PRIVATUS: Wallet-Friendly Privacy Protection for Smart Meters

Jinkyu Koo, Xiaojun Lin, and Saurabh Bagchi

Purdue University, West Lafayette, IN 47907, USA
{kooj,linx,sbagchi}@purdue.edu

Abstract. In smart power grids, a smart meter placed at a consumer-end point reports fine-grained usage information to utility providers. Based on this information, the providers can perform demand prediction and set on-demand pricing. However, this also threatens user privacy, since users' specific activity or behavior patterns can be deduced from the finely granular meter readings. To resolve this issue, we design PRIVATUS, a privacy-protection mechanism that uses a rechargeable battery. In PRIVATUS, the meter reading reported to the utility is probabilistically independent of the actual usage at any given time instant. PRIVATUS also considerably reduces the correlation between the meter readings and the actual usage pattern over time windows. Further, using stochastic dynamic programming, PRIVATUS charges/discharges the battery in the optimal way to maximize savings in the energy cost, given prior knowledge of time periods for the various price zones.

Keywords: smart grid, smart meter, privacy, cost saving, dynamic programming, battery.

1 Introduction

A smart grid is a type of the electrical grid in which electricity delivery systems are equipped with computer-based remote control and automation, which can revolutionize the way that energy is generated and consumed. A key component of the smart grid is the use of the smart meters, which measure energy usage at a fine granularity (e.g., once in a few minutes). However, by gathering hundreds of data points even in a day via the smart meter, the utility companies and third parties may learn a lot about our daily lives, *e.g.*, when we wake up, when we go out for work, and when we come back after work. In an industrial setting, this may be used to reveal details of the industrial process being used, or when a new process is adopted (which is achievable if the new machinery has electricity usage very distinct from prior machinery). Because of this privacy concern, there have been lawsuits to stop the installation of smart meters [1]. As a result, such privacy concerns have delayed the wide and quick deployment of smart grids.

There are a number of possible threat models for the above privacy risks. Given that we do need to report our energy usage profile to the utility company, the most important threat is that the metering data may be unwittingly disclosed from the utility company to third-party vendors. This problem is well illustrated

S. Foresti, M. Yung, and F. Martinelli (Eds.): ESORICS 2012, LNCS 7459, pp. 343–360, 2012.

in an article in MSNBC RedTape [2]. This article introduces a possible scenario with the smart grid that you get a discount with your power company at the cost that your auto insurance company may learn when you are home from the utility company. Additionally, due to possibly poor implementation of cryptography mechanisms, an eavesdropper on the wireless channel between the consumer's premises and the wireless network collection point may also determine the usage.

To resolve this issue, *the first objective of this paper is to make it difficult for an adversary to infer, based on the energy usage profile reported to the utilities, what is going on inside the house.* We achieve this objective by putting a rechargeable battery at the user-end point (*e.g.*, a home). The rechargeable battery acts like a buffer between the power grid and the end user in such a way that the actual energy usage pattern looks different from the energy usage pattern reported to the utility.

Additionally, the rechargeable battery provides us with an opportunity to lower the energy bill, by exploiting the time-of-use (TOU) pricing feature of smart grid, whereby electricity price varies according to pre-established time zones during a day. Basically, the cost-saving will be accomplished by charging the battery when the price is low and using the saved energy from the battery when the price is high. However, the two goals of privacy protection and cost saving are not always compatible with each other. *Our goal is therefore to achieve as much energy cost savings as possible, subject to privacy protection constraints.* To the best of our knowledge, we are the first to propose a mechanism that considers both privacy protection and cost saving simultaneously.

In this paper, we present Privatus, our solution that *guarantees* that instantaneous values of the actual usage and the energy draw visible outside the home are independent in an information-theoretic sense. Further, the patterns of both of these variables are also designed to look dissimilar. We set up a dynamic programming problem that minimizes the energy cost while preserving the privacy guarantee mentioned above.

We evaluate our solution in terms of both the privacy information leakage and the cost saving, and compare it to a previous solution that masked high frequency variation in energy usage [3]. In our simulation environment, Privatus can preserve at least 83% of the uncertainty of the actual usage sequences. In addition, Privatus can achieve 72% of the theoretically-possible maximum cost saving with a 6.43kWh battery. This translates to a saving of $16 per month in a typical residential pricing plan [4], assuming the average daily usage of 30kWh. We believe that this saving could provide an extra and significant incentive for users to invest in our solution in addition to privacy protection. The interested reader is referred to Appendix A for further discussion about this incentive.

2 Related Work

There has been extensive research about privacy protection in the area of database systems, where the goal is to provide statistical information (such as sum, average, or maximum) without revealing sensitive information about

individuals. The common approach to achieve this goal is data perturbation [5, 6]. However, none of methods in this area is directly applicable to hide the privacy information in the meter readings from the smart meters, because the utility companies do have to know precise meter reading records for billing purpose.

Recently, many studies raised the privacy concern in the smart grid both from a technical perspective and from a legal perspective [7–9]. However, only a few works have been proposed so far on the design of technical solutions to handle the privacy issue in the smart grid. Rial *et. al.* [10] proposed a privacy-preserving metering system, where the energy bill for a specific period is calculated by the user and then sent to the utility company. This system allows the user not to report the fine-granularity meter readings. However, it limits the power grid operator's capability such as demand prediction. Kalogridis *et. al.* [3] used a rechargeable battery to perform low-pass filtering over the load profile. Their algorithm forces the battery to charge (or discharge) a certain amount of energy if possible, when the required load is smaller (or larger) than the previously metered load. Thus, the high-frequency variation on energy usage profile is not visible to the smart meter. This approach can help eliminate load signatures that indicate which appliance is being used. However, the low-frequency components of a load profile are still revealed without any protection. Further, the proposed solution did not consider the cost-saving opportunity of using the rechargeable battery. Another work using the rechargeable battery is proposed by Varodayan *et. al.* [11]. They considered a simple binary-state battery model, where the battery is probabilistically charged by drawing the energy from the grid and discharged to feed the appliances. However, in their model, the charging and discharging processes at a given time instant are not independent of each other. This leads to a high level of information leakage (at least 0.5 bit for one-bit information). The authors also failed to consider the possible saving in the electricity cost by using the rechargeable battery.

Our work also adopts the rechargeable battery to protect the user privacy, but we design a mechanism by which the charging and discharging processes are guaranteed to be independent of each other at a given time instant. Further, our design also considers to reduce the correlation between the sequences of the charging and discharging processes over multiple time instants (instead of just for a single time instant). This makes it difficult for the adversary to make a meaningful guess on the user behavior by observing the sequence of meter readings. In addition, our design ensures that the way of charging the battery is optimal in the sense that we can maximize the average saving in the energy cost. This is achieved by controlling the charging process by dynamic programming.

3 System Model

Suppose that the smart meter measures the energy consumption once in every fixed interval (*e.g.*, 15 minutes), which we call the *measurement interval*. We denote by $X(n)$ the amount of energy consumed in the n-th measurement interval. We call $X(n)$ the *use process*. Denote the amount of energy that we draw

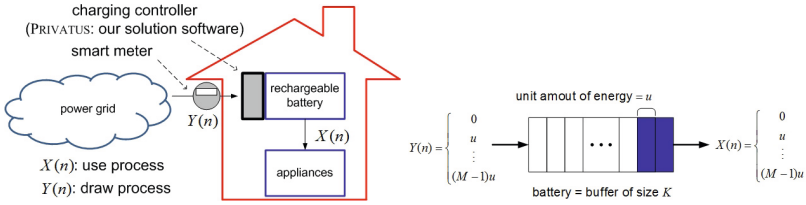

(a) Abstract model to draw and use energy.

(b) The battery as a buffer.

Fig. 1. System model

from the power grid in the n-th measurement interval by $Y(n)$, which we call the *draw process*. The smart meter measures $Y(n)$ and reports it to the utility. Without any special technique, *i.e.*, as it happens today, the draw process $Y(n)$ is the same as the use process $X(n)$. What we want to achieve in this paper is to de-correlate $X(n)$ and $Y(n)$ so that even if an adversary can observe $Y(n)$, no information is leaked about the use process $X(n)$. Toward this end, we put a rechargeable battery at the user-end as shown in Figure 1(a). The rechargeable battery acts as a buffer between $X(n)$ and $Y(n)$: instead of directly feeding $X(n)$ by $Y(n)$, we charge the battery by $Y(n)$, and use the saved energy in the battery to supply $X(n)$. We will design an algorithm in the charging controller, which will choose the value of $Y(n)$ carefully to ensure that the battery always has the appropriate level of energy (*i.e.*, no shortage to feed $X(n)$ or no overflow), and that $X(n)$ looks independent of $Y(n)$.

We assume that the values of $X(n)$ and $Y(n)$ may take any of the M different levels $\{0, u, 2u, \ldots, (M-1)u\}$, where u represents a unit amount of energy. We denote by $B(n)$ the energy level remaining in the battery at the *end* of the n-th measurement interval. Assuming for simplicity that there is no energy loss when charging and discharging the battery (for extension to the case with energy loss, see Appendix A), the value of $B(n)$ can be expressed as

$$B(n) = B(0) + \sum_{m=1}^{n} D(m), \tag{1}$$

where $D(m) = Y(m) - X(m)$ and $B(0)$ is the initial energy level of the battery that is also a multiple of u. Note that $D(n)$ also takes its value as a multiple of u, which is over the range $[-(M-1)u, (M-1)u]$. We model the battery as a buffer of size K as illustrated in Figure 1(b), which implies that the battery capacity is Ku, *i.e.*, the range of $B(n)$ is $0 \le B(n) \le Ku$.

The probability distributions of $X(n)$ and $Y(n)$ are described by $p_X(i;n)$ and $p_Y(i;n)$, respectively, where $p_X(i;n) = P(X(n) = iu)$ and $p_Y(i;n) = P(Y(n) = iu)$. Define the distribution vectors of $X(n)$ and $Y(n)$ as $P_X(n)=[p_X(0;n),p_X(1;n),\ldots,p_X(M-1;n)]$ and $P_Y(n)=[p_Y(0;n),p_Y(1;n),\ldots,p_Y(M-1;n)]$, respectively. We assume that $P_X(n)$ is known to the user (*i.e.*, the home owner). We also assume that $X(n)$ is independent, but does not need to be identically distributed across the measurement interval index n. This means that for instance,

$X(5)$ is independent of $X(11)$, but $P_X(5)$ can be different from $P_X(11)$. As we will see later, $P_Y(n)$ is our control parameter.

We are interested in the case where the electricity price per unit amount of energy varies from time to time. More specifically, we first focus on the case where there exist two time zones within a day, one of which has a low rate R_L (dollars/u) and the other has a high rate R_H (dollars/u). The zone with a low rate is called the *low-price zone* and the other is called the *high-price zone*. For ease of exposition, we assume that the measurement intervals from $n = 1$ to $n = n_L$ fall into the low-price zone, and the measurement intervals from $n = n_L + 1$ to $n = n_H$ correspond to the high-price zone. We treat the initial point $n = 0$ as the beginning of a day and the end of the measurement interval of $n = n_H$ as the end of the day. In Appendix A, we will discuss how we can generalize the solution to handle the case with more than two price zones in a day, and the case when the low-price and high-price zones are interleaved.

Because of the page limit, this paper assumes that the total amount of energy usage per day is the same over days on average. Appendix A introduces a way to release this assumption and generalize our solution.

4 Solution Approach I: Basic Formulation

4.1 Mapping between $X(n)$ and $Y(n)$

In order to hide $X(n)$ from an external adversary (*i.e.*, an adversary outside the home), we make $Y(n)$ be independent of $X(n)$. This implies that observing $Y(n)$ gives no meaningful information about $X(n)$. This is achieved when we map $X(n)$ to $Y(n)$ in such a way that $p_Y(i;n) \equiv P(Y(n) = iu) = P(Y(n) = iu|X(n) = ju)$ for any possible i and j. Practically, we achieve this by probabilistically choosing the value of $Y(n)$ according to $P_Y(n)$, which is decided before the n-th measurement interval starts, *without considering what the value of* $X(n)$ *will be*.

However, selecting $Y(n)$ randomly without being aware of $X(n)$ may cause energy shortage or overflow in the battery. For example, when $B(n-1) = 0$ (*i.e.*, there is no energy remaining in the battery before the n-th measurement interval starts), if $Y(n)$ is chosen to be zero, we cannot feed any non-zero value of $X(n)$. This means that sometimes we cannot use the appliances when we want. Similarly, when $B(n-1) = Ku$ (*i.e.*, the battery is full), a non-zero value of $Y(n)$ does not make sense if $X(n) = 0$, since we cannot draw the energy from the power grid unless we throw it away.

To handle this issue, we put a restriction on $P_Y(n)$ when the energy left in the battery is smaller than $(M-1)u$ (near-empty) or larger than $(K-(M-1))u$ (near-full), which we call the *corner cases*. More specifically, when $B(n-1) = ju$ for $j < (M-1)$, we choose $P_Y(n)$ such that $p_Y(i;n) = 0$ for $i < (M-1)-j$. Similarly, when $B(n-1) = (K-j)u$ for $j < (M-1)$, we choose $P_Y(n)$ such that $p_Y(i;n) = 0$ for $i > j$. We refer the readers to [12] for more detailed explanation of what this restriction means. The rationale behind this restriction on $P_Y(n)$

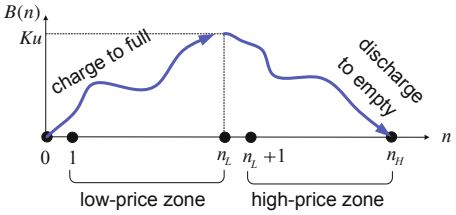

Fig. 2. Desired battery state profile

is that the battery must always have enough amount of energy to feed $X(n)$ even at the near-empty case, and that we never charge the battery more than its capacity whatever $X(n)$ is.

4.2 Strategy for Charging/Discharging the Battery

The only way to achieve cost saving by exploiting the time-of-use pricing policy is to charge the battery in the low-price zone and use the stored energy in the high-price zone. If we charge iu amount of energy in the low-price zone and use it in the high price zone, we can save $(R_H - R_L)i$ (dollars). For this reason, the maximum possible cost saving is $(R_H - R_L)K$ (dollars) per day, which is obtained when we charge the battery from empty to full in the low-price zone and discharge the battery to zero by feeding $X(n)$ in the high-price zone. Note that the maximum cost saving is proportional to the battery capacity Ku.

Therefore, our strategy to achieve the saving in the energy bill is to force the battery state to follow the trend shown in Figure 2. We achieve this by changing $P_Y(n)$ for every n, which is discussed in detail in the following subsection.

4.3 Basic Approach

We first define the distribution vector space $\mathcal{P}$ as follows.

$$\mathcal{P} = \left\{ [p_0, p_1, \ldots, p_{(M-1)}] : \sum_{i=0}^{M-1} p_i = 1, 0 \le p_i \le 1 \right\}, \tag{2}$$

where we limit the value of p_i to be a multiple of a constant c $(0 < c < 1)$, in order to make $\mathcal{P}$ be a finite set. For example, when $c = 0.1$ and $M = 4$, the distribution vector space $\mathcal{P}$ contains $[0.1, 0.2, 0.3, 0.4]$ and $[0.5, 0.5, 0, 0]$ as two of its elements. Then, $P_Y(n)$ is assigned one element in $\mathcal{P}$ in the n-th measurement interval. Recall that we force some elements of $P_Y(n)$ to be zero, depending on the battery level (Section 4.1). Therefore, the possible choice set in the n-th measurement interval is dependent on $B(n-1)$ and we denote it by $\mathcal{P}_{B(n-1)}$. Now, the key question for us is "*what would be the best choice for* $P_Y(n) \in \mathcal{P}_{B(n-1)}$ *for each* n *to maximize the cost saving?*" This question is answered by solving the following stochastic optimal control problems:

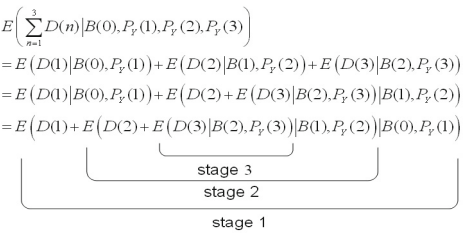

Fig. 3. An example to derive the dynamic programming framework

$$\max_{\substack{P_Y(n)\in\mathcal{P}_{B(n-1)} \\ 0<n\le n_L}} E\left(B(n_L)|B(0),P_Y(1),P_Y(2),\ldots,P_Y(n_L)\right) \tag{3}$$

in the low-price zone, and

$$\min_{\substack{P_Y(n)\in\mathcal{P}_{B(n-1)} \\ n_L<n\le n_H}} E\left(B(n_H)|B(n_L),P_Y(n_L+1),P_Y(n_L+2),\ldots,P_Y(n_H)\right) \tag{4}$$

in the high-price zone. Namely, we maximize (or minimize) the expected amount of the energy in the battery when each zone ends, given the battery level at the beginning of the zone and the distribution vectors $P_Y(1)$ through $P_Y(n_L)$ (or $P_Y(n_L+1)$ through $P_Y(n_H)$). We solve these optimization problems using dynamic programming [13].

To see how we use dynamic programming, let us first consider the following simple example in the low-price zone, where $n_L = 3$. Then, the optimization objective is to maximize $E\left(B(3)|B(0),P_Y(1),P_Y(2),P_Y(3)\right)$, which is equal to $B(0) + E\left(\sum_{n=1}^{3} D(n)|B(0),P_Y(1),P_Y(2),P_Y(3)\right)$, where $D(n) = Y(n) - X(n)$ as introduced earlier. Since $B(0)$ is given, we only need to focus on maximizing $E\left(\sum_{n=1}^{3} D(n)|B(0),P_Y(1),P_Y(2),P_Y(3)\right)$, which can be re-written as shown in Figure 3. Note in the figure that the calculations can be done recursively. Stage 2 calculations are based on stage 3, stage 1 only on stage 2. Thus, the optimal solution can be performed by maximizing the stage 3, stage 2, and stage 1 in this order. In this manner, we first compute the optimal value of $P_Y(3)$ given $B(2)$, then we compute the optimal value of $P_Y(2)$ given $B(1)$ until we reach and compute the optimal value of $P_Y(1)$. In the general case, $P_Y(n_L)$ is computed first and then other $P_Y(n)$'s are computed in a backward direction (time-wise) till $P_Y(1)$ is computed.

Namely, the optimal solution for (3) is obtained by a backward-directional computation procedure. In general, this procedure can be described by the following recursive equation, called the Bellman equation:

$$\begin{aligned} &J(n_L+1,B(n_L)) = 0, \\ &J(n,B(n-1)) = \max_{P_Y(n)\in\mathcal{P}_{B(n-1)}} E\left(D(n) + J(n+1,B(n))|B(n-1),P_Y(n)\right), \end{aligned} \tag{5}$$

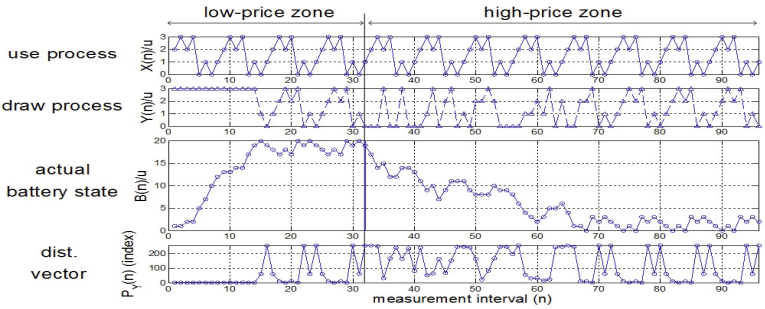

Fig. 4. Simulation results for the basic approach

for $n = n_L, (n_L - 1), \ldots, 1$. Solving (5) results in the optimal decision for $P_Y(n)$ when the value of $B(n-1)$ is given, in the sense that $P_Y(n)$ will maximize $E(B(n_L))$. Refer to [12] for further detail to solve (5). The optimal solution for (4) can also be obtained in a similar way.

In summary, what we have done is to calculate a decision table. Each entry in the decision table maps the given values of n and $B(n-1)$ to the optimal vector $P_Y(n)$ at the state. Note that the decision table can be pre-calculated before the run-time. During the run-time, we just look up the decision table for a given state, *i.e.*, n and $B(n-1)$, and probabilistically choose the value of $Y(n)$ via the distribution specified by the decision table entry. The size of this table can be large in practice if K and n_H are large. Thus, calculating the decision table can be computationally expensive. However, note that the table can be reused from one day to another till the distribution of the use process $X(n)$ changes significantly. Discussion about table complexity can be found at [12].

4.4 Simulation Study for the Basic Approach

We now present simulation results for our basic solution approach. By this simulation study, we will identify the issues with the basic approach, which will motivate us to improve our solution in Section 5.1 and propose PRIVATUS.

In the simulation, we choose $M = 4$, $K = 20$, and $c = 0.1$. We fix each measurement interval to be 15 minutes and thus we have 96 measurement intervals a day. Thus, the value of n_H becomes $n_H = 96$ and we set $n_L = 32$. In order to see more clearly what $Y(n)$ looks like compared to $X(n)$, we make $X(n)$ as a known repeated pattern, instead of generating it randomly (Figure 4).

A sample result of the simulation is shown in Figure 4, where "$P_Y(n)$ (index)" in the bottom graph means the index number of the element in $\mathcal{P}$ selected as $P_Y(n)$. We can see that at each measurement interval, the values of $X(n)$ and $Y(n)$ are mapped to each other in a random fashion. Further, the battery level indeed moves according to the trend that it is charged to the full level in the low-price zone and fully discharged in the high-price zone. However, we also observe that there exist similar patterns for the sequences of $X(n)$ and $Y(n)$ for

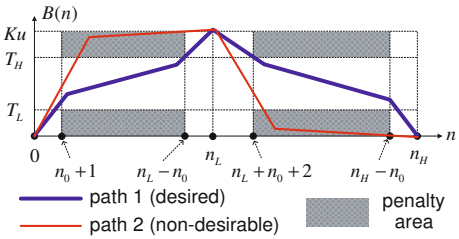

Fig. 5. Penalty areas

the measurement intervals of $16 \leq n \leq 32$ and $70 \leq n \leq 96$. More precisely, we see that the value of $X(n)$ highly likely reappears as the value of $Y(n+1)$ when the battery is at the corner cases. This is an undesirable behavior because if the adversary learns this characteristic, he or she may infer the original values of $X(n)$ with high accuracy by observing the values of $Y(n+1)$. Through this, we realize that our point-by-point de-correlation between $X(n)$ and $Y(n)$ leaves an obvious vulnerability in practice.

After more careful study, we find that this issue occurs because of two reasons: **(R1)** The first reason is that we charge/discharge the battery too fast. In the low-price zone, the battery reaches the full state much earlier than the end of the zone. Once at the full state, the battery stays close to the near-full states, since there is no benefit to bring the energy level down to a lower one according to our optimization objective in (3). The near-constant energy level of the battery implies that whatever the value of $X(n)$ is, the draw process $Y(n)$ should somehow compensate for it. Since the value of $Y(n)$ is chosen before the value of $X(n)$, we see this compensation effect in $Y(n+1)$. Similar logic applies to the high-price zone; **(R2)** The second reason is that we have too much freedom when choosing $P_Y(n)$. As a result, the draw process can take a specific symbol with a very high probability to compensate the use process. For example, if $X(n) = 3u$ and the draw process needs to compensate it (due to the first reason), the basic approach will likely choose $P_Y(n+1) = [0, 0, 0, 1]$. This implies that we will charge with the current value of $3u$ with probability 1 at the $(n+1)$-th measurement interval. In other words, due to the high degree of freedom to choose $P_Y(n)$, $Y(n)$ is chosen to be very similar to $X(n-1)$ in the corner cases.

In the next section, we will propose PRIVATUS that suppresses these undesirable effects **(R1)** and **(R2)**.

5 Solution Approach II: Advanced Formulation

5.1 Advanced Approach: PRIVATUS

In order to fix **(R1)**, we introduce penalty areas for when the battery level gets too close to empty or too close to full as shown in Figure 5. The penalty areas correspond to the battery states higher than the upper threshold T_H or

lower than the lower threshold T_L. In each zone (low-priced or high-priced), the penalty areas begin after n_0 measurement intervals, and end n_0 measurement intervals before the end of the zone. We modify our optimization objective in such a way that we incur some penalty, whenever the battery state $B(n)$ falls into the penalty areas. Hence, the optimal decision for $P_Y(n)$ would be changed to the one that still charges or discharges the battery according to the trend in Figure 2, but does not hit the penalty areas in the middle of the zones. In this sense, the modified optimization objective would result in "path 1"-like battery profile rather than "path 2"-like one in Figure 5. The "path 2"-like battery profile is what we have seen in the basic approach.

We consider the *effective battery state* $B_e(n)$ in the optimization objective function, instead of the actual battery state $B(n)$. The effective battery state $B_e(n)$ is designed to increase as the actual battery state $B(n)$ increases in the low-price zone (or $B(n)$ decreases in the high-price zone). However, every time $B(n)$ goes into a penalty area, $B_e(n)$ is deducted by some penalty amount. Denote by $[x]^+$ the projection of x to non-negative values, *i.e.*, $[x]^+ = x$ if $x > 0$, and $[x]^+ = 0$ if $x \leq 0$. Then, the effective battery state $B_e(n)$ in the low-price zone is defined as $B_e(n) = B_e(0) + \sum_{m=1}^{n} D_e(m)$. Here, $B_e(0) = \alpha B(0)$ and $D_e(m)$ is given as, if $m \leq n_0$ or $m > n_L - n_0$ (*i.e.*, in near-beginning or near-end of the low-price zone), $D_e(m) = \alpha D(m)$, and if $m > n_0$ and $m \leq n_L - n_0$,

$$D_e(m) = \alpha D(m) - \beta \left([B(m) - T_H]^+ + [T_L - B(m)]^+ \right), \tag{6}$$

where α and β are positive integers, $T_L = (M-1)u$, and $T_H = (K-(M-1))u$. In the high-price zone, we define $B_e(n)$ as $B_e(n) = B_e(n_L) + \sum_{m=n_L+1}^{n} D_e(m)$, where $B_e(n_L) = \alpha(Ku - B(n_L))$, and further, if $m \leq n_L + n_0$ or $m > n_H - n_0$, $D_e(m) = -\alpha D(m)$, and if $m > n_L + n_0$ and $m \leq n_H - n_0$,

$$D_e(m) = -\alpha D(m) - \beta \left([B(m) - T_H]^+ + [T_L - B(m)]^+ \right). \tag{7}$$

Note that if we ignore the second terms in (6) and (7), we simply have $B_e(n) = \alpha B(n)$ in the low-price zone, and $B_e(n) = \alpha(Ku - B(n))$ in the high-price zone. That is, $B_e(n)$ increases from zero to the maximum αKu in both zones as $B(n)$ moves like in Figure 2. Thus, our optimization objective for achieving the maximal cost saving is to maximize $E(B_e(n_L))$ in the low-price zone and $E(B_e(n_H))$ in the high-price zone, given initial conditions. On the other hands, the terms leading by β in (6) and (7) take into account the penalty. Whenever $D(n)$ causes $B(n)$ to fall into a penalty area, we subtract $\beta[B(n) - T_H]^+$ or $\beta[T_L - B(n)]^+$ from $B_e(n)$. Hence, we will expect that in the optimal decision for $P_Y(n)$, $B(n)$ would avoid hitting the penalty area, or $B(n)$ would attempt to get out of a penalty area if $B(n-1)$ was already in the penalty area. The relative magnitudes of α and β determines how sensitive we are to the penalty. If β is very large compared to α, $B(n)$ may not even go close to the penalty area to avoid any chance of incurring a high penalty score. Refer to [12] to see more detail about the choices for α and β.

On the other hand, to address **(R2)**, we adopt two strategies. *First*, we put the restriction on $\mathcal{P}_{B(n-1)}$ that it only contains the vectors $v \in \mathcal{P}$ such that

$\|v - V_k\| < T_k$. Here, T_k is a threshold at $B(n-1) = ku$, and V_k is the distribution vector of $Y(n)$ for which the possible values of $Y(n)$ at $B(n-1) = ku$ are selected equi-probably. For instance, when $M = 4$ and $K = 10$, we have $V_5 = [0.25, 0.25, 0.25, 0.25]$ when $B(n-1) = 5u$, and $V_1 = [0.5, 0.5, 0, 0]$ when $B(n-1) = u$. With this strategy, we are forcing the different elements of $P_Y(n)$ to be more or less equal, thus eliminating the possibility that $Y(n)$ is chosen deterministically (or with a high probability). By controlling the threshold T_k, we can control how close to equal probability we want. If T_k is low, then the choices are close to equally probable, but we also lose controllability in forcing $B(n)$ to the desired state according to the trend in Figure 2.

Second, we add one more restriction on $P_Y(n)$ in non-corner cases (*i.e.*, battery neither empty nor full) such that it does not differ significantly from $P_Y(n-1)$. If the two differ significantly, then $Y(n)$ may try compensating for the use value in the previous measurement interval and will hence track $X(n-1)$. Therefore, our strategy is that $\|P_Y(n) - P_Y(n-1)\| < T_D$, where T_D is called the distance threshold. We enforce this restriction to be applied only when the actual battery state stays in non-corner cases for two consecutive measurement intervals, *i.e.*, $T_L \le B(n-2) \le T_H$ and $T_L \le B(n-1) \le T_H$. Our intention behind this is to quickly get out of the corner cases (which hits the penalty areas). In the extreme case, with this strategy, $P_Y(n-1) = P_Y(n)$ implying that $Y(n)$ is independent of $X(n-1)$.

Reflecting all the changes, the optimal choice for $P_Y(n)$ in the low-price zone is obtained by solving the following Bellman equation.

$$\begin{aligned} &J(S(n_L + 1)) = 0, \\ &J(S(n)) = \max_{P_Y(n) \in \mathcal{P}^*_{B(n-1)}} E\left(D_e(n) + J(S(n+1)) | S(n)\right), \end{aligned} \tag{8}$$

for $n = n_L, (n_L - 1), \ldots, 1$. Here, $S(n)$ represents the state vector defined as $S(n) = [n, B(n-1), B_e(n-1), P_Y(n-1)]$. $\mathcal{P}^*_{B(n-1)}$ is defined as a subset of $\mathcal{P}$ whose element v is such that the two restrictions described above are satisfied, *i.e.*, $v \in \mathcal{P}_{B(n-1)}$, and if $T_L \le B(n-2) \le T_H$ and $T_L \le B(n-1) \le T_H$, $\|v - P_Y(n-1)\| < T_D$. The optimal choice for $P_Y(n)$ in the high-price zone can also be decided in a similar way.

5.2 Simulation Study for PRIVATUS

Now, we conduct a simulation test for PRIVATUS. In order to see the difference from the basic approach, we use the same simulation environment as in Section 4.4. We choose $T_k = 0.3$ for $k = 3, 4, \ldots, 17$; $T_k = 0.25$ for $k = 2, 18$; $T_k = 0.2$ for $k = 1, 19$; $T_k = 0.1$ for $k = 0, 20$. With these threshold values, $\mathcal{P}_k$ only contains $[0, 0, 0.4, 0.6]$, $[0, 0, 0.5, 0.5]$, and $[0, 0, 0.6, 0.4]$ for $k = 1, 19$, for instance. For the remaining parameters, we set $\alpha = 2$, $\beta = 1$, $n_0 = 3$, and $T_D = 0.2$.

Figure 6 shows a sample result for the simulation, where the solid red lines in the "$B(n)/u$" graph indicate the energy levels corresponding to the penalty area thresholds T_H and T_L. First, we can see that $B(n)$ follows the trend in

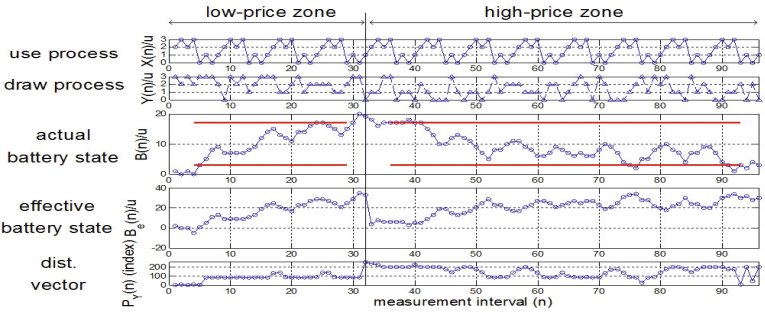

Fig. 6. Simulation results for PRIVATUS

Figure 2, and it seldom hits the penalty area as we desired. Although $B(n)$ enters the penalty area at around $n = 39, 76, 92$, we can also see that $B(n)$ tries to get out of penalty area quickly. As a result, the battery neither goes to the full-state too quickly in the low-price zone, nor goes to the empty-state too quickly in the high-price zone. Second, in the "$P_Y(n)$(index)" graph, we observe that for many times, the decision for $P_Y(n)$ remains the same, or the speed of changing a decision becomes much slower (compared to the result in Figure 4). By these two fixes, we see that the correlation between the use process and the draw process is significantly reduced. We can no longer find similar patterns between the two. The point-by-point comparison of $X(n)$ and $Y(n)$ still gives no meaningful clue from $Y(n)$ to $X(n)$, as this is by design that is maintained in the basic approach and PRIVATUS. Of course, this might be seen as a subjective interpretation of the result. Thus, in the experiment section, we will consider a metric to quantitatively measure how well we are protecting the privacy and re-visit these results.

6 Experiment

6.1 Metrics and Simulation Parameters

First, we define the metric of information leakage from the use process to the draw process as follows: for a positive integer m,

$$L^s_{(n,m)} = I(\bar{X}_{(n,m)}; \bar{Y}^s_{(n,m)})/H(\bar{X}_{(n,m)}), \tag{9}$$

where $\bar{X}_{(n,m)}=[X(n-m+1),X(n-m),...,X(n)]$, and $\bar{Y}^s_{(n,m)}=[Y(n-m+1+s),Y(n-m+s),...,Y(n+s)]$, and s is a non-negative integer called the timeshift offset. Here, $H(\mathcal{X})$ denotes the *uncertainty* of $\mathcal{X}$, and $I(\mathcal{X};\mathcal{Y})$ is the mutual information between $\mathcal{X}$ and $\mathcal{Y}$. Namely, $H(\mathcal{X}) = -\sum_i P(\mathcal{X}=i)\log P(\mathcal{X}=i)$ and

$$I(\mathcal{X};\mathcal{Y})=\sum_i \sum_j P(\mathcal{X}=i,\mathcal{Y}=j)\log \frac{P(\mathcal{X}=i,\mathcal{Y}=j)}{P(\mathcal{X}=i)P(\mathcal{Y}=j)} \tag{10}$$

Note that $\bar{X}_{(n,m)}$ and $\bar{Y}^s_{(n,m)}$ represent sequences of length m in the use process and the draw process, respectively, with the draw process being time delayed by s

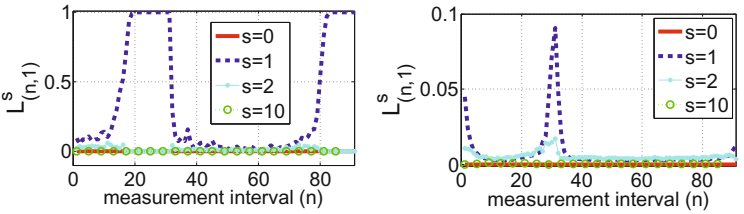

(a) Basic approach. (b) PRIVATUS ($\alpha = 2$; $\beta = 1$).

Fig. 7. Information leakage when $K = 20$ and $m = 1$

measurement intervals. Since $I(\bar{X}_{(n,m)}; \bar{Y}^s_{(n,m)}) = H(\bar{X}_{(n,m)}) - H(\bar{X}_{(n,m)}|\bar{Y}^s_{(n,m)})$, the metric $L^s_{(n,m)}$ can be interpreted as a measure of the *uncertainty reduction* in $\bar{X}_{(n,m)}$ by observing $\bar{Y}^s_{(n,m)}$, normalized to the uncertainty of $\bar{X}_{(n,m)}$. Thus, by this metric, we can quantify *how uncertain the adversary is when he attempts to guess the sequence* $\bar{X}_{(n,m)}$ *of the use process, based on the observed sequence* $\bar{Y}^s_{(n,m)}$ *of the draw process.* For example, the adversary knows that $\bar{X}_{(n,m)}$ is surely the same as $\bar{Y}^s_{(n,m)}$, when $L^s_{(n,m)} = 1$. In contrast, $L^s_{(n,m)} = 0$ means that $\bar{Y}^s_{(n,m)}$ gives no clue about $\bar{X}_{(n,m)}$ at all.

Second, given that the battery capacity is Ku, we define the metric for the cost saving for a day as

$$S_{(r,K)} = E\left(-\sum_{m=1}^{n_L} rR_H D(m) - \sum_{m=n_L+1}^{n_H} R_H D(m)\right), \tag{11}$$

where r denotes the ratio of R_L to R_H. The term $S_{(r,K)}$ is *the expected difference between the original cost for what the user actually consumes* ($\sum_{m=1}^{n_L} rR_H X(m) + \sum_{m=n_L+1}^{n_H} R_H X(m)$)*, and the money that a user pays to the utility company* ($\sum_{m=1}^{n_L} rR_H Y(m) + \sum_{m=n_L+1}^{n_H} R_H Y(m)$). A positive value of $S_{(r,K)}$ means that we achieve cost saving. If $S_{(r,K)}$ is negative, it means that we have to pay more compared to the baseline no-privacy-protection scheme.

To be consistent with the previous simulations (in Figures 4 and 6), we use the same parameters as before (*i.e.*, $M = 4$; $K = 20$; $n_L = 32$; $n_H = 96$; $\alpha = 2$; $\beta = 1$; $n_0 = 3$; $c = 0.1$) throughout the whole experiments, unless otherwise stated. However, we randomly generate $X(n)$ through $P_X(n) = [0.5, 0.2, 0.2, 0.1]$ in the low-price zone and $P_X(n) = [0.1, 0.3, 0.4, 0.2]$ in the high-price zone. This setting results in about $138u$ for the expected daily usage $E(\sum_{n=1}^{n_H} X(n))$. To get the results, we run 100,000 days in such a way that the remaining energy in the battery at the end of a day becomes the initial energy level of the battery in the next day.

6.2 Information Leakage and Cost Saving

General Performance Trend: Figure 7 shows the general performance trend of our solution approaches (for $m = 1$). We can see that when $s = 0$, $X(n)$

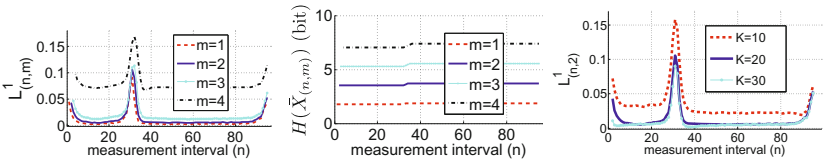

(a) Information leakage according to m ($s = 1$; $K = 20$).

(b) Uncertainty of $\bar{X}_{(n,m)}$ ($K = 20$).

(c) Information leakage according to K.

Fig. 8. Effects of sequence length m and capacity K in PRIVATUS ($\alpha = 2$; $\beta = 1$)

and $Y(n)$ are indeed independent in both the basic approach and PRIVATUS. We can also see that information leakage is the highest when $s = 1$, *i.e.*, $X(n)$ and $Y(n + 1)$ has the highest dependency in our solution approaches. This is due to our solution's inherent nature that $Y(n)$ is chosen to change the current battery state resulting from $X(n - 1)$ and the previous battery state. Figure 7(a) confirms again that in the basic approach, this issue can be quite significant because $Y(n)$ perfectly compensates $X(n-1)$ and reveals all information about $X(n - 1)$ (*i.e.*, $L^1_{(n,1)} = 1$) when the battery is in the corner cases. However, we see in Figure 7(b) that this compensation effect is greatly reduced. That is, in PRIVATUS, $Y(n)$ results in mostly near-zero uncertainty reduction about $X(n - 1)$. In even the worst case (for some measurement intervals, with delay of 1 measurement interval), the uncertainty reduction is less than 10%. We see that the worst-case information leakage in the advanced approach occurs around the price zone boundaries. We suspect that this is because around the price zone boundaries, there is no penalty defined and thus the battery state has a relatively higher chance to remain costant, which again makes it more likely that $Y(n)$ tries to compensate for $X(n - 1)$. On the other hand, we can see from the case when $s = 10$ that, with higher delays (*i.e.*, larger values of s), the sequences of the use process and the draw process become independent.

Effect of Sequence Length: In Figure 8(a), we see that in PRIVATUS, the information leakage increases as the sequence length m increases. This seems to imply that the adversary gains more information when he observes longer sequences. However, note from Figure 8(b) that the uncertainty of the use-process sequence $H(\bar{X}_{(n,m)})$ also grows as m increases. In Figure 8(b), x-bit uncertainty can be understood in such a way that approximately the use-process sequence has 2^x possible realizations with equal probability $1/2^x$. Since $M = 4$, the uncertainty of the use-process sequence becomes larger by a factor close to $\log_2 4$ (more precisely, $\log_2 2^{1.7}$ in our simulation setting) as m increases by 1. Thus, the minor increment in percentage-wise uncertainty reduction does not make it easier for the adversary to make guesses about the use-process sequence. For example, when $m = 3$ and $n = 32$, the uncertainty of the use-process sequence is 5.3 bits and uncertainty reduction is 11%. This implies that the remaining uncertainty of the use-process sequence after observing the draw-process sequence is $5.3(1-0.11) = 4.72$ bits, *i.e.*, the adversary faces the uncertainty to pick one out

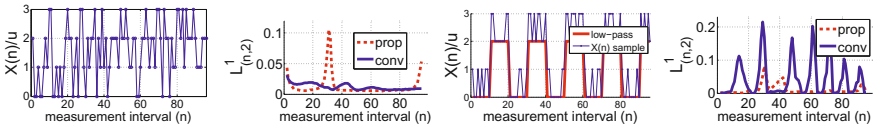

(a) $X(n)$ w/o a significant low-pass component. (b) Information leakage for (a). (c) $X(n)$ w/ a significant low-pass component. (d) Information leakage for (c).

Fig. 9. Information leakage comparison between PRIVATUS with $\alpha = 2$ and $\beta = 1$ (legend: 'prop') and an existing scheme [3] (legend: 'conv'), when $K = 20$ and $m = 2$. The higher is $L^1_{(n,2)}$, the worse is the information leakage.

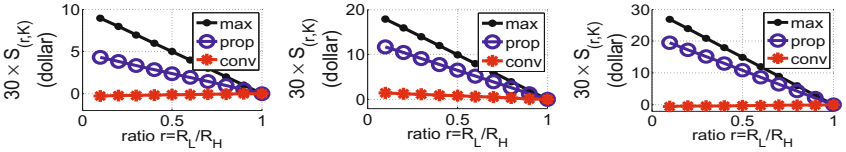

(a) $K = 10$ (2.15kWh). (b) $K = 20$ (4.3kWh). (c) $K = 30$ (6.43kWh).

Fig. 10. Cost saving comparison between PRIVATUS and an existing scheme [3]. Here, we set $u = 0.2143$kWh and $R_H = \$0.033/u = \$0.155$/kWh. This results in the average daily usage (*i.e.*, $E(\sum_{n=1}^{n_H} X(n))$) equal to 30kWh.

of $2^{5.3(1-0.11)} = 26.3$ possible sequences, in order to make a guess about the use-process sequence. On the other hand, when $m = 4$ and $n = 32$, the uncertainty is 7.0 bits and the uncertainty reduction is 17%. This results in $2^{7.0(1-0.17)} = 56.1$ possible sequences as candidates for the use-process sequence. Therefore, we conclude that the adversary has no advantage in observing a longer sequence in the draw process.

Effect of Battery Capacity: Figure 8(c) shows how PRIVATUS acts when the battery capacity varies. We can infer from the figure that when the battery capacity is too small, information leakage may be significant. This can be explained again by the compensation effect of our solution. If the battery capacity is too small, there is not much room for the battery state to fluctuate between the two penalty area thresholds T_L and T_H (see Figure 6). This means that the battery state remains relatively constant, which makes the compensation effect prominent. On the other hand, once the battery capacity is above a threshold, further increasing the battery capacity leads to little benefit in terms of further reducing the information leakage.

Comparison to Prior Work: In Figures 9 and 10, we compare PRIVATUS ('prop' in the figures) with an existing scheme ('conv' in the figures) proposed by Kalogridis *et. al.* [3]. Kalogridis' scheme performs a simple low-pass filtering over the use process in a best-effort manner without considering the energy cost factor. Thus, it reduces the high frequency variations in the resulting draw

process. Kalogridis' scheme needs to estimate the value of $X(n)$ beforehand (refer to [3] for detail). We assume in the simulation that the estimation is perfect (*i.e.*, without errors). Figure 9(a) shows a sample realization of $X(n)$, obtained from $P_X(n)$ given in Section 6.1. Note that since $X(n)$ is randomly chosen among M possible values from $P_X(n)$, which is the same within each price-zone, there is not a significant low-frequency component in $X(n)$. In this case, we can see from Figure 9(b) that PRIVATUS performs slightly better than Kalogridis' to keep the privacy information, except at the price zone boundaries. However, if there is a significant low-pass component in $X(n)$, PRIVATUS will provide much better privacy protection than Kalogridis'. This is because Kalogridis' scheme still allows the low-pass component of load profile to be revealed. To see this, we generate $X(n)$ by adding a random value 0 or u to a rectangular pulse whose period is 20 measurement intervals, as shown in Figure 9(c). Comparison result in such a case is given in Figure 9(d). Indeed, PRIVATUS results in better lower information leakage than Kalogridis' when there exists a considerable low-frequency component in $X(n)$. Meanwhile, Figure 10 shows that from the cost saving point of view, PRIVATUS has a huge advantage against Kalogridis'. In all of the cases studied, Kalogridis' scheme does not achieve a significant cost saving. On the other hand, compared to the maximum possible cost saving, computed according to Section 4.2 ('max' in the figures), PRIVATUS achieves the saving of 48% of the maximum when $K = 10$, 66% of the maximum when $K = 20$, and 72% of the maximum when $K = 30$. Thus, PRIVATUS strikes a desirable balance between privacy and cost saving. Considering that the average electricity consumption for a U.S. residential customer was 30kWh per day [14], Figure 10(c) shows that a typical home can achieve about $16 saving for a month with a 6.43kWh battery, based on the following tariff example: $R_L = 0.04$/kWh and $R_H = 0.15$/kWh [4].

7 Conclusion and Future Work

In order to resolve the privacy issue in smart grid, we proposed PRIVATUS to de-correlate the meter reading information from user behavior. PRIVATUS uses a rechargeable battery to make the meter reading reported to the utilities look independent of the actual usage at any given measurement interval. The correlation between the meter readings and the actual usage pattern over multiple measurement intervals is also reduced by changing the probability distribution of charging the battery in each interval through careful design. PRIVATUS is also geared to the future of time-of-use pricing of electricity and it ensures that the battery is charged to achieve the maximal savings in the energy cost. We formulate the problem rigorously and use stochastic dynamic programming to devise our solution. The experiment results show that PRIVATUS is successfully able to hide the actual usage from what is drawn from the grid, and achieves considerable amount of saving in the energy cost, subject to the availability of a reasonable-sized battery. Compared to prior work, we achieve much better privacy when there is a conspicuous low-frequency component in load profile, and significantly higher cost savings.

Our future work will focus on generalizing PRIVATUS under more dynamic scenarios, *e.g.*, where the price zones are dynamically changed from one day to the next, or the price varies over time in a demand-driven and adaptive manner.

Acknowledgments. This work has been partially supported by the National Science Foundation through award CNS-0831999. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

1. Beckman, H.: Lawsuit filed to stop installaton of smart meters, http://napervillesun.suntimes.com/news/9723766-418/lawsuit-filed-to-stop-smart-meter-installation.html
2. Sullivan, B.: What will talking power meters say about you?, http://redtape.msnbc.msn.com
3. Kalogridis, G., Efthymiou, C., Denic, S.Z., Lewis, T.A., Cepeda, R.: Privacy for smart meters: Towards undetectable appliance load signatures. In: 2010 First IEEE International Conference on Smart Grid Communications (2010)
4. Tucson electric power: Residential time-of-use pricing plan, https://www.tep.com/doc/customer/rates/R-21F.pdf
5. Agrawal, D., Aggarwal, C.C.: On the design and quantification of privacy preserving data mining algorithms. In: PODS 2001, pp. 247–255. ACM, New York (2001)
6. Sweeney, L.: k-anonymity: a model for protecting privacy. Int. J. Uncertain. Fuzziness Knowl.-Based Syst. 10, 557–570 (2002)
7. Stallman, R.: Is digital inclusion a good thing? How can we make sure it is? Comm. Mag. 48, 112–118 (2010)
8. Khurana, H., Hadley, M., Lu, N., Frincke, D.A.: Smart-grid security issues. IEEE Security and Privacy 8, 81–85 (2010)
9. Quinn, E.L.: Privacy and the new energy infrastructure. In: SSRN (2009)
10. Rial, A., Danezis, G.: Privacy-preserving smart metering. In: Proceedings of the 10th Annual ACM Workshop on Privacy in the Electronic Society, WPES 2011. ACM (2011)
11. Varodayan, D.P., Khisti, A.: Smart meter privacy using a rechargeable battery: Minimizing the rate of information leakage. In: ICASSP (2011)
12. Koo, J., Lin, X., Bagchi, S.: Technical Report: Wallet-Friendly Privacy Protection for Smart Meters, https://engineering.purdue.edu/~linx/papers.html
13. Bertsekas, D.P., Shreve, S.E.: Stochastic Optimal Control: The Discrete-Time Case. Athena Scientific (2007)
14. Administration, U.E.I.: Average electricity consumption for a us residential customer, http://www.eia.gov/tools/faqs/faq.cfm?id=97&t=3
15. U.S. Department of Energy: Battery power for your residentialsolar electric system, http://www.nrel.gov/docs/fy02osti/31689.pdf

A Discussion

Battery Cost: In Section 6, we showed that a 6.43kWh battery can achieve $16 saving per month, assuming 30kWh use in a day. People may argue that this is the relatively small savings compared to the high battery cost. Indeed, initial costs for residential batteries range from $80 to $200 per kWh [15], and thus the battery cost of 6.43kWh may range from $514 to $1,280. However, note that people buy a hybrid car to save the fuel-cost and the environment, although it requires a considerable initial cost due to the battery. Even though the fuel saving of the hybrid cars does not completely offset its high cost, the fuel saving serves as a significant incentive for consumers (who may only be mildly environment-conscious) to buy hybrid cars. Similarly, in our case, the cost savings will encourage privacy-conscious customers to buy our solution. In addition, given a 6.43kWh battery and $16 saving per month, the battery cost may be balanced out by the saving in 2.6 to 6.6 years. We think that this is similar to the period to recover the additional cost of a hybrid car compared to a normal car.

Energy Loss in a Battery: By multiplying coefficients (< 1) by $X(n)$ and $Y(n)$ in (1), our model can be easily extended to include the energy loss in the battery that occurs when charging and discharging.

More Than Two Price Zones: Once we know the rates of energy usage and the boundaries of each price zone, we can calculate the desired pattern of battery charge and discharge—akin to that in Figure 2. Namely, what we need to do is to calculate to what level the battery can be charged or discharged in each zone. Then, the solution approach outlined earlier applies directly to the case with more than two price zones.

Interleaved Low-Price and High-Price Zones: This situation is equivalent to the case where there are multiple price zones, one group of which have a low price, and the other group have a high price. Thus, this case can be treated in the same way as the above.

The Amount of Energy Usage Per Day Varying Over Days: This paper focuses on hiding the energy consumption pattern within a day. Across days, the total usage per day can still be revealed to the adversary (by which the adversary may know whether you are home or not for a given day). The other part of PRIVATUS, which is not presented in this paper due to the page limit, handles this issue. At the high level, the solution is to flatten the energy use across days, by charging more in days with less usage and by using the saved energy in days with more usage. The solution does not affect the current randomization framework within each day; it only modifies the total use in each day.