Privacy-Cost Trade-offs in Demand-Side Management with Storage

Onur Tan*, Jesús Gómez-Vilardebó*, Deniz Gündüz†
*Centre Tecnològic de Telecomunicacions de Catalunya (CTTC/CERCA), Barcelona, Spain.
†Imperial College London, London, UK.
Emails: {onur.tan@cttc.es, jesus.gomez@cttc.es, d.gunduz@imperial.ac.uk }

***Abstract*—Demand-side energy management (EM) is studied from a *privacy-cost trade-off* perspective, considering time-of-use pricing and the presence of an energy storage unit. Privacy is measured as the variation of the power withdrawn from the grid from a fixed target value. Assuming non-causal knowledge of the household's aggregate power demand profile and the electricity prices at the energy management unit (EMU), the privacy-cost trade-off is formulated as a convex optimization problem, and a low-complexity *backward water-filling algorithm* is proposed to compute the optimal EM policy. The problem is studied also in the online setting assuming that the power demand profile is known to the EMU only causally, and the optimal EM policy is obtained numerically through dynamic programming (DP). Due to the high computational cost of DP, a low-complexity heuristic EM policy with a performance close to the optimal online solution is also proposed, exploiting the water-filling algorithm obtained in the offline setting. As an alternative, information theoretic leakage rate is also evaluated, and shown to follow a similar trend as the load variance, which supports the validity of the load variance as a measure of privacy. Finally, the privacy-cost trade-off, and the impact of the size of the storage unit on this trade-off are studied through numerical simulations using real smart meter data in both the offline and online settings.**

***Index Terms*—Smart meter, privacy, demand-side management, energy storage, home energy management.**

I. Introduction

Smart meters (SMs) are key components for demand-side management in smart grids. SMs measure power consumption of users and transmit their readings to the utility provider (UP) in almost real-time. This allows the UPs to closely monitor the grid, improving its reliability, robustness and efficiency [1]. For example, the UPs can support time-of-use electricity pricing based on fine-grained SM readings and encourage consumers to shift their demands to off-peak hours with the promise of reduced energy costs. Despite many potential benefits, the possible misuse of SM data by third parties, as well as the UP, raises serious privacy concerns [2]. Intruders can analyze SM readings [3], [4], and extract private information regarding user activities, such as residential occupancy, sleep schedule, meal time [5], and appliance usage patterns [6].

Privacy can be achieved by modifying SM readings before being reported to the UP. For example, by compressing SM readings [7], by adding random noise [8], or by sending the aggregated energy consumption of a group of users [9]. However, tampering with SM readings also reduces their relevance and value as control signals. Alternatively, privacy can also be provided by demand-side management utilizing storage units, such as rechargeable batteries (RBs) [10]–[19], and alternative energy sources, such as a renewable energy source like a solar panel [10], [11], [20]. In [10] user's privacy is protected by using a RB and a renewable energy source from an information theoretic perspective. Heuristic algorithms are proposed in [13]–[15]. The joint optimization of privacy and energy cost with a RB is addressed in [17]–[19]. The authors in [17] and [18] propose online algorithms based on stochastic dynamic programming (DP) and Lyapunov optimization techniques, respectively. The authors in [19], [21] and [22] study a stochastic control model, formulated as a partially observeable Markov decision process. Characterizing the optimal strategy is computationally challenging due to the continuous state-action space; while approximate solutions can be obtained numerically through discretization, or upper and lower bounds can be derived.

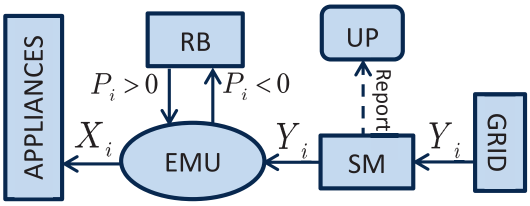

Fig. 1. A smart meter (SM) system diagram with an energy management unit (EMU) and a rechargeable battery (RB) at the user's household. The EMU manages the power flows (solid lines) among the power grid, the appliances and the RB. SM reports its readings to the UP at certain time instants, illustrated by the dashed line.

In this paper, we consider the SM system depicted in Fig. 1. The energy flow is managed by the energy management unit (EMU), which satisfies the power demands of the appliances, X_i, from the power grid and the RB. We do not allow any outages or shifting of user demands. The SM measures the power withdrawn from the grid, Y_i, and reports it to the UP at certain time instants without any tampering. Assuming that the electricity price is time-dependent, the EMU utilizes the RB both to reduce the energy cost, and to mask the energy consumption profile of the user. One can argue that perfect privacy is achieved if a constant SM reading is reported to the

The first two authors received support from the Catalan Government under grant SGR2014-1567, the Spanish Government under grant TEC2013-44591-P (INTENSYV), and the European Unions Horizon 2020 Research and Innovation Programme under grant agreement No 646469, while the last author's research has been funded by the U.K. Engineering and Physical Sciences Research Council (EPSRC) (EP/NO21738/1).

UP over time [13]. Consequently, we measure user privacy in terms of the variation of the power withdrawn from the grid, Y_i, from a constant target consumption value over the period of interest [12]. In addition to the load variance, we also evaluate the information leakage rate, which is defined as the mutual information rate between the aggregated power demands of the appliances and the SM readings. Mutual information, which takes into account the statistics of the user's demand has been previously considered as a measure of privacy for SM systems [7], [10], [11], [14]–[16], [19]–[21], [23]. On the other hand, the energy cost is measured with a time-varying time-of-use electricity pricing model. Our goal here is to design energy management (EM) policies that jointly increase the privacy of the user and reduce the energy cost over a given period of time under a RB capacity constraint.

Building upon our previous work [12], we first characterize the optimal *offline* EM policy, assuming that the energy demands and electricity prices are known non-causally by the EMU over the period of interest. We formulate the privacy-cost trade-off as a convex optimization problem, and identify the structure of the optimal policy. Exploiting this structure, we provide a *backward water-filling algorithm*, which efficiently finds the optimal EM policy.

Next, we study the online optimization problem considering only causal knowledge of the energy demands at the EMU, that is, the EMU learns only the energy demand at the current time slot (TS). We characterize the optimal online policy using DP. Due to the continuous state space, DP algorithms with good approximation to the optimal solution are prohibitively complex; and therefore, we propose a simple heuristic online algorithm, which exploits the backward water-filling algorithm obtained in the offline setting. Finally, we numerically evaluate the load variance as well as the information leakage rate, and characterize the trade-off between the privacy and cost. The operating points on this trade-off can be chosen based on user's preferences. We also investigate the impact of the RB capacity on the performance. Our main contributions can be summarized as follows:

- We consider the SM system illustrated in Fig. 1, and study the design of EM policies that aim at minimizing a joint privacy-cost objective.
- We formulate the optimal privacy-cost trade-off in the offline setting as a convex optimization problem. We identify the structure of the optimal solution, and provide a low-complexity backward water-filling algorithm for computing it.
- We solve the online optimization problem by first discretizing the continuous state space, and then applying DP. Alternatively, we provide an efficient heuristic algorithm that exploits the optimal offline algorithm to solve a particular subproblem constructed at each iteration.
- The information leakage rate between the user's demand profile and the SM readings is also evaluated. Comparison with the load variance indicate that the two privacy measures exhibit similar trends.
- Finally, the performances of the proposed offline and online EM policies are assessed through numerical simulations, using a real SM data set. The privacy-cost trade-off, and the impact of the RB capacity on this trade-off are analyzed for the proposed policies.

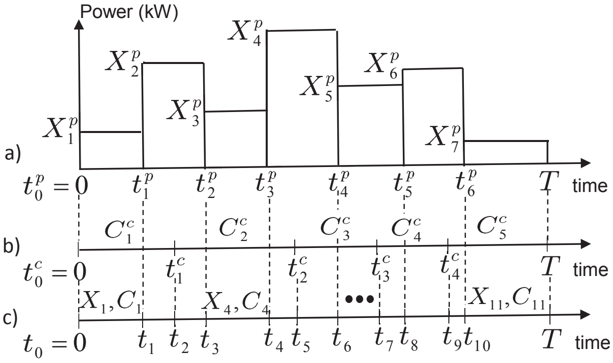

Fig. 2. Illustration of the timelines for the total power demand of the household, and the cost per unit energy. The total power demand changes at time instants $t_1^p, t_2^p, \ldots, t_6^p$, while the price of energy changes at time instants $t_1^c, t_2^c, t_3^c, t_4^c$.

The remainder of the paper is structured as follows. In Section II, we describe the system model. In Section III, we characterize the optimal offline EM policy, and provide the backward water-filling algorithm. The optimal and heuristic online EM policies are presented in Section IV. In Section V, we explain how to characterize the information leakage rate. In Section VI, extensive numerical results are presented. Finally, we conclude our paper in Section VII.

II. System Model

We consider a discrete-time power consumption model in a household (see Fig. 2(a)). In this model, each appliance consumes a constant power for an arbitrary duration when it is active. Appliances can be in the active or inactive state at any time. Let $t_0^p = 0 < t_1^p < \cdots < t_{(K-1)}^p < T$ be the time instants at which there is a change in the total power consumption. We denote the total power consumption within $[t_{(k-1)}^p, t_k^p]$ by X_k^p (kW) for $k = 1, \ldots, K$, where $X_k^p \geq 0$.

We also consider a time-varying electricity pricing model in which the cost per unit energy changes over time at certain time instants, and remains constant in between (see Fig. 2(b)). Let $t_0^c = 0 < t_1^c < \cdots < t_{(M-1)}^c < T$ be the time instants at which the cost of energy changes. We denote the cost per unit energy within $[t_{(m-1)}^c, t_m^c]$ by C_m^c (cent/kWh) for $m = 1, \ldots, M$. We can combine the time instants at which the power consumption or the cost per unit energy changes into a single time series $t_0 = 0 < t_1 < \cdots < t_{N-1} < t_N = T$ (see Fig. 2(c)). The duration of the TS between two consecutive time instants is denoted by $\tau_i \triangleq t_i - t_{i-1}$ (min), for $i = 1, \ldots, N$. TSs do not necessarily have the same duration. We denote the total power consumption and the cost per unit energy within TS i as X_i (kW) and C_i (cent/kWh), respectively, where $X_i \geq 0$. Note that for any two consecutive TSs, either the power demand X_i, or the cost C_i, or both may change, whereas they remain constant within a TS.

We study the SM model depicted in Fig. 1, where X_i (kW) denotes the aggregated real power consumption in TS i, while Y_i (kW) denotes the real power drawn from the grid. SM reports $\{Y_i\}_{i=1}^N$ to the UP without any tampering. We assume that Y_i remains constant within each TS. We show later that this assumption is without loss of optimality. Accordingly, there is no loss of information by the SM reporting its measurement once per TS. We integrate a RB with finite capacity B_{max} (kWh), and an EMU that manages the energy flow in the system. The EMU can use both the grid and the RB to satisfy the energy demand X_i, such that $Y_i = X_i - P_i$, where P_i (kW) is the power charged to ($P_i < 0$), or discharged from ($P_i > 0$) the RB in TS i, and $Y_i \in \mathbb{R}^+$, where $\mathbb{R}^+$ denotes the set of nonnegative real numbers. By constraining $Y_i \geq 0$ $\forall i$, we do not allow the user to sell his excess energy back to the grid.

We can argue that "perfect privacy" is achieved if the power withdrawn from the grid, Y_i, takes a constant value $\bar{E}$, $\forall i$. Ideally, if a user has a flat power demand from the grid, all individual appliance signatures are filtered out from the aggregated energy consumption data, and we can assume that the UP cannot learn anything about her energy consumption behaviour [13]. Accordingly, the privacy of an EM policy is measured by the *load variance*, defined as:

$$\mathcal{V} \triangleq \frac{1}{T}\sum_{i=1}^{N} \tau_i \cdot (Y_i - \bar{E})^2. \tag{1}$$

Perfect privacy is achieved when $\mathcal{V} = 0$. The target power demand $\bar{E}$ is a constant parameter in our model, which is selected by the user in advance[1].

The *average energy cost* of an EM policy is defined as:

$$\mathcal{C} \triangleq \frac{1}{T}\sum_{i=1}^{N} \tau_i \cdot Y_i \cdot C_i. \tag{2}$$

In our model all the energy demands must be satisfied at the time of request, i.e., outages or demand shifting are not allowed. Hence, assuming that the RB is empty at $t = 0$, Y_i have to satisfy the following cumulative constraints:

$$\sum_{j=1}^{i} \tau_j \cdot X_j \leq \sum_{j=1}^{i} \tau_j \cdot Y_j, \quad i = 1, \ldots, N, \tag{3}$$

which assure that a sufficient amount of energy is drawn from the grid to satisfy the energy demands of the appliances at each TS. We allow drawing more energy from the grid than that is requested by the appliances, which is then stored in the RB. Since the RB capacity is finite, the battery energy at TS i must satisfy:

$$B_i \triangleq \sum_{j=1}^{i} \tau_j \cdot (Y_j - X_j) \leq B_{max}, \quad i = 1, \ldots, N, \tag{4}$$

which assure that the difference between the cumulative energy drawn from the grid and the cumulative energy demand of the appliances up to each TS is not more than the capacity of the RB. This guarantees that the extra energy drawn from the grid can be stored for future use, and no energy is wasted due to battery overflows. The constraint in (4) assumes that energy cannot be drawn from the grid to be wasted for the sake of privacy. However, we do not constrain the final battery state to be empty; hence, more energy than requested by the appliances can be drawn to be left in the battery at the end of TS N.

Since both objective functions (1) and (2) are convex, and the constraints are linear, achievable pairs of $(\mathcal{V}, \mathcal{C})$ under constraints (3) and (4) form a convex region, and the optimal operating points are characterized by the Pareto boundary of this region [24]. Hence, we use the weighted average of $\mathcal{V}$ and $\mathcal{C}$ to identify all the points on the Pareto boundary [24]. The convex optimization problem can be written as:

$$\begin{aligned} &\min_{Y_i \geq 0} \sum_{i=1}^{N} \Big[\theta \cdot \tau_i \cdot \big(Y_i - \bar{E}\big)^2 + (1-\theta) \cdot \tau_i \cdot Y_i \cdot C_i\Big] \\ &\text{s.t.} \quad (3) \text{ and } (4), \end{aligned} \tag{5}$$

where $0 < \theta \leq 1$ is the parameter that adjusts the trade-off between the privacy and cost. The value of θ is set in advance by the user. If $\theta = 1$, the user is interested only in privacy; while if $\theta = 0$, only in the cost. Since the cost per unit energy and the user's total load remain constant over each TS, it follows from the convexity of the objective function in (5) that the optimal power drawn from the grid must remain constant within a TS [25]. Hence, the assumption of having the SM report only once per TS does not lead to any loss of information.

In Section III, we identify the *optimal offline EM policy* that minimizes (5), where all the demands and prices are known by the EMU in advance at $t_0 = 0$. While non-causal knowledge of the user's energy consumption may not be realistic for certain appliances, activity patterns of majority of appliances, such as refrigerators, heating, programmable washing machines and dish washers, electrical vehicles, are highly predictable during their operation periods [26]. Alternatively, we will study the online optimization in Section IV.

III. Optimal Offline EM Policy

To obtain the optimal offline EM policy for the problem in (5), we define the Lagrangian function:

$$\begin{aligned} \mathcal{L} = &\sum_{i=1}^{N} \Big[\theta\tau_i\big(Y_i - \bar{E}\big)^2 + (1-\theta)\tau_i Y_i C_i\Big] \\ &+ \sum_{i=1}^{N} \lambda_i\Big(\sum_{j=1}^{i} \tau_j (X_j - Y_j)\Big) \\ &+ \sum_{i=1}^{N} \mu_i\bigg(\Big(\sum_{j=1}^{i} \tau_j (Y_j - X_j)\Big) - B_{max}\bigg) - \sum_{i=1}^{N} v_i Y_i, \end{aligned} \tag{6}$$

where $\lambda_i \geq 0$, $\mu_i \geq 0$ and $v_i \geq 0$, $i = 1, \ldots, N$, are the Lagrange multipliers, and the complementary slackness conditions are:

$$\lambda_i\Big(\sum_{j=1}^{i} \tau_j (X_j - Y_j)\Big) = 0, \quad i = 1, \ldots, N, \tag{7}$$

[1] Our framework can be easily adapted to consider a time-varying target energy profile, $\{\bar{E}_i\}_{i=1}^N$. This more general model could allow the user to emulate a completely different energy consumption profile to confuse an intruder. In the paper we have not specified how the target value $\bar{E}$ is chosen by the user, and for the simulations we have considered $\bar{E}$ as the average power demand of the appliances. While we have observed through numerical simulations that, this value provides sufficient flexibility to the EMU, we think that the determination of the target value $\bar{E}$, or the target sequence $\{\bar{E}_i\}_{i=1}^N$, is an interesting future research problem.

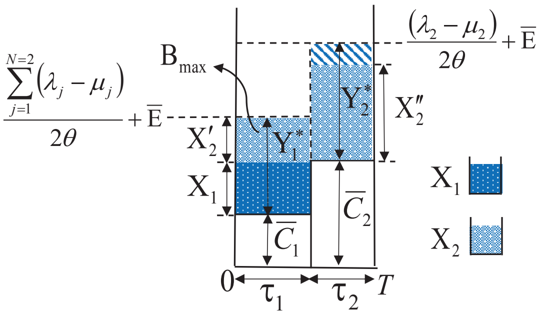
$\frac{\sum_{j=1}^{N=2}(\lambda_j-\mu_j)}{2\theta}+\bar{E}$
B_{max}
$\frac{(\lambda_2-\mu_2)}{2\theta}+\bar{E}$
X_2'
Y_1^*
Y_2^*
X_2''
X_1
$\bar{C}_1$
$\bar{C}_2$
X_1
X_2
0
τ_1
τ_2
T

Fig. 3. Illustration of an example of the optimal EM policy through water-filling for $N = 2$.

$$\mu_i\left(\left(\sum_{j=1}^{i}\tau_j(Y_j-X_j)\right)-B_{max}\right)=0,\quad i=1,\ldots,N, \tag{8}$$

$$v_iY_i=0,\quad i=1,\ldots,N. \tag{9}$$

Applying the Karush Kuhn Tucker (KKT) necessary conditions, we obtain, for $i = 1,\ldots,N$:

$$\frac{\partial\mathcal{L}}{\partial Y_i}=2\theta\tau_i\big(Y_i^*-\bar{E}\big)+(1-\theta)\tau_iC_i+\tau_i\sum_{j=i}^{N}(\mu_j-\lambda_j)-v_i,$$
$$=0. \tag{10}$$

Then the optimal values for Y_i are found as:

$$Y_i^*=\left[\left(\frac{\sum_{j=i}^{N}(\lambda_j-\mu_j)}{2\theta}+\bar{E}\right)-\bar{C}_i\right]^+,\quad\forall i, \tag{11}$$

where $[x]^+$ is equal to x if $x \geq 0$, and 0 otherwise, and the *weighted cost level*, $\bar{C}_i$ is defined as:

$$\bar{C}_i\triangleq\frac{(1-\theta)C_i}{2\theta},\quad\forall i. \tag{12}$$

We note that the solution in (11) resembles the classical water-filling solution [27], where $Y_i^*+\bar{C}_i$ corresponds to the *water level* in TS i, and one can interpret the optimal EM policy as pouring water over TSs. The classical water-filling solution is encountered in various problems in the literature, most notably, the optimal power allocation among parallel Gaussian channels under a sum power constraint so as to maximize the capacity [27]. In the classical water-filling solution, the water level is constant, and adjusted through a constant Lagrangian multiplier, which is chosen to satisfy the sum power constraint in the above example.

In our problem, “water” corresponds to the energy allocated to each TS. To clarify how it differs from the classical water-filling solution, and to provide some intuition for the constraints in (3) and (4), we next provide an example of the optimal EM policy through water-filling for $N = 2$ in Fig. 3. The heights of the white rectangles correspond to the weighted cost levels, $\bar{C}_i$'s, while their widths correspond to the TS durations, τ_i's, for $i = 1, 2$. The first power demand X_1 is given as the height of the corresponding filled area on top of the first white rectangle, while the second power demand X_2 is given as $X_2=\frac{\tau_1X_2'+\tau_2X_2''}{\tau_2}$. The optimal values of the power withdrawn from the grid, Y_i^*, are illustrated as the height of the filled areas above $\bar{C}_i$, leading to the water levels $Y_i^*+\bar{C}_i$, $\forall i$.

We can see in Fig. 3 that unlike the classical water-filling solution, in our model the water level does not have to be constant. Instead, it changes from one TS to the next. This is because we have multiple constraints in (3) and (4), which should be satisfied at each TS as opposed to the classical water-filling problem, in which there is a single constraint for TS N. As seen in Fig. 3, Y_1^* is sufficient to satisfy the first power demand X_1 as well as part of X_2, denoted by X_2', which is first stored in the RB. Following (11), the water level in the first TS is found to be $Y_1^*+\bar{C}_1=\left(\sum_{j=1}^{2}(\lambda_j-\mu_j)\right)/2\theta+\bar{E}$. Y_2^* satisfies part of demand X_2, i.e., X_2'', and the rest is stored in the RB, i.e., $Y_2^*-X_2''$. From (11), the water level in the second TS is given by $Y_2^*+\bar{C}_2=(\lambda_2-\mu_2)/2\theta+\bar{E}$, different from the water level in the first TS.

Next, we identify some properties of the optimal EM policy based on the KKT conditions in (7)-(10), which are both necessary and sufficient due to the convexity of the optimization problem in (5).

Lemma 1. *In the optimal EM policy, given $Y_i > 0$ $\forall i$, whenever the water level, i.e., $Y_i+\bar{C}_i$, increases (decreases) from TS i to TS $i+1$, i.e., $Y_i+\bar{C}_i < Y_{i+1}+\bar{C}_{i+1}$ ($Y_i+\bar{C}_i > Y_{i+1}+\bar{C}_{i+1}$), the RB must be full (empty) at TS i, i.e., $B_i = B_{max}$ ($B_i = 0$). Moreover, if the RB is neither empty nor full at TS i, i.e., $0 < B_i < B_{max}$, then the water level does not change from TS i to TS $i+1$, i.e., $Y_i+\bar{C}_i = Y_{i+1}+\bar{C}_{i+1}$.*

Proof. From the slackness conditions in (7) and (8), we can argue that the RB is full whenever $\lambda_i = 0$ and $\mu_i > 0$, and the RB is empty whenever $\lambda_i > 0$ and $\mu_i = 0$. Note that λ_i and μ_i cannot be positive simultaneously. From (11), we see that $Y_i+\bar{C}_i < Y_{i+1}+\bar{C}_{i+1}$ implies $\lambda_i = 0$ and $\mu_i > 0$, and $Y_i+\bar{C}_i > Y_{i+1}+\bar{C}_{i+1}$ implies $\lambda_i > 0$ and $\mu_i = 0$. Therefore, we can conclude that whenever the water level increases (decreases) from TS i to TS $i+1$, the RB must be full (empty) at TS i. Moreover, if the RB is neither empty nor full at TS i, i.e., $0 < B_i < B_{max}$, the i-th constraints in (3) and (4) are satisfied with strict inequality. This implies from the slackness conditions in (7) and (8) that $\lambda_i = 0$ and $\mu_i = 0$. From (11), we can conclude that, if the RB is neither empty nor full at TS i, the water level does not change from TS i to TS $i+1$, i.e., $Y_i+\bar{C}_i = Y_{i+1}+\bar{C}_{i+1}$. □

Lemma 2. *In the optimal EM policy, given $Y_i^* > 0$ $\forall i$, if the RB is never full from TS i to TS N, i.e., $B_j < B_{max}$ for $j = i, i+1, \ldots, N$, then the optimum water levels from TS i to TS N, i.e., $Y_j^*+\bar{C}_j$, for $j = i, i+1, \ldots, N$, must satisfy $Y_j^*+\bar{C}_j \geq \bar{E}$. If the RB is neither empty nor full from TS i to TS N, i.e., $0 < B_j < B_{max}$, for $j = i, i+1, \ldots, N$, then the optimum water levels from TS i to TS N should be equal to $\bar{E}$, i.e., $Y_j^*+\bar{C}_j = \bar{E}$, for $j = i, i+1, \ldots, N$.*

Proof. If the RB is never full from TS i to TS N, i.e., $B_j < B_{max}$ for $j = i, i{+}1, \ldots, N$, the constraints in (4) are satisfied with strict inequality. It follows from the slackness conditions in (8) that $\mu_j = 0$, for $j = i, i+1, \ldots, N$. From (11), this implies that $Y_j^* + \bar{C}_j \geq \bar{E}$, and we can conclude that, if the RB is never full from TS i to TS N, the optimum water levels from TS i to TS N should satisfy $Y_j^* + \bar{C}_j \geq \bar{E}$, for $j = i, i+1, \ldots, N$. If the RB is neither empty nor full from TS i to TS N, i.e., $0 < B_j < B_{max}$, for $j = i, i+1, \ldots, N$, the constraints in (3) and (4) are satisfied with strict inequality. It follows from (7) and (8) that $\lambda_j = \mu_j = 0$, for $j = i, i+1, \ldots, N$. From (11), this implies that $Y_j^* + \bar{C}_j = \bar{E}$, and we can conclude that, if the RB is neither empty nor full from TS i to TS N, the optimum water levels from TS i to TS N should be equal to $\bar{E}$, i.e., $Y_j^* + \bar{C}_j = \bar{E}$, for $j = i, i{+}1, \ldots, N$. □

A. Implications of the Lemmas

For clarity, we first consider the solution for an infinite RB. If B_{max} is infinite, the RB is never full and the constraints in (4) are never satisfied with equality, i.e., $\mu_i = 0$, $\forall i$. Then, it follows from Lemma 1 that the water level is monotonically decreasing from one TS to the next. This is because the water (energy) can only flow backwards in our model, i.e., energy requested at a TS can be requested from the grid over earlier TSs, but not the future ones. Accordingly, whenever the constraint in (3) is not satisfied with equality at TS i, i.e., $\lambda_i = 0$, then some energy for future use is drawn in advance within current TS i. Hence, in the optimal EM policy, if some drawn power is transferred from future TSs to the current one, the water level remains the same from the current TS to the next. Conversely, when there is a water level decrease from the current TS to the next, that is, if $\lambda_i > 0$, no drawn power is allocated from future TSs to the current, i.e., the RB is empty at TS i, as argued in Lemma 1. Moreover, from Lemma 2, we can conclude that all optimal water levels must satisfy $Y_i^* + \bar{C}_i \geq \bar{E}$, $\forall i$, since the RB is never full.

If B_{max} is finite, the amount of energy drawn for future use within TS i is limited by the remaining RB capacity at TS i, i.e., $B_{max} - B_i$. When the energy demand of future TSs that is requested from the grid in the current one is less than $B_{max} - B_i$, the constraints in (3) and (4) are satisfied with strict inequality, i.e., $\lambda_i = \mu_i = 0$, and the water level does not change from TS i to TS $i+1$, as argued in Lemma 1. Conversely, when there is a water level increase from TS i to TS $i+1$, that is, if $\lambda_i = 0$ and $\mu_i > 0$, the amount of energy demand for future TSs satisfied in the current one is equal to $B_{max} - B_i$, which implies that the RB is full at TS i. If the RB is full in the current TS, then no future energy demand can be satisfied in the current and previous TSs anymore due to the RB capacity limitation. When there is a water level decrease from TS i to TS $i+1$, that is, if $\lambda_i > 0$ and $\mu_i = 0$, no future energy demand is satisfied in the current TS, i.e., the RB is empty at TS i, as argued in Lemma 1. If the RB is never full from TS i to TS N, i.e., $B_j < B_{max}$ for $j = i, i+1, \ldots, N$, we can conclude from Lemma 2 that the optimum water levels from TS i to TS N, must satisfy $Y_j^* + \bar{C}_j \geq \bar{E}$, for $j = i, i+1, \ldots, N$.

B. Backward Water-Filling Algorithm

Based on the aforementioned implications of Lemma 1 and Lemma 2, we next describe the backward water-filling algorithm through an example in Fig. 4. The heights of the white rectangles correspond to the weighted cost levels, $\bar{C}_i$'s, while their widths correspond to the TS durations, τ_i's, for $i = 1, 2, 3$. We also fix a target consumption value $\bar{E}$ illustrated in Fig. 4. Fig. 4(a) depicts the power demands, X_i, as the heights of the filled areas on top of the white rectangles. Thus, the initial water levels are given by $X_i + \bar{C}_i$, $\forall i$. Observe that the RB is initially empty at every TS. Considering the example in Fig. 4(a), in Fig. 4(b) and Fig. 4(c) we illustrate the optimal offline EM policy in the presence of an infinite and a finite capacity RB, respectively.

In the infinite RB case, the first demand X_1 is satisfied from the grid within the first TS, as seen in the first plot in Fig. 4(b). The demand in the second TS, X_2, can be satisfied during the first and second TSs. The algorithm decides how much power to draw in the first and second TSs, applying the water-filling solution in (11), where we use $\bar{C}_1 + X_1$ instead of $\bar{C}_1$.

Since the electricity price is more expensive in the second TS, part of X_2 is drawn within the first TS, and stored in the RB (see the second plot in Fig. 4(b)). The rest of X_2 is drawn from the grid within the second TS. Hence, X_2 is fulfilled from both the RB and the grid. Observe that the RB is not empty at the end of the first TS; and hence, the water level does not change from the first TS to the second as argued in Lemma 1.

The demand in the third TS can be drawn from the grid in the first three TSs (see the third plot in Fig. 4(b)). Observe that the RB is not empty at the end of first and second TSs; and hence, all water levels are equalized as argued in Lemma 1. On the other hand, the RB is empty at the end of the third TS. If the current water levels satisfy the conditions argued in Lemma 2 in this step, the algorithm leads to the optimal solution. As depicted in the third plot in Fig. 4(b), all water levels are smaller than the target value $\bar{E}$; and hence, Lemma 2 is not satisfied. To remedy this, the algorithm allocates further grid energy to all three TSs. Accordingly, all water levels are raised up to $\bar{E}$ as seen in the fourth plot in Fig. 4(b), leading to the optimal values of the power drawn from the grid, Y_i^*, as the height of the filled areas above $\bar{C}_i$, $\forall i$. Observe that the optimal power withdrawn from the grid in the first TS, Y_1^*, depends on the user's demand and the weighted cost levels in the following TSs as well. For N TSs, the optimal power values withdrawn from the grid can be obtained by $N + 1$ iterations of the backward water-filling algorithm.

Fig. 4(c) depicts the optimal backward water-filling solution and the optimal values of the power withdrawn from the grid, Y_i^*, in the presence of a finite capacity RB. X_1 is satisfied from the grid within the first TS, and the RB is empty at the end of the first TS (see the first plot in Fig. 4(c)). In contrast to the infinite capacity RB case, the portion of the user's total load, X_2, drawn in advance within the first TS is limited by the RB capacity B_{max} (see the second plot in Fig. 4(c)). In other words, the part of the demand in the second TS that is drawn from the grid in the first TS is equal to B_{max}, and the

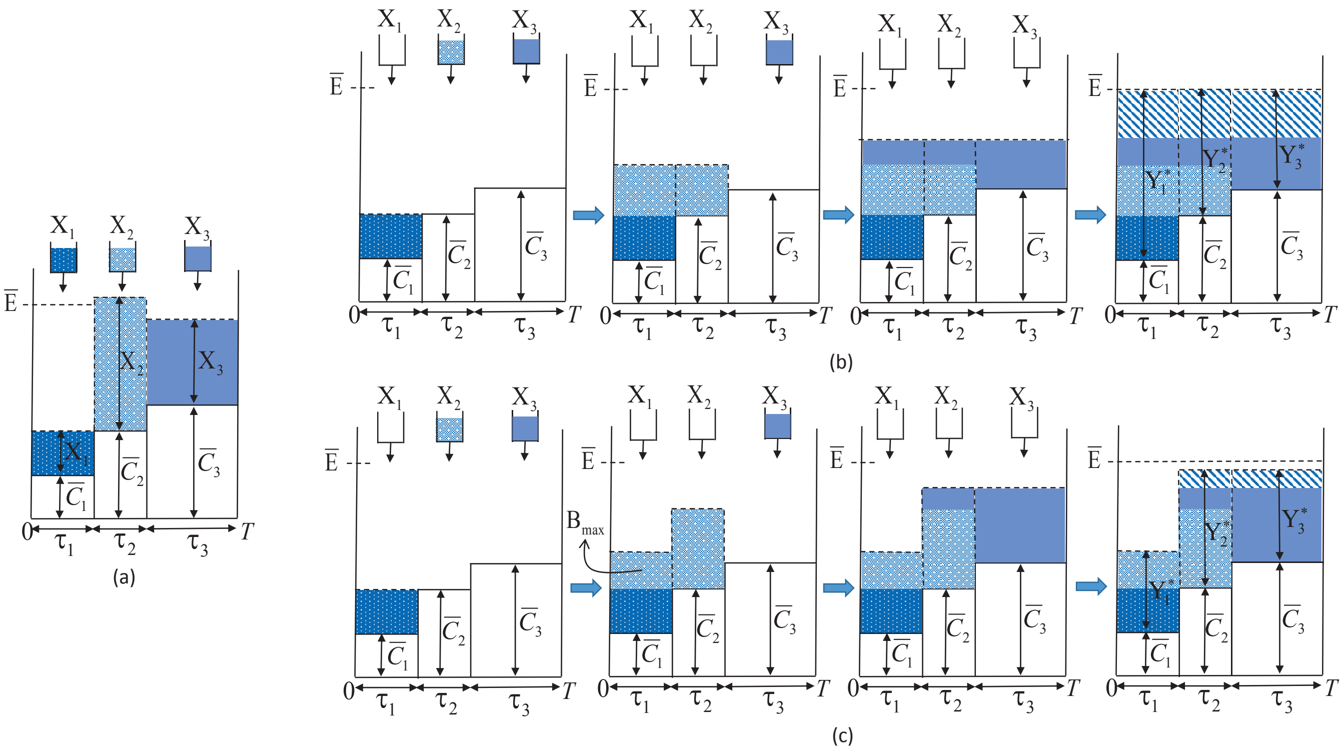

Fig. 4. Depiction of the user's total loads and initial water levels (a), and the optimal backward water-filling algorithm in the presence of (b) infinite, and (c) finite capacity RBs, respectively.

RB is full at the end of the first TS. This explains the water level increase from the first TS to the second as argued in Lemma 1. X_3 can only be satisfied over the second and third TSs since the RB is full at the end of the first TS (see the third plot in Fig. 4(c)). Part of X_3 is drawn within the second TS and stored in the RB. Observe that the RB is neither empty nor full at the end of the second TS; and hence, the water level does not change from the second TS to the third as argued in Lemma 1. Similarly to the infinite RB capacity case, if the current water levels satisfy the conditions argued in Lemma 2 in this step, the algorithm leads to the optimal solution. As seen in the third plot in Fig. 4(c), all water levels are smaller than $\bar{E}$. Observe that the RB is full at the end of the first TS; and hence, the water level at the first TS cannot be raised further due to the RB capacity limitation. On the other hand, the water levels at the second and third TSs do not satisfy the optimality conditions argued in Lemma 2 as the RB is neither empty nor full at the second TS, and empty at the third TS. Therefore, the algorithm needs to draw further grid energy in the second and third TSs. As depicted in the fourth plot in Fig. 4(c), the algorithm allocates the same amount of energy to the second and third TSs, and raises water levels, leading to the optimal values of the power withdrawn from the grid, Y_i^*, $\forall i$. Observe that the water levels at the second and third TSs are raised in accordance with Lemma 1 satisfying the RB capacity constraint. Since the RB is neither empty nor full at the end of the second TS, the water level does not change from the second TS to the third. The water levels at the second and third TSs do not reach $\bar{E}$, since the RB gets full at the end of the third TS.

IV. Online EM Policies

In this section, we consider causal (online) knowledge of the user's total load at the EMU, while the electricity prices are still known in advance[2]. First we provide the optimal online EM policy by solving the associated DP problem [28]. As DP quickly becomes computationally intractable with the increasing size of the state space of the problem, we also propose an efficient heuristic online policy that iteratively uses the offline backward water-filling algorithm developed in the previous section. For simplicity in this case, we assume unit TS durations, i.e., $\tau_i = 1$, $\forall i$. Similarly to the offline setting, we assume that the target value $\bar{E}$ is a constant parameter and is specified in advance.

A. *Optimal Online Policy*

The state of the system at the beginning of TS i is determined by the energy demand, $X_i \in \mathcal{X}$, and the battery state, $B_{i-1} \in \mathcal{B}$. While the energy demand and battery state are continuous bounded variables, we discretize both in order to apply DP. Hence, $\mathcal{X}$ and $\mathcal{B}$ are finite discrete sets generated by discretizing the feasible state spaces of the energy demand and battery state with particular energy quantizers, which are detailed in Section VI. Let $|\mathcal{X}|$ and $|\mathcal{B}|$ denote the cardinality of sets $\mathcal{X}$ and $\mathcal{B}$, respectively. We assume that the discrete energy demand follows a stationary first-order Markov relation, with transition probabilities q_{mn} between energy demand states x_m

[2] Non-causal knowledge of the energy prices by the consumers is realistic for the current smart grids, where the energy prices change very slowly. Morever, the DP formulation and the proposed heuristic algorithm can be easily extended to the scenario with causal knowledge of the energy prices.

and x_n, i.e., $q_{mn} = \Pr\{X_{i+1} = x_n | X_i = x_m\}$. The online EM policy at TS i, i.e., $\pi_i(X_i, B_{i-1})$, maps each state to a power value to be drawn from the grid, Y_i, that is selected from the finite discrete set $\mathcal{Y}_i$, i.e., $\pi_i : \mathcal{X} \times \mathcal{B} \rightarrow \mathcal{Y}_i$. The battery state at the end of TS i, B_i, is given by:

$$B_i = B_{i-1} + Y_i - X_i. \tag{13}$$

Following (13), $\mathcal{Y}_i$ can be defined as the set of feasible decisions for energy demand X_i and the battery state B_{i-1} at the beginning of TS i, i.e., $\mathcal{Y}_i = \{Y_i \in \mathbb{R}^+ | Y_i = B_i - B_{i-1} + X_i, B_i \in \mathcal{B}\}$. The EMU is not allowed to waste any energy by limiting the battery state B_i to be lower than B_{max}. Following the objective function in (5), we can write the cost function for decision Y_i as follows:

$$g_i(Y_i) \triangleq \Big[\theta \cdot \big(Y_i - \bar{E}\big)^2 + (1-\theta) \cdot Y_i \cdot C_i\Big]. \tag{14}$$

We aim at minimizing the average cost over N TSs. The optimal online policy is a collection of decision functions, i.e., $\pi^* = \{\pi_1^*, \pi_2^*, \ldots, \pi_N^*\}$, which leads to the optimal power values to be drawn from the grid $Y_i^* = \pi_i^*(X_i, B_{i-1})$, and is found as the solution to the following optimization problem:

$$\begin{aligned} &\min_{\pi_1,\ldots,\pi_N} \sum_{i=1}^{N} \mathrm{E}\Big[g_i(\pi_i(X_i, B_{i-1}))\Big] \\ &\text{s.t. } \pi_i(X_i, B_{i-1}) \geq 0,\ i = 1, \ldots, N, \\ &0 \leq B_{i-1} + \pi_i(X_i, B_{i-1}) - X_i \leq B_{max},\ i = 1, \ldots, N, \end{aligned} \tag{15}$$

where the expectation is taken with respect to the statistics of the user's demand distribution. The optimal online policy, $\pi_i^*(X_i, B_{i-1})$, can be obtained through DP by proceeding backwards from the N-th TS to the first, namely, backward induction, as follows:

$$\begin{aligned} J_N^*(X_N, B_{N-1}) &\triangleq \min_{Y_N \in \pi_N(X_N, B_{N-1})} g_N(Y_N), \\ J_i^*(X_i, B_{i-1}) &\triangleq \min_{Y_i \in \pi_i(X_i, B_{i-1})} \mathrm{E}\Big[g_i(Y_i) + J_{i+1}^*(X_{i+1}, B_i)\Big], \\ &= \min_{Y_i} \Big\{g_i(Y_i) + \sum_n q_{mn} J_{i+1}^*(x_n, B_{i-1} + Y_i - x_m)\Big\}, \\ &\qquad i = N-1, \ldots, 1, \end{aligned} \tag{16}$$

where J_i^* denotes the optimal cost function at TS i that assigns to the energy demand, X_i, and the battery state, B_{i-1}, the minimum cost $J_i^*(X_i, B_{i-1})$. We recursively solve (16) to obtain the optimal policy $\pi_i^*(X_i, B_{i-1})$, $\forall i$. The EMU records this function as a $|\mathcal{X}| \times |\mathcal{B}|$ look-up table. Each entry of the table corresponds to the optimal decision for the power withdrawn from the grid Y_i^*, for state (X_i, B_{i-1}). At the end of backward induction, we obtain a look-up table for each TS i. Then, proceeding forwards from the first TS to the N-th, namely, through forward induction, and using the corresponding look-up tables, the optimal power withdrawn from the grid sequence can be obtained for a particular energy demand realization.

Algorithm 1 Heuristic Online Policy

```
$B_0 \leftarrow 0$                              ▷ Initially battery is empty
for i = 1 to N do                                ▷ TS $i$
    1. Subproblem Construction:
    Set the power demands for two TSs
    $\hat{X}_1 \leftarrow [X_i - B_{i-1}]^+$, $\hat{X}_2 \leftarrow 3\bar{E}$
    Set the battery energies for two TSs
    $\hat{B}_1 \leftarrow [B_{i-1} - X_i]^+$, $\hat{B}_2 \leftarrow \hat{B}_1$
    Set the electricity prices for two TSs
    $\hat{C}_1 \leftarrow C_i$, $\hat{C}_2 \leftarrow \frac{1}{N}\sum_{i=1}^{N} C_i$
    2. Subproblem Solution:
    Solve the constructed subproblem by using the backward water-filling algorithm.
    Feed the optimal power withdrawn from the grid, $\hat{Y}_1^*$, into the real timeline.
    3. Power Decision:
    $Y_i \leftarrow \hat{Y}_1^*$        ▷ Set the power drawn from the grid at TS $i$
    $B_i \leftarrow B_{i-1} + (Y_i - X_i)$        ▷ Update the battery energy
end for
```

Observe that the size of the look-up tables, and equivalently, the complexity of the DP algorithm, grows very quickly with the increasing number of states, which depends on the quantizer precision. Hence, DP algorithms can easily become computationally intractable with the increasing quantizer precision, which is needed to approximate the optimal solution for the original problem with a continuous state space [28].

B. Heuristic Online Policy

Due to the high computational complexity of DP solutions, here we propose a low complexity heuristic online algorithm, detailed in Algorithm 1. At each TS i, this algorithm creates a two-TS subproblem. Accordingly, each subproblem consists of the power demands, the electricity prices and the battery states for two TSs, which are denoted as $(\hat{X}_1, \hat{X}_2)$, $(\hat{C}_1, \hat{C}_2)$ and $(\hat{B}_1, \hat{B}_2)$, respectively. In each subproblem, the first TS is representative for the past and present information, while the second TS is representative for future information. The parameters for the first TS of the subproblem, i.e., $\hat{X}_1$, $\hat{C}_1$, $\hat{B}_1$, are set based on the current information available at the EMU, such as, the current power demand, X_i, the current electricity price, C_i, and the battery state, B_{i-1}. The algorithm sets $\hat{X}_1$ as the part of the current power demand, X_i, which can not be satisfied from the available energy in the battery, $[X_i - B_{i-1}]^+$, $\hat{B}_1$ as the remaining energy in the battery after satisfying part of the current power demand, $[B_{i-1} - X_i]^+$, and $\hat{C}_1$ as the current electricity price, C_i. The parameters for the second TS of the subproblem, i.e., $\hat{X}_2$, $\hat{C}_2$, $\hat{B}_2$, are set as follows. The algorithm sets $\hat{X}_2$ as three times the target power demand[3] $\bar{E}$, $\hat{C}_2$ as the mean electricity price, and $\hat{B}_2$ as $\hat{B}_1$. At each step, the algorithm optimally solves the constructed subproblem using the backward water-filling algorithm developed in Section III-B. The power values arising from the optimal solution for the first and second TSs are

[3] We set $\hat{X}_2$ more than the target power demand $\bar{E}$ in order to compensate for future rise in demand. This allows the algorithm to charge the RB further, so that a possible peak in demand in future TSs can be tackled without diverging much from the target $\bar{E}$.

denoted by $\hat{Y}_1^*$ and $\hat{Y}_2^*$, respectively. The algorithm is only interested in the optimal solution for the first TS, i.e., $\hat{Y}_1^*$. Therefore, the algorithm sets the decision for the power to be withdrawn from the grid Y_i at TS i as $\hat{Y}_1^*$. Finally, it updates the battery state, B_i, by using B_{i-1}, X_i and Y_i.

The heuristic policy mimics the backward water-filling algorithm in an online setting. Moreover, the choice of a two-TS subproblem at each iteration allows to reduce the complexity of the heuristic algorithm, and there is no need to quantize the state space. While, mimicking the optimal offline policy as mentioned above is a reasonable and low-complexity heuristic to replace the optimal DP solution, supported also by the numerical results in Section VI, we are not able to provide any theoretical performance guarantees.

V. Information leakage rate

In the previous sections, we have considered the load variance as the privacy measure. An alternative information theoretic privacy measure is the information leakage rate [10], which is defined as the average mutual information between the sequences of the user's total load and the power drawn from the grid:

$$I_p \triangleq \frac{1}{N} I(X^N; Y^N). \tag{17}$$

The information leakage rate can be argued to be a more accurate privacy measure as it takes into account the statistical behaviour of the user's total load. Note that the information leakage rate measures the reduction in the UP's uncertainty (entropy) about user's energy consumption, X^N, after receiving meter readings, Y^N, as we have $I_p = \frac{1}{N}\left[H(X^N) - H(X^N|Y^N)\right]$. As an information theoretic privacy measure, the information leakage rate provides privacy guarantees regardless of the computational power of the attacker. However, the optimal decision policy in terms of the information leakage rate is significantly harder to characterize [21]. Moreover, even when $I_p = 0$, for large N, it is possible that the SM readings completely reveal certain portions of the energy consumption profile. Here, we first provide a computational expression for the information leakage rate. In the next section, we will numerically evaluate and compare the load variance and the information leakage rate privacy measures, and demonstrate that the two follow similar trends.

As a first step towards computing the information leakage rate, we quantize the vectors of the user's demand and the power drawn from the grid. Let $\tilde{X}^N$ and $\tilde{Y}^N$ denote the quantized versions of X^N and Y^N, respectively. The samples of $\tilde{X}^N$ and $\tilde{Y}^N$ take values from finite discrete sets $\tilde{\mathcal{X}}$ and $\tilde{\mathcal{Y}}$, respectively. For simplicity, we assume that the samples of $\tilde{X}^N$ and the joint samples of $(\tilde{X}^N, \tilde{Y}^N)$ follow stationary first-order Markov distributions. Then we can write the distributions of $\tilde{X}^N$ and $(\tilde{X}^N, \tilde{Y}^N)$ as follows:

$$p(\tilde{X}^N) = p(\tilde{X}_1) \prod_{i=2}^{N} p(\tilde{X}_i | \tilde{X}_{i-1}), \tag{18a}$$

$$p(\tilde{X}^N, \tilde{Y}^N) = p(\tilde{X}_1, \tilde{Y}_1) \prod_{i=2}^{N} p(\tilde{X}_i, \tilde{Y}_i | \tilde{X}_{i-1}, \tilde{Y}_{i-1}). \tag{18b}$$

Note that the condition (18b) holds when (18a) holds, and the power withdrawn from the grid at TS i, Y_i, depends only on the current load X_i, and the previous values of the load and the power withdrawn from the grid, (X_{i-1}, Y_{i-1}). Under these assumptions, we derive an upper bound on the information leakage rate, I_p, as shown in Appendix. The obtained information leakage rate upper bound will be evaluated numerically as a measure of information theoretical privacy leakage for the EM policies derived in Section III and Section IV.

We note that the condition (b) in Appendix holds with equality if we assume that the sequence of the power withdrawn from the grid, $\tilde{Y}^N$, is also a stationary first-order Markov process. This assumption has been made in [15] for the computation of the information leakage rate; however, adding this extra Markov assumption together with the initial ones may not lead to any realistic model or non-trivial EM strategy.

To numerically evaluate the mutual information expressions for given sequences of the user's total load and the power drawn from the grid, we can explicitly write the information leakage rate I_p, as follows:

$$\frac{1}{N}\left(\sum_{i=2}^{N} \sum_{\substack{\tilde{x}_{i-1} \in \tilde{\mathcal{X}} \\ \tilde{y}_{i-1} \in \tilde{\mathcal{Y}}}} \sum_{\substack{\tilde{x}_{i} \in \tilde{\mathcal{X}} \\ \tilde{y}_{i} \in \tilde{\mathcal{Y}}}} p(\tilde{x}_i, \tilde{x}_{i-1}, \tilde{y}_i, \tilde{y}_{i-1}) \log \tfrac{p(\tilde{x}_i, \tilde{x}_{i-1}, \tilde{y}_i, \tilde{y}_{i-1})}{p(\tilde{x}_i, \tilde{x}_{i-1}) p(\tilde{y}_i, \tilde{y}_{i-1})} \right.$$
$$\left. - \sum_{i=3}^{N} \sum_{\substack{\tilde{x}_{i-1} \in \tilde{\mathcal{X}} \\ \tilde{y}_{i-1} \in \tilde{\mathcal{Y}}}} p(\tilde{x}_{i-1}, \tilde{y}_{i-1}) \log \frac{p(\tilde{x}_{i-1}, \tilde{y}_{i-1})}{p(\tilde{x}_{i-1}) p(\tilde{y}_{i-1})} \right). \tag{19}$$

We can compute I_p by estimating all the joint and marginal distributions in (19). We use empirical distributions as the estimates for these distributions, i.e., we count the number of joint or single appearances over all realizations, and normalize them to obtain the corresponding probabilities.

VI. Numerical results

In this section, we provide further insights into the proposed offline and online EM policies through numerical simulations. We analyze the trade-off between the user's privacy and energy cost as well as the effect of the RB capacity on this trade-off. We consider the real SM readings obtained from [29] with a time resolution on the order of three seconds. For our simulations we convert the readings obtained from one household for a period of one month to a time resolution of one-minute, and use as the load profile. To be consistent with our power consumption model, we assume that the discrete time instants in Fig. 2(a) correspond to the sampling times of the SM. We set the electricity price in our simulations based on the real pricing tariffs [30]: the off-peak price is 5 cent per kWh during 00:00 to 12:00, the on-peak price is 20 cent per kWh during 12:00 to 20:00, and the medium-peak price is 10 cent per kWh during 20:00 to 00:00. For the simulations, we consider the target value $\bar{E}$ as the average power demand of the user, i.e., $\bar{E} = \frac{1}{T} \sum_{i=1}^{N} \tau_i \cdot X_i$, which is assumed to be known in advance. To discretize the state space for the online problem, we use a 4-bit non-uniform mu-law quantizer [31] for the

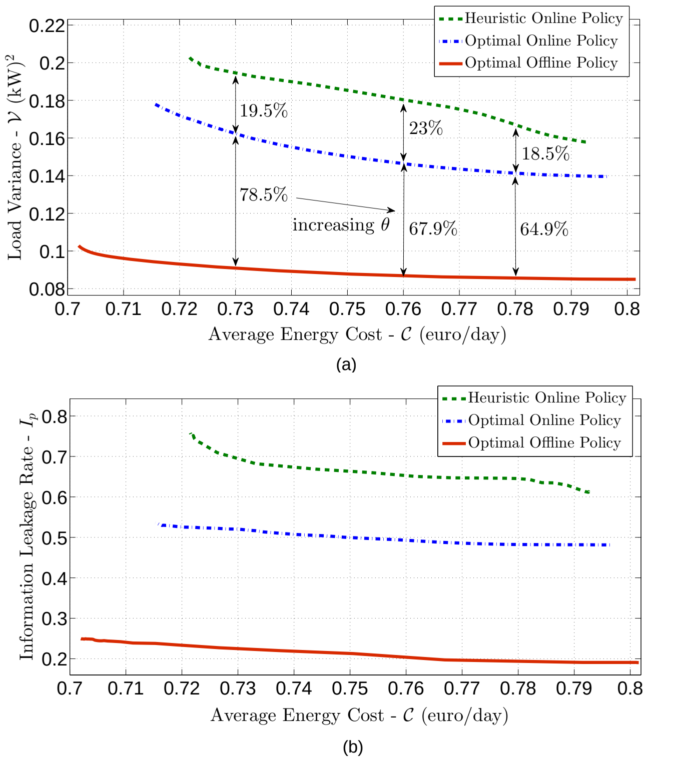

(a)

(b)

Fig. 5. (a) The load variance, $\mathcal{V}$, versus the average energy cost, $\mathcal{C}$, and (b) the information leakage rate, I_p, versus the average energy cost, $\mathcal{C}$, resulting from the proposed offline and online EM policies under the RB capacity, $B_{max} = 0.5$ kWh.

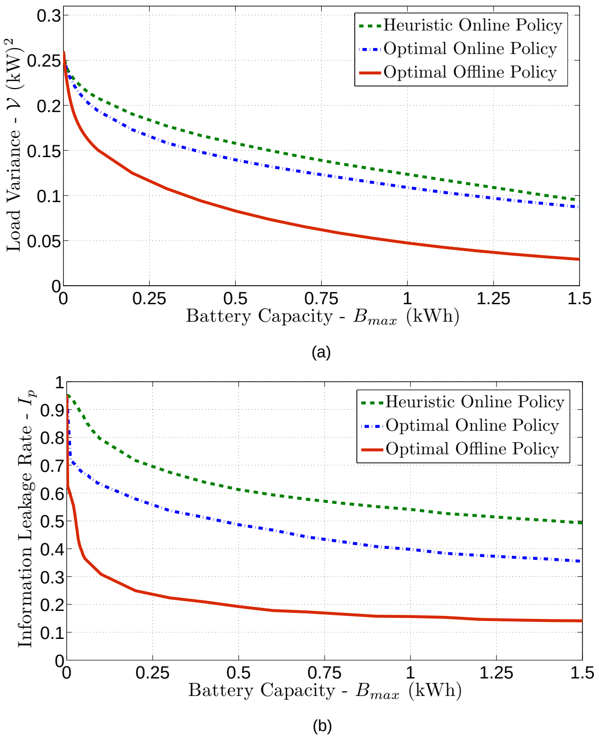

(a)

(b)

Fig. 6. (a) The load variance, $\mathcal{V}$, versus battery capacity, B_{max}, and (b) the information leakage rate, I_p, versus battery capacity, B_{max}, for the proposed offline and online EM policies under $\theta = 1$.

energy demand, and a 2-bit uniform quantizer for the battery state, respectively. For the characterization of the information leakage rate, we discretize the user's total load sequence and the sequence of the power drawn from the grid that results from the proposed policies by using a 5-bit non-uniform mu-law quantizer. Note that due to the non-uniformly distributed characteristics of the readings used in the simulations, the non-uniform mu-law quantizer allows us to reduce the quantization noise [31].

TABLE I
CORNER POINTS OF THE TRADE-OFF CURVES IN FIG. 5

	Heuristic Policy	Optimal Online Policy	Optimal Offline Policy
$\min \mathcal{V}$	0.157	0.139	0.085
$\min I_p$	0.612	0.481	0.19
$\mathcal{C}$	0.792	0.796	0.801
$\mathcal{V}$	0.204	0.178	0.103
I_p	0.758	0.536	0.249
$\min \mathcal{C}$	0.721	0.715	0.702

In Fig. 5, we illustrate the trade-offs between the privacy and cost resulting from the proposed EM policies with a RB capacity $B_{max} = 0.5$ kWh. The Pareto optimal trade-off curves between the load variance, $\mathcal{V}$, and the average energy cost, $\mathcal{C}$, in Fig. 5(a), and the trade-off curves between the information leakage rate, I_p, and the average energy cost, $\mathcal{C}$, in Fig. 5(b), are formed by varying θ values. For all the proposed policies, the average energy cost increases, while the load variance and the information leakage rate diminish as θ increases. According to the user preferences or requirements of the system, the operating point can be chosen anywhere on the trade-off curve. We observe that the load variance and the information leakage rate behave similarly for all the policies. Based on this observation, we can argue that the load variance can be used as a reliable privacy measure for SM systems. The corner points of the trade-off curves for the proposed policies in Fig. 5 are given in Table I. Observe that the heuristic online policy performs close to the optimal online policy both at the maximum privacy and the minimum cost corner points, while the optimal offline policy outperforms both of them as expected. Observe in Fig. 5(a) that for average energy costs $\mathcal{C} = \{0.73, 0.76, 0.78\}$ (euro/day), reductions in the load variance from the heuristic to the optimal online policy are found to be 19.5%, 23% and 18.5%, respectively, while reductions in the load variance from the optimal online policy to the optimal offline policy are found to be 78.5%, 67.9% and 64.9%, respectively.

Next, we investigate the effect of the battery capacity on the maximum privacy and the minimum cost. We plot the load variance, $\mathcal{V}$, versus B_{max} in Fig. 6(a), and the information leakage rate, I_p, versus B_{max} in Fig. 6(b) for $\theta = 1$. Observe that both the load variance and the information leakage rate diminish as RB capacity increases. Similar behaviours of the load variance and the information leakage rate with respect to the RB capacity further consolidates the argument that the load

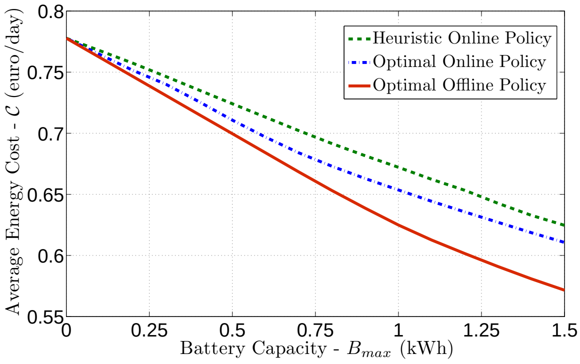

Fig. 7. The average energy cost, $\mathcal{C}$, versus battery capacity, B_{max}, resulting from the proposed offline and online EM policies under $\theta = 0.001$.

variance can be used as a proxy for the information leakage rate in SM systems. When there is no RB in the system, i.e., $B_{max} = 0$, the proposed offline and online policies behave identically. This also corresponds to a naive algorithm that ignores the RB, and draws the necessary energy to satisfy the demand at each TS i from the grid, i.e, $Y_i^* = X_i$, $\forall i$. In this case, the UP knows the user's total load sequence perfectly, and the load variance is found to be $\mathcal{V} = 0.259$, while the information leakage rate reduces to the entropy rate of the user's total load sequence, which is found to be $I_p = 0.952$.

Observe that the information leakage rates achieved by the optimal offline and online policies drop very quickly with even a small RB capacity. While the information leakage rate achieved by the optimal offline policy quickly saturates to its minimum value, those achieved by the optimal and heuristic online policies decrease smoothly with the RB capacity. Observe that the heuristic online policy performs reasonably close the optimal online policy for both privacy measures. The gain on the performances of the proposed policies can be achieved by virtue of the degree-of-freedom provided by the RB. For example, the heuristic online policy outperforms the naive algorithm that does not exploit the battery for any non-zero capacity. A battery with capacity 1 (kWh) provides 52.4% reduction in the load variance of the heuristic online policy, and 58% and 81.7% reductions in that of the optimal online and offline policies, respectively. When $B_{max} = 1.5$ (kWh), the information leakage rate of the heuristic online policy is found to be $I_p = 0.49$, and that of the optimal online and offline policies are found to be $I_p = 0.354$ and $I_p = 0.141$, respectively. These results show that a moderate RB capacity leads to a significant reduction in the information leakage rate. For RB capacities beyond 1.5 kWh, we do not expect a significant privacy gain. We also expect that the information leakage rate of the heuristic policy approaches to the optimal online policy as the RB capacity becomes sufficiently large.

Fig. 7 illustrates the average cost, $\mathcal{C}$, versus B_{max} for $\theta = 0.001$, which corresponds to the scenario in which the consumer is more interested in minimizing the cost of energy rather than privacy. The highest value for the average energy cost is $\mathcal{C} = 0.778$ (euro/day), achieved for $B_{max} = 0$. The heuristic online policy again outperforms the naive algorithm for any RB capacity, and performs very close to the optimal online policy. A battery with capacity 1 (kWh) provides 13.5% reduction in the average energy cost of the heuristic online policy, and 15.9% and 19.6% reductions in that of the optimal online and offline policies, respectively. When $B_{max} = 1.5$ (kWh), the average energy cost of the heuristic online policy is found to be $\mathcal{C} = 0.624$ (euro/day), and that of the optimal online and offline policies are found to be $\mathcal{C} = 0.61$ (euro/day) and $\mathcal{C} = 0.57$ (euro/day), respectively. We see that the user can reduce his/her energy consumption cost significantly with the proposed policies in the presence of a moderate capacity RB.

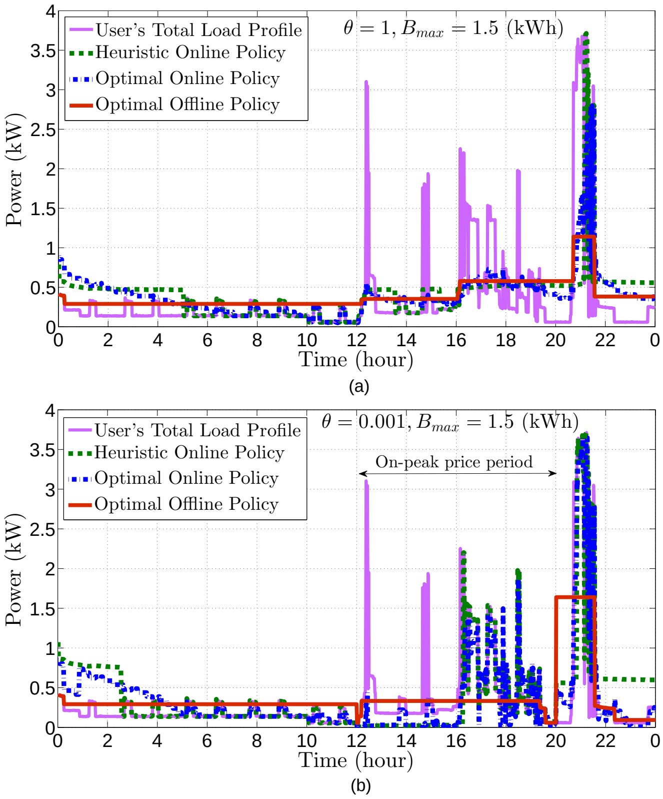

Fig. 8. Comparison of the user's total load profile with the profile of the power withdrawn from the grid resulting from the proposed offline and online EM policies with a RB of capacity, $B_{max} = 1.5$ kWh, for (a) $\theta = 1$, (b) $\theta = 0.001$, respectively.

Finally, we compare the user's total load profile with the profile of the power withdrawn from the grid resulting from the proposed offline and online EM policies with $B_{max} = 1.5$ kWh, for $\theta = 1$ and $\theta = 0.001$, in Fig. 8(a) and (b), respectively. When $\theta = 1$, the goal is to maximize the privacy, that is, to generate a smooth power profile for the energy withdrawn from the grid. Observe in Fig. 8(a) that the optimal offline policy generates quite a smooth profile showing off most of the peaks in the demand profile. Particularly, if we focus on the peak power of the original load profile between 20.00 and 22.00, we can see that the optimal offline policy masks most of the peak signal, while the optimal and heuristic online policies still retain significant peaks. On the other hand, they both perform well in masking the peak values at other times of the day.

When $\theta = 0.001$, the proposed policies intend to minimize the energy cost of the user. As seen in Fig. 8(b), the proposed policies store extra energy in the RB during the off-peak price period, and satisfy the demand of the peak period from the RB in order to reduce the cost. In the peak period between 12.00 and 20.00, the optimal offline policy draws nearly constant power from the grid, and satisfies the rest of the demand from the RB; on the other hand, the optimal and heuristic online policies satisfy the demand more from the RB between 12.00 and 16.00, and more from the grid between 16.00 and 20.00. We can envision that as the RB capacity increases, the optimal and heuristic online policies can store more energy in the battery to be used in the peak period, which would reduce the average costs.

VII. Conclusions

We studied demand-side EM policies from a privacy-energy cost trade-off perspective for a SM system with a finite-capacity energy storage unit. We considered a discrete-time energy consumption model, in which both the power consumption of the consumer and the electricity prices vary over time. We considered the variance of the power withdrawn from the grid around a predetermined constant target value as a measure of privacy for the consumer. First, assuming that the user's energy demand profile and the electricity prices are known non-causally, we formulated the optimal offline privacy-cost trade-off as a convex optimization problem, and characterized various properties of the optimal policy. Then, we proposed a low-complexity backward water-filling algorithm which efficiently computes the optimal offline EM policy.

Next, assuming that the user's power consumption profile is known only causally, we characterized the optimal online policy using DP. We also proposed a low-complexity heuristic online algorithm, and showed through numerical simulations that it performs close to the optimal online solution. In addition to the load variance, we also characterized the information leakage rate between the sequences of the user's total load and the power drawn from the grid.

Extensive numerical simulations have been presented using real SM consumption data to illustrate the trade-offs between privacy and energy cost resulting from the proposed offline and online policies. Our results indicate that the privacy-cost trade-offs for the load variance and the information leakage rate have similar behaviours. We also showed that most of the privacy gains can be obtained with a relatively small capacity RB.

As future extensions of the current work, various practical issues, such as storage inefficiencies, battery leakages and peak power constraints, can be considered with a slightly more complex model and analysis. Also, in this work, we consider SMs that only report the real power consumption of the user; whereas, SMs can report other relevant information, such as the reactive power, the power factor or various harmonics, which can also be used to make deductions about the users' energy consumption behaviours. These variables can be included in the analysis. Another interesting extension would be to allow the user to sell the surplus energy back to the UP, and assess its impact on the achieved privacy-cost trade-off.

Appendix

Here we derive an upper bound on the information leakage rate.

$$
\begin{aligned}
I_p &= \frac{1}{N} I(\tilde{X}^N; \tilde{Y}^N), \\
&= \frac{1}{N}\Big(H(\tilde{X}^N) + H(\tilde{Y}^N) - H(\tilde{X}^N, \tilde{Y}^N) \Big), \\
&\overset{(a)}{=} \frac{1}{N} \sum_{i=1}^{N} \Big(H(\tilde{X}_i|\tilde{X}^{i-1}) + H(\tilde{Y}_i|\tilde{Y}^{i-1}) \\
&\qquad\qquad - H(\tilde{X}_i, \tilde{Y}_i|\tilde{X}^{i-1}, \tilde{Y}^{i-1}) \Big), \\
&\overset{(b)}{\leq} \frac{1}{N} \sum_{i=1}^{N} \Big(H(\tilde{X}_i|\tilde{X}_{i-1}) + H(\tilde{Y}_i|\tilde{Y}_{i-1}) \\
&\qquad\qquad - H(\tilde{X}_i, \tilde{Y}_i|\tilde{X}_{i-1}, \tilde{Y}_{i-1}) \Big), \\
&\overset{(c)}{=} \frac{1}{N} \sum_{i=1}^{N} \Big(H(\tilde{X}_i|\tilde{X}_{i-1}) - H(\tilde{X}_i, \tilde{X}_{i-1}|\tilde{Y}_i, \tilde{Y}_{i-1}) \\
&\qquad\qquad + H(\tilde{X}_{i-1}|\tilde{Y}_{i-1}) \Big), \\
&\overset{(d)}{=} \frac{1}{N} \Bigg(\sum_{i=1}^{N} \Big(H(\tilde{X}_i, \tilde{X}_{i-1}) - H(\tilde{X}_i, \tilde{X}_{i-1}|\tilde{Y}_i, \tilde{Y}_{i-1}) \Big) \\
&\qquad - \sum_{i=1}^{N} \Big(H(\tilde{X}_{i-1}) - H(\tilde{X}_{i-1}|\tilde{Y}_{i-1}) \Big) \Bigg), \\
&\overset{(e)}{=} \frac{1}{N} \Bigg(\Big(H(\tilde{X}_1) - H(\tilde{X}_1|\tilde{Y}_1) \Big) \\
&\qquad + \sum_{i=2}^{N} \Big(H(\tilde{X}_i, \tilde{X}_{i-1}) - H(\tilde{X}_i, \tilde{X}_{i-1}|\tilde{Y}_i, \tilde{Y}_{i-1}) \Big) \\
&\qquad - \Big(H(\tilde{X}_1) - H(\tilde{X}_1|\tilde{Y}_1) \Big) \\
&\qquad - \sum_{i=3}^{N} \Big(H(\tilde{X}_{i-1}) - H(\tilde{X}_{i-1}|\tilde{Y}_{i-1}) \Big) \Bigg), \\
&\overset{(f)}{=} \frac{1}{N} \Bigg(\sum_{i=2}^{N} I(\tilde{X}_i, \tilde{X}_{i-1}; \tilde{Y}_i, \tilde{Y}_{i-1}) - \sum_{i=3}^{N} I(\tilde{X}_{i-1}; \tilde{Y}_{i-1}) \Bigg),
\end{aligned}
$$

where (a) follows from the chain rule of entropy; (b) follows from the first-order Markov assumption for $\tilde{X}_i$ and $(\tilde{X}_i, \tilde{Y}_i)$ in (18) and the fact that conditioning reduces entropy; (c), (d) and (e) also follow from the chain rule of entropy, and applying the necessary cancellations; and (f) follows from the definition of the mutual information.

References

[1] A. Ipakchi and F. Albuyeh, "Grid of the future," *IEEE Power Energy Mag.*, vol. 7, no. 2, pp. 52–62, Mar.-Apr. 2009.

[2] P. McDaniel and S. McLaughlin, "Security and privacy challenges in the smart grid," *IEEE Security Privacy*, vol. 7, no. 3, pp. 75–77, May-Jun. 2009.

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

[4] G. Kalogridis and S. Z. Denic, "Data mining and privacy of personal behaviour types in smart grid," in *Proc. IEEE Int. Conf. Data Mining Wkshp.*, Vancouver, Canada, Dec. 2011, pp. 636–642.

[5] A. Molina-Markham, P. Shenoy, K. Fu, E. Cecchet, and D. Irwin, “Private memoirs of a smart meter,” in *Proc. 2nd ACM Workshop Embedded Sens. Syst. Energy-Efficiency Building (BuildSys)*, Zurich, Switzerland, 2010, pp. 61–66.

[6] U. Greveler, P. Glosekotter, B. Justus, and D. Loehr, “Multimedia content identification through smart meter power usage profiles,” in *Proc. Int. Conf. Inform. and Knowledge Eng.*, Las Vegas, NV, July 2012.

[7] L. Sankar, S. R. Rajagopalan, S. Mohajer, and H. V. Poor, “Smart meter privacy: A theoretical framework,” *IEEE Trans. Smart Grid*, vol. 4, no. 2, pp. 837–846, Jun. 2013.

[8] 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. Smart Grid*, vol. 3, no. 3, pp. 1317–1324, Sep. 2012.

[9] J.-M. Bohli, C. Sorge, and O. Ugus, “A privacy model for smart metering,” in *Proc. IEEE Int. Comm. Conf.*, Capetown, South Africa, May 2010, pp. 1–5.

[10] O. Tan, D. Gündüz, and H. V. Poor, “Increasing smart meter privacy through energy harvesting and storage devices,” *IEEE Journal on Selected Areas in Comm. (J-SAC)*, vol. 31, no. 7, pp. 1331–1341, Jul. 2013.

[11] D. Gündüz, J. Gómez-Vilardebó, O. Tan, and H. V. Poor, “Information theoretic privacy for smart meters,” in *Proc. Inform. Theory and Applications Wkshp. (ITA)*, San Diego, CA, Feb. 2013, pp. 1–7.

[12] O. Tan, D. Gündüz, and J. Gómez-Vilardebó, “Optimal privacy-cost trade-off in demand-side management with storage,” in *Proc. IEEE Int. Wrkshp. Signal Process. Advances in Wireless Comm. (SPAWC)*, Stockholm, Sweden, Jun.-Jul. 2015, pp. 370–374.

[13] G. Kalogridis, C. Efthymiou, S. Denic, T. A. Lewis, and R. Cepeda, “Privacy for smart meters: Towards undetectable appliance load signatures,” in *Proc. IEEE Smart Grid Comm. Conf.*, Gaithersburg, MD, Oct. 2010, pp. 232–237.

[14] S. McLaughlin, P. McDaniel, and W. Aiello, “Protecting consumer privacy from electric load monitoring,” in *Proc. 18th ACM Conf. on Computer and Commun. Security*, Chicago, IL, Oct. 2011, pp. 87–98.

[15] W. Yang, N. Li, Y. Qi, W. Qardaji, S. McLaughlin, and P. McDaniel, “Minimizing private data disclosures in the smart grid,” in *Proc. ACM Conf. Comput. Commun. Secur.*, Raleigh, NC, Oct. 2012, pp. 415–427.

[16] D. Varodayan and A. Khisti, “Smart meter privacy using a rechargeable battery: Minimizing the rate of information leakage,” in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, Prague, Czech Republic, May 2011, pp. 1932–1935.

[17] J. Koo, X. Lin, and S. Bagchi, “PRIVATUS: Wallet-friendly privacy protection for smart meters,” in *Proc. 17th Eur. Symp. Res. Comp. Security*, Pisa, Italy, Sept. 2012, pp. 343–360.

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

[19] J. Yao and P. Venkitasubramaniam, “On the privacy-cost tradeoff of an in-home power storage mechanism,” in *Proc. Allerton Conf. Communication, Control, and Computing*, Monticello, IL, Oct. 2013, pp. 115–122.

[20] J. Gómez-Vilardebó and D. Gündüz, “Smart meter privacy for multiple users in the presence of an alternative energy source,” *IEEE Trans. Inform. Forensics and Security*, vol. 10, no. 1, pp. 132–141, Jan. 2015.

[21] S. Li, A. Khisti, and A. Mahajan, “Privacy-optimal strategies for smart metering systems with a rechargeable battery,” available at arxiv “http://bit.ly/20m8RxN”.

[22] G. Giaconi and D. Gündüz, “Smart meter privacy with renewable energy and a finite capacity battery,” in *Proc. IEEE Int. Wrkshp. Signal Process. Advances in Wireless Comm. (SPAWC)*, Edinburgh, UK, Jul. 2016.

[23] L. Sankar, S. R. Rajagopalan, and H. V. Poor, “Utility-privacy tradeoffs in databases: An information-theoretic approach,” *IEEE Trans. Inform. Forensics and Security*, vol. 8, no. 6, pp. 838–852, Jun. 2013.

[24] O. Grodzevich and O. Romanko, “Normalization and other topics in multi-objective optimization,” in *Proc. Fields MITACS Ind. Prob. Wrkshp.*, Toronto, Canada, Aug. 2006, pp. 89–101.

[25] M. A. Zafer and E. Modiano, “A calculus approach to energy-efficient data transmission with quality-of-service constraints,” *IEEE/ACM Trans. Networking*, vol. 17, no. 3, pp. 898–911, Jun. 2009.

[26] A. Reinhardt, D. Christin, and S. S. Kanhere, “Predicting the power consumption of electric appliances through time series pattern matching,” in *Proc. 5th ACM Wkshp. Embedded Syst. Energy-Efficient Buildings (BuildSys)*, Rome, Italy, Nov. 2013, pp. 1–2.

[27] T. Cover and J. Thomas, *Elements of Information Theory*. Hoboken, NJ, USA: Wiley, 1991.

[28] D. P. Bertsekas, *Dynamic programming and optimal control*. Athena Scientific, 2007.

[29] J. Z. Kolter and M. J. Johnson, “REDD: A public data set for energy disaggregation research,” in *Proc. Wkshp. Data Mining Applications Sustainability (SustKDD)*, San Diego, CA, Aug. 2011.

[30] Eurostat, “European Comission energy price statistics (2013),” available at “http://bit.ly/1AigKSR”.

[31] C. Brokish and M. Lewis, *A-Law and mu-Law Companding Implementations Using the TMS320C54x*. Digital Signal Processing Solutions, Texas Instruments, 1997.

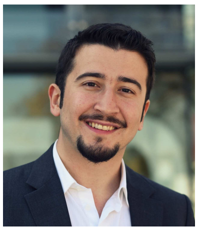

Onur Tan received the B.S. degree in Electrical and Electronics Engineering from the Middle East Technical University (METU), Turkey in June 2007, the M.S. degree in Electrical and Electronics Engineering from Bilkent University, Turkey in June 2010, and the Ph.D. degree, with ”Excellent Cum Laude” distinction, in Signal Theory and Communications Department from the Universitat Politècnica de Catalunya (UPC), Spain in July 2016, respectively. From November 2011 to July 2016, he was a researcher at the Centre Tecnològic de Telecomunicacions de Catalunya (CTTC), Barcelona, Spain. From September 2013 to February 2014, and from October 2014 to January 2015, he was a visiting researcher at the Department of Electrical and Electronic Engineering, the Intelligent Systems and Networks Research Group, Imperial College London, UK. The current research interests of Dr. Tan lie in the areas of machine learning, signal processing, data privacy, data analytics and big data problems.

Jesús Gómez-Vilardebó received his M.Sc. and Ph.D. degrees in Telecommunication Engineering from the Universitat Politècnica de Catalunya (UPC) in October 2003 and July 2009, respectively. In September 2005, he was granted by the Centre Tecnològic de Telecomunicacions de Catalunya (CTTC) to obtain the Ph.D. on Signal Theory and Communications at the UPC. He is now with the CTTC holding a Research Associate position. His current research interests include information theory, stochastic signal processing, and their applications in wireless multi-user communications and information privacy.

Deniz Gündüz [S’03-M’08-SM’13] received the B.S. degree in electrical and electronics engineering from METU, Turkey in 2002, and the M.S. and Ph.D. degrees in electrical engineering from NYU Polytechnic School of Engineering in 2004 and 2007, respectively. After his PhD, he served as a postdoctoral research associate at Princeton University, and as a consulting assistant professor at Stanford University. He was a research associate at CTTC in Barcelona, Spain until September 2012, when he joined the Electrical and Electronic Engineering Department of Imperial College London, UK, as a Lecturer. Currently he is a Reader in the same department.

His research interests lie in the areas of communications and information theory with special emphasis on joint source-channel coding, multi-user networks, energy efficient communications and privacy in cyber-physical systems. Dr. Gündüz is an Editor of the IEEE TRANSACTIONS ON COMMUNICATIONS, and the IEEE TRANSACTIONS ON GREEN COMMUNICATIONS AND NETWORKING. He is the recipient of a Starting Grant of the European Research Council (ERC) in 2016, IEEE Communications Society Best Young Researcher Award for the Europe, Middle East, and Africa Region in 2014, Best Paper Award at the 2016 IEEE Wireless Communications and Networking Conference (WCNC), and the Best Student Paper Award at the 2007 IEEE International Symposium on Information Theory (ISIT). He served as the General Co-chair of the 2016 IEEE Information Theory Workshop, and a Co-chair of the PHY and Fundamentals Track of the 2017 IEEE Wireless Communications and Networking Conference.