PRIVACY-PRESERVING ENERGY FLOW CONTROL IN SMART GRIDS

Zuxing Li, Tobias J. Oechtering, and Mikael Skoglund

School of Electrical Engineering and the ACCESS Linnaeus Centre
KTH Royal Institute of Technology, Stockholm, Sweden

ABSTRACT

In this paper, an energy flow control strategy to reduce the smart meter privacy leakage is studied. The considered smart grid is equipped with an energy storage device. The privacy leakage is modeled as optimal Bayesian detections on the behaviors of the consumer made by an authorized adversary. To evaluate the privacy risk, a Bayesian detection-operational privacy leakage metric is proposed. The design of an optimal privacy-preserving energy control strategy can be formulated as a belief state MDP problem. Therefore, standard methods and algorithms can be utilized to obtain or to approximate the optimal control strategy. A simplified problem to design an instantaneous optimal privacy-preserving control strategy is also considered. It is shown that the problem of the instantaneous optimal control strategy design can be formulated as a set of linear programmings.

***Index Terms*—** Bayesian detection, linear programming, MDP, smart meter privacy

1. INTRODUCTION

A smart grid is an energy network which manages the energy generation and distribution more efficiently following the real-time consumer's energy demand through control and communication technologies [1]. As benefits from the smart grid, energy efficiency can be improved; reliability and robustness can be increased; and costs of the energy provider and consumer can be reduced. However, these benefits come with privacy/secrecy challenges [2, 3]. In a smart grid, the smart meter provides the real-time information of energy supplies from the energy provider on the demands of the consumer, which can be utilized to infer on the privacy of the consumer [4, 5]. Regarding the smart meter privacy problem, different privacy-preserving approaches have been developed. An encryption method was proposed in [6] to protect privacy of individual consumers through the neighborhood-level data aggregation. In [7], a privacy scheme was devised to schedule the usage of delay-tolerable appliances to hide the energy consumption characteristics of other appliances.

The work has been supported by the Swedish Research Council (VR) within the CHIST-ERA project COPES under Grant 2015-06815.

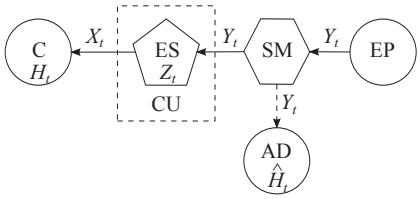

Fig. 1. The studied model of smart meter privacy leakage to an authorized adversary where energy and information flows are represented by solid and dashed arrows respectively.

The above methods work well in most cases except for having compromised authorized people. Taking into account the latter threat, some works utilize alternative energy sources or energy storage devices to hide the energy demand characteristics. In [1, 8–10], they modeled and solved such problems by using information theoretic methods. In [11], this idea led to a problem to minimize the variance of energy supplies from the energy provider and therefore to "flatten" the smart meter readings. In [12–14], privacy problems in information theoretic formulation were shown to be reformulated as problems of Markov decision process (MDP).

In this paper, we consider the smart meter privacy leakage to an informed, greedy, authorized adversary and design the optimal privacy-preserving energy control strategy in the presence of an energy storage device. Different from the abstract privacy leakage interpretations used in [1, 8–14], we model the smart meter privacy leakage from a Bayesian detection-operational perspective. Some works have been done to relate the hypothesis detection and privacy risk [15–18]. Following our previous work [18], we propose a privacy leakage metric in terms of instantaneous minimal Bayesian risk of the adversary and then identify a belief state MDP problem to optimize the privacy-preserving energy control strategy. In addition, a lower-complexity problem of instantaneous optimal energy control strategy design is discussed later.

In the following, we will denote a random variable by a capital letter, its realization by the lower-case letter, and its definition domain by the calligraphic letter. Let X_t^{t+k} and x_t^{t+k} denote a random sequence $(X_t, \ldots, X_{t+k})$ and its realization $(x_t, \ldots, x_{t+k})$. Particularly, X^k, x^k are used when

$t = 0$; X_t^∞, x_t^∞ are used when $k = \infty$; and X^∞, x^∞ are used when $t = 0$ and $k = \infty$.

2. SMART GRID MODEL

In this paper, we consider the smart grid model shown in Fig. 1. In the infinite time horizon, during each time slot $t \in \{0, 1, \dots\}$, the accumulated energy demand X_t of the consumer (C), energy storage state Z_t of the energy storage (ES) device, and energy supply Y_t from the energy provider (EP) are defined on $\mathcal{X} = \{0, e, \dots, ue\}$, $\mathcal{Z} = \{0, e, \dots, me\}$, and $\mathcal{Y} = \{0, e, \dots, (u+m)e\}$ where e is the energy measuring precision. For the consumer, let H_t defined on $\mathcal{H}$ denote an n-ary hypothesis of his behavior in the slot t. Assume that the initial state (H_0, X_0, Z_0) is generated following the p.m.f. p_{H_0,X_0,Z_0}; H_{t+1} is generated depending on H_t only and following the time-invariant[1] p.m.f. $p_{H_{t+1}|H_t}$; the energy demand X_{t+1} is generated depending on (H_{t+1}, X_t) only and following the time-invariant[2] p.m.f. $p_{X_{t+1}|H_{t+1},X_t}$; and the control unit (CU) requests an energy supply Y_t depending on (X_t, Z_t) only and following a control strategy characterized by the p.m.f. $p_{Y_t|X_t,Z_t}$. These settings imply the following Markov chains.

$$H_{t+1} - H_t - (H^{t-1}, H_{t+2}^\infty, X^\infty, Z^\infty, Y^\infty)$$
$$X_{t+1} - (H_{t+1}, X_t) - (H^t, H_{t+2}^\infty, X^{t-1}, X_{t+2}^\infty, Z^\infty, Y^\infty)$$
$$Y_t - (X_t, Z_t) - (H^\infty, X^{t-1}, X_{t+1}^\infty, Z^{t-1}, Z_{t+1}^\infty, Y^{t-1}, Y_{t+1}^\infty)$$

In our model, the instantaneous energy demand x_t is always satisfied and no energy is wasted. Then, we have

$$z_t + y_t - x_t = z_{t+1}$$

which leads to

$$p_{Z_{t+1}|X_t,Z_t}(z_t + y_t - x_t | x_t, z_t) = p_{Y_t|X_t,Z_t}(y_t|x_t, z_t). \quad (1)$$

Thus, Z_{t+1} depends on (X_t, Z_t) only and a control strategy can be equivalently characterized by the p.m.f. $p_{Z_{t+1}|X_t,Z_t}$. In addition, a valid control strategy satisfies the following property:

$$p_{Y_t|X_t,Z_t}(y_t|x_t, z_t) = 0 \text{ if } \begin{cases} y_t < \max\{0, x_t - z_t\} \\ \text{or} \\ y_t > me + x_t - z_t \end{cases}, \quad (2)$$

where the first condition indicates that the lower bound of the energy supply y_t is $x_t - z_t$ to provide the rest energy when the energy storage z_t cannot satisfy the energy demand x_t solely; and the second condition indicates that the upper bound of y_t is constrained by the maximum energy storage limitation.

We consider a privacy leakage problem that the smart meter (SM) readings Y^∞ are utilized by an authorized adversary (AD) to infer on the behaviors H^∞ by making guesses $\hat{H}^\infty$. We assume that a decision $\hat{H}_t$ of the adversary is made depending on Y_t only and following a decision strategy characterized by the p.m.f. $p_{\hat{H}_t|Y_t}$. This assumption implies the following Markov chain.

$$\hat{H}_t - Y_t - (\hat{H}^{t-1}, \hat{H}_{t+1}^\infty, H^\infty, X^\infty, Z^\infty, Y^{t-1}, Y_{t+1}^\infty)$$

In addition, we assume the adversary is informed and greedy such that the optimal decision strategies $\{p^*_{\hat{H}_t|Y_t}\}_{t=0}^\infty$ are used based on his knowledge of the p.m.f.s p_{H_0,X_0,Z_0}, $\{p_{H_{t+1}|H_t}\}_{t=0}^\infty$, $\{p_{X_{t+1}|H_{t+1},X_t}\}_{t=0}^\infty$, and used control strategies $\{p_{Y_t|X_t,Z_t}\}_{t=0}^\infty$. In practice, such an informed, greedy, authorized adversary can be a compromised manager of the energy provider.

3. ENERGY CONTROL AGAINST PRIVACY LEAKAGE

3.1. Bayesian Detection Model of Privacy Leakage

In the previous works, the smart meter privacy leakage was measured by the uncertainty of the adversary about energy demand profile, e.g., entropy and variance. Here, we model the smart meter privacy problem as a Bayesian hypothesis detection, i.e., the privacy leakage has an operational interpretation.

Let $c(\hat{h}_t, h_t)$ denote the non-negative cost of the adversary to make a decision $\hat{h}_t$ when the true behavior is h_t. The detection costs are assigned following our own privacy-preserving design interest. That will lead to an energy control strategy guaranteeing the desired privacy-preserving performance of us (the designers). According to [19], the optimal decision strategy $p^*_{\hat{H}_t|Y_t}$ used by the informed and greedy adversary is a deterministic likelihood-ratio test (LRT) and the minimal Bayesian risk r_t^* is a function of p_{H_t,X_t,Z_t} and the control strategy $p_{Y_t|X_t,Z_t}$ as

$$\begin{aligned} r_t^* &= \min \sum_{\hat{h}_t, h_t \in \mathcal{H}^2} c(\hat{h}_t, h_t) p_{\hat{H}_t,H_t}(\hat{h}_t, h_t) \\ &= \sum_{y_t \in \mathcal{Y}} \min_{\hat{h}_t \in \mathcal{H}} \{ \sum_{h_t,x_t,z_t \in \mathcal{H}\times\mathcal{X}\times\mathcal{Z}} c(\hat{h}_t, h_t) \\ &\qquad p_{Y_t|X_t,Z_t}(y_t|x_t, z_t) p_{H_t,X_t,Z_t}(h_t, x_t, z_t) \}. \end{aligned} \quad (3)$$

For the considered privacy leakage problem in the infinite time horizon, we propose to use the accumulated discounted minimal Bayesian risk V as the privacy leakage metric as

$$V = \sum_{t=0}^\infty \beta^t r_t^*, \quad (4)$$

where $0 \le \beta < 1$ is the discount factor. It is obvious that V is a function of $\{p_{H_t,X_t,Z_t}\}_{t=0}^\infty$ and $\{p_{Y_t|X_t,Z_t}\}_{t=0}^\infty$. From the operation perspective, the proposed privacy leakage metric V is applicable in the scenarios where privacy-preserving concern degrades as time goes on.

[1] $\forall t \ge 0, k \ge -t, p_{H_{t+1}|H_t} = p_{H_{t+k+1}|H_{t+k}}$.

[2] $\forall t \ge 0, k \ge -t, p_{X_{t+1}|H_{t+1},X_t} = p_{X_{t+k+1}|H_{t+k+1},X_{t+k}}$.

3.2. Optimal Privacy-Preserving Control Strategy Design in Belief State MDP Formulation

Using the proposed privacy leakage metric, our objective is to design the optimal control strategies to suppress the privacy leakage risk, i.e., to maximize V:

$$\{p^*_{Y_t|X_t,Z_t}\}_{t=0}^{\infty} = \underset{p_{Y_t|X_t,Z_t}\in\mathcal{A},\forall t\geq 0}{\operatorname{argmax}} V, \tag{5}$$

where $\mathcal{A}$ is the valid control strategy set and consists of all p.m.f.s satisfying the condition (2).

Based on the settings in Section 2, we have

$$\begin{aligned}
&p_{H_{t+1},X_{t+1},Z_{t+1}|H_t,X_t,Z_t}\\
=&p_{Z_{t+1}|H_t^{t+1},X_t^{t+1},Z_t}\cdot p_{X_{t+1}|H_t^{t+1},X_t,Z_t}\cdot p_{H_{t+1}|H_t,X_t,Z_t}\\
=&p_{Z_{t+1}|X_t,Z_t}\cdot p_{X_{t+1}|H_{t+1},X_t}\cdot p_{H_{t+1}|H_t}.
\end{aligned}$$

Since the p.m.f. $p_{Z_{t+1}|X_t,Z_t}$ represents a control strategy, the transition from (H_t, X_t, Z_t) to $(H_{t+1}, X_{t+1}, Z_{t+1})$ depends on the used control strategy. On observing p_{H_t,X_t,Z_t} and using control-strategy-based $p_{H_{t+1},X_{t+1},Z_{t+1}|H_t,X_t,Z_t}$, the control unit can determine (observe) $p_{H_{t+1},X_{t+1},Z_{t+1}}$. The minimal Bayesian risk of the adversary r^*_t can be seen as the reward to use a control strategy $p_{Y_t|X_t,Z_t}$ given the p.m.f. p_{H_t,X_t,Z_t}. Based on these observations, we intuitively have the following proposition.

Proposition 1. *The privacy-preserving energy flow control design can be formulated as a belief state MDP problem.*

Explicit identification of elements in the belief state MDP problem is still required and shown in the following constructive proof of Proposition 1.

Proof. Elements of the belief state MDP problem are identified as:

- State: $s_t = (h_t, x_t, z_t) \in \mathcal{S} = \mathcal{H}\times\mathcal{X}\times\mathcal{Z}$.
- Belief state: $b_t = p_{H_t,X_t,Z_t} \in \mathcal{B}$.
- Action: $a_t = p_{Y_t|X_t,Z_t} \in \mathcal{A}$.
- Belief state transition: $p_{B_{t+1}|B_t,A_t}(b_{t+1}|b_t,a_t) =$
$$\begin{cases} 1, \text{ if } b_{t+1}(s_{t+1}) = \displaystyle\sum_{s_t\in\mathcal{S}} p_{S_{t+1}|S_t}(s_{t+1}|s_t)b_t(s_t) \\ \quad \text{for all } s_{t+1}\in\mathcal{S} \\ 0, \text{ otherwise} \end{cases}.$$
Note that $p_{S_{t+1}|S_t} = p_{Z_{t+1}|X_t,Z_t}p_{X_{t+1}|H_{t+1},X_t}p_{H_{t+1}|H_t}$ where $p_{Z_{t+1}|X_t,Z_t}$ can be substituted by the used control strategy $a_t = p_{Y_t|X_t,Z_t}$ according to (1).
- Belief state reward: $r^*_t(b_t, a_t)$.
- Policy: $\delta_t : \mathcal{B}\to\mathcal{A}$ which maps a belief state b_t to a control strategy a_t.

□

By formulating the privacy-preserving energy flow control design as a belief state MDP problem, standard methods can be used to obtain or to approximate the optimal design as shown in the following.

3.3. Optimal Energy Flow Control

For the belief state MDP problem, define $\Delta = \{\delta_0, \delta_1, \dots\}$. Regarding the objective to suppress the privacy leakage risk to the lowest, $\Delta^* = \{\delta^*_0, \delta^*_1, \dots\}$ is the solution to the optimization problem as:

$$\Delta^* = \underset{\Delta}{\operatorname{argmax}}\, V(\Delta, b_0), \text{ for all } b_0 \in \mathcal{B}. \tag{6}$$

Bellman's principle of optimality [20] indicates that for all $t \geq 0$ and a given $b \in \mathcal{B}$

$$V(\Delta^*, b) = \max_{a\in\mathcal{A}} \{r^*_t(b,a) + \beta V(\Delta^*, b')\} \tag{7}$$

where b' satisfies $p_{B_{t+1}|B_t,A_t}(b'|b,a) = 1$. If the optimal argument a^* of the Bellman equation (7) exists, a stationary solution Δ^* exists and satisfies

$$a^* = \delta^*_t(b), \text{ for all } t \geq 0 \text{ and a given } b \in \mathcal{B}. \tag{8}$$

Generally, solving the Bellman equation (7) of the belief state MDP problem is computationally complex due to the infinite belief state and action domains. An approximation idea is to use some discretization procedure, e.g., using α-vector algorithm [21] by discretizing the action domain, using value iteration algorithm by discretizing both the belief state and action domains.

3.4. Instantaneous Optimal Control

Due to the computational complexity of the formulated belief state MDP problem, we consider an instantaneous optimal control strategy. In each time slot, a control strategy is used to suppress the instantaneous privacy leakage risk by maximizing the instantaneous minimal Bayesian risk of the authorized adversary. Denote the instantaneous optimal policies as $\Delta^\# = \{\delta^\#_0, \delta^\#_1, \dots\}$. Then,

$$\delta^\#_t(b_t) = \underset{p_{Y_t|X_t,Z_t}\in\mathcal{A}}{\operatorname{argmax}} r^*_t(b_t, p_{Y_t|X_t,Z_t}). \tag{9}$$

The problem (9) is not a convex optimization. However, it can be rewritten as a set of linear programming problems as discussed in the following.

It is known that the optimal decision strategy for the adversary in a slot t is deterministic. There are $l = |\hat{\mathcal{H}}|^{|\mathcal{Y}|} = n^{|\mathcal{Y}|}$ deterministic mappings from $\mathcal{Y}$ to $\hat{\mathcal{H}}$, i.e., there are l deterministic candidate decision strategies for the adversary. Let $\phi_j : \mathcal{Y}\to\hat{\mathcal{H}}$ with $j \in \{1, 2, \dots, l\}$ denote the j-th deterministic candidate decision strategy. Given a belief state

$b_t = p_{H_t,X_t,Z_t}$ and a deterministic candidate decision strategy ϕ_j, define a subset of $\mathcal{A}$ as

$$\mathcal{A}_j(b_t) = \left\{ p_{Y_t|X_t,Z_t} : \begin{array}{l} p_{Y_t|X_t,Z_t} \in \mathcal{A}; \\ \forall (y_t, \hat{h}_t) \in \mathcal{Y} \times \mathcal{H}, \\ \sum_{s_t \in \mathcal{S}} \{(c(\phi_j(y_t), h_t) - c(\hat{h}_t, h_t)) \\ p_{Y_t|X_t,Z_t}(y_t|x_t, z_t) b_t(s_t)\} \leq 0. \end{array} \right\}.$$

It is obvious that a subset $\mathcal{A}_j(b_t)$ is defined by linear constraints on $p_{Y_t|X_t,Z_t}$. In addition, it can be easily verified that $\bigcup_{j=1}^{l} \mathcal{A}_j(b_t) = \mathcal{A}$.

Proposition 2. *The non-convex optimization problem in (9) can be rewritten as*

$$\max_{j \in \{1,2,\ldots,l\}} \left\{ \max_{p_{Y_t|X_t,Z_t} \in \mathcal{A}_j(b_t)} r_t^*(b_t, p_{Y_t|X_t,Z_t}) \right\} \tag{10}$$

to design the instantaneous optimal control strategy $\delta_t^{\#}(b_t)$. In (10), the inner optimizations are all linear programmings.

Proof. For the j-th inner optimization in (10), the objective $r_t^*(b_t, p_{Y_t|X_t,Z_t})$ is maximized over a subset $\mathcal{A}_j(b_t)$. The definition of $\mathcal{A}_j(b_t)$ makes the objective in the j-th inner optimization reduce to

$$\sum_{y_t \in \mathcal{Y}} \sum_{s_t \in \mathcal{S}} \{c(\phi_j(y_t), h_t) p_{Y_t|X_t,Z_t}(y_t|x_t, z_t) b_t(s_t)\}.$$

It is obvious the objective is a linear function of $p_{Y_t|X_t,Z_t}$. In addition, the subset $\mathcal{A}_j(b_t)$ is defined by a set of linear constraints of $p_{Y_t|X_t,Z_t}$. Therefore, each inner optimization in (10) is a linear programming problem to maximize a linear objective of $p_{Y_t|X_t,Z_t}$ subject to a set of linear constraints on $p_{Y_t|X_t,Z_t}$. □

Using standard methods, the inner linear programmings can be efficiently solved. The outer optimization of (10) simply compares the results of the inner optimizations to determine the maximum instantaneous minimal Bayesian risk of the adversary and instantaneous optimal strategy $\delta_t^{\#}(b_t)$.

Remark 1. *The instantaneous optimal policies $\Delta^{\#} = \{\delta_0^{\#}, \delta_1^{\#}, \ldots\}$ are stationary such that $\delta_t^{\#} = \delta^{\#}$, $\forall t \geq 0$.*

4. NUMERICAL STUDY

Here, we illustrate a simple numerical example. The system is set as: a binary hypothesis $n = 2$, binary energy demand $u = 1$, binary energy storage state $m = 1$. The p.m.f.s are set as:

$$\begin{aligned} p_{H_0,X_0,Z_0}(h_a, 0, 0) &= p_{H_0,X_0,Z_0}(h_a, e, e) = 0.1, \\ p_{H_0,X_0,Z_0}(h_a, 0, e) &= p_{H_0,X_0,Z_0}(h_a, e, 0) = 0.2, \\ p_{H_0,X_0,Z_0}(h_b, 0, 0) &= p_{H_0,X_0,Z_0}(h_b, 0, e) = 0.1, \\ p_{H_0,X_0,Z_0}(h_b, e, 0) &= p_{H_0,X_0,Z_0}(h_b, e, e) = 0.1, \end{aligned}$$

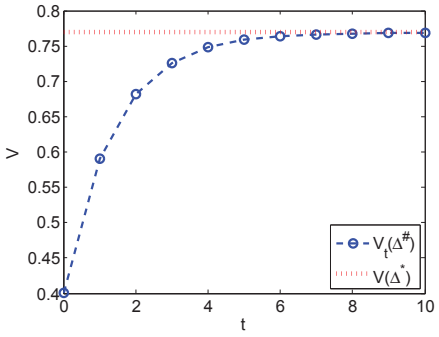

Fig. 2. Privacy-preserving performance comparison between using optimal control strategies and using instantaneous optimal control strategies.

$$\begin{aligned} p_{H_{t+1}|H_t}(h_a|h_a) &= 0.9, \\ p_{H_{t+1}|H_t}(h_b|h_b) &= 0.8, \\ p_{X_{t+1}|H_{t+1},X_t}(0|h_a, 0) &= p_{X_{t+1}|H_{t+1},X_t}(e|h_a, e) = 0.7, \\ p_{X_{t+1}|H_{t+1},X_t}(0|h_b, 0) &= p_{X_{t+1}|H_{t+1},X_t}(e|h_b, e) = 0.7. \end{aligned}$$

The costs are set as: $c(h_a, h_a) = c(h_b, h_b) = 0$ and $c(h_a, h_b) = c(h_b, h_a) = 1$. The discount factor is $\beta = 0.5$.

In Fig. 2, we use a blue circle to represent $V_t(\Delta^{\#}) = \sum_{k=0}^{t} \beta^k r_k^*(\delta_k^{\#})$ which denotes an accumulated discounted minimal Bayesian risk of the adversary until time slot t when instantaneous optimal policies are used. The red dot line represents $V(\Delta^*)$ when the optimal privacy-preserving energy control strategies are used in the infinite time horizon. For this example, the numerical results indicate that the instantaneous optimal policies $\Delta^{\#}$ can approach the privacy-preserving performance of the optimal privacy-preserving policies Δ^* asymptotically.

5. CONCLUSION

In this paper, we consider the smart meter privacy leakage to an informed, greedy, authorized adversary. We model the privacy leakage as an optimal Bayesian detection on the behavior hypothesis of the consumer and measure the instantaneous privacy leakage risk by the minimal Bayesian risk of the adversary in each time slot. We identify the design of privacy-preserving energy control strategies as a belief state MDP problem. Therefore, established standard methods and algorithms can be used to solve the optimal energy control policies which maximize the accumulated discounted minimal Bayesian risk of the adversary. When we focus on suppressing the instantaneous privacy leakage risk, the design problem of an instantaneous optimal control policy can be described by a set of linear programmings. Thereof, it can be solved efficiently.

6. REFERENCES

[1] O. Tan, D. Gunduz, and H. V. Poor, "Increasing smart meter privacy through energy harvesting and storage devices," *IEEE Journal on Selected Areas in Communications*, vol. 31, no. 7, pp. 1331–1341, 2013.

[2] E. L. Quinn, "Privacy and the new energy infrastructure," *SSRN*, 2009.

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

[4] F. Sultanem, "Using appliance signatures for monitoring residential loads at meter panel level," *IEEE Transactions on Power Delivery*, vol. 6, no. 4, pp. 1380–1385, 1991.

[5] A. Molina-Markham, P. Shenoy, K. Fu, E. Cecchet, and D. Irwin, "Private memoirs of a smart meter," in *Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building*, 2010, pp. 61–66.

[6] F. Li, B. Luo, and P. Liu, "Secure information aggregation for smart grids using homomorphic encryption," in *Proceedings of SmartGridComm 2010*, 2010, pp. 327–332.

[7] G. Kalogridis, C. Efthymiou, S. Z. Denic, T. A. Lewis, and R. Cepeda, "Privacy for smart meters: Towards undetectable appliance load signatures," in *Proceedings of SmartGridComm 2010*, 2010, pp. 232–237.

[8] D. Varodayan and A. Khisti, "Smart meter privacy using a rechargeable battery: Minimizing the rate of information leakage," in *Proceedings of ICASSP 2011*, 2011, pp. 1932–1935.

[9] D. Gunduz and J. Gomez-Vilardebo, "Smart meter privacy in the presence of an alternative energy source," in *Proceedings of ICC 2013*, 2013, pp. 2027–2031.

[10] Z. Li and T. J. Oechtering, "Privacy on hypothesis testing in smart grids," in *Proceedings of ITW 2015 Fall*, 2015, pp. 337–341.

[11] L. Yang, X. Chen, J. Zhang, and H. V. Poor, "Optimal privacy-preserving energy management for smart meters," in *Proceedings of INFOCOM 2014*, 2014, pp. 513–521.

[12] J. Yao and P. Venkitasubramaniam, "On the privacy-cost tradeoff of an in-home power storage mechanism," in *Proceedings of Allerton 2013*, 2013, pp. 115–122.

[13] J. Yao and P. Venkitasubramaniam, "The privacy analysis of battery control mechanisms in demand response: Revealing state approach and rate distortion bounds," in *Proceedings of CDC 2014*, 2014, pp. 1377–1382.

[14] S. Li, A. Khisti, and A. Mahajan, "Structure of optimal privacy-preserving policies in smart-metered systems with a rechargeable battery," in *Proceedings of SPAWC 2015*, 2015, pp. 375–379.

[15] J. C. Duchi, M. I. Jordan, and M. J. Wainwright, "Local privacy and statistical minimax rates," in *Proceedings of FOCS 2013*, 2013, pp. 429–438.

[16] R. F. Barber and J. Duchi, "Privacy: A few definitional aspects and consequences for minimax mean-squared error," in *Proceedings of CDC 2014*, 2014, pp. 1365–1369.

[17] P. Kairouz, S. Oh, and P. Viswanath, "Extremal mechanisms for local differential privacy," eprint arXiv:1407.1338.

[18] Z. Li and T. J. Oechtering, "Privacy-aware distributed Bayesian detection," *IEEE Journal of Selected Topics in Signal Processing*, vol. 9, no. 7, pp. 1345–1357, 2015.

[19] P. K. Varshney, *Distributed Detection and Data Fusion*, Springer-Verlag, 1996.

[20] R. Bellman, "The theory of dynamic programming," *Bull. Amer. Math. Soc.*, vol. 60, pp. 503–516, 1954.

[21] R. D. Smallwood and E. J. Sondik, "The optimal control of partially observable Markov processes over a finite horizon," *Operations Research*, vol. 21, pp. 1071–1088, 1973.