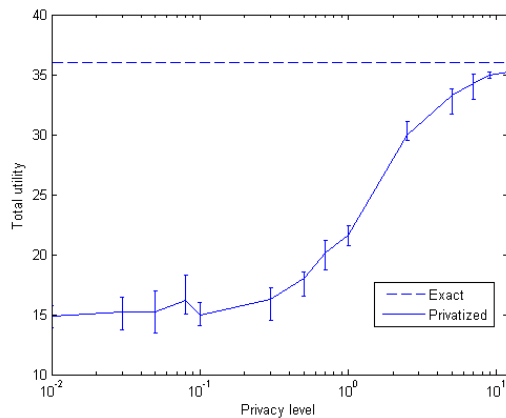


Fig. 5. Utility degradation as a function of the privacy parameter ϵ with $\alpha = \frac{\epsilon}{5}$.

The performance of the system is evaluated according to the equations (22)-(23). The total utility that the nodes enjoy degrades with increased privacy. This is expected, as privacy constraints do not allow computing the optimal prices and lead to a loss in efficiency of the power system. According to Figure 5, for moderate privacy levels ($\epsilon \approx 3$) there is only a slight loss of performance. However, in the high privacy regime, performance degrades significantly.

VI. CONCLUSIONS

In the present work, we highlighted privacy concerns that are raised from employing a pricing scheme to achieve demand response. Using the notion of differential privacy, we proposed a perturbation scheme, where each node reports noisy versions of their utility function. Formal guarantees regarding privacy aspects of interest were given while the performance-privacy trade-off was numerically explored. While performance degrades dramatically in the high privacy regime, it is possible to provide formal privacy guarantees while slightly degrading performance.

Finally, some noteworthy issues, including possible future extensions are the following:

- Merging multiple adjacency relations in a single relation is a non-automated process. However, incorporating multiple adjacency relations is a promising direction of research. For example, in the case of the PHEV charging schedule depicted in Figure 1, the user may be interested in providing privacy with respect to the time driving or the number of trips. The ability to adjust the privacy-aware pricing scheme to incorporate new guarantees is an important aspect.
- A proxy utility function is generated once every time the optimization problem is solved. Each time that the optimization problem is solved, each node can choose either to keep the same proxy utility function or use its privacy budget to generate a new one. In the first case, every realization is going to be biased in the sense that the temporal 'mean' value of the proxy utility function will not be the private one. On the other hand, a mechanism that regenerates privacy budget is required, in order to refresh the proxy function in regular intervals.
- A stochastic treatment of pricing is interesting. First, it can accommodate the random nature of renewable energy sources. Second, as shown in [6], it is possible to redesign the price computing mechanism to account for the added variability which can significantly improve the performance compared to a naive implementation.

REFERENCES

- [1] Michael Backes and Sebastian Meiser, *Differentially private smart metering with battery recharging*, IACR Cryptology ePrint Archive, 2012.
- [2] Cynthia Dwork, *Differential privacy*, Automata, languages and programming, 2006.
- [3] ———, *Differential privacy: A survey of results*, Theory and Applications of Models of Computation, 2008.
- [4] Marisa B Figueiredo, Ana De Almeida, and Bernardete Ribeiro, *An experimental study on electrical signature identification of non-intrusive load monitoring (nilm) systems*, Adaptive and Natural Computing Algorithms, 2011.
- [5] Zhenqi Huang, Sayan Mitra, and Nitin Vaidya, *Differentially private distributed optimization*, arXiv preprint arXiv:1401.2596, 2014.
- [6] Jerome Le Ny and George J Pappas, *Differentially private filtering*, IEEE 51st Annual Conference on Decision and Control, 2012.
- [7] Na Li, Lijun Chen, and Steven H Low, *Optimal demand response based on utility maximization in power networks*, IEEE Power and Energy Society General Meeting, 2011.
- [8] Patrick McDaniel and Stephen McLaughlin, *Security and privacy challenges in the smart grid*, IEEE Security and Privacy, 2009.
- [9] Frank McSherry and Kunal Talwar, *Mechanism design via differential privacy*, IEEE Symposium on Foundations of Computer Science, 2007.
- [10] Darakhshan J Mir, Sibren Isaacman, Ramón Cáceres, Margaret Martonosi, and Rebecca N Wright, *Dp-where: Differentially private modeling of human mobility*, IEEE International Conference on Big Data, 2013.
- [11] Jason Reed and Benjamin C Pierce, *Distance makes the types grow stronger: a calculus for differential privacy*, ACM Sigplan Notices, 2010.
- [12] Alfredo Rial and George Danezis, *Privacy-preserving smart metering*, Proceedings of the 10th annual ACM workshop on Privacy in the electronic society, 2011.
- [13] Pedram Samadi, A Mohsenian-Rad, Robert Schober, Vincent WS Wong, and Juri Jatskevich, *Optimal real-time pricing algorithm based on utility maximization for smart grid*, IEEE International Conference on Smart Grid Communications, 2010.
- [14] Lalitha Sankar, S Raj Rajagopalan, Soheil Mohajer, and H Vincent Poor, *Smart meter privacy: A theoretical framework*, IEEE Transactions on Smart Grid, 2013.