Smart Meter Privacy: Exploiting the Potential of Household Energy Storage Units

Yanan Sun, Lutz Lampe, and Vincent W.S. Wong
Department of Electrical and Computer Engineering
The University of British Columbia, Vancouver, Canada
e-mail: {ynsun, lampe, vincentw}@ece.ubc.ca

***Abstract*—The Internet of Things (IoT) extends network connectivity and computing capability to physical devices. However, data from IoT devices may increase the risk of privacy violations. In this paper, we consider smart meters as a prominent early instance of the IoT, and we investigate their privacy protection solutions at customer premises. In particular, we design a load hiding approach that obscures household consumption with the help of energy storage units. For this purpose, we leverage the opportunistic use of existing household energy storage units to render load hiding less costly. We propose combining the use of electric vehicles (EVs) and heating, ventilating, and air conditioning (HVAC) systems to reduce or eliminate the reliance on local rechargeable batteries for load hiding. To this end, we formulate a Markov decision process to account for the stochastic nature of customer demand and use a Q-learning algorithm to adapt the control policies for the energy storage units. We also provide an idealized benchmark system by formulating a deterministic optimization problem and deriving its equivalent convex form. We evaluate the performance of our approach for different combinations of storage units and with different benchmark methods. Our results show that the opportunistic joint use of EV and HVAC units can reduce the need of dedicated large-capacity or fast-charging-cycle batteries for load hiding.**

***Index Terms*—Internet of Things, privacy, smart metering, Markov decision process, Q-learning, electric vehicle.**

I. Introduction

The continuous evolution of pervasive computation, communication and control is creating a new world populated by intelligent connectivity on physical devices. The emerging *Internet of Things (IoT)* provides technology-enabled solutions for physical assets, which covers a broad spectrum of use cases that are driven by the ability to connect, monitor, exchange information, and take autonomous actions [1]. An important and rapidly growing realization of IoT is the smart grid. By connecting different grid devices to a communication network for real-time monitoring and surveillance, the smart grid can extend computing intelligence into power system infrastructures, which creates an *Internet of Energy* [2].

A key element of the smart grid is the advanced metering infrastructure (AMI), which relies on smart meters for bi-directional power flow and communication capabilities. The mass rollout of smart meters provides significant benefits for managing energy supply and demand for power utilities and customers alike. However, this industrial IoT solution increases the risk of privacy violations. For example, by using non-intrusive load monitoring (NILM) data analysis techniques, the fine-grained smart meter readings can be disaggregated to reveal customer usage patterns, personal routines, and behavioral preferences [3]–[6]. In general, privacy concerns are amplified by the presence of IoT devices, which expand the reach of tracking, monitoring, and surveillance. In this paper, we focus on privacy-preserving mechanisms for user consumption data reported by smart meters.

Data protection techniques, such as data anonymization and data aggregation, can be applied to mitigate the potential privacy leakage of smart meters. Data anonymization removes any attribute information from the meter readings to obscure their relationship with the customers. It relies on an escrow company as an intermediary to pseudonymize meter measurements [7], [8]. However, those pseudonymized traces can still be associated with the households that produced them [9]. Data aggregation, on the other hand, aims to reduce the amount of sensitive information that can be leaked. The privacy-preserving aggregation relies on the homomorphic features of cryptographic computation [10]–[13]. These methods are limited by their required computation or communication complexity, and their performance depends on specific model design with respect to privacy protection [14]. Furthermore, aggregation is primarily designed to tackle the computation overhead and resource management issues caused by storing and analyzing the huge volumes of measurement data.

The above solutions are designed from the utility companies' perspective. In general, these solutions can be applied to other IoT applications that collect private information, e.g., e-healthcare data management. However, IoT privacy challenges go beyond these conventional methods in that different participants (e.g., utility companies and customers) in the IoT marketplace may have unaligned interests in collecting and using the data. Therefore, customer-oriented solutions are required to accommodate individual privacy preferences.

One approach to achieve customer-oriented privacy protection is through data perturbation. For example, smart meter measurements can be perturbed by injecting random noise or applying data compression techniques [15]–[17]. The trade-off between the measurement precision and perturbation noise is analyzed in [18]. However, tampering with smart meter readings before transmission to the utility company may

This work was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC).

reduce their relevance for billing purposes. An alternative approach applies data obfuscation to alter the actual energy consumption that is measured and reported by the smart meter. The customer energy profile can be distorted by integrating an alternative energy source such as a renewable energy generation unit [19]–[22] or with the help of energy storage units at the customer premises. In this paper, we pursue a privacy-preserving approach that builds on the use of household energy storage units, which we refer to as *household load hiding*.

A. Related Work

Prior load hiding schemes obfuscate the smart meter measurements by using either a local rechargeable battery or a user controllable load. The former and latter are commonly referred to as *battery-based load hiding* (BLH) and *load-based load hiding* (LLH).

BLH methods hide the energy consumption variations through controlling the battery charging and discharging process. The variances of household load profile can be offset to maintain a constant output load with the presence of a rechargeable battery [23]–[25]. The physical constraints of the battery, however, make it often impractical to flatten out the actual power variations. BLH schemes can also achieve hiding by randomizing the load profile to preserve a certain differential privacy [26], or using stochastic battery control strategy to minimize information leakage [27]. The trade-off between the smart meter privacy and electricity cost in the context of home energy management and demand response by using a rechargeable battery is addressed in [28]–[32]. Considering the battery purchase expenses, however, BLH methods can be costly to implement. Different from these BLH methods, this paper aims to exploit alternative energy storage that has already been in place in the household to reduce the cost incurred by batteries, and thus to render load hiding less costly in the implementation phase.

LLH methods achieve hiding by shifting the energy demand of user controllable loads. The proposed schemes either flatten the energy consumption or inject artificial power signatures to hide the load profile and prevent attack cases such as occupancy detection [33]–[36]. Considering the restrictions of using appliances that are interruptible and can store energy, the available options are limited for LLH schemes. Different from these LLH schemes, this paper aims to combine BLH and LLH by exploiting the potential use of household controllable loads with local dedicated batteries, and thus to render load hiding more practical in the implementation phase.

The idea of using assistive battery to achieve hiding was first proposed in [37], [38] through the integration of electric vehicles (EVs). A convex optimization problem was formulated to exploit the combined use of EVs with local dedicated rechargeable batteries to disguise the household load profile [37]. The use of EVs as assistive batteries to replace local dedicated large-capacity or fast-charging-cycle batteries for load hiding was discussed in [38], where the stochastic nature of both the EV charging process and household energy demand are captured by the Markov decision process (MDP). The use of cascaded rechargeable batteries to alleviate the privacy leakage from smart meter measurement was discussed in [39]. In this paper, we extend the approach from [37], [38] to the combined use of existing household energy storage units (e.g., controllable loads, EVs) to design more cost-effective privacy-preserving solutions for customers. Our proposed solution is different from [37]–[39] in that we aim to exploit the existing household energy storage units for load hiding. In particular, we consider their use cases as both assistive and alternative energy storage solution to a dedicated battery to achieve hiding by leveraging the combination of BLH and LLH. Therefore, our proposed solution can address smart meter privacy concerns and accommodate individual privacy preference completely at the customer premises in a cost-effective manner.

B. Contributions

Achieving our proposed combined load hiding is a challenging problem, mainly due to the limited predictability of household demand, physical charging and discharging constraints of energy storage units, and the fact that thermal appliances first and foremost need to satisfy the customer's comfort requirements. This is aggravated by the interdependencies of scheduling decisions, i.e., current scheduling decisions affect the availability of energy storage and demand in the future. Our main contributions can be summarized as follows:

- We formulate an MDP problem to capture the uncertainties in household demand and customer behavior. The MDP addresses the difficulties in how to schedule the energy storage units considering the charging and discharging constraints as well as the limited predictability of customer demand.
- We propose a model-free learning method that can adapt to the optimal control policies for scheduling the energy storage units, which addresses the difficulty added by the interdependencies of scheduling decisions. We also derive a benchmark system by formulating a deterministic optimization problem, which can be used to evaluate the effectiveness of our learning method.
- We evaluate our approach by simulation using different combinations of energy storage with a local dedicated battery, an EV, and a heating, ventilating, and air conditioning (HVAC) unit. Our results show that the latter can be effectively used as assistive batteries, trading off the level of privacy achieved through load hiding and the aggregate cost for energy consumption. By comparing with different benchmark methods, we also validate our idea that the combination of BLH and LLH as well as the opportunistic use of typical household energy storage units (e.g., EV, HVAC system) can greatly alleviate the need for local dedicated large-capacity or fast-charging-cycle batteries.

The rest of the paper is organized as follows. Section II introduces the system model. Section III presents an MDP formulation. A model-free learning method for solving the MDP problem is proposed in Section IV, together with a benchmark by formulating a deterministic optimization problem.

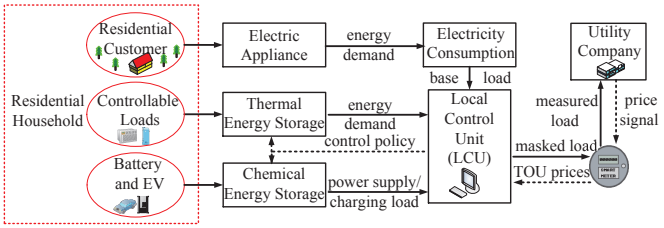

Fig. 1. A general household load hiding framework.

Simulation results are presented and discussed in Section V. Section VI concludes the paper.

II. System Model

Household load hiding is a customer-oriented approach to avoid possible privacy leakage from the consumption data reported by the smart meters. Fig. 1 shows the general framework for a single residential household that has both chemical and thermal energy storage units available which can be used to alter the energy drawn from the power grid. A smart meter is installed to measure the household power consumption. It interacts with the utility company to exchange electricity pricing information and household load measurements for billing purposes. A local control unit (LCU) is designed to optimize the operation schedules of the thermal and chemical energy storage units based on the time of use (TOU) pricing information from the smart meter and the current energy demand of the customer, while satisfying a set of customer preferences and requirements. The household base load and the usage of those energy storage units are combined when monitored and reported by the smart meter. Thus, the measured consumption data transmitted to the utility company has been distorted to disguise the original household load profile.

We discretize time according to the typical smart meter sampling intervals, denoted by Δt, to derive the proposed load hiding scheme. Each load scheduling period then consists of a set of discrete time slots, denoted by $\mathcal{T} = \{1, 2, \ldots, T\}$ and $T \leq \infty$. We now introduce the models for thermal and chemical energy storage units, which are characterized by different system dynamics and physical constraints.

A. Thermal Energy Storage Unit Model

Thermal energy storage units include thermostatically controllable loads such as HVAC systems and electric water tanks, which we refer to as thermal appliances in this paper and denote their set by Ψ. Thermal appliances can convert electricity to heat. Their storage and load-shifting behavior is affected by thermal dynamics and customer activities.

Let $T_{i,t}^{\text{in}}$ and $T_{i,t}^{\text{amb}}$ denote the temperature inside and outside the space of appliance $i \in \Psi$ at time $t \in \mathcal{T}$, respectively. Given its heat rated power $q_{i,t}$, we further introduce θ_i and $\sigma_{i,t}$ to specify the thermal coefficient (rate-of-heat flow) of appliance i and the heat transfer between appliance i and its surrounding environment at time t, respectively. The system thermal dynamics can be expressed by [40]

$$T_{i,t}^{\text{in}} = T_{i,t-1}^{\text{in}} + \sigma_{i,t}\left(T_{i,t}^{\text{amb}} - T_{i,t-1}^{\text{in}}\right) + \theta_i q_{i,t}. \tag{1}$$

Note that $\theta_i > 0$ if the thermal appliance is heating, and $\theta_i < 0$ if it is cooling. $\sigma_{i,t}$ can be calculated based on [41] and [42] for different types of thermal appliances.

Furthermore, we have constraints which account for the maximum heat rated power that thermal appliance i can provide and the comfort zone required by the customer, denoted by $q_i^{\max}$ and $\left[T_i^{\min}, T_i^{\max}\right]$, respectively, as given by

$$0 \leq q_{i,t} \leq q_i^{\max}, \qquad i \in \Psi, t \in \mathcal{T}, \tag{2}$$

$$T_i^{\min} \leq T_{i,t}^{\text{in}} \leq T_i^{\max}, \qquad i \in \Psi, t \in \mathcal{T}. \tag{3}$$

B. Chemical Energy Storage Unit Model

Local rechargeable batteries and batteries in EVs are chemical energy storage units that can change the household power consumption through controlling their charging or discharging process. Let C_i, $q_i^{\min}$ and $q_i^{\max}$, SOC_i^{init}, $SOC_i^{\min}$ and $SOC_i^{\max}$, and e_i denote the capacity, the minimum and maximum charging rate, the initial state of charge (SOC), the lower and upper limit of the SOC, and the charging efficiency factor of a local rechargeable battery ($i = \text{B}$) or an EV ($i = \text{EV}$), respectively. While a local rechargeable battery can always be scheduled, an EV can only be scheduled when plugged-in at the household. Denoting the set of all time slots for which the EV is plugged-in during one arrival and departure event by $\mathcal{T}_{\text{EV}} = \{t_a, t_{a+1}, \ldots, t_d\} \subset \mathcal{T}$, we have the following constraints:

$$q_{\text{B}}^{\min} \leq q_{\text{B},t} \leq q_{\text{B}}^{\max}, \quad t \in \mathcal{T}, \tag{4}$$

$$SOC_{\text{B},t} = SOC_{\text{B},t-1} + \frac{q_{\text{B},t-1}\Delta t\, e_{\text{B}}}{C_{\text{B}}}, \quad t \in \mathcal{T}, \tag{5}$$

$$q_{\text{EV}}^{\min} \leq q_{\text{EV},t} \leq q_{\text{EV}}^{\max}, \quad t \in \mathcal{T}_{\text{EV}}, \tag{6}$$

$$q_{\text{EV},t} = 0, \quad t \in \mathcal{T} \setminus \mathcal{T}_{\text{EV}}, \tag{7}$$

$$SOC_{\text{EV},t+1} = SOC_{\text{EV},t} + \frac{q_{\text{EV},t}\Delta t\, e_{\text{EV}}}{C_{\text{EV}}}, \quad t \in \mathcal{T}_{\text{EV}} \setminus \{t_{\text{d}}\}, \tag{8}$$

$$SOC_i^{\min} \leq SOC_{i,t} \leq SOC_i^{\max}, \quad i \in \{\text{B}, \text{EV}\}, t \in \mathcal{T}, \tag{9}$$

where $SOC_{\text{B},0} = SOC_{\text{B}}^{\text{init}}$ and $SOC_{\text{EV},t_{\text{a}}} = SOC_{\text{EV}}^{\text{init}}$.

Moreover, customers require a certain SOC level when the EV departs, denoted by $SOC_{\text{EV}}^{\text{req}}$, and thus

$$SOC_{\text{EV},t_d} = SOC_{\text{EV}}^{\text{req}}. \tag{10}$$

Notice that the EV can have multiple arrival and departure events during one scheduling period and it will follow the above constraints whenever such an event occurs.

C. Combining Energy Storage Units for Load Hiding

We consider three specific combinations of different energy storage units for load hiding purposes.

1) A dedicated local rechargeable battery with thermal appliances: The thermal appliances demand can be shifted to alleviate the need for a large-capacity or fast-charging-cycle battery. Customers can choose to use those loads during off-peak hours instead of peak hours based on their comfort requirements. Customers may also be given an incentive by the

utility company to apply such load shifting through demand response program, unrelated to privacy protection.

2) A dedicated local rechargeable battery with an EV: The use of an EV can also reduce the reliance on a dedicated large-capacity or fast-charging-cycle battery. Although the EV can only be scheduled when parked at home, its plugged-in time often overlaps with residential peak demand periods. Therefore, the EV can be exploited as an assistive battery which is beneficial in both load hiding and peak reduction.

3) An EV and thermal appliances: An EV and thermal appliances can jointly be used for load hiding. The EV serves as the battery when plugged-in at home, while the thermal appliances are assistive energy storage units when the EV departs. By using those existing energy storage units, no extra expenses for battery purchase and installation will be incurred.

III. Combining Energy Storage Units: An Optimization Framework

In this section, we propose an optimization framework to combine household energy storage units to achieve load hiding. Our primary objective is to protect customer privacy in a cost-effective manner. Specifically, we are motivated by the fact that household electricity consumption essentially comes from the inhabitants' activities. This indicates that the household occupancy patterns can be associated with the finite set of on-off household appliance loads. Markovian models have been shown to be effective in simulating active occupancy patterns [43]. To this end, we formulate an MDP problem that captures the uncertainties of household demand and energy storage availability as well as the interdependencies of scheduling decisions.

A. State Space and Action Space

We consider a finite horizon time model where the decision epochs are indexed by $t \in \mathcal{T}$. There are a total of W energy storage units used for load hiding, which are selected from the appliance set Ψ, the local dedicated battery, and the EV. Let $\phi_{i,t}$ and ϕ_t denote the state vector of the ith energy storage unit and all the storage units used at time t, respectively, such that $\phi_t = [\phi_{1,t}, \ldots, \phi_{W,t}]$. Denoting the state vector of the household base load at time t by l_t, we define the system state at time t as $s_t = [l_t, \phi_t] \in \mathcal{S}$, where $\mathcal{S}$ is the state space.

For thermal energy storage units, we have $\phi_{i,t} = \left(T^{\text{in}}_{i,t}, T^{\text{amb}}_{i,t}, \sigma_{i,t}\right)$, and the system state update follows equation (1). For chemical energy storage units, we differentiate between a local dedicated rechargeable battery and an EV. We introduce a binary indicator $a_{\text{EV},t}$ to denote the EV plugged-in status, where $a_{\text{EV},t} = 1$ when the EV is plugged-in at time t and can be scheduled, and $a_{\text{EV},t} = 0$ otherwise. We also introduce $d_{\text{EV},t} \in \mathbb{N}$ to denote the remaining plugged-in time of the EV at time t. We have $d_{\text{EV},t} = 0$ when $a_{\text{EV},t} = 0$. When $a_{\text{EV},t} = 1$, i.e., $t \in \mathcal{T}_{\text{EV}}$, we have $d_{\text{EV},t_{\text{a}}} = t_{\text{d}} - t_{\text{a}}$ and $d_{\text{EV},t} = d_{\text{EV},t-1} - 1$ for all $t \in \mathcal{T}_{\text{EV}} \setminus \{t_{\text{a}}\}$. Thus, the state vectors of the local dedicated battery and the EV are given by $\phi_{\text{B},t} = SOC_{\text{B},t}$ and $\phi_{\text{EV},t} = (SOC_{\text{EV},t}, a_{\text{EV},t}, d_{\text{EV},t})$, respectively.

At each decision epoch, the LCU needs to choose an action to schedule the energy storage units. Let u_t denote the control action vector at time t, where $u_t = (q_{1,t}, \ldots, q_{W,t})$. Due to the constraints of the energy storage units, not all the actions can be chosen at a given state. We thus introduce $\mathcal{U}(s)$ to denote the feasible set of all possible actions given state s. $\mathcal{U}(s)$ satisfies equations (1) – (3) for thermal appliances, and equations (4) – (10) for batteries.

B. System Dynamics

Given the system state s_t and control action u_t, the evolution of the MDP can be described by the system state transition probability $\mathbb{P}(s_{t+1} \,|\, s_t, u_t)$, for some $s_t, s_{t+1} \in \mathcal{S}$, $u_t \in \mathcal{U}(s_t)$, and $t \in \mathcal{T}$. We assume the household base load evolves according to a Markov chain, with its transition probabilities dependent only on occupancy patterns, where the evolution of the household base load is not affected by the control actions of the energy storage units. Therefore, the system transition probabilities between the composite states s_t and $s_{t+1} \in \mathcal{S}$ when action u_t is taken can be expressed as

$$\mathbb{P}(s_{t+1} \mid s_t, u_t) = \mathbb{P}(l_{t+1} \mid l_t)\mathbb{P}(\phi_{t+1} \mid \phi_t, u_t). \tag{11}$$

Since the state of a specific energy storage unit evolves with transition probabilities independent of other energy storage units, we have

$$\mathbb{P}(\phi_{t+1} \mid \phi_t, u_t) = \prod_{i=1}^{W} \mathbb{P}(\phi_{i,t+1} \mid \phi_{i,t}, q_{i,t}). \tag{12}$$

We now describe how to determine $\mathbb{P}(\phi_{i,t+1} \mid \phi_{i,t}, q_{i,t})$.

For thermal appliances, we assume that heat transfer between appliance i and its surrounding environment again evolves according to a Markov chain, independent of temperature as well as the control action. To simplify the discussion, we also assume $T^{\text{amb}}_{i,t}$ to be known, which can be obtained from the day-ahead forecast information. Therefore,

$$\mathbb{P}(\phi_{i,t+1} \mid \phi_{i,t}, q_{i,t}) = \begin{cases} \mathbb{P}(\sigma_{i,t+1} \mid \sigma_{i,t}), & \text{if (1) and (3) hold,} \\ 0, & \text{otherwise.} \end{cases} \tag{13}$$

The plugged-in status and remaining plugged-in time of the EV are mainly affected by the customer driving preference. Here we assume that their transition probabilities are independent of the control actions and energy level of the EV. Therefore, we have

$$\mathbb{P}(\phi_{\text{EV},t+1} \mid \phi_{\text{EV},t}, q_{\text{EV},t}) = \begin{cases} \mathbb{P}(a_{\text{EV},t+1}, d_{\text{EV},t+1} \mid a_{\text{EV},t}, d_{\text{EV},t}), & \text{if (8)} - \text{(10) hold,} \\ 0, & \text{otherwise.} \end{cases} \tag{14}$$

As shown in Fig. 2, there are only two states for the EV arrival status. When $a_{\text{EV},t} = 0$, the EV has not arrived at home. Consequently $d_{\text{EV},t} = 0$. Then it is with probability $p_{t,0}$ that the EV still does not arrive at time $t+1$. If the EV arrives at time $t+1$, and assuming we have N different states for the possible remaining parking duration $d_{\text{EV},t+1}$, the transitions

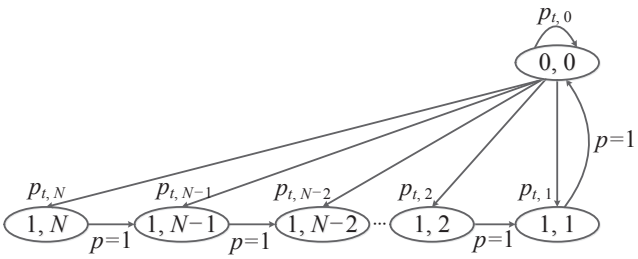

Fig. 2. Illustration of the Markovian process for the EV plugged-in status and remaining plugged-in time. The circles stand for the EV state $(a_{\mathrm{EV},t}, d_{\mathrm{EV},t})$ at time t.

to state $(1, k')$ occurs with probability $p_{t,k'}$. Note that we have $\sum_{k'=0}^{N} p_{t,k'} = 1$. When the EV is parking at home, its arrival and remaining parking duration states will move to $(a_{\mathrm{EV},t+2}, d_{\mathrm{EV},t+2}) = (1, d_{\mathrm{EV},t+1} - 1)$ with probability $p = 1$ in the next time slot. The process is repeated until the remaining parking duration becomes zero and the EV departs. Therefore, the evolution of the EV plugged-in status and remaining plugged-in time can be expressed as

$$\mathbb{P}(a_{\mathrm{EV},t+1} = j', d_{\mathrm{EV},t+1} = k' \mid a_{\mathrm{EV},t} = j, d_{\mathrm{EV},t} = k)$$
$$= \begin{cases} 1, & j' = j = 1, \ \ k' = k - 1, \\ 1, & j' = k' = 0, \ j = k = 1, \\ p_{t,0}, & j' = j = 0, \ \ k' = k = 0, \\ p_{t,k'}, & j' = 1, j = k = 0, k' \in \{1, \ldots, N\}, \\ 0, & \text{otherwise.} \end{cases} \tag{15}$$

For the local rechargeable battery, we have $\mathbb{P}(\phi_{\mathrm{B},t+1} \mid \phi_{\mathrm{B},t}, q_{\mathrm{B},t}) = 1$ if (5) and (9) hold, and $\mathbb{P}(\phi_{\mathrm{B},t+1} \mid \phi_{\mathrm{B},t}, q_{\mathrm{B},t}) = 0$ otherwise.

C. Objective Function

We consider both privacy leakage and electricity cost in the objective of the proposed load hiding scheme, since the customers may only be willing to tolerate certain additional consumption costs for the sake of privacy. Intuitively, a constant reported electricity usage can eliminate the possibility of inferring individual appliances usages. We thus aim to maintain a constant load and quantify the reduction in privacy leakage as to what extent the scheduled operations can reduce the energy fluctuations. Specifically, denoting the constant load we want to maintain (i.e., the average consumption over a scheduling period) by l_{c}, the privacy leakage $f(s_t, u_t)$ given system state s_t and action taken u_t can be expressed as

$$f(s_t, u_t) = \left\| \frac{l_{\mathrm{c}} - l_t - \sum_{i=1}^{W} q_{i,t} \Delta t}{l_{\mathrm{c}}} \right\|_2, \ \ t \in \mathcal{T}. \tag{16}$$

The household purchases and sells electricity at prices h_t^{c} and h_t^{d} at time t, respectively, which are assumed to be perfectly known, as the utility companies generally use fixed price plans as the TOU prices at the residential sector. The electricity cost given system state s_t and action taken u_t can be written as

$$g(s_t, u_t) = h_t^{\mathrm{c}} \left[\sum_{i=1}^{W} q_{i,t} \Delta t + l_t \right]^+ - h_t^{\mathrm{d}} \left[- \sum_{i=1}^{W} q_{i,t} \Delta t - l_t \right]^+, \tag{17}$$

where $t \in \mathcal{T}$ and $[x]^+ = \max(x, 0)$.

The one-step cost function at time t given system state s_t and action taken u_t is then defined as

$$c(s_t, u_t) = \lambda g(s_t, u_t) + (1 - \lambda) f(s_t, u_t), \ \ t \in \mathcal{T}, \tag{18}$$

where the weighting parameter $\lambda \in [0, 1]$.

At each decision epoch, the LCU implements a policy π, which is a mapping from each state $s \in \mathcal{S}$ to an action $u \in \mathcal{U}(s)$. Given the initial state s_1, following π yields a random path which consists of a sequence of states, actions, and costs over the entire scheduling period. For each decision epoch t, we have $u_t = \pi(s_t)$ and s_{t+1} is reached with probability $\mathbb{P}(s_{t+1} | s_t, u_t)$. We refer to such a sequence as an episode. Our goal is to minimize the total costs of the episode, which is a random variable dependent on the state transition probabilities. Thus, we define the objective function as the expected total cost and can be given as

$$J_{s_1}(\pi) = \mathbb{E}_{s_1}^{\pi} \left\{ \sum_{t=1}^{T} c(s_t, u_t) \mid s_1 \right\}. \tag{19}$$

The objective of the LCU is to find an optimal policy π^* that minimizes the expected total cost (19).

IV. Solution Methodology

In this section, we present the solution for the MDP problem. We devise a model-free learning method to minimize the objective function (19). In addition, we present a benchmark for the performance achieved with this practical stochastic problem framework by formulating a deterministic problem that considers household demand and the availability of storage as perfectly known for the duration of a scheduling period. We show how to solve the deterministic problem by transforming it into an equivalent convex form.

A. Model-Free Learning Using the Q-learning Algorithm

To determine the optimal policy of our MDP problem, we need to know the state transition probabilities $\mathbb{P}(l_{t+1} \mid l_t)$ and $\mathbb{P}(\phi_{t+1} \mid \phi_t, u_t)$. Unfortunately, the non-stationary transition probabilities of household base load and energy storage are difficult to approximate. The usage of energy storage units largely depends on customer preferences, and the household base load can vary drastically from peak to off-peak hours. Therefore, we use a model-free learning method. Specifically, we adopt the Q-learning algorithm for its simplicity [44]. The algorithm learns the action-value function that captures the expected cost associated with a state-action pair (s, u), which is defined as their Q value, by repeatedly choosing an action, yielding some costs, and obtaining information about outcome states. The optimal scheduling policy can then be constructed by selecting the action with the least cost in each state once the action-value function is learned.

The Q-learning algorithm uses an exploration and exploitation policy to choose actions given the current state. In particular, an η-greedy policy is used, which explores a random action with probability η and exploits a greedy action with probability $1 - \eta$. At each time t, given the current state s_t,

it chooses the greedy action u^* which minimizes the current action value $Q(s_t, u)$ most of the time, but with probability η, it instead selects an action randomly. Thus, the policy $\pi(s_t, u)$, $\forall u \in \mathcal{U}(s_t)$ can be given as [44, p. 122]

$$\pi(s_t, u) = \begin{cases} 1 - \eta + \frac{\eta}{|\mathcal{U}(s_t)|}, & \text{if } \ u = u^* \\ \frac{\eta}{|\mathcal{U}(s_t)|}, & \text{if } \ u \neq u^*. \end{cases} \tag{20}$$

After taking action u_t, the system incurs an instantaneous cost $c(s_t, u_t)$ and the system state transits to s_{t+1}. The update rule can be given by [44, p. 148]

$$\begin{aligned} Q(s_t, u_t) \leftarrow Q(s_t, u_t) + \alpha(c(s_t, u_t) + \\ \min_{u \in \mathcal{U}(s_{t+1})} Q(s_{t+1}, u) - Q(s_t, u_t)), \end{aligned} \tag{21}$$

where $\alpha \in [0, 1]$ is the learning rate. It allows us to interpolate between the old information (current Q value of the state-action pair (s_t, u_t)) and the new information (the observed state s_{t+1} and the one-step cost $c(s_t, u_t)$). The true expected value of the Q function will eventually be learned as long as this process continues, as we will observe the possible succeeding states that could have occurred and aggregate over their outcomes, even though we never directly learn the transition probabilities. Thus, the learnt Q function directly approximates the optimal action-value function Q^* independent of the policy being followed thereafter.

A description of the Q-learning algorithm is presented in Algorithm 1. It is executed by the LCU to learn the optimal policy based on historical household consumption data until the Q function converges. To start learning, the LCU initializes the Q function and sets the learning and weighting parameter (step 1). During the learning phase, it follows the η-greedy policy and keeps updating the Q function over all the decision epochs of multiple episodes. At the beginning of each episode, the LCU receives the pricing information of the entire scheduling period from the utility company and sets the desired constant load (step 3). It sets all the parameters for the energy storage units (step 4) and the initial state (step 5). For each decision epoch, the LCU observes the current state (step 8) and determines the control action (step 10). It then updates the Q function (step 13) after taking the action (step 11) and observing the outcome state (step 12). Steps 8 - 14 are repeated until the end of the episode. Once the LCU has finished learning, it determines the optimal action $u^* \in \mathcal{U}(s)$ based on the optimal policy π^* by observing the current system state $s \in \mathcal{S}$ for all decision epochs $t \in \mathcal{T}$ over the entire scheduling period.

Algorithm 1 Q-learning algorithm

1: Initialize Q function value, learning parameters α, η and weighting parameter λ.
2: **Repeat** (for each episode):
3: Obtain electricity pricing information h_t^{c} and h_t^{d}, $t \in \mathcal{T}$. Set desired household constant load output l_c.
4: Initialize energy storage units parameters.
5: Set the initial state: initialize energy storage units status, and observe the initial household base load.
6: $t := 1$.
7: **Repeat** (for each decision epoch of episode):
8: Observe the current state s_t.
9: $u^* = \arg\min_{u \in \mathcal{U}(s_t)} Q(s_t, u)$.
10: Determine the action u_t from (20).
11: Take the action u_t and calculate the cost from (18).
12: Observe the next state s_{t+1}: observe the energy storage units status update ϕ_{t+1} and the household base load l_{t+1}.
13: Update the Q function from (21).
14: $t := t + 1$.
15: **Until** the end of the episode.
16: **Until** the end of the learning phase.

B. Benchmark Problem Formulation

We now formulate a deterministic problem for which we assume future household demand and customer behavior are known, using the same privacy and cost measures defined in (16) – (18). This idealized formulation serves as a benchmark for the practical scenario considered in the previous subsection. Since the future household demand and customer behavior are assumed to be known, the privacy leakage index and the electricity cost can be expressed as

$$f_T = \sum_{t=1}^{T} f(s_t, u_t) \ \text{ and } \ g_T = \sum_{t=1}^{T} g(s_t, u_t),$$

respectively. Thus, the privacy-cost minimization problem is given by

$$\underset{q_{i,t},\, i \in \Psi \cup \{\mathrm{B}, \mathrm{EV}\},\, t \in \mathcal{T}}{\text{minimize}} \quad (1 - \lambda) f_T + \lambda g_T \tag{22a}$$

$$\text{subject to} \quad \text{constraints (1) – (10).} \tag{22b}$$

Problem (22) is non-convex due to the definition of $g(s_t, u_t)$, which is the difference between two convex functions. However, we can transform the original problem into its equivalent convex form as follows. We assume that $h_t^{\mathrm{c}} \geq h_t^{\mathrm{d}}$, as the utility companies will sell electricity to the customers at a higher price than purchasing back in order to make profit. By introducing two non-negative auxiliary variables κ_t^+ and κ_t^-, we can rewrite $g(s_t, u_t)$ as

$$g(s_t, u_t) = h_t^{\mathrm{c}} \kappa_t^+ - h_t^{\mathrm{d}} \kappa_t^-. \tag{23}$$

Substituting (23) into problem (22), the equivalent convex form can be given by

$$\underset{q_{i,t}, \kappa_t^+, \kappa_t^-,\, i \in \Psi \cup \{\mathrm{B}, \mathrm{EV}\},\, t \in \mathcal{T}}{\text{minimize}} \quad (1 - \lambda) f_T + \lambda g_T \tag{24a}$$

$$\text{subject to} \quad \kappa_t^+ - \kappa_t^- = \sum_{i=1}^{W} q_{i,t} \Delta t + l_t, \ t \in \mathcal{T}, \tag{24b}$$

$$\kappa_t^+, \kappa_t^- \geq 0, \quad t \in \mathcal{T}, \tag{24c}$$

$$\text{constraints (1) – (10).} \tag{24d}$$

Problem (24) can be solved using solvers such as CVX [45].

V. Performance Evaluation

In this section, we evaluate the performance of our proposed household load hiding method. We consider a 24 hour scheduling period in our simulation. The AMI currently deployed in, for example, Ontario, Canada supports smart meter data transmission periodically every 5 to 60 minutes [46]. The next generation of smart meters can provide AMI with up to one minute interval measurement data [47]. In our sim-

ulation, we set the smart meter sampling interval Δt to be 2 minutes. Therefore, we have $T = 720$ decision epochs in total for each scheduling period. The EV related parameters are set according to the specification for Chevy Volt model [48]. The maximum EV charging and discharging rate is $q_{\mathrm{EV}}^{\max} = -q_{\mathrm{EV}}^{\min} = 1.44$ kW, with a capacity $C_{\mathrm{EV}} = 8$ kWh. For simplicity, we assume that the EV arrives and departs only once during each scheduling period. Following the National Household Travel Survey 2009 [49], the EV is simulated to arrive within [5 pm, 7 pm], and depart within [6 am, 8 am] the next day with uniform probability. The rationale is that the customers will generally park the EV during the night to fully charge it to their desired SOC level $SOC_{\mathrm{EV}}^{\mathrm{req}}$ for use during the next day. Moreover, we assume that the customer will inform the LCU about $SOC_{\mathrm{EV}}^{\mathrm{req}}$ and the parking duration $d_{\mathrm{EV},t}$ upon the EV arrival to ensure that constraint (10) on EV departure SOC level is satisfied. The initial SOC of the EV, $SOC_{\mathrm{EV}}^{\mathrm{init}}$, when plugged-in is a random variable uniformly distributed in [0.1, 0.9]. The desired SOC of the EV at departure, $SOC_{\mathrm{EV}}^{\mathrm{req}}$, is set to 0.9. The initial SOC of the battery, $SOC_{\mathrm{B}}^{\mathrm{init}}$, is also selected uniformly at random from [0.1, 0.9]. We set $q_{\mathrm{B}}^{\max} = -q_{\mathrm{B}}^{\min} = C_{\mathrm{B}}/4$ in our experiments, which means it takes 4 hours to charge or discharge the battery. By using this definition, batteries with large-capacity indicate that they have fast-charging-cycle. Similarly, batteries with small capacity indicate that they have slow-charging-cycle. The lower and upper SOC limits of the EV and the battery are set as $SOC_{\mathrm{EV}}^{\min} = SOC_{\mathrm{B}}^{\min} = 0.1$ and $SOC_{\mathrm{EV}}^{\max} = SOC_{\mathrm{B}}^{\max} = 0.9$, respectively, and the charging efficiencies of the battery and the EV are $e_{\mathrm{B}} = e_{\mathrm{EV}} = 1$.

We consider two commonly used HVAC models, with their thermal characteristics set according to [41] and [50], respectively. Note that the HVAC system in [41] is more energy efficient compared to that in [50], which means it consumes less energy to achieve the same heating or cooling performance. We calculate the temperature settings when the HVAC system is not used for load hiding, and apply these settings with a four degree temperature band (plus or minus two degrees deviation) as the customer desired comfort zone that must be satisfied when the HVAC system is used as a controllable load. The household base load is simulated from [51], where the energy demand from either a weekday or weekend can be specified. In our simulation, the learning process is carried out separately for weekday and weekend profiles to tackle their differences in terms of consumption patterns. The electricity pricing information is obtained from [52]. The parameters used in the learning algorithm are set as $\alpha = 0.05$ and $\eta = 0.02$. Unless otherwise stated, the trade-off parameter is set as $\lambda = 0$. The learning process is repeated for 20000 episodes.

A. *Hiding Effect with Different Energy Storage Combinations*

We first illustrate how our method can disguise the household load by combining different energy storage units. The original load and the masked load profiles of a weekday using different combinations are shown in Fig. 3, with respect to both the Q-learning algorithm and the benchmark solution.

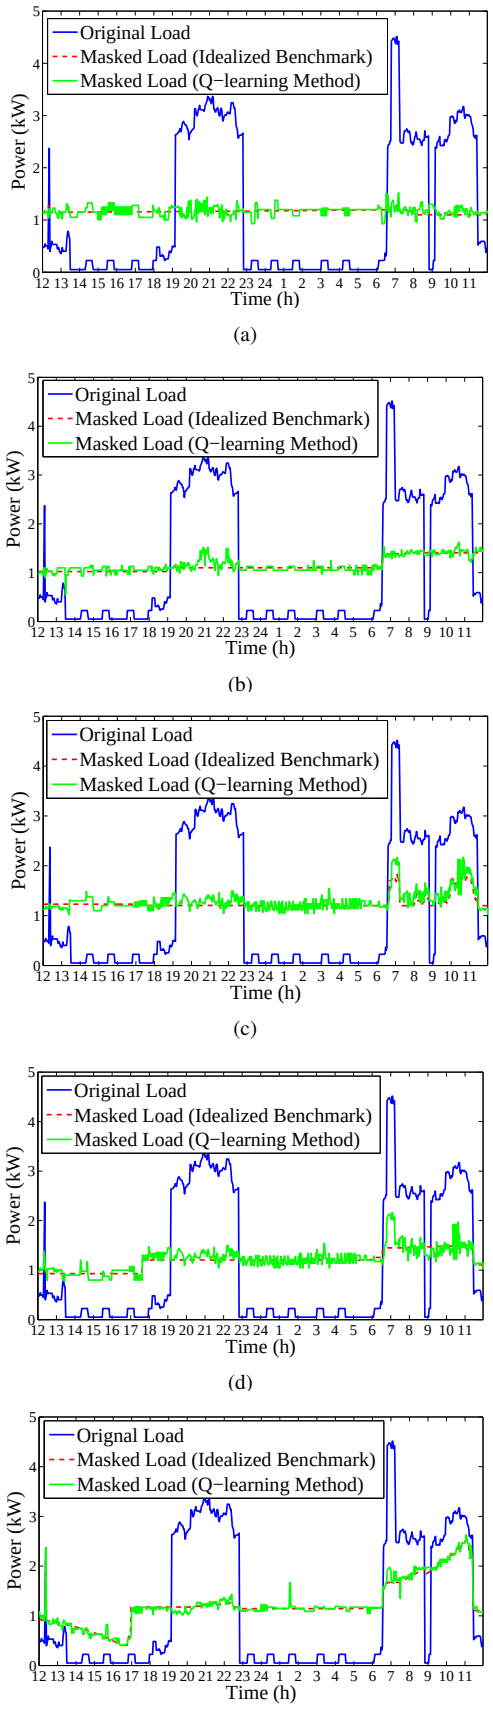

Fig. 3. The observed household electricity consumption using combinations of (a) Battery-HVAC 1, (b) Battery-HVAC 2, (c) a local rechargeable battery and an EV (Battery-EV), (d) EV-HVAC 1, and (e) EV-HVAC 2 for load hiding. The size of the local rechargeable battery is 10 kWh for all cases.

To differentiate between different HVAC systems, we refer to the combination of a local rechargeable battery and the HVAC system with its thermal settings from [41] and [50] as Battery-HVAC 1 and Battery-HVAC 2, respectively. Similarly, the combination of the EV and the HVAC system with its thermal settings from [41] and [50] are referred to as EV-HVAC 1 and EV-HVAC 2, respectively.

In Fig. 3(a), we can observe an almost flattened load curve using Battery-HVAC 1 for load hiding. Since the HVAC system contributes to a large portion of the household load peaks, shifting its demand effectively eliminates load variations. Fig. 3(b) shows the hiding results of Battery-HVAC 2. We still observe a similar trend in the flattened load curve, but to a lesser extent as the HVAC system in this setting is less efficient in terms of converting electricity to heat. Overall, the results validate our approach of using HVAC systems to assist a local rechargeable battery for load hiding, which can be generalized to other existing energy-intensive household thermal appliances (e.g., electric water tanks). Fig. 3(c) shows that using Battery-EV for load hiding produces a similarly flattened load curve, except for the period after 7 am when the EV has departed. This means that the desired output cannot always be maintained due to the EV availability and the physical constraints for the battery. We observe from Fig. 3(d) that using an EV and an HVAC system (without the extra battery) produces a piecewise flattened load curve. However, if an EV and a less efficient HVAC system are used for hiding, notable load variations remain, as shown in Fig. 3(e). In this case, the flexibility of using the HVAC system for load hiding is limited by the need to first accommodate the customer's comfort requirements.

The results in Fig. 3 support the idea of using EV and HVAC as opportunistic energy storage units for load hiding. However, their hiding performance is limited by the EV availability, requirements on the EV SOC level at departure, the thermal efficiency of the HVAC system, and restrictions on the flexibility of the HVAC use. We further observe from the figures that the load curves of the Q-learning method exhibit small fluctuations around the flat lines obtained with the idealized benchmark optimization. However, these fluctuations do not allow energy usage information to be inferred about the customer. Hence, we conclude that the learning-based method can successfully schedule the energy storage units to disguise the household load, using only the available system state information.

B. Cost and Privacy Performance Evaluation

We now quantitatively evaluate the privacy preserving and cost reduction performance of our Q-learning method and of the benchmark solution. We adopt the mutual information between the original household load and the masked load as the privacy protection measure [25]. For both the masked load time series as well as the original load time series, we compute the mutual information based on [25, Eq. (1)] using the load differences at each consecutive time slot pair (i.e., t, $t+1 \in \mathcal{T}$) during one scheduling period, applying the quantization interval $m = 100$ (see [25] for more details).

Fig. 4 shows the electricity cost of the customer per scheduling period as a function of mutual information when different energy storage unit combinations are used for hiding. As more and more households are equipped with high energy efficient HVAC systems, we choose the HVAC system from [41] for performance evaluation here. The dedicated battery has a storage capacity of 10 kWh. The results are obtained using different parameters $\lambda \in [0, 1]$, enabling a trade-off between the cost of energy consumption and privacy leakage (smaller mutual information corresponds to better privacy protection).

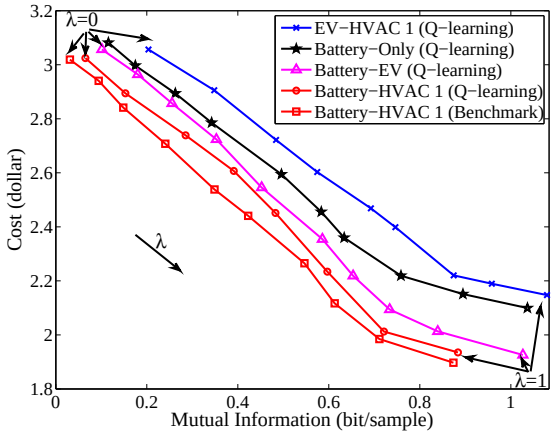

Fig. 4. Cost versus mutual information for Q-learning method using Battery-HVAC 1, EV-HVAC 1, Battery-EV, and Battery-only combinations for hiding. Also included are the results for the deterministic benchmark formulation for Battery-HVAC 1 case.

We observe that the Battery-HVAC 1 combination for load hiding achieves the best cost-privacy performance, followed by the Battery-EV, Battery-only, and the EV-HVAC 1 combination. The results substantiate that EV and HVAC can both be exploited as an alternative or an assistive energy storage solution to a dedicated battery for load hiding. The battery-assisted use case is more effective in terms of electricity cost though, mainly due to the comfort and EV charging requirements of the customer. We further note that the proposed Q-learning method provides a cost-effective trade-off close to that obtained for the idealized benchmark case, which assumes perfect knowledge of future household consumption and EV arrival and departure times. We show the Battery-HVAC 1 case in Fig. 4, but the same trend has been observed for the other storage unit combinations.

We next compare the performance of the Q-learning approach with that of the online control algorithm from [29]. Since this method is only applicable to a dedicated battery case, Fig. 5 shows the cost-privacy trade-off when both methods make use of a 10 kWh battery for load hiding. We observe that both methods achieve fairly similar cost-privacy performances for this scenario. This corroborates the effectiveness of our learning framework, as it does not incur a performance penalty for the battery-only case, while its actual utility lies in the applicability to load hiding with a combination of dedicated and opportunistic storage units.

Finally, we evaluate the possible reduction of required capacity for the dedicated battery when assisted by EV and H-VAC units. For this, we consider the cost-privacy performance for different combinations of storage units with batteries of different sizes in Fig. 6. For concreteness, we focus on the case of maximal privacy (i.e., $\lambda = 0$). We observe that using a single assistive storage unit can reduce the battery size by

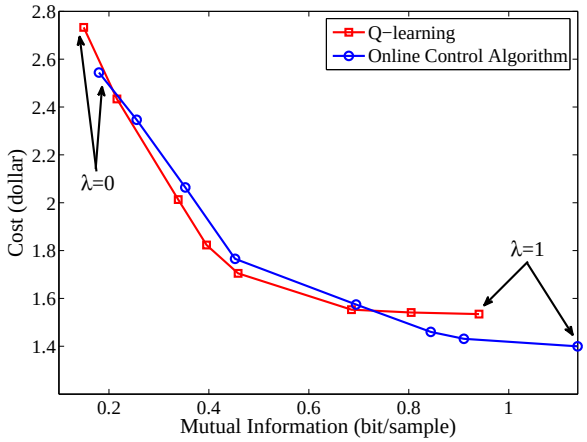

Fig. 5. Cost versus mutual information for proposed Q-learning method and the online control algorithm from [29] using different trade-off parameters. In both cases, a dedicated rechargeable battery is used for load hiding.

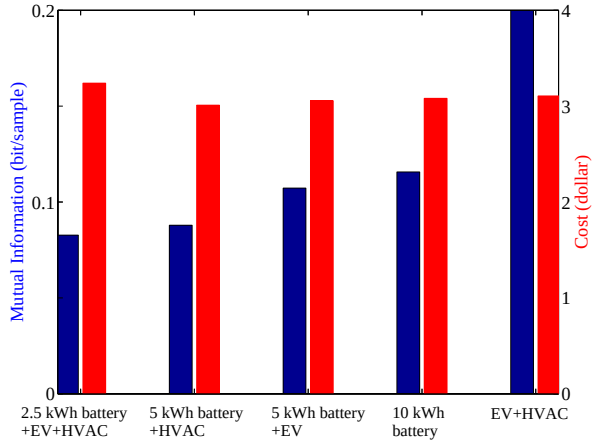

Fig. 6. Cost and mutual information for different battery sizes with and without the use of EV or HVAC systems for load hiding.

50%. Furthermore, the joint use of EV and HVAC achieves a better privacy protection at about the same energy consumption cost but with a battery of only one quarter the capacity of the battery-only case.

VI. Conclusion

In this paper, we have investigated the potential of existing household energy storage units to tackle the potential privacy leakage from smart meter readings. We have proposed a cost-effective privacy preserving solution from the customers' perspective by leveraging the combined use of existing thermal appliances and energy storage units for household load hiding. To accomplish this, we have formulated an MDP problem that captures the uncertainties in both the household power demand and customer usage behavior. A model-free learning framework has been designed to solve the MDP problem without prior information on that state transition probabilities. We have also provided a deterministic optimization problem, whose solution can be considered as a benchmark for solution obtained with the proposed method. Simulation results validated our approach and showed its effectiveness in achieving a favorable privacy-cost trade-off with a relatively small-size storage battery. The idea of designing alternative solutions that can be deployed at customer premise has merits in addressing the new privacy challenges raised by the use of IoT devices.

References

[1] L. Atzori, A. Iera, and G. Morabito, "The Internet of Things: A survey," *Computer Networks*, vol. 54, no. 15, pp. 2787–2805, Oct. 2010.

[2] N. Bui, A. P. Castellani, P. Casari, and M. Zorzi, "The Internet of Energy: A web-enabled smart grid system," *IEEE Network*, vol. 26, no. 4, pp. 39–45, Jul. 2012.

[3] G. Hart, "Nonintrusive appliance load monitoring," *Proc. of the IEEE*, vol. 80, no. 12, pp. 1870–1891, Dec. 1992.

[4] M. Zeifman and K. Roth, "Nonintrusive appliance load monitoring: Review and outlook," *IEEE Trans. on Consumer Electronics*, vol. 57, no. 1, pp. 76–84, Feb. 2011.

[5] J. Kelly and W. Knottenbelt, "Neural NILM: Deep neural networks applied to energy disaggregation," in *Proc. of ACM Int'l Conf. on Embedded Systems for Energy-Efficient Built Environments*, Seoul, South Korea, Nov. 2015.

[6] M. Baker, G. Hicks, S. Rodriguez, and M. Fuller, "A test of commercially available products for estimating end uses from smart meter data," in *Proc. of Int'l Workshop on Non-Intrusive Load Monitoring*, Vancouver, Canada, May 2016.

[7] C. Efthymiou and G. Kalogridis, "Smart grid privacy via anonymization of smart metering data," in *Proc. of IEEE SmartGridComm*, Gaithersburg, MD, Oct. 2010.

[8] C. Rottondi, G. Mauri, and G. Verticale, "A data pseudonymization protocol for smart grids," in *Proc. of IEEE GreenCom*, Piscataway, NJ, Sep. 2012.

[9] M. Jawurek, M. Johns, and K. Rieck, "Smart metering de-pseudonymization," in *Proc. of Computer Security Applications Conf.*, Orlando, FL, Dec. 2011.

[10] S. Ruj and A. Nayak, "A decentralized security framework for data aggregation and access control in smart grids," *IEEE Trans. on Smart Grid*, vol. 4, no. 1, pp. 196–205, Mar. 2013.

[11] C. Rottondi, G. Verticale, and C. Krauss, "Distributed privacy-preserving aggregation of metering data in smart grids," *IEEE J. on Sel. Areas in Commun.*, vol. 31, no. 7, pp. 1342–1354, Jul. 2013.

[12] C. I. Fan, S. Y. Huang, and Y. L. Lai, "Privacy-enhanced data aggregation scheme against internal attackers in smart grid," *IEEE Trans. on Industrial Informatics*, vol. 10, no. 1, pp. 666–675, Feb. 2014.

[13] H. Shen, M. Zhang, and J. Shen, "Efficient privacy-preserving cube-data aggregation scheme for smart grids," *IEEE Trans. on Information Forensics and Security*, vol. 12, no. 6, pp. 1369–1381, Jun. 2017.

[14] Z. Erkin, J. R. Troncoso-Pastoriza, R. L. Lagendijk, and F. Perez-Gonzalez, "Privacy-preserving data aggregation in smart metering systems: An overview," *IEEE Signal Processing Mag.*, vol. 30, no. 2, pp. 75–86, Mar. 2013.

[15] J. M. Bohli, C. Sorge, and O. Ugus, "A privacy model for smart metering," in *Proc. of IEEE ICC Workshops*, Capetown, South Africa, May 2010.

[16] S. Wang, L. Cui, J. Que, D. H. Choi, X. Jiang, S. Cheng, and L. Xie, "A randomized response model for privacy preserving smart metering," *IEEE Trans. on Smart Grid*, vol. 3, no. 3, pp. 1317–1324, Sept. 2012.

[17] L. Sankar, S. R. Rajagopalan, S. Mohajer, and S. Mohajer, "Smart meter privacy: A theoretical framework," *IEEE Trans. on Smart Grid*, vol. 4, no. 2, pp. 837–846, Jun. 2013.

[18] M. Savi, C. Rottondi, and G. Verticale, "Evaluation of the precision-privacy tradeoff of data perturbation for smart metering," *IEEE Trans. on Smart Grid*, vol. 6, no. 5, pp. 2409–2416, Sept. 2015.

[19] D. Günüz and J. Gómez-Vilardebó, "Smart meter privacy in the presence of an alternative energy source," in *Proc. of IEEE ICC*, Budapest, Hungary, Jun. 2013.

[20] O. Tan, D. Günüz, and H. V. Poor, "Increasing smart meter privacy through energy harvesting and storage devices," *IEEE J. on Sel. Areas in Commun.*, vol. 31, no. 7, pp. 1331–1341, Jul. 2013.

[21] G. Giaconi, D. Günüz, and H. V. Poor, "Smart meter privacy with an energy harvesting device and instantaneous power constraints," in *Proc. of IEEE ICC*, London, UK, Jun. 2015.

[22] G. Giaconi and D. Günüz and H. V. Poor, "Smart meter privacy with renewable energy and an energy storage device," accepted for publication in *IEEE Trans. on Information Forensics and Security*, 2017.

[23] G. Kalogridis, C. Efthymiou, S. Denic, T. Lewis, and R. Cepeda, "Privacy for smart meters: Towards undetectable appliance load signatures," in *Proc. of IEEE SmartGridComm*, Gaithersburg, MD, Oct. 2010.

[24] S. McLaughlin, P. McDaniel, and W. Aiello, "Protecting consumer privacy from electric load monitoring," in *Proc. of ACM Conf. on Computer and Communications Security*, Chicago, IL, Oct. 2011.

[25] W. Yang, N. Li, Y. Qi, W. Qardaji, S. McLaughlin, and P. McDaniel, "Minimizing private data disclosures in the smart grid," in *Proc. of ACM Conf. on Computer and Commun. Security*, Raleigh, NC, Oct. 2012.

[26] J. Zhao, T. Jung, Y. Wang, and X. Li, "Achieving differential privacy of data disclosure in the smart grid," in *Proc. of IEEE INFOCOM*, Toronto, Canada, Apr. 2014.

[27] S. Li, A. Khisti, and A. Mahajan, "Privacy-optimal strategies for smart metering systems with a rechargeable battery," in *Proc. of American Control Conf. (ACC)*, Boston, MA, July 2016.

[28] J. Yao and P. Venkitasubramaniam, in *Proc. of Allerton Conf. on Communication, Control, and Computing*, Monticello, IL, Oct. 2013.

[29] L. Yang, X. Chen, J. Zhang, and H.V. Poor, "Cost-effective and privacy-preserving energy management for smart meters," *IEEE Trans. on Smart Grid*, vol. 6, no. 1, pp. 486–495, Jan. 2015.

[30] O. Tan, J. Gmez-Vilardeb, and D. Gndz, "Privacy-cost trade-offs in demand-side management with storage," *IEEE Trans. on Information Forensics and Security*, vol. 12, no. 6, pp. 1458–1469, Jun. 2017.

[31] J. Koo, X. Lin, and S. Bagchi, "RL-BLH: Learning-based battery control for cost savings and privacy preservation for smart meters," in *Proc. of IEEE/IFIP Int'l Conf. on Dependable Systems and Networks (DSN)*, Denver, CO, Jun. 2017.

[32] Z. Zhang, Z. Qin, L. Zhu, J. Weng, and K. Ren, "Cost-friendly differential privacy for smart meters: Exploiting the dual roles of the noise," *IEEE Trans. on Smart Grid*, vol. 8, no. 2, pp. 619–626, Mar. 2017.

[33] D. Egarter, C. Prokop, and W. Elmenreich, "Load hiding of household's power demand," in *Proc. of IEEE SmartGridComm*, Venice, Italy, Nov. 2014.

[34] D. Chen, D. Irwin, P. Shenoy, and J. Albrecht, "Combined heat and privacy: Preventing occupancy detection from smart meters," in *Proc. of IEEE PerCom*, Budapest, Hungary, Mar. 2014.

[35] D. Chen, S. Kalra, D. Irwin, P. Shenoy, and J. Albrecht, "Preventing occupancy detection from smart meters," *IEEE Trans. on Smart Grid*, vol. 6, no. 5, pp. 2426–2434, Aug. 2015.

[36] A. Reinhardt, G. Konstantinou, D. Egarter, and D. Christin, "Worried about privacy? Let your PV converter cover your electricity consumption fingerprints," in *Proc. of IEEE SmartGridComm*, Miami, FL, Nov. 2015.

[37] Y. Sun, L. Lampe, and V.W.S. Wong, "Combining electric vehicle and rechargeable battery for household load hiding," in *Proc. of IEEE SmartGridComm*, Miami, FL, Nov. 2015.

[38] Y. Sun, L. Lampe, and V.W.S. Wong, "EV-assisted battery load hiding: A Markov decision approach," in *Proc. of IEEE SmartGridComm*, Sydney, Australia, Nov. 2016.

[39] Y. H. Liu, S. H. Lee, and A. Khisti, "Information-theoretic privacy in smart metering systems using cascaded rechargeable batteries," *IEEE Signal Processing Letters*, vol. 24, no. 3, pp. 314–318, Mar. 2017.

[40] B. Ramanathan and V. Vittal, "A framework for evaluation of advanced direct load control with minimum disruption," *IEEE Trans. on Power Systems*, vol. 23, no. 4, pp. 1681–1688, Nov. 2008.

[41] M. Ilic, J. W. Black, and J. L. Watz, "Potential benefits of implementing load control," in *Proc. of IEEE Power Engineering Society Winter Meeting*, New York, NY, Jan. 2002.

[42] L. Paull, H. Li, and L. Chang, "A novel domestic electric water heater model for a multi-objective demand side management program," *Electric Power Systems Research*, vol. 80, no. 12, pp. 1446–1451, Dec. 2010.

[43] I. Richardson, M. Thomson, and D. Infield, "A high-resolution domestic building occupancy model for energy demand simulations," *Energy and Buildings*, vol. 40, no. 8, pp. 1560–1566, 2008.

[44] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. MIT Press, 1998.

[45] CVX Research, Inc., "CVX: Matlab software for disciplined convex programming, version 2.1," Mar. 2017. [Online]. Available: http://cvxr.com/cvx.

[46] IESO, "Building privacy into Ontario's smart meter data management system: A control framework," May 2012.

[47] Sensus, "FlexNet AMI System." [Online]. Available: https://sensus.com/solutions/advanced-metering-infrastructure-ami/

[48] J. Lyle, "Latest Chevy Volt battery pack and generator details and clarications," Aug. 2007. [Online]. Available: http://gm-volt.com/2007/08/29/latest-chevy-volt-battery-pack-and-generator-details-and-clarications/.

[49] U.S. Department of Transporation, "Summary of travel trends, 2009 national household travel survey," Washington, DC, 2009.

[50] J. Black, "Integrating demand into the US electric power system: Technical, economic, and regulatory frameworks for responsive load." Ph.D. Dissertation, Engineering Systems Division. Cambridge, MA, Massachusetts Institute of Technology, 2005.

[51] I. Richardson, M. Thomson, D. Infield, and C. Clifford, "Domestic electricity use: A high-resolution energy demand model," *Energy Buildings*, vol. 42, no. 10, pp. 1878–1887, Oct. 2010.

[52] "Customer generation price plan." [Online]. Available: http://www.srpnet.com/prices/home/customergenerated.aspx

Yanan Sun (S'14) received the B.S. and M.S. degrees both from Xi'an Jiaotong University, Xi'an, China, in 2011 and 2014, respectively. She is currently a Ph.D. candidate in the Department of Electrical and Computer Engineering, The University of British Columbia (UBC), Vancouver, BC, Canada. Her research interests lie in the broad area of cyber-physical system. Her current work focus on smart meter privacy and energy storage applications in smart grid.

Lutz Lampe (M'02-SM'08) received the Dipl.-Ing. and Dr.-Ing. degrees in electrical engineering from the University of Erlangen, Germany, in 1998 and 2002, respectively. Since 2003, he has been with the Department of Electrical and Computer Engineering, The University of British Columbia, Vancouver, BC, Canada, where he is a Full Professor. His research interests are broadly in theory and application of wireless, power line, optical wireless and optical fibre communications. Dr. Lampe is currently an Associate Editor for the IEEE TRANSACTIONS ON COMMUNICATIONS, the IEEE COMMUNICATIONS LETTERS, and the IEEE COMMUNICATIONS SURVEYS AND TUTORIALS. He was a (co-)recipient of a number of best paper awards, including awards at the 2006 IEEE International Conference on Ultra-Wideband (ICUWB), the 2010 IEEE International Communications Conference (ICC), and the 2011 and 2017 IEEE International Symposium on Power Line Communications and Its Applications (ISPLC). He was the General (Co-)Chair for the 2005 IEEE ISPLC, the 2009 IEEE ICUWB and the 2013 IEEE International Conference on Smart Grid Communications (SmartGridComm). He is a co-editor of the book Power Line Communications: Principles, Standards and Applications from Multimedia to Smart Grid, published by John Wiley & Sons in its 2nd edition in 2016.

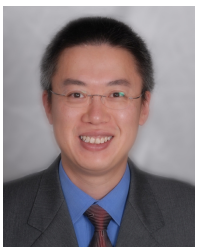

Vincent W.S. Wong (S'94, M'00, SM'07, F'16) received the B.Sc. degree from the University of Manitoba, Winnipeg, MB, Canada, in 1994, the M.A.Sc. degree from the University of Waterloo, Waterloo, ON, Canada, in 1996, and the Ph.D. degree from the University of British Columbia (UBC), Vancouver, BC, Canada, in 2000. From 2000 to 2001, he worked as a systems engineer at PMC-Sierra Inc. (now Microsemi). He joined the Department of Electrical and Computer Engineering at UBC in 2002 and is currently a Professor. His research areas include protocol design, optimization, and resource management of communication networks, with applications to wireless networks, smart grid, mobile cloud computing, and Internet of Things. Dr. Wong is an Editor of the *IEEE Transactions on Communications*. He has served as a Guest Editor of *IEEE Journal on Selected Areas in Communications* and *IEEE Wireless Communications*. He has also served on the editorial boards of *IEEE Transactions on Vehicular Technology* and *Journal of Communications and Networks*. He was a Technical Program Co-chair of *IEEE SmartGridComm*'14, as well as a Symposium Co-chair of *IEEE SmartGridComm* ('13, '17) and *IEEE Globecom*'13. He is the Chair of the IEEE Vancouver Joint Communications Chapter and has served as the Chair of the IEEE Communications Society Emerging Technical Sub-Committee on Smart Grid Communications. Dr. Wong received the 2014 UBC Killam Faculty Research Fellowship.