Optimal Privacy-Cost Trade-off
in Demand-Side Management with Storage

Onur Tan*, Deniz Gündüz†, Jesús Gómez Vilardebó*

*Centre Tecnològic de Telecomunicacions de Catalunya (CTTC), Barcelona, Spain.
†Department of Electrical and Electronic Engineering, Imperial College London, London, UK.

Abstract—Demand-side energy storage management is studied from a joint privacy-energy cost optimization perspective. Assuming that the user's power demand profile as well as the electricity prices are known non-causally, the optimal energy management (EM) policy that jointly increases the privacy of the user and reduces his energy cost is characterized. The *backward water-filling* interpretation is provided for the optimal EM policy. While the energy cost is reduced by requesting more energy when the prices are lower, energy consumption privacy is achieved by a smoother output load. It is shown that both gains can be achieved by using a limited size storage unit. The optimal trade-off between the user's privacy and energy cost is characterized, and the impact of the size of the storage unit and the resolution of the smart meter readings on this trade-off is studied.

I. Introduction

Smart meters (SMs) measure the power consumption of the users connected to the power grid and transmit their readings to the utility provider (UP) in almost real-time. This allows the UPs to closely monitor the grid and provide potential benefits in reliability, robustness and efficiency [1]. For example, the UPs can support dynamic electricity pricing based on SM readings and encourage the users to dynamically shift their demands to off-peak hours with the promise of reducing their energy costs. However, the possible misuse of these fine-grained readings by the UP or other third parties raise serious privacy and security concerns for the consumers [2].

Various techniques have been studied in the literature to provide a certain level of privacy to SM users. On the one hand, privacy can be provided by tampering the SM readings before being reported to the UP. Following this approach, [3] proposes the compression of SM data before being transmitted to the UP, and [4] considers adding random noise to SM readings to protect user's privacy. On the other hand, without tampering the SM readings, privacy can also be achieved by demand-side management with the utilization of storage units, such as rechargeable batteries (RBs) [5]–[10], and alternative energy sources [7], [10], [11]. In [5], a heuristic algorithm is proposed with the utilization of an RB to protect privacy, while in [6] and [7], user's privacy is protected by using an RB and an energy harvesting device, respectively, from an information theoretic perspective. The joint optimization of

Fig. 1. A smart-meter (SM) system diagram.

privacy and energy cost for SMs is addressed in [8] and [9]. Authors propose an online control algorithm and a dynamic programming with the utilization of an RB, respectively.

We consider the SM system depicted in Fig. 1. The energy management unit (EMU) satisfies the power demands of the appliances from the power grid and the RB. We do not allow outages or shifting of user demands. The SM measures the output load, Y_i, and reports its readings to the UP at certain time instants. Assuming that the electricity price is time-varying, the EMU utilizes the RB to reduce the user's energy consumption cost, as well as to mask the power consumption profile from the UP and other third parties. We assume that the future power demands as well as the electricity prices are known non-causally by the EMU. Exploiting this information the EMU can store extra energy into the RB in advance in order to achieve these gains. We assume that perfect privacy can be achieved if a constant SM reading is reported to the UP over time [5]. Consequently, we measure user privacy in terms of the deviation of the output load, Y_i, from the average power demand over the period of interest. On the other hand, the average energy cost is measured with a time-varying electricity pricing model. Our goal here is to characterize the optimal energy management (EM) policy that jointly optimizes privacy and energy cost over a given period of time under an RB capacity constraint. Note that an EM policy corresponds to power values requested by the EMU from the grid over the given time window.

We first formulate the joint privacy-energy cost minimization as a convex optimization problem. The optimal solution is characterized as the *backward water-filling algorithm*, in which the energy received from the grid can only be shifted to earlier time slots (TSs), and the water levels can be equalized to the extend the RB capacity allows. We characterize the

Fig. 2. Illustration of the timelines with variations in the total power consumption and the cost per unit energy.

trade-off between the user's privacy and energy cost for the optimal EMU operation. The operating point on this trade-off can be chosen based on the user's requirements on privacy and energy cost. We also investigate the impact of the RB capacity and the resolution of the SM readings on this trade-off.

II. SYSTEM MODEL

We consider a discrete-time power consumption model in a household (see Fig. 2(a)). In this model, each appliance consumes constant power for an arbitrary duration when it is active. Appliances can be in 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 state of at least one appliance. We denote the total power consumption within $[t_{(k-1)}^p, t_k^p]$ by X_k^p (kW) for $k = \{1, 2, \ldots, K\}$.

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 $m = \{1, 2, \ldots, M\}$. 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 $\tau_i \triangleq t_i - t_{i-1}$ (sec), for $i = 1, 2, \ldots, N$. 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. Note that for any two consecutive TSs, either the power demand or the cost per unit energy or both may change, whereas they remain constant within each TS. In our model, TSs do not necessarily have the same duration.

Hinged on the discrete-time power consumption and pricing model illustrated in Fig. 2, we study the power input/output system depicted in Fig. 1. We consider an SM that reports the output load, Y_i (kW), to the UP at each TS i[1]. We integrate an RB with a finite capacity B (kWh), and an EMU which

manages the power flow. The EMU can use both the power grid and the RB to satisfy the user's power demand, i.e., the input load X_i, as $X_i = Y_i + P_i$, where P_i (kW) is the power charged to ($P_i < 0$), or discharged from ($P_i > 0$) the RB during TS i, and $Y_i \geq 0$. We consider an EM policy that jointly optimizes the privacy and energy cost of the user within the time frame $[0, T]$ by utilizing the RB. Note that an energy management policy corresponds to the vector of output loads $[Y_1, Y_2, \ldots, Y_N]$. We are interested in *offline optimization*, that is, we assume that the EMU knows the power demand, X_i, and the cost per unit energy, C_i, for all TSs within $[0, T]$ in advance at $t_0 = 0$.

We assume that perfect privacy is achieved if the output load Y_i at each TS is equal to the user's average power demand within $[0, T]$. Ideally, if the user has a flat power demand from the grid at all times, we assume that the UP can not learn anything about the user's energy consumption behaviour [5]. Accordingly, we define the average power demand of the user as $\bar{E} \triangleq \frac{1}{T} \sum_{i=1}^{N} \tau_i \cdot X_i$. Then, the privacy provided by an EM policy is measured by the *load variance*, which is defined as:

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

Observe that perfect privacy is achieved when $\mathcal{V} = 0$, in which case $Y_i = \bar{E}$ for all TSs.

The energy consumption cost obtained by an EM policy is measured by the *average energy cost*, which is defined as:

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

We assume that all the power demands of appliances must be satisfied at the time that they are requested, i.e., we guarantee that the appliances do not incur any outages and we do not allow rescheduling; hence, assuming that the RB is empty at $t = 0$, the output load values have to satisfy the following constraints:

$$\sum_{i=1}^{n} \tau_i \cdot X_i \leq \sum_{i=1}^{n} \tau_i \cdot Y_i, \quad n = 1, \ldots, N. \tag{3}$$

On the other hand, the energy that has been drawn prior to the demand of the appliances needs to be stored in the RB. Since the RB capacity is finite, we require:

$$\sum_{i=1}^{n} \tau_i \cdot (Y_i - X_i) \leq B, \quad n = 1, \ldots, N. \tag{4}$$

It is possible to show that the set of all achievable $(\mathcal{V}, \mathcal{C})$ pairs under constraints (3) and (4) form a convex region. Then the optimal operating points are characterized by the Pareto boundary of this region. Hence, we use the weighted average

---

[1]Here, we assume that Y_i remains constant within each TS i. In the sequel, we will show that this assumption is indeed optimal. Accordingly, there is no loss of information on the UP side when SM reports once per each TS.

of $\mathcal{V}$ and $\mathcal{C}$ to identify all the points on the Pareto boundary. The optimization problem can be written as follows:

$$\min_{Y_i \geq 0} \sum_{i=1}^{N} \left[\theta \cdot \tau_i \cdot \left(Y_i - \bar{E}\right)^2 + (1-\theta) \cdot \tau_i \cdot Y_i \cdot C_i \right] \quad (5)$$

$$\text{s.t. } (3) \text{ and } (4),$$

where $0 \leq \theta \leq 1$ is the parameter that adjust the trade-off between privacy and energy cost. The value of θ can be set in advance by the user. If $\theta = 1$, then the user is interested only in maximizing the privacy; if $\theta = 0$, the user intends only to minimize the energy cost. Since the cost per unit energy and the input load remain constant over each TS, it follows from the convexity of the objective function that the optimal output load must remain constant within a TS [12]. Hence, the assumption of having the SM report only once per TS does not lead to any loss of information on the UP side. Since (5) is a convex optimization problem, it can be solved by the classical Lagrangian methods [13]. In the following section, we will provide some specifics of the optimal solution along with a water-filling interpretation for $0 < \theta \leq 1$. For $\theta = 0$, we obtain the optimal solution by using classical linear programming techniques since the objective function in (5) becomes linear in this case.

III. OPTIMAL ENERGY MANAGEMENT (EM) POLICY

Here, we provide the optimal EM policy that minimizes the privacy-energy cost function in (5). We define the Lagrangian function with the Lagrangian multipliers $\lambda_i \geq 0$, $\mu_i \geq 0$ and $v_i \geq 0$, as follows:

$$\mathcal{L} = \sum_{i=1}^{N} \left[\theta\tau_i\left(Y_i - \bar{E}\right)^2 + (1-\theta)\tau_i Y_i C_i \right]$$
$$+ \sum_{j=1}^{N} \lambda_j \left(\sum_{i=1}^{j} \tau_i (X_i - Y_i) \right)$$
$$+ \sum_{j=1}^{N} \mu_j \left(\left(\sum_{i=1}^{j} \tau_i (Y_i - X_i) \right) - B \right) - \sum_{j=1}^{N} v_j Y_j. \quad (6)$$

Corresponding complementary slackness conditions are:

$$\lambda_j \left(\sum_{i=1}^{j} \tau_i (X_i - Y_i) \right) = 0, \quad j = 1, \dots, N, \quad (7)$$

$$\mu_j \left(\left(\sum_{i=1}^{j} \tau_i (Y_i - X_i) \right) - B \right) = 0, \quad j = 1, \dots, N, \quad (8)$$

$$v_j Y_j = 0, \quad j = 1, \dots, N. \quad (9)$$

We apply the KKT conditions on the Lagrangian function:

$$\frac{\partial \mathcal{L}}{\partial Y_i} = 2\theta\tau_i\left(Y_i - \bar{E}\right) + (1-\theta)\tau_i C_i$$
$$+ \tau_i \sum_{j=i}^{N} (\mu_j - \lambda_j) - v_i = 0. \quad (10)$$

Fig. 3. Depiction of the backward water-filling algorithm for the optimal output loads with (a) infinite, and (b) finite capacity RBs, respectively, and the trade-off parameter $\theta = 1/3$.

Then the optimal output load in TS i, Y_i^*, is found in terms of the Lagrange multipliers, the cost per unit energy, and the trade-off parameter θ, as follows:

$$Y_i^* = \left[\alpha_i - \frac{(1-\theta)C_i}{2\theta} \right]^+, \quad 0 < \theta \leq 1, \forall i. \quad (11)$$

where $[x]^+$ is equal to x if $x \geq 0$, and 0 otherwise, and the water level in TS i is defined as:

$$\alpha_i \triangleq \frac{\sum_{j=i}^{N}(\lambda_j - \mu_j)}{2\theta} + \bar{E}, \quad 0 < \theta \leq 1, \forall i. \quad (12)$$

We first consider the special case when the RB capacity is infinite, i.e., $B \to \infty$. For this case, the constraints in (4) are satisfied without equality; and thereby, we have $\mu_j = 0$ for $\forall j$, from the slackness conditions in (8). Since $\lambda_i \geq 0$ for $\forall i$, it follows from (12) that the water level is monotonically decreasing with time, i.e., $\alpha_i \geq \alpha_{i+1}$. This implies that the water (power) can only flow backwards in our model, because the input load in a TS has to be satisfied within that TS, i.e., the output load can only be assigned to previous TSs, rather than to the future ones. If the RB is not empty after satisfying all input loads up to TS i, this implies that the i-th constraint in (3) is satisfied with strict inequality. It follows from the slackness condition in (7) that $\lambda_i = 0$; which gives rise to the fact from (12) that the water level remains constant, i.e., $\alpha_i = \alpha_{i+1}$.

In Fig. 3(a), we present a graphical interpretation of the optimal EM policy for three TSs in the presence of an infinite capacity RB. Selecting $\theta = 1/3$, the height of the rectangles correspond to the costs per unit energy, C_i's, while their widths correspond to the TS durations, τ_i's, for TSs $i = 1, 2, 3$. Fig. 3(a) depicts the optimal water-filling solution and the optimal output load values, Y_i^*, which is given by the height of the total dashed areas below the water level and above C_i. The input load X_i at TS i is given by the height of the corresponding dashed areas. Accordingly, the first power demand X_1 corresponds to the height of the first dashed area above C_1, and is satisfied from the grid within that TS, as seen in the figure. For the input load X_2, the output load is allocated by using the water-filling algorithm in reverse direction as

seen in Fig. 3(a). Since the cost per unit energy is relatively expensive in the second TS, part of X_2 is drawn in advance within the first TS, i.e., the height of the second dashed area above C_1, and stored into the RB. The rest of X_2 is drawn from the grid within the second TS, i.e., the height of the dashed area above C_2. Hence, X_2 is satisfied both from the RB and the grid. The power demand in the third TS, X_3, is satisfied from the grid within that TS. Thus, the optimal output load in the first TS, Y_1^*, depends on the input loads and the costs per unit energy in the following TSs. For N TSs, the optimal output load values can be obtained by N iterations of the water-filling algorithm. Since each input load can be satisfied by backward power allocation over the current and the previous TSs, we call this algorithm as *backward water-filling*.

Next we consider the general case when the RB capacity is finite. For this case, since the constraints in (4) can be satisfied with or without equality, we also need to consider the Lagrangian multipliers $\mu_j \geq 0$ in (12). The optimal solution is similar to the backward water-filling solution; however, the amount of water that can be poured into the previous TSs is now bounded by the RB capacity. If the RB is full in a TS, the excess drawn power would be wasted. Therefore, the RB capacity introduces an upper bound on the output load at each TS. Since $\lambda_i \geq 0$ and $\mu_i \geq 0 \; \forall i$, it follows from (12) that the water level is not necessarily decreasing with time, and can increase or decrease among two consecutive TSs. For example, if the RB is full after satisfying all the input load demands up to TS i, then this implies that the i-th constraint in (3) is satisfied with strict inequality while the i-th constraint in (4) is satisfied with strict equality. It follows from the slackness conditions in (7) and (8) that $\lambda_i = 0$ and $\mu_i \geq 0$, respectively; which lead to the fact from (12) that the water level increases, i.e., $\alpha_i \leq \alpha_{i+1}$.

In Fig. 3(b), we depict the graphical interpretation of the optimal EM policy for three TSs in the presence of a finite capacity RB. Differently from the infinite capacity RB case in Fig. 3(a), the portion of the input load in the second TS drawn in advance within the first TS is limited by the RB capacity. In other words, the amount of water that can be poured from the second TS to the first is bounded by the RB capacity. Therefore, the water levels can be equalized to the extend the water-filling direction and the RB constraints allow. Observe in Fig. 3(b) that the water level increases from the first TS to the second, which implies that the RB is full at the end of first TS.

IV. NUMERICAL RESULTS

In this section, we provide further insights about the optimal EM policy through some numerical results. We numerically analyze the trade-off between the user's privacy and energy cost as well as the effect of the RB capacity and the SM resolution on this trade-off. We consider the real power consumption data obtained from [14] with a time resolution varying on the order of three seconds. For our simulations we randomly take a whole-day power consumption data of one household, and


Fig. 4. The load variance, $\mathcal{V}$, versus the average energy cost, $\mathcal{C}$, for the RB capacities, $B = 1$ kWh, $B = 1.5$ kWh and $B = 2$ kWh, respectively.

convert the load profile to a time resolution of one-minute. To be consistent with our power consumption model, we assume that the sampling times of the original power data correspond to the discrete time-instants in Fig. 2(a). We set the electricity price by considering real pricing tariffs [15]. We assume that the electricity price can only change at the sampling times of the original power data. Accordingly, we set the off-peak price as 5 cent per kWh during $00:00$ to $12:00$, the on-peak price as 20 cent per kWh during $12:00$ to $20:00$ and the medium-peak price as 10 cent per kWh during $20:00$ to $00:00$.

In Fig. 4, we characterize the trade-off between the user's privacy and energy cost for RB capacities $B = \{1, 1.5, 2\}$ kWhs, respectively. The Pareto optimal trade-off curves between the load variance and the average energy cost are formed by varying θ from 0 to 1. The average energy cost increases and the load variance diminishes as $\theta \rightarrow 1$, and vice versa, as $\theta \rightarrow 0$. When $\theta = 1$, the load variance achieves its minimum value; on the other hand, the average energy cost achieves its minimum value before θ reaches to 0. Observe in Fig. 4 that, while the average energy cost cannot be reduced further after a particular θ value (as $\theta \rightarrow 0$), the load variance continues to increase. However, the user does not operate in this regime, and the operating point can be chosen elsewhere on the trade-off curve according to the requirements of the user. In Fig. 4, we also investigate the effect of the RB capacity on this trade-off. Observe that the Pareto optimal trade-off curve moves towards the origin as the RB capacity increases. This implies that with increasing RB capacity, the load variance can be reduced further under a fixed average energy cost, and the average energy cost can be reduced further under a fixed load variance. Both gains can be achieved by virtue of the degree-of-freedom provided by the RB.

Next, we compare the original load profile with the load profiles resulting from the optimal EM policy under the RB capacity, $B = 2$ kWh, and $\theta = \{0, 0.002, 1\}$, in Fig. 5. When $\theta = 0$, the EM policy minimizes only the energy cost of the user. As seen in Fig. 5, the EM policy stores extra energy in the RB in the off-peak price period and satisfies the demand of the on-peak price period from the RB as much as possible in order to reduce the cost. When $\theta = 1$, the EM policy maximizes
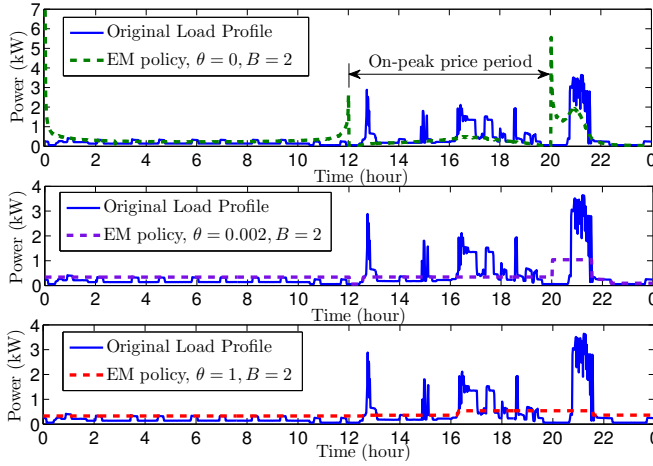
Fig. 5. Comparison of the original load profile with the load profiles resulting from the optimal EM policy under the RB capacity, $B = 2$ kWh, and $\theta = 0$, $\theta = 0.002$ and $\theta = 1$, respectively.

Fig. 6. The total load variance, $N\mathcal{V}$, versus the average energy cost, $\mathcal{C}$, for the RB capacity, $B = 1.5$ kWh, and the load profiles with a time resolution varying on the order of 5, 10, 15 minutes, and 1 hour, respectively.

only the privacy of the user. We can see that the EM policy generates a smooth load profile with which the peaks arising from on-off switching of appliances in the original load profile are masked. When $\theta = 0.002$, the EM policy jointly optimizes the user's privacy and energy cost.

Finally, we investigate the impact of the SM resolution on the trade-off between the user's privacy and energy cost. To that end, we modify the original load profile into new load profiles with lower resolutions. Accordingly, the new load profiles have time resolutions varying on the order of 5, 10, 15 minutes, and 1 hour, respectively. We then characterize the Pareto optimal trade-off between the total load variance, $N\mathcal{V}$, and the average energy cost, $\mathcal{C}$, for the load profiles with given resolutions and the RB capacity $B = 1.5$ kWh in Fig. 6. We see that the Pareto optimal trade-off curve moves downwards as the SM resolution gets lower. This implies that with a decreasing resolution, the EM policy can provide higher energy consumption privacy under a fixed average energy cost. This is due to the fact that a load sampled at a lower-resolution is smoother, and has a smaller variance compared to the same load sampled at a higher-resolution.

V. CONCLUSIONS

We studied the optimal demand-side EM policy that minimizes the joint privacy-energy cost measure in an SM system in the presence of a finite-capacity energy storage unit. We considered a discrete-time power demand profile, i.e., input load, for the user as well as time-varying electricity prices. We assumed that the user's input load profile along with the electricity prices are known non-causally at the EMU, and the power demands of the appliances have to be satisfied without any outages or rescheduling. We characterized the optimal EM policy that jointly increases the user's privacy and reduces the energy cost. The RB is utilized in order to achieve both gains with an adaptive EM policy. We showed that the optimal solution is characterized as the backward water-filling

algorithm. We characterized the optimal trade-off between the user's privacy and the energy cost, and investigated the impact of the RB capacity and the SM resolution on this trade-off.

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] 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.

[4] 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.

[5] 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.

[6] 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. (ICASSP)*, Prague, Czech Republic, May 2011.

[7] 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.: Smart Grid Comm.*, vol. 31, no. 7, pp. 1331–1341, Jul. 2013.

[8] 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.

[9] 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.

[10] 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.

[11] 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.

[12] 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.

[13] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.

[14] 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.

[15] Online, "European comission energy price statistics (2013)," available at "http://bit.ly/1AigKSR".