

Achieving Differential Privacy of Data Disclosure in the Smart Grid

Jing Zhao* Taeho Jung* Yu Wang† Xiangyang Li*

* Department of Computer Science, Illinois Institute of Technology, Chicago, IL 60616, USA.

† Department of Computer Science, University of North Carolina at Charlotte, Charlotte, NC 28223, USA.

Abstract— The smart grid introduces new privacy implications to individuals and their family due to the fine-grained usage data collection. For example, smart metering data could reveal highly accurate real-time home appliance energy load, which may be used to infer the human activities inside the houses. One effective way to hide actual appliance loads from the outsiders is *Battery-based Load Hiding* (BLH), in which a battery is installed for each household and smartly controlled to store and supply power to the appliances. Even though such technique has been demonstrated useful and can prevent certain types of attacks, none of existing BLH works can provide probably privacy-preserving mechanisms. In this paper, we investigate the privacy of smart meters via differential privacy. We first analyze the current existing BLH methods and show that they cannot guarantee differential privacy in the BLH problem. We then propose a novel randomized BLH algorithm which successfully assure differential privacy without considering real-world constraints, and further propose the *Multitasking-BLH-Exp3* algorithm which adaptively update the BLH algorithm based on the context and the constraints. Results from extensive simulations show the efficiency and effectiveness of the proposed method over existing BLH methods.

Index Terms—Smart Grid, Smart Meter, Privacy, Differential Privacy, Data Disclosure

I. INTRODUCTION

With the rapid development of the advanced meter infrastructure (AMI) [1] as part of a move to smart grids, the privacy issues regarding the electricity usage information are receiving more and more attention recently. AMI is composed of networked smart meters, but these smart meters not only collect register reads, the monthly electricity consumption information for billing purposes, but also collect interval data (typically the minute-level or second-level electricity usage profile) for controlling purposes. On one hand, this fine-grained information enables trending, forecasting and fault detection analysis, which leads to a more efficient and robust grid system; on the other hand, this information reveals important personal information – human behaviours. For example, by applying Non-Intrusive Load Monitoring (NILM) techniques [2]–[6], attackers can efficiently derive the appliance usage patterns of the residents from the fine-grained energy usage profile.

The concept of NILM is proposed as opposed to intrusive load monitoring (ILM). In ILM, there is an individual monitor for each appliance while in NILM there is only one monitor to acquire the aggregate energy consumption of all the appliances. The target of NILM is to derive the

energy usage profile of each appliance from this aggregated information. The techniques to realize NILM include edge detection, pattern recognition, quadratic integer programming, etc. NILM is originally designed to support the construction of smart homes [4], [5], which learns the lifestyle of the residents, monitors ageing and problematic appliances, and consequently provides safe environment for the alone elder people. However, the NILM technique also enables malicious third parties to acquire the residents behavior patterns, which will reveal the vacant times, the number and the location of the residents inside a house, or even the ages and brands of the appliances. This may cause severe security hazards. For example, if a burglar acquires this information, he immediately knows when and where to break in. In fact, Rouf *et al.* [7] have shown that they can spoof the energy usage information from real world deployed meter systems as a third party and realize the analysis mentioned above (identifying unoccupied residences or people’s routines).

Due to these privacy implications, the deployment of smart meters have encountered obstacles from the public outcry [8], [9]. Some parts in North America and Europe have already banned the deployment of the smart meters [10]. Furthermore, the disputes over the law aspects of AMI is also ongoing [11]. Considering the smart meter system’s great benefits, addressing the privacy issues of smart metering data is crucial to the deployment of smart grid systems.

One effective way to deal with the privacy leakage from the smart meters is *Battery-based Load Hiding* (BLH). The main idea of BLH is to install a battery for each household and use the energy provided (discharge of the battery) and energy consumed (charge of the battery) by the battery to perturb the real energy consumed by the household appliances. By doing so, the real energy consumption of the appliances is hidden in the energy consumption reported by the smart meter. The main challenges faced by the BLH schemes come from the constraints on the capacity and the maximum charging/discharging rate of the battery.

Existing BLH schemes generally try to flatten the energy consumption observed by the smart meters. The representative BLH schemes include the Best Effort (BE) scheme [12], the Non-intrusive Load Levelling (NILL) scheme [13] and the Stepping Framework (SF) [14]. These schemes share the same principle: try to maximize the “distance” between the energy consumed by the appliances and the energy consumption reported by the smart meter. The definition of “distance”

differs from scheme to scheme, but it is based on the entropy theory in general. Other metrics measuring the privacy also include the number of events detected, the cluster classification based metrics, regression based metrics, etc.

Even though current BLH schemes have been demonstrated useful in privacy leakage prevention, they all have certain weaknesses. First of all, their privacy evaluation metric is steamed from the general information theory. The relationship between these metrics and the real privacy is unclear, i.e., there lacks a rigorous definition of the privacy which leads to provably privacy-preserving mechanisms. Moreover, the attacks considered by current BLH schemes are generally limited to edge detection based NILM methods. Thus, it is possible that they are potentially vulnerable for other kinds of attacks. Therefore, there is a need to formally define privacy in BLH problem and design new provably privacy-preserving BLH schemes.

In this paper, we investigate the privacy issues of the smart meters in the *differential privacy* context, which is originally proposed by Dwork *et al.* [15]–[17] as a privacy measure for database queries. Differential privacy mainly captures the increased risk to one’s privacy incurred by participating in a database. It measures the difference of the output distribution before and after an item is put into the dataset. The mechanisms developed under this definition achieve provable privacy in statistic queries, machine learning, and pricing. The most common way to achieve differential privacy is to add noise to the real query result. We study and analyze the BLH problem in the differential privacy context, and formally define the privacy of it.

The BLH problem is not directly solved by a simple perturbation because the noise (i.e., the energy provided or consumed by the battery) in a real smart grid is constrained by the features of the battery, such as capacity, maximum charge/discharge rates. Taking these fundamental constraints into consideration, we further model the BLH problem as a multiple armed bandit (MAB) problem [18], [19], which is an online sequential decision problem. We utilize the Exp3 algorithm ([20]) for MAB to adaptively update the noise distribution in BLH.

The rest of this paper is organized as follows. We briefly review related works on NILM and BLH in Section II and introduce backgrounds of differential privacy and MAB in Section III. We then introduce our system model and the formal BLH problem under differential privacy setting in Section IV. In Section V, we analyze current BLH schemes from the perspective of differential privacy. We then propose a randomization algorithm which can assure differential privacy in Section VI. Section VII presents results from simulation evaluations and Section VIII concludes this paper with some possible future work.

II. RELATED WORK

A. NILM Techniques

The most important category of the NILM technique is the edge detection based mechanism [2]–[4]. These detection

methods aim to capture the event when an appliance is turned on or turned off. By analyzing the sharp changes in the aggregated energy usage profile from the smart meter, these mechanisms could also efficiently derive which appliance is turned on/off. The common features used for analysis include the shape, the amount, the duration, and the time constraint of the changes.

Other methods try to capture the steady state features of the energy usage. Inagaki *et al.* [5] propose a quadratic integer programming based method which tries to find the combination of appliances whose composite current is closest to the observed current. Another technique, ElectriSense [6], is based on the fact that most modern electronics and fluorescent lighting employ switch mode power supplies, which will generate high frequency electromagnetic interference (EMI). By analyzing the features of the EMI, ElectriSense can derive which appliances are in operation. As most smart meters do not have the ability to measure EMI, we do not consider this kind of privacy attacks.

All above mechanisms are generally proposed in the cooperative context, i.e., the NILM algorithm could have all the needed aggregated information. Rouf *et al.* [7] show how to acquire the energy consumption data as a malicious third party in the real world. As the meter system is non-cooperative, they can only acquire partial information of the data using various methods. However, they have shown it is possible to derive the appliance usage information utilizing this partial information.

B. BLH Mechanisms

Current BLH mechanisms generally aim to flatten the energy consumption observed by the smart meter. Mechanisms of this kind try to maintain a constant external load seen by the smart meter. The main difference among these mechanisms is how to react when the battery is too low or too high. In the Best Effort (BE) scheme [12], when the energy level of the battery reaches the minimum level or the maximum level, it requires the battery to charge/discharge at the maximum rate. In the Non-Intrusive Load Levelling scheme (NILM) [13], instead of charging or discharging the battery at the maximum rate, the system chooses a charging/discharging rate that is related to the energy consumption of the appliances.

Yang *et al.* [14] analyze the above two mechanisms and show that these two mechanisms will disclose the true energy consumption when the battery is too low or too high. In addition, they propose a stepping framework (SF) for the BLH system. In this framework, instead of trying to maintain a single constant external load, the BLH system can choose a load to be seen by the smart meter from a set of predefined values according to the current energy consumption level of the appliances.

The problem with all above BLH mechanisms is that they lack a theoretical discussion and evaluation of their system. The BE system is evaluated in terms of relative entropy, clustering classification and correlation/regression while the NILM and the stepping framework mainly evaluate their system in terms of entropy. However, there is no clear evidence to

show how these metrics are directly related to the privacy. Even if a system achieves a high score according to these metrics, how private the system is and whether the system is safe to the attacks other than the attacks mentioned in their papers remain to be questioned.

III. BACKGROUNDS

A. Differential Privacy

The concept of *differential privacy* is originally introduced by Dwork [17]. In [17], Dwork prove that it is infeasible to achieve the universal privacy in the database, i.e., it is impossible to achieve that “access to a statistical database should not enable one to learn anything about an individual that could not be learned without access”. In contrast, he proposes a weaker definition: differential privacy, which captures the increased risk to one’s privacy incurred by participating a database. Ever since differential privacy is introduced, a bunch of privacy mechanisms [21]–[24], which can achieve provable privacy by the rigorous definition, have been proposed in the areas of data mining, statistical query, and many related areas.

In this paper, by leveraging the powerful usage of differential privacy, we will use it to study smart metering data privacy. Here, we exploits the following definition of (δ, ϵ) -differential privacy.

Definition 1. Given an n -dimension datasets D^n , a randomized algorithm is (δ, ϵ) -differentially private if $\forall x, y \in D^n$ that differs only in one element and all $S \in \text{range}(A)$,

$$\Pr[A(x) \in S] \leq e^\epsilon \times \Pr[A(y) \in S] + \delta,$$

where $\text{range}(A)$ denotes the output range of A .

Informally, this definition says if two datasets differs only in one element, the outcome of the query A over these two datasets should be indistinguishable. The closer ϵ and δ are to zero, the more private A is. One common way to achieve such A is to add special noise to the original queries. In addition, suppose $f(\cdot)$ is some query function over the dataset. The global sensitivity Δf is defined as follows:

$$\Delta f = \max_{x,y} |f(x) - f(y)|,$$

for all x, y that differing in at most one element. Then, when $\Delta f = 1$, the function $f(\cdot)$ achieves (ϵ, δ) -differential privacy if noises from the binomial distribution $B(\frac{1}{2}, n) - \frac{n}{2}$ is added to it, where n is number of queries, and n satisfies $n \geq -64 \ln(\delta)/\epsilon^2$ [16].

B. MAB Problems and Relative Solutions

The multi-armed bandit problem [18], [19] is a sequential decision problem defined by a set of actions. At each step, the system can choose an action from the action set and some payoff is observed. The fundamental issues in the MAB problem is to handle the trade off between exploration and exploitation in sequential experiments to maximize the payoffs. According to the nature of the payoff rewarding process, the MAB problem can be categorized as stochastic,

adversarial and Markovian [18]. In the stochastic MAB, the reward follows some distribution; in the adversarial MAB, the reward is given in an arbitrary manner; in the Markovian MAB, the reward is given according to the state of the arm. Different MAB problems have different solutions. The stochastic MAB is generally solved by upper confidence bound (UCB) based schemes [25]. The adversarial MAB is mainly solved by the Exp3 algorithm [20] and its variations. The Markovian MAB is usually tackled with Gittins indices [18], [19].

We define a contextual bandit problem for our BLH problem, and use the \mathcal{S} -Exp3 algorithm [18] to guarantee the lower bound of the reward (defined later) in this paper. The generic contextual bandit problem is shown as follows.

Contextual Bandit Problem

Known parameters: K arms to choose & number of rounds $n \geq K$

For each round $t = 1, 2, \dots$

- (1) Forecaster chooses $I_t \in \{1, \dots, K\}$
- (2) Adversary chooses a gain vector $g_t = (g_{1,t}, \dots, g_{K,t}) \in [0, 1]^K$
- (3) Forecaster receives $g_{I_t,t}$ and learns nothing else.

Typical objective function to maximize in MAB problem is the pseudo-regret $\bar{R}_n = \max_i E [\sum_{t=1}^n g_{i,t} - \sum_{t=1}^n g_{I_t,t}]$.

IV. PROBLEM FORMATION

A. System Model

The battery with capacity C is connected to the original house hold electricity network, and charging or discharging it (whose maximum rates are both β) adds noises (i.e., the battery energy $b(t)$) to the real load of the appliances $d(t)$, which constitutes the smart meter’s reading $s(t) = b(t) + d(t)$, and the $c(t)$ is the energy stored at the battery at time t . Table I summarizes the notations and Figure 1 shows the BLH system we have. In addition, we use d_{max} to denote the upper bound of the $d(t)$ over the entire time domain. We assume a discrete time domain having equal-length intervals (e.g., smart meter’s data collection cycle) in this paper, then we have the following constraints for any $t \in \{0, 1, 2, \dots\}$:

$$\begin{cases} c(t) = c(0) + \sum_{i=0}^t b(i) & \text{Accumulated Usage} \\ 0 \leq c(t) \leq C & \text{Capacity Requirement} \\ -\beta \leq b(t) \leq \beta & \text{Charging/Discharging Rate} \\ s(t) \geq 0 & \text{Households Cannot Emit Energy} \end{cases}$$

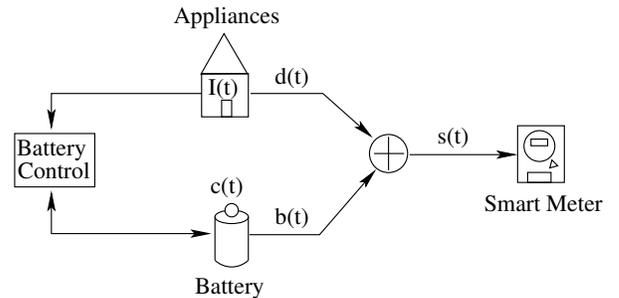


Fig. 1. The BLH system.

TABLE I
IMPORTANT NOTATIONS

C	Capacity of the battery
β	Maximum charging/discharging rate
$s(t)$	Smart meter's reading
$b(t)$	Battery's energy consumption $b(t) > 0$ means battery is being charged
$c(t)$	Energy stored at the battery
$d(t)$	Appliances' energy consumption
$I(t)$	Set of running appliances

B. Adversary Model

The battery is out of adversaries' control, and they can only access the smart meters' reading. Adversaries are interested in the real load of household's appliances $d(t)$, and they try to analyze the time-series data $s(0), s(1), \dots$ to infer the appliances usage pattern (i.e., distribution of the usage).

Adversaries are assumed to be *ignorant*, in the sense that they may have some general 'common sense' about the energy consumption patterns (high-usage in evenings and low-usage in midnights), but they do not have detailed and targeted prior information on specific household's appliances usage pattern.

C. Problem Formulation

Given the definition of (δ, ϵ) -differential privacy in Section III, we can formulate the BLH problem as follows.

Problem 1 (BLH Problem). *Given the dataset $I(t)$ and the query function $f(\cdot)$ over $I(t)$ which returns $d(t)$ (i.e., $f(I(t)) = d(t)$), we aim to devise a randomization algorithm \mathcal{A} which adds $b(t)$ as a noise to the query result to hide $d(t)$ from the adversaries so that (δ, ϵ) -differential privacy is guaranteed with good δ, ϵ .*

It seems the solution is already given in [16], but we cannot directly use it because Δf is larger than 1 and n may be smaller than $-64 \ln(\delta)/\epsilon^2$ in the smart grid context. Thus, we need to find a good randomization algorithm which is differentially private in the smart grid context.

V. ANALYSIS OF EXISTING BLH SCHEMES WITH DIFFERENTIAL PRIVACY

In this section, we first analyze the existing BLH schemes from the perspective of differential privacy. More specifically, we mainly analyze existing BLH schemes by showing the best δ (the smallest δ) they can achieve. We use δ as the metric to measure the existing BLH scheme since if δ is large enough, we can achieve arbitrarily small ϵ . If $\delta = 1$, $Pr[A(x) \in S] \leq 1$ is for sure, i.e. any scheme can achieve a $(0, 1)$ -differential privacy.

A. The BE scheme

The Best Effort (BE) scheme [12] tries to maintain a steady state where the system exposes a constant external load $s(t) = K_c$ to the smart meter. However, there are four cases that BE could not maintain this steady. These cases are listed as follows:

$$\begin{cases} c(t-1) - d(t) + K_c < 0 & 1) \text{ Battery energy is too low} \\ c(t-1) - d(t) + K_c > C & 2) \text{ Battery energy is too high} \\ d(t) - K_c > \beta & 3) \text{ Charging is too fast} \\ K_c - d(t) > \beta & 4) \text{ Discharging is too fast} \end{cases}$$

If case 1) or case 2) occurs, BE switches its target output load to the current true demand i.e. $s(t) = d(t)$. If case 3) occurs, BE charges at the maximum rate. If case 4) occurs BE discharges at the maximum rate.

Now we analyze BE scheme from the perspective of differential privacy. Recall that in differential privacy, we mainly focus on two datasets that differ only in one item. Correspondingly, in BLH, we mainly focus on two time points t_1 and t_2 where the sets of operating appliances $I(t_1)$ and $I(t_2)$ only differ in one appliance. If BE can maintain a constant external load, it obvious that $s(t_1) = s(t_2)$. If BE cannot maintain the constant external load, i.e. the system encounters case 1), 2), 3) or 4) as described above, the system will charge or discharge at the maximum rate. Then, we have:

$$\begin{cases} s(t_1) = d(t_1) \& s(t_2) = d(t_2) & \text{Case 1) and 2)} \\ s(t_1) = d(t_1) + \beta \& s(t_2) = d(t_2) + \beta & \text{Case 3)} \\ s(t_1) = d(t_1) - \beta \& s(t_2) = d(t_2) - \beta & \text{Case 4)} \end{cases}$$

Since $I(t_1)$ and $I(t_2)$ differ in one element, as long as $|d(t_1) - d(t_2)| \neq \beta$, which is a trivial case, we have $d(t_1) \neq d(t_2)$. This indicates $s(t_1) \neq s(t_2)$ in all cases. Let $S_1 = \{d(t_1), d(t_1) + \beta, d(t_1) - \beta\}$ and $S_2 = \{d(t_2), d(t_2) + \beta, d(t_2) - \beta\}$, then we have

$$\begin{aligned} Pr[s(t_1) \in S_2] &= Pr[s(t_2) \in S_1] = 0 \\ \Rightarrow \begin{cases} Pr[s(t_1) \in S_1] - e^\epsilon * Pr[s(t_2) \in S_1] \leq \delta \\ Pr[s(t_2) \in S_2] - e^\epsilon * Pr[s(t_1) \in S_2] \leq \delta \end{cases} \\ \Rightarrow \begin{cases} \delta \geq Pr[s(t_1) \in S_1] \\ \delta \geq Pr[s(t_2) \in S_2] \end{cases} \\ \Rightarrow \delta &\geq Pr[BE \text{ is not in steady state}] \end{aligned}$$

which means the δ that BE can achieve is larger than the probability that the system encounters cases 1), 2), 3) and 4).

The probability of the system enters the non-steady state in BE is mostly determined by the load patterns. If the load is constantly above or below predetermined K_c , it is likely that the system will enter the non-steady state. Thus, the differential privacy of BE could not be guaranteed.

Theorem 1. *The BE cannot guarantee differential privacy for BLH problem.*

B. The NILL scheme

The Non-Intrusive Load Levelling scheme (NILL) [13] defines three states: (1) the stable state where the residue energy in the battery is neither too low or too high; (2) the low recovery state where maintaining the K_c will deplete the residue battery energy; (3) the high recovery state where maintaining the K_c will overcharge the battery.

In the stable states, there are two sub states: (S1) the battery charge rate could maintain K_c ; (S2) the battery charge rate

could not maintain K_c . In the first sub state, $s(t) = K_c$; in the second sub state $s(t) = d(t) + \beta$ or $s(t) = d(t) - \beta$. If the battery is too low, the system will enter the low recovery state; if the battery is too high, the system will enter the high recovery state.

In the low recovery state, there are also two sub states: $L1$) $d(t) \leq \beta$ and $L2$) $d(t) > \beta$. In the previous case, $s(t)$ is set to β ; in the later state, the external load $s(t)$ is set to be $d(t)$, i.e the real load. The system will return to the stable state if the residue energy in the battery is larger than $0.8C$ i.e. $c(t) > 0.8C$, where C is the capacity of the battery.

In the high recovery state, the external load $s(t)$ is set to be $d(t) - 0.5AMP$. The system will return to the steady state if $c(t) > 0.5C$ or $d(t) > 4.5AMP + d(t - 1)$.

Now we analyze the NILL algorithm from the perspective of differential privacy. In the state $S1$, $S2$, $L2$ and the high recovery state, two neighbouring sets of appliances $I(t_1)$ and $I(t_2)$ that differ in only one appliance will have distinct external load. Using a similar analysis as that used in BE , we can derive that δ that NILL can achieve should be larger than the probability that the system is in state $S1$, $S2$, $L2$ and the high recovery state.

Again, the chance that NILL encounters $S1$, $S2$, $L2$ and the high recovery state is determined by the pattern of $d(t)$. There is no guarantee of the δ .

Theorem 2. *The NILL cannot guarantee differential privacy for BLH problem.*

C. The Stepping Framework

The schemes under the stepping framework (SF) [14] try to maintain the external load of the algorithm to be multiples of β . If the real load satisfies $(k-1)\beta \leq d(t) \leq k\beta$, the external load $s(t)$ will be set to $(k-1)\beta$ or $k\beta$. Suppose the largest energy consumption of an appliance is b_{max} . If $b_{max} > 2\beta$, clearly a data set containing b_{max} and a data set that doesn't contain b_{max} will never output the same external load (as a matter of fact, no scheme can). In this case, $\delta = 1$. Therefore, we focus on the case where $b_{max} < 2\beta$.

We will give an instance to derive the lower bound of δ for the stepping frameworks. We create two neighbouring appliance sets $I(t_1)$ and $I(t_2)$ that differ only in one appliance with energy consumption b_{max} . Without loss of generality, suppose $d(t_1) = k\beta + b'$, $d(t_2) = k\beta + b' - b_{max}$ and $d(t_1) < (k+1)\beta$. Suppose $b' < b_{max}$, $s(t_1)$ is chosen from $\{k\beta, (k+1)\beta\}$ and $s(t_2)$ is chosen from $\{k\beta, (k-1)\beta\}$. Then, we require 1) $\delta \geq Pr[s(t_1) = (k+1)\beta]$ 2) $\delta \geq Pr[s(t_2) = (k-1)\beta]$ We also require that $I(t_1)$ appear consecutively for N times and $I(t_2)$ appear for another N times consecutively. Let $Pr[s(t_1) = k\beta] = p$ and $Pr[s(t_2) = k\beta] = p'$, then due to the constraint of the battery, we have

$$\begin{cases} -C \leq Npb' - N(1-p)(\beta - b') \leq C \\ -C \leq N(1-p')(\beta + b' - b_{max}) - Np'(b_{max} - b') \leq C \end{cases} \Rightarrow \begin{cases} pb' - (1-p)(\beta - b') = 0 \\ (1-p')(\beta + b' - b_{max}) - p'(b_{max} - b') = 0 \end{cases}$$

when $N \gg C$

$$\Rightarrow \begin{cases} p = 1 - \frac{b'}{\beta} \\ p' = 1 - \frac{b_{max} - b'}{\beta} \end{cases} \Rightarrow \begin{cases} Pr[s(t_1) = (k+1)\beta] = 1 - p = \frac{b'}{\beta} \\ Pr[s(t_2) = (k-1)\beta] = 1 - p' = \frac{b_{max} - b'}{\beta} \end{cases} \Rightarrow \delta \geq \frac{b_{max}}{\beta}$$

when b' is close to 0

$$\Rightarrow \epsilon = \frac{p}{p'} = \frac{\beta - b'}{\beta + b' - b_{max}} = \frac{\beta}{\beta - b_{max}}$$

Therefore, we can also conclude with the following theorem:

Theorem 3. *The SF schemes cannot guarantee differential privacy for BLH problem.*

VI. NOVEL BLH SCHEMES ACHIEVING DIFFERENTIAL PRIVACY

In this section, we first give two randomization algorithms which generate noises to assure differential privacy, and then present which noise to choose from the candidate set given by the randomization algorithms in a real life scenario by considering the constraints from the battery and system over time period.

A. Randomization Algorithms

We now present two randomization algorithms to generate noises to assure differential privacy.

1) *Coarse-grained Noise:* As a first step, we present the first randomization algorithm as Algorithm 1.

Algorithm 1 BLH based on Coarse-grained Noise

Input: $I(t)$ and $f(\cdot)$, s.t., $f(I(t)) = d(t)$

Output: $\mathcal{A}_C(f(\cdot)) = s(t)$.

1: $\mathcal{A}_C(f(\cdot)) = f(\cdot) + r \cdot b_{max}$, where r is generated from $B(\frac{1}{2}, n) - \frac{n}{2}$ and $n \geq \frac{-64 \ln(\delta)}{\epsilon^2}$. In other words, the battery chooses a noise $r \cdot b_{max}$ from the noises set $\{-\frac{n}{2}\beta, (1 - \frac{n}{2})\beta, \dots, (\frac{n}{2} - 1)\beta, \frac{n}{2}\beta\}$ and charges or discharges accordingly.

2: **return** $\mathcal{A}_C(f(\cdot))$

Theorem 4. *Algorithm 1 ensures (δ, ϵ) -differential privacy as long as n satisfies $n \geq \frac{-64 \ln(\delta)}{\epsilon^2}$.*

Proof: Recall that if we add a binomial noise $B(\frac{1}{2}, n) - \frac{n}{2}$ that satisfies $n \geq -64 \ln(\delta)/\epsilon^2$, we can achieve (δ, ϵ) -differential privacy for a query $f(\cdot)$ with global sensitivity 1 (Section III).

We have defined the query function $f(I(t)) = d(t)$ for the set of appliances at time t . However, the global sensitivity Δf in our BLH problem is the energy consumption of the most energy consuming appliance, which is assumed to be b_{max} . Therefore, we first define a new query $f'(\cdot)$ over the set of $I(t)$ as

$$f'(I(t)) = d(t)/b_{max}$$

In this case, the global sensitivity $\Delta f' = 1$ and $s(t) = f(I(t)) \cdot d_{max}$. Then, the following randomization algorithm \mathcal{A} for the function f' achieves (δ, ϵ) -differential privacy:

$$\mathcal{A}(f'(\cdot)) = f'(\cdot) + r, \quad r \sim B\left(\frac{1}{2}, n\right) - \frac{n}{2} \quad \text{s.t.} \quad n \geq \frac{-64 \ln(\delta)}{\epsilon^2}$$

Finally, our randomization algorithm \mathcal{A}_C for the query function f in our BLH problem is:

$$\mathcal{A}_C(f(\cdot)) = f(\cdot) + r \cdot b_{max}, \quad r \sim B\left(\frac{1}{2}, n\right) - \frac{n}{2} \quad \text{s.t.} \quad n \geq \frac{-64 \ln(\delta)}{\epsilon^2}$$

which is (δ, ϵ) -differentially private. \blacksquare

2) *Fine-grained Noise*: The above algorithm is coarse-grained in the sense that the smallest unit of noise is b_{max} . Now, we introduce a fine-grained randomization algorithm whose noise unit is 1 (Algorithm 2).

Algorithm 2 BLH based on Fine-grained Noise

Input: $I(t)$ and $f(\cdot)$, s.t., $f(I(t)) = d(t)$

Output: $\mathcal{A}_F(f(\cdot)) = s(t)$.

- 1: $\mathcal{A}_F(f(\cdot)) = f(\cdot) + r$, where r is generated from $B(\frac{1}{2}, n) - \frac{n}{2}$ and $n \geq \frac{-3 \ln \theta}{\Theta^2(\epsilon)}$. In other words, the battery chooses a noise from $\{-\frac{n}{2}, -\frac{n}{2} + 1, \dots, \frac{n}{2} - 1, \frac{n}{2}\}$ and charges/discharges accordingly.
 - 2: **return** $\mathcal{A}_F(f(\cdot))$
-

Theorem 5. *Algorithm 2 ensures (δ, ϵ) -differential privacy as long as n satisfies $n \geq \frac{-3 \ln \theta}{\Theta^2(\epsilon)}$, where $\Theta(\epsilon) = \frac{(1-l)e^{\frac{\epsilon}{b_{max}}} - 1}{2(1+e^{\frac{\epsilon}{b_{max}}})}$*

Proof: Suppose there are two neighbouring appliance sets $I(t_1)$ and $I(t_2)$ and $d(t_1) - d(t_2) = b_j$, where b_j denotes the difference of energy consumption of two sets. Then we have $s(t_1) = d(t_1) + \text{noise}$ and $s(t_2) = d(t_1) + \text{noise} + b_j$. Finding the bound of the probability that $s(t_1)$ and $s(t_2)$ give the same value is equivalent to find the bound of the probability that generated noise is x and $x + b_j$. The larger b_j is, the greater the possible probability gap will be. As $b_j \leq b_{max}$. We only need to find a bound of the probability that an arbitrary noise is x and $x + b_{max}$.

Now we consider adding a binomial noise to the real demand. Suppose the noise is drawn from $B(\frac{1}{2}, n) - \frac{n}{2}$, then the probability of adding a noise x is $Pr[x + \frac{n}{2}] = \binom{n}{n/2+x} \frac{1}{2^n}$. The probability of adding a noise $x + b_{max}$ is $Pr[x + \frac{n}{2} + b_{max}] = \binom{n}{n/2+x+b_{max}} \frac{1}{2^n}$. Then we have

$$\begin{aligned} & \frac{Pr[x + \frac{n}{2}]}{Pr[x + \frac{n}{2} + b_{max}]} \\ &= \frac{\binom{n}{\frac{n}{2} + x + 1} \binom{n}{\frac{n}{2} + x + 2} \dots \binom{n}{\frac{n}{2} + x + b_{max}}}{\binom{n}{\frac{n}{2} - x - b_{max} + 1} \binom{n}{\frac{n}{2} - x - b_{max} + 2} \dots \binom{n}{\frac{n}{2} - x}} \\ &\leq \left(\frac{\frac{n}{2} + x + 1}{\frac{n}{2} - x - b_{max} + 1}\right)^{b_{max}} \end{aligned}$$

Suppose $b_{max} \leq \frac{l}{2}n$ where $0 \leq l \leq 1$. Then, we have:

$$x \leq \frac{(1-l)e^{\frac{\epsilon}{b_{max}}} - 1}{2(1+e^{\frac{\epsilon}{b_{max}}})} * n \rightarrow \frac{Pr[x + \frac{n}{2}]}{Pr[x + \frac{n}{2} + b_{max}]} \leq \epsilon$$

Since b_{max} and l are known parameters, we further define

$$\Theta(\epsilon) = \frac{(1-l)e^{\frac{\epsilon}{b_{max}}} - 1}{2(1+e^{\frac{\epsilon}{b_{max}}})}$$

According to the Chernoff bound, we have

$$Pr[y > \frac{n}{2} + \Theta(\epsilon)n] = Pr[y > \frac{n}{2}(1 + 2\Theta(\epsilon))] \leq e^{(-\frac{n\Theta^2(\epsilon)}{3})}$$

$$\text{and} \quad n \geq \frac{-3 \ln \delta}{\Theta^2(\epsilon)} \rightarrow e^{(-\frac{n\Theta^2(\epsilon)}{3})} < \delta$$

Then, as long as $n \geq \frac{-3 \ln \delta}{\Theta^2(\epsilon)}$ the following randomization algorithm \mathcal{A}_F guarantees (δ, ϵ) -differential privacy.

$$\mathcal{A}_F(f(\cdot)) = f(\cdot) + r, \quad r \sim B\left(\frac{1}{2}, n\right) - \frac{n}{2} \quad \text{s.t.} \quad n \geq \frac{-3 \ln \theta}{\Theta^2(\epsilon)}$$

B. Noise Selection under Constraints

In the aforementioned algorithms, the battery can choose a noise to add to $d(t)$ by charging or discharging itself. The battery draws the noise from a binomial distribution, but it is not always feasible to add the desired noise owing to several constraints on the battery and the system over the time period as summarized as follows: $\forall t$: 1) $0 \leq c(t) \leq C$; 2) $-\beta \leq b(t) \leq \beta$; 3) $s(t) \geq 0$.

Therefore, we need to update the distribution according to the context $(c(t), b(t), d(t))$. Since we cannot foresee the future's usage pattern, this forms an online selection problem with constraints. We solve this non-trivial problem by solving the following contextual multi-armed bandit (MAB) problem, which is defined as follows.

Problem 2 (Contextual Multi-Armed Bandit Problem For Noise Selection). *Given a set of arms $\{k_1, \dots, k_m\}$ which satisfy $\frac{-\beta}{b_{max}} < k_1 < k_2 < \dots < k_{m-1} < k_m < \frac{\beta}{b_{max}}$. Our randomization algorithm \mathcal{A}_C or \mathcal{A}_F need to choose an arm k_i at t and modify $q(t)$ as*

$$q(t) = \min\left\{\frac{2(\beta - |k_i|)}{b_{max}}, \frac{2(c(t) + k_i)}{b_{max}}, \frac{2(C - c(t) - k_i)}{b_{max}}, \frac{2(d(t) + k_i)}{b_{max}}\right\}$$

Then, the randomization algorithm (\mathcal{A}_C or \mathcal{A}_F) chooses $r_{i,t}$ from the binomial distribution $B(\frac{1}{2}, q(t)) - \frac{q(t)}{2} + k_i$.

If we deem δ as a predefined value and set $q(t) = \frac{-64 \ln(\delta)}{\epsilon^2}$ for \mathcal{A}_C or $q(t) = \frac{-3 \ln \theta}{\Theta^2(\epsilon)}$ for \mathcal{A}_F , we can acquire ϵ as a function of the arm k_i and the time t , which we denote as $\epsilon(i, t)$.

Then, the loss of each arm k_i at time t is defined as:

$$L_{i,t} = (1 - \alpha) \cdot \left| \frac{1}{2} - \frac{c(t-1) + r_{i,t}}{C} \right| + \alpha \cdot e^{-\epsilon(i,t)}$$

The battery needs to choose arms in a online manner such that following pseudo-regret is minimized:

$$\bar{R}_n = \max_i E \left[\sum_{t=1}^n L_{I_t,t} - \sum_{t=1}^n L_{i,t} \right]$$

The ‘‘Context’’ comes from different $(c(t), d(t))$ pairs. $q(t)$ is updated every time $(c(t), d(t))$ is changed, and thus the randomization algorithm (\mathcal{A}_C or \mathcal{A}_F) have a different binomial algorithm to choose the noise from.

The reason to use $L_{i,t}$ as the loss is: 1) the closer ϵ is to zero the more private the algorithm will be; 2) we need to consider the potential danger of exhausting or overcharging the battery and try keep the residue energy of the battery to be around $\frac{C}{2}$.

Given the above MAB problem, we use the following BLH-Exp3 algorithm (Algorithm 3, a fast variant of Exp3 [20]), as a building block to choose the noises at each t .

Algorithm 3 BLH Exp3

Parameters: $\eta = \sqrt{\frac{2 \ln K}{nK}}$, $\mathbf{p}_1 = (p_{1,1}, p_{2,1}, \dots, p_{m,1}) = (\frac{1}{m}, \dots, \frac{1}{m})$, $\forall i: \hat{L}_{i,1} = 0$

- 1: Find the m arms between $(\frac{-\beta}{b_{max}}, \frac{\beta}{b_{max}})$ which are uniform randomly distributed.
- 2: **for all** round $t' = 1, 2, \dots, n$ **do**
- 3: Choose a noise $k_{I_{t'}}$ where $I_{t'} \sim \mathbf{p}_{t'}$
- 4: For the noise $k_{I_{t'}}$, compute the estimated loss $l_{I_{t'}, t'} = \frac{L_{I_{t'}, t'}}{p_{I_{t'}, t'}}$. For other noises, set the estimated loss as 0.
- 5: Update every k_i 's cumulative loss

$$\hat{L}_{i,t'} = \hat{L}_{i,t'-1} + l_{i,t'}$$

- 6: Compute the new distribution

$$\mathbf{p}_{t'+1} = (p_{1,t'+1}, \dots, p_{m,t'+1}),$$

$$\text{where } p_{i,t'+1} = \frac{\exp(-\eta \hat{L}_{i,t'})}{\sum_{k=1}^m \exp(-\eta \hat{L}_{k,t'})}$$

Every time the battery is faced with the new context $(c(t), d(t))$, it runs a separated new instance of BLH-Exp3 where every instance owns its own clock time t' . The clock time t' increases only when the battery encounters the same context and thus recalling the corresponding BLH-Exp3 instance (similar to a CPU's multitasking). We denote this algorithm as **Multitasking-BLH-Exp3**.

Lemma 1. *The Multitasking-BLH-Exp3 algorithm guarantees the following upper bound of the pseudo-regret until $t = n$:*

$$\bar{R}_n = \max_i E \left[\sum_{t=1}^n L_{I_t, t} - \sum_{t=1}^n L_{i, t} \right] \leq \sqrt{2n|S| \ln m}$$

where I_t is the battery's arm selection (i.e., k_{I_t}) at time t , and S is the universe set of all contexts

The lemma is derived directly from the corresponding proof in [18], which is omitted due to space limit.

VII. EVALUATION

In this section we will evaluate our proposed BLH schemes based on real world electricity usage trace. The dataset we use is the MIT REDD dataset [26] which provides second level power consumption information of six houses for roughly one month. Of these six houses, the data traces of two houses are too sparse, thus we mainly use the traces of the other four. An example of power consumption of a house is shown in

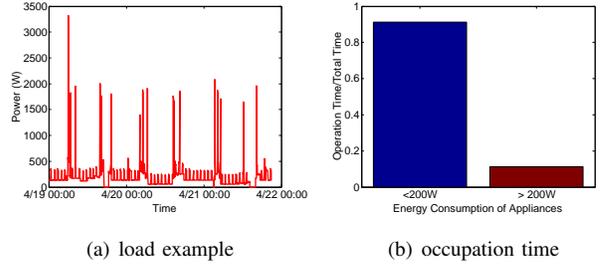


Fig. 2. (a) An example of power consumption of a single house in the MIT REDD data traces. (b) Time occupations of the appliances.

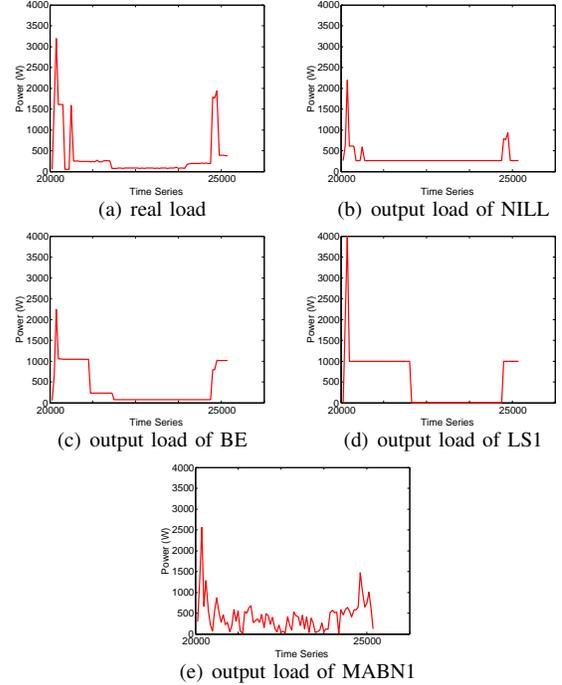


Fig. 3. Examples of the real load (a) and the corresponding outputs from different algorithms (b-e).

Figure 2(a). We mainly analyze our and existing BLH schemes from three aspects: 1) the differential privacy, i.e. (δ, ϵ) ; 2) the mutual information metric used by Yang *et al* [14]; 3) the event detection accuracy. We compare our schemes with BE, NILL and the schemes under the stepping framework, namely LS1, LS2, LC, RC. For our proposed schemes, we test the following two: MAB using coarse noise generation scheme (Algorithms 1 and 3 together, denoted by MABN1) and MAB using fine grained noise generation scheme (Algorithms 2 and 3 together, denoted by MABN2). An example of the outcomes of these algorithms is given in Figure 3.

As we discussed in Section VI, the differential privacy we can get is directly related to the maximum energy consumption of a single appliance and the maximum charge/discharge rate of the battery. The greatest energy consumption of household appliances is typically 3KW (such as washers and driers). If we try to protect the differential privacy of those appliances, the resulting maximum charge/discharge rate could be too large.

However, as Figure 2(b) shows, in the power usage traces, the energy consuming appliances (here we refer to the appliances with power larger than 200W) only operates for about 11% of the time and those works under 200W operates for more than 90% of the time. This indicates that we can treat those energy consuming appliances as outliers and only provide the most commonly used appliances with high ranked privacy. In our evaluations, we set the maximum energy consumption of the appliances to be 200W and the maximum charge/discharge rate we use for the battery is 1000W. For δ , we set its value to be 0.2. The bound of ϵ for coarse grained noise is 3.2 and the ϵ 's bound for fine grained noise (with a granularity of 50W) is 20.43 in theory. However, as we will show later, the real epsilon we acquired from real word data traces shows that the actual ϵ we can get is better than these theoretical bound (the bound is not tight for specific data traces).

Note that though we only provide differential privacy for part of the appliances, we by no means provide less protection for the appliances than existing BLH methods. Therefore, we also measure the metric used by Yang *et al.* [14] for comparison in our evaluations.

A. Differential Privacy

For differential privacy, we choose a granularity of 50W, i.e., the minimum set we consider for differential privacy is with a range of 50. Here, we consider two cases: 1) consider the energy consuming appliances; 2) do not consider the energy consuming appliances. We denote our schemes under case one as MABN1(LA) and MABN2(LA). For this case, we set the energy consumption of the largest appliance to be 3000W. For this specific setting, we temporarily set the maximum discharge rate to 15KW to make it possible to achieve a δ of 0.2. From Figure 4, we can see that under this setting, it is infeasible to achieve desirable δ and ϵ . The main reason is that under this setting, the added noise will quickly deplete the energy of the battery or overcharge the battery. This will make it hard to add the desirable noise in the long run.

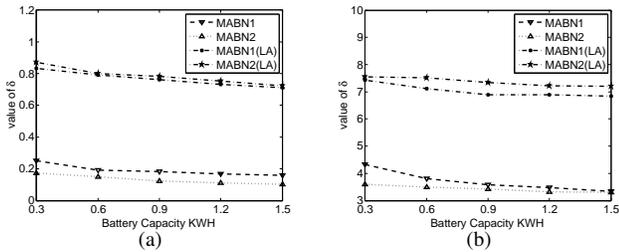


Fig. 4. The changes of δ (a) or ϵ (b) with respect to battery capacity.

For the second case, we only compare the distribution of neighbouring appliances sets I_{t_1} and I_{t_2} when $|d(t_1) - d(t_2)| < 200$ (as the greatest energy consumption we consider here is 200W). The changes of the observed δ and ϵ with respect to the change of battery capacity are shown in Figure 4 too. We can observe that MABN1 performs worse than MABN2 in terms of δ and ϵ . The reason is that the theoretical

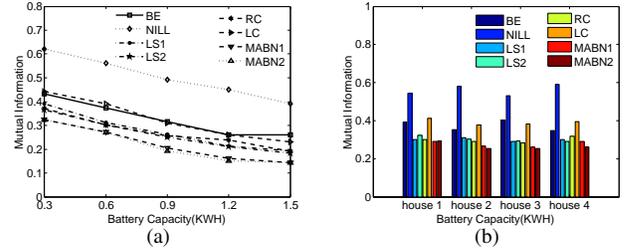


Fig. 5. (a) The changes of the mutual information between $s'(t)$ and $d'(t)$ with respect to battery capacity. (b) The mutual information between $s'(t)$ and $d'(t)$ for different houses when the battery capacity is set to 0.6KWH.

bound of MABN1 is given based a granularity of the maximum energy consumption of the appliances (which is 200W). In this evaluation, however, we evaluate the algorithm in a granularity of 50W. Thus, MABN1 doesn't perform better than MABN2 as the theoretical bound shows.

B. Mutual Information

The mutual information used by Yang *et al.* [14] is defined as follows. Given $d(t)$ and $e(t)$ over time series $t = 1, 2, 3, \dots, n$, the mutual information between $s(t)$ and $d(t)$ is defined as

$$M(s, d) = \sum_t \sum_{s(t)} \sum_{d(t)} \log \frac{p(s(t), d(t))}{p(s(t))p(d(t))}.$$

The values that Yang *et al.* used are slightly different from ours as the mutual information they evaluate is $s'(t) = s(t) - s(t-1)$ and $d'(t) = d(t) - d(t-1)$, i.e. the mutual information between the change of values. They mainly evaluate the how robust their scheme is against edge detection. However, we believe that the protection of the mutual information of the absolute values is also important. It has been shown that the operating appliances can be inferred purely based on $d(t)$ [5]. The simulation results of the mutual information for load change and for absolute value are shown in Figure 5 and Figure 6 respectively. We can see that the protections of load change of MABN1 and MABN2 are better than all the other algorithms. As for the protection of absolute load value, MABN1 and MABN2 performs better than BE and NILL, but worse than the schemes of the stepping framework. The main reason is that the schemes under stepping framework tries to maintain a small set of discrete values. This will hide more information for the pure load. However, the schemes under stepping framework is not aware of differential privacy and could not provide the differential bound as our schemes do.

C. Events Detection Accuracy

Here we define the events based on the changes of the overall energy consumption. We deem the change of demand greater than 50W as an occurred event. The events detection accuracy is defined as the ratio between the accurately occurred events and the total detected events from the load output. Results are reported in Table II. It is obviously that both our algorithm and stepping framework based methods outperforms BE and NILL.

TABLE II

EVENT DETECTION PRECISION. THE PRECISION IS DENOTED BY a/b , WHERE a IS THE AVERAGE NUMBER DAILY EVENTS THAT COULD ACCURATELY DETECTED (THE DIFFERENCE OF LOAD CHANGE SHOULD NOT BE LESS THAN 10%) AND b IS THE AVERAGE NUMBER OF DAILY EVENTS THAT COULD BE DETECTED FROM THE OUTCOME OF THE SCHEME.

Battery	NILL	BE	LS1	LS2	LC	RC	MABN1	MABN2
0.3KWH	45.1/52.9(85.28%)	11.2/12.0(93.23%)	1.8/19.8(9.26%)	2.6/23.2(11.21%)	3.6/43.4(8.23%)	1.5/60.6(2.42%)	2.2/113.5(1.92%)	3.0/123.8(2.43%)
0.6KWH	42.5/50.4(84.36%)	7.8/8.9(87.16%)	1.8/19.4(9.31%)	2.3/21.4(10.83%)	3.3/40.2(8.12%)	1.6/60.2(2.71%)	2.1/110.1(1.87%)	2.6/122.5(2.13%)
0.9KWH	41.3/49.7(83.17%)	6.8/8.1(83.35%)	1.4/18.3(7.43%)	2.0/19.7(10.35%)	2.4/38.5(6.23%)	2.3/59.3(3.92%)	2.1/109.1(1.93%)	2.3/121.6(1.93%)
1.2KWH	40.2/49.4(81.45%)	6.5/7.9(82.21%)	1.3/18.3(7.11%)	1.8/19.0(9.26%)	2.8/38.0(7.38%)	2.4/58.7(4.13%)	2.3/108.2(2.13%)	1.5/121.8(1.27%)
1.5KWH	40.5/49.2(82.33%)	6.2/7.5(82.63%)	1.3/17.7(7.23%)	1.5/18.1(8.17%)	2.4/37.7(6.32%)	2.4/57.9(4.21%)	2.4/107.5(2.26%)	2.1/121.1(1.73%)

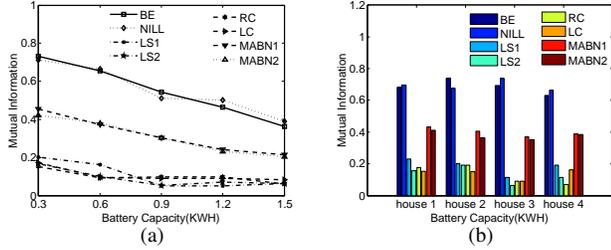


Fig. 6. (a) The changes of the mutual information between $s(t)$ and $d(t)$ with respect to battery capacity. (b) The mutual information between $s(t)$ and $d(t)$ for different houses when the battery capacity is set to 0.6KWH.

VIII. CONCLUSION AND FUTURE WORK

Recent studies have reveal the privacy concerns on smart metering data. Though current BLH solutions have been demonstrated to be useful for certain kinds of attacks, the information leakage risk of these solutions is much unknown. There lacks a definition of privacy for the BLH solutions they cannot guarantee differential privacy. We then propose novel randomized BLH algorithms which can indeed achieve certain differential privacy bound while not validating the battery constraint. Results from extensive simulations demonstrate the efficiency and effectiveness of the proposed method over existing BLH methods.

There are several interesting challenges left for further investigation. First, in this paper we formulate the BLH problem as an online optimization problem. If we assume the load of a day could be predicted, what performance, from the perspective of differential privacy, could an offline scheme achieve? This will give a general bound on the BLH algorithms. Second, right now we have not considered the economic cost/benefit of noise generation. On one hand, the charge/discharge of the battery will decrease the lifetime of the battery. On the other hand, considering the real time pricing used in smart grid, where the electricity prices of different hours are different, the battery could charge in low-price hours and discharge in high-price hours to gain economic benefit. To acquire economical benefit while satisfying certain privacy requirement is still a challenging issue for the BLH solutions.

REFERENCES

[1] "Understanding the potential of smart grid data analytics," GTM Research Report, 2012.

[2] G. Hart, "Nonintrusive appliance load monitoring," *Proceedings of the IEEE*, 18(12):1870-1891, 1992.

[3] C. Laughman, K. Lee, R. Cox, S. Shaw, S. Leeb, et al., "Power signature analysis," *Power and Energy Magazine*, 1(2):56-63, 2003.

[4] M. Marceau and R. Zmeureanu, "Nonintrusive load disaggregation computer program to estimate the energy consumption of major end uses in residential buildings," *Energy Conversion and Management*, 41(13):1389-1403, 2000.

[5] S. Inagaki, T. Egami, T. Suzuki, H. Nakamura, and K. Ito, "Nonintrusive appliance load monitoring based on integer programming," *Electrical Engineering in Japan*, 173(2):18-25, 2011.

[6] S. Gupta, M. S. Reynolds, and S. N. Patel, "Electrisense: Single-point sensing using emi for electrical event detection and classification in the home," in *Proc. of UbiComp 2010*, 2010.

[7] I. Rouf, H. Mustafa, et al., "Neighborhood watch: Security and privacy analysis of automatic meter reading systems," in *Proc. ACM CCS*, 2012.

[8] K. Fehrenbacher, "Smart meter worm could spread like a virus," available at <http://earth2tech.com/2009/07/31/smart-meter-worm-could-spread-like-a-virus/>, 2009.

[9] U.S. Dept. of Energy, "Smart grid privacy workshop summary report," in *Proc. of Smart Grid Privacy Workshop*, 2012.

[10] G. P. Zachary, "Saving smart meters from a backlash," *IEEE Spectrum*, 2011.

[11] U.S. Dept. of Energy, "Data access and privacy issues related to smart grid technologies," Report, 2010.

[12] G. Kalogridis, C. Efthymous, S. Denic, T. Lewis, and R. Cepeda, "Privacy for smart meters: Towards undetectable appliance load signatures," in *Proc. of SmartGridComm*, 2010.

[13] S. McLaughlin, P. McDaniel, and W. Aiello, "Protecting consumer privacy from electric load monitoring," in *Proc. of ACM CCS*, 2011.

[14] W. Yang, N. Li, Y. Qi, W. Qardaji, S. McLaughlin, et al., "Minimizing private data disclosures in the smart grid," in *Proc. of ACM CCS*, 2012.

[15] C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and M. Naor, "Our data, ourselves: Privacy via distributed noise generation," in *Proc. of EUROCRYPT*, 2006.

[16] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Proc. of TCC*, 2006.

[17] C. Dwork, "Differential privacy," in *Proc. of ICALP*, 2006.

[18] S. Bubeck and N. Cesa-Bianchi, "Regret analysis of stochastic and nonstochastic multi-armed bandit problems," *Foundations and Trends in Machine Learning*, 5(2):1-122, 2012.

[19] J. Gittins, K. Glazebrook, and R. Weber, "Multi-Armed Bandit Allocation Indices," Wiley, 2011.

[20] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire, "The non-stochastic multiarmed bandit problem," *SIAM Journal of Computing*, 32(1):48-77, 2002.

[21] K. Nissim, S. Raskhodnikova, and A. Smith, "Smooth sensitivity and sampling in private data analysis," in *Proc. of ACM STOC*, 2007.

[22] A. Ghosh, T. Roughgarden, and M. Sundararajan, "Universally utility-maximizing privacy mechanisms," in *Proc. of ACM STOC*, 2009.

[23] A. Blum, K. Ligett, and A. Roth, "learning theory approach to non-interactive database privacy," in *Proc. of ACM STOC*, 2008.

[24] F. McSherry and K. Talwar, "Mechanism design via differential privacy," in *Proc. of IEEE FOCS'07*, 2007.

[25] P. Auer, N. Cesa-Bianchi, and P. Fischer, "Finite-time analysis of the multiarmed bandit problem," *Machine Learning*, 47(2-3):235-256, 2002.

[26] J. Z. Kolter and M. J. Johnson, "REDD: A Public Data Set for Energy Disaggregation Research," in *Proc. of ACM SustKDD*, 2011.