

Information-Theoretic Privacy for Smart Metering Systems with a Rechargeable Battery

Simon Li, Ashish Khisti, and Aditya Mahajan

Abstract—Smart-metering systems report electricity usage of a user to the utility provider on almost real-time basis. This could leak private information about the user to the utility provider. In this work we investigate the use of a rechargeable battery in order to provide privacy to the user.

We assume that the user load sequence is a first-order Markov process, the battery satisfies ideal charge conservation, and that privacy is measured using normalized mutual information (leakage rate) between the user load and the battery output. We consider battery charging policies in this setup that satisfy the feasibility constraints. We propose a series reductions on the original problem and ultimately recast it as a Markov Decision Process (MDP) that can be solved using a dynamic program.

In the special case of i.i.d. demand, we explicitly characterize the optimal policy and show that the associated leakage rate can be expressed as a single-letter mutual information expression. In this case we show that the optimal charging policy admits an intuitive interpretation of preserving a certain invariance property of the state. Interestingly an alternative proof of optimality can be provided that does not rely on the MDP approach, but is based on purely information theoretic reductions.

I. INTRODUCTION

Smart meters are a critical part of modern power distribution systems because they provide fine-grained power consumption measurements to utility providers. These fine-grained measurements improve the efficiency of the power grid by enabling services such as time-of-use pricing and demand response [1]. However, this promise of improved efficiency is accompanied by a risk of privacy loss. It is possible for the utility provider—or an eavesdropper—to infer private information including load taxonomy from the fine-grained measurements provided by smart meters [2]–[4]. Such private information could be exploited by third parties for the purpose of targeted advertisement or surveillance. Traditional techniques in which an intermediary anonymizes the data [5] are also prone privacy loss to an eavesdropper. One possible solution is to partially obscure the load profile by using a rechargeable battery [6]. As the cost of rechargeable batteries decreases (for example, due to proliferation of electric vehicles), using them for improving privacy is becoming economically viable.

In a smart metering system with a rechargeable battery, the energy consumed from the power grid may either be less than

the user’s demand—the rest being supplied by the battery; or may be more than the user’s demand—the excess being stored in the battery. A rechargeable battery provides privacy because the power consumed from the grid (rather than the user’s demand) gets reported to the electricity utility (and potentially observed by an eavesdropper). In this paper, we focus on the mutual information between the user’s demand and consumption (i.e., the information leakage) as the privacy metric. Mutual Information is a widely used metric in the literature on information theoretic security, as it is often analytically tractable and provides a fundamental bound on the probability of detecting the true load sequence from the observation [7]. Our objective is to identify a battery management policy (which determine how much energy to store or discharge from the battery) to minimize the information leakage rate.

We briefly review the relevant literature. The use of a rechargeable battery for providing user privacy has been studied in several recent works, e.g., [6], [8]–[11]. Most of the existing literature has focused on evaluating the information leakage rate of specific battery management policies. These include the “best-effort” policy [6], which tries to maintain a constant consumption level, whenever possible; and battery conditioned stochastic charging policies [8], in which the conditional distribution on the current consumption depends only on the current battery state (or on the current battery state and the current demand). In [6], the information leakage rate was estimated using Monte-Carlo simulations; in [8], it was calculated using the BCJR algorithm [12]. The methodology of [8] was extended by [9] to include models with energy harvesting and allowing for a certain amount of energy waste. Bounds on the performance of the best-effort policy and hide-and-store policy for models with energy harvesting and infinite battery capacity were obtained in [10]. The performance of the best effort algorithm for an individual privacy metric was considered in [11]. None of these papers address the question of choosing the optimal battery management policy.

Rate-distortion type approaches have also been used to study privacy-utility trade-off [13]–[15]. These models allow the user to report a distorted version of the load to the utility provider, subject to a certain average distortion constraint. Our setup differs from these works as we impose a constraint on the *instantaneous* energy stored in the battery due to its limited capacity. Both our techniques and the qualitative nature of the results are different from these papers.

Our contributions are two-fold. First, when the demand is Markov, we show that the minimum information leakage rate and optimal battery management policies can be obtained by solving an appropriate dynamic program. These results are

Simon Li was with the University of Toronto when this work was done. Ashish Khisti (akhisti@ece.utoronto.ca) is with the University of Toronto, Toronto, ON, Canada. Aditya Mahajan (aditya.mahajan@mcgill.ca) is with McGill University, Montreal, QC, Canada.

Preliminary version of this work was presented at the Signal Processing for Advanced Wireless Communications (SPAWC), Jul. 2015, Stockholm, Sweden, International Zurich Seminar (IZS), Mar. 2-4, 2016, Zurich, Switzerland and the American Control Conference (ACC), July 6-8, 2016, Boston, MA, USA.

similar in spirit to the dynamic programs obtained to compute capacity of channels with memory [16]–[18]; however, the specific details are different due to the constraint on the battery state. Second, when the demand is i.i.d., we obtain a single letter characterization of the minimum information leakage rate; this expression also gives the optimal battery management policy. We prove the single letter expression in two steps. On the achievability side we propose a class of policies with a specific structure that enables a considerable simplification of the leakage-rate expression. We find a policy that minimizes the leakage-rate within this restricted class. On the converse side, we obtain lower bounds on the minimal leakage rate and show that these lower bound match the performance of the best structured policy. We provide two proofs. One is based on the dynamic program and the other is based purely on information theoretic arguments.

After the present work was completed, we became aware of [21], where a similar dynamic programming framework is presented for the infinite horizon case. However, no explicit solutions of the dynamic program are derived in [21]. To the best of our knowledge, the present paper is the first work that provides an explicit characterization of the optimal leakage rate and the associated policy for i.i.d. demand.

A. Notation

Random variables are denoted by uppercase letters (X, Y , etc.), their realization by corresponding lowercase letters (x, y , etc.), and their state space by corresponding script letters (\mathcal{X}, \mathcal{Y} , etc.). \mathcal{P}_X denotes the space of probability distributions on \mathcal{X} ; $\mathcal{P}_{X|Y}$ denotes the space of stochastic kernels from \mathcal{Y} to \mathcal{X} . X_a^b is a short hand for $(X_a, X_{a+1}, \dots, X_b)$ and $X^b = X_1^b$. For a set \mathcal{A} , $\mathbb{1}_{\mathcal{A}}(x)$ denotes the indicator function of the set that equals 1 if $x \in \mathcal{A}$ and zero otherwise. If \mathcal{A} is a singleton set $\{a\}$, we use $\mathbb{1}_a(x)$ instead of $\mathbb{1}_{\{a\}}(x)$.

Given random variables (X, Y) with joint distribution $P_{X,Y}(x, y) = P_X(x)q(y|x)$, $H(X)$ and $H(P_X)$ denote the entropy of X , $H(Y|X)$ and $H(q|P_X)$ denote conditional entropy of Y given X and $I(X; Y)$ and $I(q; P_X)$ denote the mutual information between X and Y .

II. PROBLEM FORMULATION AND MAIN RESULTS

A. Model and problem formulation

Consider a smart metering system as shown in Fig. 1. At each time, the energy consumed from the power grid must equal the user's demand plus the additional energy that is either stored in or drawn from the battery. Let $\{X_t\}_{t \geq 1}$, $X_t \in \mathcal{X}$, denote the user's demand; $\{Y_t\}_{t \geq 1}$, $Y_t \in \mathcal{Y}$, denote the energy drawn from the grid; and $\{S_t\}_{t \geq 1}$, $S_t \in \mathcal{S}$, denote the energy stored in the battery. All alphabets are finite. For convenience, we assume $\mathcal{X} := \{0, 1, \dots, m_x\}$, $\mathcal{Y} := \{0, 1, \dots, m_y\}$, and $\mathcal{S} = \{0, 1, \dots, m_s\}$. We note that such a restriction is for simplicity of presentation; the results generalize even when \mathcal{X} and \mathcal{Y} are not necessarily contiguous intervals or integer valued. To guarantee that user's demand is always satisfied, we assume $m_x \leq m_y$ or that $\mathcal{X} \subseteq \mathcal{Y}$ holds more generally.




Fig. 1: A smart metering system.

The demand $\{X_t\}_{t \geq 1}$ is a first-order time-homogeneous Markov chain¹ with transition probability Q . We assume that Q is irreducible and aperiodic. The initial state X_1 is distributed according to probability mass function P_{X_1} . The initial charge S_1 of the battery is independent of $\{X_t\}_{t \geq 1}$ and distributed according to probability mass function P_{S_1} .

The battery is assumed to be ideal and has no conversion losses or other inefficiencies. Therefore, the following conservation equation must be satisfied at all times:

$$S_{t+1} = S_t + Y_t - X_t. \quad (1)$$

Given the history of demand, battery charge, and consumption, a randomized *battery charging policy* $\mathbf{q} = (q_1, q_2, \dots)$ determines the energy consumed from the grid. In particular, given the histories (x^t, s^t, y^{t-1}) of demand, battery charge, and consumption at time t , the probability that current consumption Y_t equals y is $q_t(y | x^t, s^t, y^{t-1})$. For a randomized charging policy to be feasible, it must satisfy the conservation equation (1). So, given the current power demand and battery charge (x_t, s_t) , the feasible values of grid consumption are defined by

$$\mathcal{V}_o(s_t - x_t) = \{y \in \mathcal{Y} : s_t - x_t + y \in \mathcal{S}\}. \quad (2)$$

Thus, we require that

$$q_t(\mathcal{V}_o(s_t - x_t) | x^t, s^t, y^{t-1}) := \sum_{y \in \mathcal{V}_o(s_t - x_t)} q_t(y | x^t, s^t, y^{t-1}) = 1.$$

The set of all such feasible policies is denoted by \mathcal{Q}_A^2 . Note that while the charging policy $q_t(\cdot)$ can be a function of the entire history, the support of $q_t(\cdot)$ only depends on the present value of x_t and s_t through the difference $s_t - x_t$. This is emphasized in the definition in (2).

The quality of a charging policy depends on the amount of information leaked under that policy. There are different notions of privacy; in this paper, we use mutual information as a measure of privacy. Intuitively speaking, given random variables (Y, Z) , the mutual information $I(Y; Z)$ measures the decrease in the uncertainty about Y given by Z (or vice-versa). Therefore, given a policy \mathbf{q} , the information about (X^T, S_1) leaked to the utility provider or eavesdropper is

¹In practice, the energy demand is periodic rather than time homogeneous. We are assuming that the total consumption may be split into a periodic predictable component and a time-homogeneous stochastic component. In this paper, we ignore the predictable component because it does not affect privacy.

²With a slight abuse of notation, we use \mathcal{Q}_A to denote the battery policy for both the infinite and finite-horizon problems

captured by $I^{\mathbf{q}}(X^T, S_1; Y^T)$, where the mutual information is evaluated according to the joint probability distribution on (X^T, S^T, Y^T) induced by the distribution \mathbf{q} as follows:

$$\begin{aligned} & \mathbb{P}^{\mathbf{q}}(S^T = s^T, X^T = x^T, Y^T = y^T) \\ &= P_{S_1}(s_1)P_{X_1}(x_1)q_1(y_1 | x_1, s_1) \prod_{t=2}^T \left[\mathbb{1}_{s_t \{s_{t-1} - x_{t-1} + y_{t-1}\}} \right. \\ & \quad \left. \times Q(x_t | x_{t-1})q_t(y_t | x^t, s^t, y^{t-1}) \right]. \end{aligned}$$

We use information leakage *rate* as a measure of the quality of a charging policy. For a finite planning horizon, the information leakage rate of a policy $\mathbf{q} = (q_1, \dots, q_T) \in \mathcal{Q}_A$ is given by

$$L_T(\mathbf{q}) := \frac{1}{T} I^{\mathbf{q}}(X^T, S_1; Y^T), \quad (3)$$

while for an infinite horizon, the worst-case information leakage rate of a policy $\mathbf{q} = (q_1, q_2, \dots) \in \mathcal{Q}_A$ is given by

$$L_{\infty}(\mathbf{q}) := \limsup_{T \rightarrow \infty} L_T(\mathbf{q}). \quad (4)$$

We are interested in the following optimization problems:

Problem A. *Given the alphabet \mathcal{X} of the demand, the initial distribution P_{X_1} and the transition matrix Q of the demand process, the alphabet \mathcal{S} of the battery, the initial distribution P_{S_1} of the battery state, and the alphabet \mathcal{Y} of the consumption:*

- 1) *For a finite planning horizon T , find a battery charging policy $\mathbf{q} = (q_1, \dots, q_T) \in \mathcal{Q}_A$ that minimizes the leakage rate $L_T(\mathbf{q})$ given by (3).*
- 2) *For an infinite planning horizon, find a battery charging policy $\mathbf{q} = (q_1, q_2, \dots) \in \mathcal{Q}_A$ that minimizes the leakage rate $L_{\infty}(\mathbf{q})$ given by (4).*

The above optimization problem is difficult because we have to optimize a multi-letter mutual information expression over the class of history dependent probability distributions. In the spirit of results for feedback capacity of channels with memory [16]–[18], we show that the above optimization problem can be reformulated as a Markov decision process where the state and action spaces are conditional probability distributions. Thus, the optimal policy and the optimal leakage rate can be computed by solving an appropriate dynamic program. We then provide an explicit solution of the dynamic program for the case of i.i.d. demand.

B. Example: Binary Model

We illustrate the special case when $\mathcal{X} = \mathcal{Y} = \mathcal{S} = \{0, 1\}$ in Fig. 2. The input, output, as well as the state, are all binary valued. When the battery is in state $s_t = 0$, there are three possible transitions. If the input $x_t = 1$ then we must select $y_t = 1$ and the state changes to $s_{t+1} = 0$. If instead $x_t = 0$, then there are two possibilities. We can select $y_t = 0$ and have $s_{t+1} = 0$ or we can select $y_t = 1$ and have $s_{t+1} = 1$. In a similar fashion there are three possible transitions from the state $s_t = 1$ as shown in Fig. 2. We will




Fig. 2: Binary System model. The battery can be either in $s = 0$ or $s = 1$. The set of feasible transitions from each state are shown in the figure.

assume that the demand (input) sequence is sampled i.i.d. from an equiprobable distribution.

Consider a simple policy that sets $y_t = x_t$ and ignores the battery state. It is clear that such a policy will lead to maximum leakage $L_T = 1$. Another feasible policy is to set $y_t = \bar{s}_t$. Thus whenever $s_t = 0$, we will set $y_t = 1$ regardless of the value of x_t , and likewise $s_t = 1$ will result in $y_t = 0$. It turns out that the leakage rate for this policy also approaches 1. To see this note that the eavesdropper having access to y^T also in turn knows s^T . Using the battery update equation (1) the sequence x_1^{T-1} is thus revealed to the eavesdropper, resulting in a leakage rate of at least $1 - 1/T$.

In reference [8] a probabilistic battery charging policy is introduced that only depends on the current state and input i.e., $q_t(y_t | x^t, s^t) \triangleq q_t(y_t | x_t, s_t)$. Furthermore the policy makes equiprobable decisions between the feasible transitions i.e.,

$$q_t(y_t = 0 | x_t, s_t) = q_t(y_t = 1 | x_t, s_t) = 1/2, \quad x_t = s_t \quad (5)$$

and $q_t(y_t | x_t, s_t) = \mathbb{1}_{x_t}(y_t)$ otherwise. The leakage rate for this policy was numerically evaluated in [8] using the BCJR algorithm and it was shown numerically that $L_{\infty} = 0.5$. Such numerical techniques seem necessary in general even for the class of memoryless policies and i.i.d. inputs, as the presence of the battery adds memory into the system.

As a consequence of our main result it follows that the above policy admits a single-letter expression for the leakage rate³ $L_{\infty} = I(S^* - X; X)$, thus circumventing the need for numerical techniques. Furthermore it also follows that this leakage rate is indeed the minimum possible one among the class of all feasible policies. Thus it is not necessary for the battery system to use more complex policies that take into account the entire history. We note that a similar result was shown in [30] for the case of finite horizon policies. However the proof in [30] is specific to the binary model. In the present paper we provide a complete single-letter solution to the case of general i.i.d. demand, and a dynamic programming method for the case of first-order Markovian demands, as discussed next.

C. Main results for Markovian demand

We identify two structural simplifications for the battery charging policies. First, we show (see Proposition 1 in Sec-

³The random variable S^* is an equiprobable binary valued random variable, independent of X .

tion III-A) that there is no loss of optimality in restricting attention to charging strategies of the form

$$q_t(y_t | x_t, s_t, y^{t-1}). \quad (6)$$

The intuition is that under such a policy, observing y^t gives partial information only about (x_t, s_t) rather than about the whole history (x^t, s^t) .

Next, we identify a sufficient statistic for y^{t-1} is the charging strategies of the form (6). For that matter, given a policy \mathbf{q} and any realization y^{t-1} of Y^{t-1} , define the belief state $\pi_t \in \mathcal{P}_{X,S}$ as follows:

$$\pi_t(x, s) = \mathbb{P}^{\mathbf{q}}(X_t = x, S_t = s | Y^{t-1} = y^{t-1}) \quad (7)$$

Then, we show (see Theorem 1 below) that there is no loss of optimality in restricting attention to charging strategies of the form

$$q_t(y_t | x_t, s_t, \pi_t). \quad (8)$$

Such a charging policy is Markovian in the belief state π_t and the optimal policies of such form can be searched using a dynamic program.

To describe such a dynamic program, we assume that there is a decision maker that observes y^{t-1} (or equivalently π_t) and chooses ‘‘actions’’ $a_t = q_t(\cdot | \cdot, \cdot, \pi_t)$ using some decision rule $a_t = f_t(\pi_t)$. We then identify a dynamic program to choose the optimal decision rules.

Note that the actions a_t take values in a subset \mathcal{A} of $\mathcal{P}_{Y|X,S}$ given by

$$\mathcal{A} = \{a \in \mathcal{P}_{Y|X,S} : a(\mathcal{Y}_o(s-x) | x, s) = 1, \forall (x, s) \in \mathcal{X} \times \mathcal{S}\}. \quad (9)$$

To succinctly write the dynamic program, for any $a \in \mathcal{A}$, we define the Bellman operator $\mathcal{B}_a : [\mathcal{P}_{X,S} \rightarrow \mathbb{R}] \rightarrow [\mathcal{P}_{X,S} \rightarrow \mathbb{R}]$ as follows: for all $\pi \in \mathcal{P}_{X,S}$,

$$\begin{aligned} [\mathcal{B}_a V](\pi) &= I(a; \pi) \\ &+ \sum_{\substack{x \in \mathcal{X}, s \in \mathcal{S}, \\ y \in \mathcal{Y}}} \pi(x, s) a(y | x, s) V(\varphi(\pi, y, a)) \end{aligned} \quad (10)$$

where the function φ is a non-linear filtering equation defined in Sec. III-C.

Our main result is the following:

Theorem 1. *In Problem A there is no loss of optimality to restrict attention to charging policies of the form (8).*

- 1) *For the finite horizon T , we can identify the optimal policy $\mathbf{q}^* = (q_1^*, \dots, q_T^*)$ of the form (8) by iteratively defining value functions $V_t : \mathcal{P}_{X,S} \rightarrow \mathbb{R}$ as follows. For any $\pi \in \mathcal{P}_{X,S}$, $V_{T+1}(\pi) = 0$, and for $t \in \{T-1, \dots, 1\}$,*

$$V_t(\pi) = \min_{a \in \mathcal{A}} [\mathcal{B}_a V_{t+1}](\pi), \quad \forall \pi \in \mathcal{P}_{X,S}. \quad (11)$$

Let $f_t^(\pi)$ denote the arg min of the right hand side of (11). Then, optimal policy $\mathbf{q}^* = (q_1^*, \dots, q_T^*)$ is given by*

$$q_t^*(y_t | x_t, s_t, \pi_t) = a_t(y_t | x_t, s_t), \text{ where } a_t = f_t^*(\pi_t).$$

Moreover, the optimal (finite horizon) leakage rate is given by $V_1(\pi_1)/T$, where $\pi_1(x, s) = P_{X_1}(x)P_{S_1}(s)$.

- 2) *For the infinite horizon, the optimal charging policy of the form (8) is time-homogeneous and is given by the following fixed point equation*

$$J + v(\pi) = \min_{a \in \mathcal{A}} [\mathcal{B}_a v](\pi), \quad \forall \pi \in \mathcal{P}_{X,S}, \quad (12)$$

where $J \in \mathbb{R}$ is a constant and $v : \mathcal{P}_{X,S} \rightarrow \mathbb{R}$. Let $f^(\pi)$ denote the arg min of the right hand side of (12). Then, the time-homogeneous optimal policy $\mathbf{q}^* = (q^*, q^*, \dots)$ given by*

$$q^*(y_t | x_t, s_t, \pi_t) = a_t(y_t | x_t, s_t), \text{ where } a_t = f^*(\pi_t)$$

is optimal. Moreover, the optimal (infinite horizon) leakage rate is given by J . \square

See Section III for proof.

The dynamic program above resembles the dynamic program for partially observable Markov decision processes (POMDP) with hidden state (X_t, S_t) , observation Y_t , and action A_t . However, in contrast to POMDPs, the expected per-step cost $I(a; \pi)$ is not linear in π . Nonetheless, one could use computational techniques from POMDPs to approximately solve the dynamic programs of Theorem 1. See Section III for a brief discussion.

D. Main result for i.i.d. demand

Assume the following:

- (A) *The demand $\{X_t\}_{t \geq 1}$ is i.i.d. with probability distribution P_X .*

We provide an explicit characterization of optimal policy and optimal leakage rate for this case.

Define an auxiliary state variable $W_t = S_t - X_t$ that takes values in $\mathcal{W} = \{s - x : s \in \mathcal{S}, x \in \mathcal{X}\}$. For any $w \in \mathcal{W}$, define:

$$\mathcal{D}(w) = \{(x, s) \in \mathcal{X} \times \mathcal{S} : s - x = w\}. \quad (13)$$

Then, we have the following.

Theorem 2. *Define*

$$J^* = \min_{\theta \in \mathcal{P}_{\mathcal{S}}} I(S - X; X) \quad (14)$$

where X and S are independent with $X \sim P_X$ and $S \sim \theta$. Let θ^ denote the arg min in (14). Define $\xi^*(w) = \sum_{(x,s) \in \mathcal{D}(w)} P_X(x)\theta^*(s)$. Then, under (A)*

- 1) *J^* is the optimal (infinite horizon) leakage rate.*
- 2) *Define $b^* \in \mathcal{P}_{Y|W}$ as follows:*

$$b^*(y|w) = \begin{cases} P_X(y) \frac{\theta^*(y+w)}{\xi^*(w)} & \text{if } y \in \mathcal{X} \cap \mathcal{Y}_o(w) \\ 0 & \text{otherwise} \end{cases} \quad (15)$$

Then, the memoryless charging policy $\mathbf{q}^ = (q_1^*, q_2^*, \dots)$ given by*

$$q_t^*(y | x_t, s_t, \pi_t) = b^*(y | s_t - x_t) \quad (16)$$

is optimal and achieves the optimal (infinite horizon) leakage rate. \square

Note that the optimal charging policy is *memoryless*, i.e., the distribution on Y_t does not depend on π_t (and, therefore on y^{t-1}).

The proof, which is presented in Section IV, is based on the standard arguments of showing achievability and a converse. On the achievability side we show that the policy in (15) belongs to a class of policies that satisfies a certain invariance property. Using this property the multi-letter mutual information expression can be reduced into a single-letter expression. For the converse we provide two proofs. The first is based on a simplification of the dynamic program of Theorem 1. The second is based on purely probabilistic and information theoretic arguments.

E. Salient features of the result for i.i.d. demand

Theorem 2 shows that even if consumption could take larger values than the demand, i.e., $\mathcal{Y} \supset \mathcal{X}$, under policy Y_t takes values only in \mathcal{X} . This agrees with the intuition that a consumption larger than m_x reveals that the battery has low charge and that the power demand is high. In extreme cases, a large consumption may completely reveal the battery and power usage thereby increasing the information leakage.

We now show some other properties of the optimal policy.

Property 1. The mutual information $I(S - X; X)$ is equal to $H(S - X) - H(S)$.

Proof. This follows from the following simplifications:

$$\begin{aligned} I(S - X; X) &= H(S - X) - H(S - X|X) \\ &= H(S - X) - H(S|X) \\ &= H(S - X) - H(S) \quad \square \end{aligned}$$

Property 2. The mutual information $I(S - X; X)$ is strictly convex in the distribution θ and, therefore, $\theta^* \in \text{int}(\mathcal{P}_S)$.

See Appendix A for proof.

As a consequence, the optimal θ^* in (14) may be obtained using the Blahut-Arimoto algorithm [19], [20].

Property 3. Under the battery charging policy specified in Theorem 2, the power consumption $\{Y_t\}_{t \geq 1}$ is i.i.d. with marginal distribution P_X . Thus, $\{Y_t\}_{t \geq 1}$ is statistically indistinguishable from $\{X_t\}_{t \geq 1}$.

See Remarks 1 and 3 in Section IV-B for proof.

Property 4. If the power demand has a symmetric PMF, i.e., for any $x \in \mathcal{X}$, $P_X(x) = P_X(m_x - x)$, then the optimal θ^* in Theorem 2 is also symmetric, i.e., for any $s \in \mathcal{S}$, $\theta^*(s) = \theta^*(m_s - s)$.

Proof. For $\theta \in \mathcal{P}_S$, define $\bar{\theta}(s) = \theta(m_s - s)$. Let $X \sim P_X$, $S \sim \theta$ and $\bar{S} \sim \bar{\theta}$. Then, by symmetry

$$I(S - X; X) = I(\bar{S} - X; X). \quad (17)$$

For any $\lambda \in (0, 1)$, let $\theta_\lambda(s) = \lambda\theta(s) + (1 - \lambda)\bar{\theta}(s)$ denote the convex combination of θ and $\bar{\theta}$. Let $S_\lambda \sim \theta_\lambda$. By Property 2, if $\theta \neq \bar{\theta}$, then

$$\begin{aligned} I(S_\lambda - X; X) &< \lambda I(S - X; X) + (1 - \lambda) I(\bar{S} - X; X) \\ &= I(S - X; X), \end{aligned}$$

where the last equation uses (17).

Thus, if $\theta \neq \bar{\theta}$, we can strictly decrease the mutual information by using θ_λ . Hence, the optimal distribution must have the property that $\theta^*(s) = \theta^*(m_s - s)$. \square

Given a distribution μ on some alphabet \mathcal{M} , we say that the distribution is *almost symmetric and unimodal* around $m^* \in \mathcal{M}$ if

$$\mu_{m^*} \geq \mu_{m^*+1} \geq \mu_{m^*-1} \geq \mu_{m^*+2} \geq \mu_{m^*-2} \geq \dots$$

where we use the interpretation that for $m \notin \mathcal{M}$, $\mu_m = 0$. Similarly, we say that the distribution is *symmetric and unimodal* around $m^* \in \mathcal{M}$ if

$$\mu_{m^*} \geq \mu_{m^*+1} = \mu_{m^*-1} \geq \mu_{m^*+2} = \mu_{m^*-2} \geq \dots$$

Note that a distribution can be symmetric and unimodal only if its support is odd.

Property 5. If the power demand is symmetric and unimodal around $\lfloor m_x/2 \rfloor$, then the optimal θ^* in Theorem 2 is *almost symmetric and unimodal* with around $\lfloor m_s/2 \rfloor$. In particular, if m_s is even, then

$$\theta_{m^*}^* \geq \theta_{m^*+1}^* = \theta_{m^*-1}^* \geq \theta_{m^*+2}^* = \theta_{m^*-2}^* \geq \dots$$

and if m_s is odd then

$$\theta_{m^*}^* = \theta_{m^*+1}^* \geq \theta_{m^*-1}^* = \theta_{m^*+2}^* \geq \theta_{m^*-2}^* = \dots$$

where $m^* = \lfloor m_s/2 \rfloor$.

Proof. Let $\bar{X} = -X$. Then, $I(S - X; S) = H(S - X) - H(S) = H(S + \bar{X}) - H(S)$. Note that $P_{\bar{X}}$ is also symmetric and unimodal around $\lfloor m_x/2 \rfloor$.

Let S° and \bar{X}° denote the random variables $S - \lfloor m_s/2 \rfloor$ and $\bar{X} - \lfloor m_x/2 \rfloor$. Then \bar{X}° is also symmetric and unimodal around origin and

$$I(S - X; X) = H(S^\circ + \bar{X}^\circ) - H(S^\circ).$$

Now given any distribution θ° of S° , let θ^+ be a permutation of θ° that is almost symmetric and unimodal with a positive bias around origin. Then by [22, Corollary III.2], $H(P_X * \theta^\circ) \geq H(P_X * \theta^+)$. Thus, the optimal distribution must have the property that $\theta^\circ = \theta^+$ or, equivalently, θ is almost unimodal and symmetric around $\lfloor m_s/2 \rfloor$.

Combining this with the result of property 4 gives the characterization of the distribution when m_s is even or odd. \square

F. Numerical Example: i.i.d. demand

Suppose there are n identical devices in the house and each is on with probability p . Thus, $X \sim \text{Binomial}(n, p)$. We derive the optimal policy and optimal leakage rate for this scenario under the assumption that $\mathcal{Y} = \mathcal{X}$. We consider two specific examples, where we numerically solve (14).

Suppose $n = 6$ and $p = 0.5$.

- 1) Consider $\mathcal{S} = [0:5]$. Then, by numerically solving (14), we get that the optimal leakage rate J^* is 0.4616 and the optimal battery charge distribution θ^* is

$$\{0.1032, 0.1747, 0.2221, 0.2221, 0.1747, 0.1032\}.$$

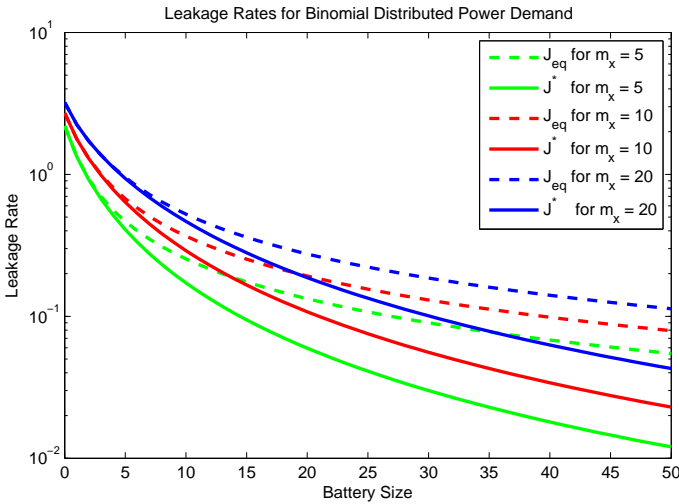


Fig. 3: A comparison of the performance of $q_{eq} \in \mathcal{Q}_B$ as defined in (18) with the optimal leakage rate for i.i.d. Binomial distributed demand $\text{Binomial}(m_x, 0.5)$ for $m_x = \{5, 10, 20\}$.

- 2) Consider $\mathcal{S} = [0:6]$. Then, by numerically solving (14), we get that the optimal leakage rate J^* is 0.3774 and the optimal battery charge distribution θ^* is

$$\{0.0773, 0.1364, 0.1847, 0.2031, 0.1847, 0.1364, 0.0773\}.$$

Note that both results are consistent with Properties 4 and 5.

We next compare the performance with the following time-homogeneous benchmark policy $q_{eq} \in \mathcal{Q}_B$: for all $y \in \mathcal{Y}, w \in \mathcal{W}$,

$$q_t(y_t|w_t) = \frac{\mathbb{1}_{\mathcal{Y}^\circ(w_t)}\{y_t\}}{|\mathcal{Y}^\circ(w_t)|}. \quad (18)$$

This benchmark policy chooses all feasible values of Y_t with equal probability. For that reason we call it *equi-probably policy* and denote its performance by J_{eq} .

In Fig. 3, we compare the performance of \mathbf{q}^* and \mathbf{q}_{eq} as a function of battery sizes for different demand alphabets.

Under \mathbf{q}_{eq} , the MDP converges to a belief state that is approximately uniform. Hence, in the low battery size regime, J_{eq} is close to optimal but its performance gradually worsens with increasing battery size.

III. PROOF OF THEOREM 1

One of the difficulties in obtaining a dynamic programming decomposition for Problem A is that the objective function is not of the form $\sum_{t=1}^T \text{cost}(\text{state}_t, \text{action}_t)$. We show that there is no loss of optimality to restrict attention to a class of policies \mathcal{Q}_B and for any policy in \mathcal{Q}_B , the mutual information may be written in an additive form.

A. Simplification of optimal charging policies

Let $\mathcal{Q}_B \subset \mathcal{Q}_A$ denote the set of charging policies that choose consumption based only on the consumption history, current demand, and battery state. Thus, for $\mathbf{q} \in \mathcal{Q}_B$, at any time t , given history (x^t, s^t, y^{t-1}) , the consumption Y_t is y

with probability $q_t(y | x_t, s_t, y^{t-1})$. Then the joint distribution on (X^T, S^T, Y^T) induced by $\mathbf{q} \in \mathcal{Q}_B$ is given by

$$\begin{aligned} & \mathbb{P}^{\mathbf{q}}(S^T = s^T, X^T = x^T, Y^T = y^T) \\ &= P_{S_1}(s_1)P_{X_1}(x_1)q_1(y_1 | x_1, s_1) \prod_{t=2}^T \left[\mathbb{1}_{s_t}\{s_{t-1}-x_{t-1}+y_{t-1}\} \right. \\ & \quad \left. \times Q(x_t|x_{t-1})q_t(y_t | x_t, s_t, y^{t-1}) \right]. \end{aligned}$$

Proposition 1. *In Problem A, there is no loss of optimality in restricting attention to charging policies in \mathcal{Q}_B . Moreover, for any $\mathbf{q} \in \mathcal{Q}_B$, the objective function takes an additive form:*

$$L_T(\mathbf{q}) = \frac{1}{T} \sum_{t=1}^T I^{\mathbf{q}}(X_t, S_t; Y_t | Y^{t-1})$$

where

$$\begin{aligned} & I^{\mathbf{q}}(X_t, S_t; Y_t | Y^{t-1}) \\ &= \sum_{\substack{x_t \in \mathcal{X}, s_t \in \mathcal{S} \\ y^t \in \mathcal{Y}^t}} \mathbb{P}^{\mathbf{q}}(X_t = x_t, S_t = s_t, Y^t = y^t) \\ & \quad \times \log \frac{q_t(y_t | x_t, s_t, y^{t-1})}{\mathbb{P}^{\mathbf{q}}(Y_t = y_t | Y^{t-1} = y^{t-1})}. \end{aligned}$$

See Appendix B for proof. The intuition behind why policies in \mathcal{Q}_B are better than those in \mathcal{Q}_A is as follows. For a policy \mathcal{Q}_A , observing the realization y^t of Y^t gives partial information about the history (x^t, s^t) while for a policy \mathcal{Q}_B , y_t gives partial information only about the current state (x_t, s_t) . The dependence on (x_t, s_t) cannot be removed because of the conservation constraint (1).

Proposition 1 shows that the total cost may be written in an additive form. Next we use an approach inspired by [16]–[18] and formulate an equivalent sequential optimization problem.

B. An equivalent sequential optimization problem

Consider a system with state process $\{X_t, S_t\}_{t \geq 1}$ where $\{X_t\}_{t \geq 1}$ is an exogenous Markov process as before and $\{S_t\}_{t \geq 1}$ is a controlled Markov process as specified below. At time t , a decision maker observes Y^{t-1} and chooses a distribution valued action $A_t \in \mathcal{A}$, where \mathcal{A} is given by (9), as follows:

$$A_t = f_t(Y^{t-1}) \quad (19)$$

where $\mathbf{f} = (f_1, f_2, \dots)$ is called the decision policy.

Based on this action, an auxiliary variable $Y_t \in \mathcal{Y}$ is chosen according to the conditional probability $a_t(\cdot | x_t, s_t)$ and the state S_{t+1} evolves according to (1).

At each stage, the system incurs a per-step cost given by

$$c_t(x_t, s_t, a_t, y^t; \mathbf{f}) := \log \frac{a_t(y_t | x_t, s_t)}{\mathbb{P}^{\mathbf{f}}(Y_t = y_t | Y^{t-1} = y^{t-1})}. \quad (20)$$

The objective is to choose a policy $\mathbf{f} = (f_1, \dots, f_T)$ to minimize the total finite horizon cost given by

$$\tilde{L}_T(\mathbf{f}) := \frac{1}{T} \mathbb{E}^{\mathbf{f}} \left[\sum_{t=1}^T c_t(x_t, s_t, a_t, y^t; \mathbf{f}) \right] \quad (21)$$

where the expectation is evaluated with respect to the probability distribution $\mathbb{P}^{\mathbf{f}}$ induced by the decision policy \mathbf{f} .

Proposition 2. *The sequential decision problem described above is equivalent to Problem A. In particular,*

1) Given $\mathbf{q} = (q_1, \dots, q_T) \in \mathcal{Q}_B$, let $\mathbf{f} = (f_1, \dots, f_T)$ be

$$f_t(y^{t-1}) = q_t(\cdot | \cdot, \cdot, y^{t-1}).$$

Then $\tilde{L}_T(\mathbf{f}) = L_T(\mathbf{q})$.

2) Given $\mathbf{f} = (f_1, \dots, f_T)$, let $\mathbf{q} = (q_1, \dots, q_T) \in \mathcal{Q}_B$ be

$$q_t(y_t | x_t, s_t, y^{t-1}) = a_t(y_t | x_t, s_t), \text{ where } a_t = f_t(y^{t-1}).$$

Then $L_T(\mathbf{q}) = \tilde{L}_T(\mathbf{f})$.

Proof. For any history (x^t, s^t, y^{t-1}) , $a^t \in \mathcal{A}$, and $s_{t+1} \in \mathcal{S}$,

$$\begin{aligned} & \mathbb{P}(S_{t+1} = s_{t+1} | X^t = x^t, S^t = s^t, Y^t = y^t, A^t = a^t) \\ &= \sum_{y_t \in \mathcal{Y}} \mathbb{1}_{s_{t+1}} \{s_t + y_t - x_t\} a_t(y_t | x_t, s_t) \\ &= \mathbb{P}(S_{t+1} = s_{t+1} | X_t = x_t, S_t = s_t, A_t = a_t). \end{aligned} \quad (22)$$

Thus, the probability distribution on (X^T, S^T, Y^T) induced by a decision policy $\mathbf{f} = (f_1, \dots, f_T)$ is given by

$$\begin{aligned} & \mathbb{P}^{\mathbf{f}}(S^T = s^T, X^T = x^T, Y^T = y^T) \\ &= P_{S_1}(s_1) P_{X_1}(x_1) q_1(y_1 | x_1, s_1) \prod_{t=2}^T \left[\mathbb{1}_{s_t} \{s_{t-1} - x_{t-1} + y_{t-1}\} \right. \\ & \quad \left. \times Q(x_t | x_{t-1}) a_t(y_t | x_t, s_t) \right]. \end{aligned}$$

where $a_t = f_t(y^{t-1})$. Under the transformations described in the Proposition, $\mathbb{P}^{\mathbf{f}}$ and $\mathbb{P}^{\mathbf{q}}$ are identical probability distributions. Consequently, $\mathbb{E}^{\mathbf{f}}[c_t(X_t, S_t, A_t, Y^t; \mathbf{f})] = I^{\mathbf{q}}(S_t, X_t; Y_t | Y^{t-1})$. Hence, $L_T(\mathbf{q})$ and $\tilde{L}_T(\mathbf{f})$ are equivalent. \square

Eq. (22) implies that $\{X_t, S_t\}_{t \geq 1}$ is a controlled Markov process with control action $\{A_t\}_{t \geq 1}$. In the next section, we obtain a dynamic programming decomposition for this problem. For the purpose of writing the dynamic program, it is more convenient to write the policy (19) as

$$A_t = f_t(Y^{t-1}, A^{t-1}). \quad (23)$$

Note that these two representations are equivalent. Any policy of the form (19) is also a policy of the form (23) (that simply ignores A^{t-1}); any policy of the form (23) can be written as a policy of the form (19) by recursively substituting A_t in terms of Y^{t-1} . Since the two forms are equivalent, in the next section we assume that the policy is of the form (23).

C. A dynamic programming decomposition

The model described in Section III-B above is similar to a POMDP (partially observable Markov decision process): the system state (X_t, S_t) is partially observed by a decision maker who chooses action A_t . However, in contrast to the standard cost model used in POMDPs, the per-step cost depends on the observation history and *past policy*. Nonetheless, if we consider the belief state as the information state, the problem can be formulated as a standard MDP.

For that matter, for any realization y^{t-1} of past observations and any choice a^{t-1} of past actions, define the belief state $\pi_t \in \mathcal{P}_{X,S}$ as follows: For $s \in \mathcal{S}$ and $x \in \mathcal{X}$,

$$\pi_t(x, s) = \mathbb{P}^{\mathbf{f}}(X_t = x, S_t = s | Y^{t-1} = y^{t-1}, A^{t-1} = a^{t-1}).$$

If Y^{t-1} and A^{t-1} are random variables, then the belief state is a $\mathcal{P}_{X,S}$ -valued random variable.

The belief state evolves in a state-like manner as follows.

Lemma 1. *For any realization y_t of Y_t and a_t of A_t , π_{t+1} is given by*

$$\pi_{t+1} = \varphi(\pi_t, y_t, a_t) \quad (24)$$

where φ is given by

$$\begin{aligned} & \varphi(\pi, y, a)(x', s') \\ &= \frac{\sum_{x \in \mathcal{X}} Q(x' | x) a(y | x, s' - x + y) \pi(x, s' - x + y)}{\sum_{(x,s) \in \mathcal{X} \times \mathcal{S}} a(y | x, s) \pi(x, s)}. \end{aligned}$$

Proof. For ease of notation, we use $\mathbb{P}(x_t, s_t | y^{t-1}, a^{t-1})$ to denote $\mathbb{P}(X_t = x_t, S_t = s_t | Y^{t-1} = y^{t-1}, A^{t-1} = a^{t-1})$. Similar interpretations hold for other expressions as well. Consider

$$\begin{aligned} & \pi_{t+1}(x_{t+1}, y_{t+1}) = \mathbb{P}(x_{t+1}, s_{t+1} | y^t, a^t) \\ &= \frac{\mathbb{P}(x_{t+1}, s_{t+1}, y_t, a_t | y^{t-1}, a^{t-1})}{\mathbb{P}(y_t, a_t | y^{t-1}, a^{t-1})} \end{aligned} \quad (25)$$

Now, consider the numerator of the right hand side.

$$\begin{aligned} & \mathbb{P}(x_{t+1}, s_{t+1}, y_t, a_t | y^{t-1}, a^{t-1}) \\ &= \mathbb{P}(x_{t+1}, s_{t+1}, y_t, a_t | y^{t-1}, a^{t-1}, \pi_t) \\ &= \sum_{(x_t, s_t) \in \mathcal{X} \times \mathcal{S}} \mathbb{P}(x_{t+1} | x_t) \mathbb{1}_{s_{t+1}}(s_t + x_t - y_t) \\ & \quad \times a_t(y_t | x_t, s_t) \mathbb{1}_{a_t}(f_t(\pi_t)) \pi_t(x_t, s_t) \end{aligned} \quad (26)$$

Substituting (26) in (25) (and observing that the denominator of the right hand side of (25) is the marginal of the numerator over (x_{t+1}, s_{t+1})), we get that π_{t+1} can be written in terms of π_t , y_t and a_t . Note that if the term $\mathbb{1}_{a_t}(f_t(\pi_t))$ is 1, it cancels from both the numerator and the denominator; if it is 0, we are conditioning on a null event in (25), so we can assign any valid distribution to the conditional probability. \square

Note that an immediate implication of the above result is that π_t depends only on (y^{t-1}, a^{t-1}) and not on the policy \mathbf{f} . This is the main reason that we are working with a policy of the form (23) rather than (19).

Lemma 2. *The cost $\tilde{L}_T(\mathbf{f})$ in (21) can be written as*

$$\tilde{L}_T(\mathbf{f}) = \frac{1}{T} \mathbb{E} \left[\sum_{t=1}^T I(a_t; \pi_t) \right]$$

where $I(a_t; \pi_t)$ does not depend on the policy \mathbf{f} and is computed according to the standard formula

$$\begin{aligned} I(a_t; \pi_t) &= \sum_{\substack{x \in \mathcal{X}, s \in \mathcal{S}, \\ y \in \mathcal{Y}}} \pi_t(x, s) a_t(y | x, s) \\ & \quad \times \log \frac{a_t(y | x, s)}{\sum_{(\tilde{x}, \tilde{s}) \in \mathcal{X} \times \mathcal{S}} \pi_t(\tilde{x}, \tilde{s}) a_t(y | \tilde{x}, \tilde{s})}. \end{aligned}$$

Proof. By the law of iterated expectations, we have

$$\tilde{L}_T(\mathbf{f}) = \frac{1}{T} \left[\sum_{t=1}^T \mathbb{E}[c_t(X_t, S_t, A_t, Y^t; \mathbf{f}) | Y^{t-1}, A^{t-1}] \right]$$

Now, from (20), each summand may be written as

$$\begin{aligned} \mathbb{E}^f[c_t(X_t, S_t, A_t, Y^t; \mathbf{f}) | Y^{t-1} = y^{t-1}, A^{t-1} = a^{t-1}] \\ = \sum_{\substack{x \in \mathcal{X}, s \in \mathcal{S}, \\ y \in \mathcal{Y}}} \pi_t(x, s) a_t(y | x, s) \\ \quad \times \log \frac{a_t(y | x, s)}{\sum_{(\tilde{x}, \tilde{s}) \in \mathcal{X} \times \mathcal{S}} \pi_t(\tilde{x}, \tilde{s}) a_t(y | \tilde{x}, \tilde{s})} \\ = I(a_t; \pi_t). \end{aligned}$$

□

Proof of Theorem 1. Lemma 1 implies that $\{\pi_t\}_{t \geq 1}$ is a controlled Markov process with control action a_t . In addition, Lemma 2 implies that the objective function can be expressed in terms of the state π_t and the action A_t . Consequently, in the equivalent optimization problem described in Section III-B, there is no loss of optimality to restrict attention to Markovian policies of the form $a_t = f_t(\pi_t)$; an optimal policy of this form is given by the dynamic programs of Theorem 1 [24]. Proposition 2 implies that this dynamic program also solves Problem A.

Note that the Markov chain $\{X_t\}_{t \geq 1}$ is irreducible and aperiodic and the per-step cost $I(a_t; \pi_t)$ is bounded. Therefore, the existence of the solution of the infinite horizon dynamic program can be shown by following the proof argument of [26]. □

D. Remarks about numerical solution

The dynamic program of Theorem 1, both state and action spaces are distribution valued (and, therefore, subsets of Euclidean space). Although, an exact solution of the dynamic program is not possible, there are two approaches to obtain an approximate solution. The first is to treat it as a dynamic program of an MDP with continuous state and action spaces and use approximate dynamic programming [23], [24]. The second is to treat it as a dynamic program for a PODMP and use point-based methods [25]. The point-based methods rely on concavity of the value function, which we establish below.

Proposition 3. *The value functions $\{V_t\}_{t=1}^T$ defined in Theorem 1 are concave.*

See Appendix C for proof.

IV. PROOF OF THEOREM 2

A. Simplification of the dynamic program

Under (A), the belief state π_t can be decomposed into product form $\pi_t(x, s) = P_X(x)\theta_t(s)$, where

$$\theta_t(s) = \mathbb{P}^f(S_t = s | Y^{t-1} = y^{t-1}, A^{t-1} = a^{t-1}).$$

Thus, in principle, we can simplify the dynamic program of Theorem 1 by using θ_t as an information state. However, for reasons that will become apparent, we provide an alternative simplification that uses an information state $\xi_t \in \mathcal{P}_W$.

Recall that $W_t = S_t - X_t$ which takes values in $\mathcal{W} = \{s - x : s \in \mathcal{S}, x \in \mathcal{X}\}$. For any realization (y^{t-1}, a^{t-1}) of past observations and actions, define $\xi_t \in \mathcal{P}_W$ as follows: for any $w \in \mathcal{W}$,

$$\xi_t(w) = \mathbb{P}^f(W_t = w | Y^{t-1} = y^{t-1}, A^{t-1} = a^{t-1}).$$

If Y^{t-1} and A^{t-1} are random variables, then ξ_t is a \mathcal{P}_W -valued random variable. As was the case for π_t , it can be shown that ξ_t does not depend on the choice of the policy \mathbf{f} .

Lemma 3. *Under (A), θ_t and ξ_t are related as follows:*

- 1) $\xi_t(w) = \sum_{(x,s) \in \mathcal{D}(w)} P_X(x)\theta_t(s)$.
- 2) $\theta_t = P_X * \xi_t$.

Proof. For part 1):

$$\begin{aligned} \xi_t(w) &= \mathbb{P}^f(W_t = w | Y^{t-1} = y^{t-1}, A^{t-1} = a^{t-1}) \\ &= \mathbb{P}^f(S_t - X_t = w | Y^{t-1} = y^{t-1}, A^{t-1} = a^{t-1}) \\ &= \sum_{(x,s) \in \mathcal{D}(w)} P_X(x)\theta_t(s). \end{aligned}$$

For part 2):

$$\begin{aligned} \theta_t(s) &= \mathbb{P}^f(S_t = s_t | Y^{t-1} = y^{t-1}, A^{t-1} = a^{t-1}) \\ &= P(W_t + X_t = s_t | y^{t-1}, a^{t-1}) \\ &= (P_X * \xi_t)(s_t). \end{aligned}$$

□

Since $\pi_t(x, s) = P_X(x)\theta_t(s)$, Lemma 3 suggests that one could simplify the dynamic program of Theorem 1 by using ξ_t as the information state instead of π_t . For such a simplification to work, we would have to use charging policies of the form $q_t(y_t | w_t, y^{t-1})$. We establish that restricting attention to such policies is without loss of optimality. For that matter, define \mathcal{B} as follows:

$$\mathcal{B} = \{b \in \mathcal{P}_{Y|W} : b(\mathcal{Y}_o(w) | w) = 1, \forall w \in \mathcal{W}\}. \quad (27)$$

Lemma 4. *Given $a \in \mathcal{A}$ and $\pi \in \mathcal{P}_{X,S}$, define the following:*

- $\xi \in \mathcal{P}_W$ as $\xi(w) = \sum_{(x,s) \in \mathcal{D}(w)} \pi(x, s)$
- $b \in \mathcal{B}$ as follows: for all $y \in \mathcal{Y}, w \in \mathcal{W}$

$$b(y | w) = \frac{\sum_{(x,s) \in \mathcal{D}(w)} a(y | x, s)\pi(x, s)}{\xi(w)};$$

- $\tilde{a} \in \mathcal{A}$ as follows: for all $y \in \mathcal{Y}, x \in \mathcal{X}, s \in \mathcal{S}$

$$\tilde{a}(y|x, s) = b(y|s - x).$$

Then under (A), we have

- 1) *Invariant Transitions:* for any $y \in \mathcal{Y}$, $\varphi(\pi, y, a) = \varphi(\pi, y, \tilde{a})$.
- 2) *Lower Cost:* $I(a; \pi) \geq I(\tilde{a}; \pi) = I(b; \xi)$.

Therefore, in the sequential problem of Sec. III-B, there is no loss of optimality in restricting attention to actions $b \in \mathcal{B}$.

Proof. 1) Suppose $(X, S) \sim \pi$ and $W = S - X$, $S_+ = W + Y$, $X_+ \sim P_X$. We will compare $\mathbb{P}(S_+ | Y)$ when

$Y \sim a(\cdot|X, S)$ with when $Y \sim \tilde{a}(\cdot|X, S)$. Given $w \in \mathcal{W}$ and $y \in \mathcal{Y}$,

$$\begin{aligned} \mathbb{P}^a(W = w, Y = y) &= \sum_{(x,s) \in \mathcal{D}(w)} a(y|x, s)\pi(x, s) \\ &= \sum_{(x,s) \in \mathcal{D}(w)} b(y|w)\pi(x, s) \\ &\stackrel{(a)}{=} \sum_{(x,s) \in \mathcal{D}(w)} \tilde{a}(y|x, s)\pi(x, s) \\ &= \mathbb{P}^{\tilde{a}}(W = w, Y = y) \end{aligned} \quad (28)$$

where (a) uses that for all $(x, s) \in \mathcal{D}(w)$, $s - x = w$. Marginalizing (28) over W , we get that $\mathbb{P}^a(Y = y) = \mathbb{P}^{\tilde{a}}(Y = y)$. Since $S_+ = W + Y$, Eq. (28) also implies $\mathbb{P}^a(S_+ = s, Y = y) = \mathbb{P}^{\tilde{a}}(S_+ = s, Y = y)$. Therefore, $\mathbb{P}^a(S_+ = s|Y = y) = \mathbb{P}^{\tilde{a}}(S_+ = s|Y = y)$.

2) Let $(X, S) \sim \pi$ and $W = S - X$. Then $W \sim \xi$. Therefore, we have

$$I(a; \pi) = I^a(X, S; Y) \geq I^a(W; Y).$$

where the last inequality is the data-processing inequality. Under \tilde{a} , $(X, S) - W - Y$, therefore,

$$I(\tilde{a}; \pi) = I^{\tilde{a}}(X, S; Y) = I^{\tilde{a}}(W; Y).$$

Now, by construction, the joint distribution of (W, Y) is the same under a , \tilde{a} , and b . Thus,

$$I^a(W; Y) = I^{\tilde{a}}(W; Y) = I^b(W; Y).$$

Note that $I^b(W; Y)$ can also be written as $I(b; \xi)$. The result follows by combining all the above relations. \square

Once attention is restricted to actions $b \in \mathcal{B}$, the update of ξ_t may be expressed in terms of $b \in \mathcal{B}$ as follows:

Lemma 5. For any realization y_t of Y_t and b_t of B_t , ξ_{t+1} is given by

$$\xi_{t+1} = \tilde{\varphi}(\xi_t, y_t, a_t) \quad (29)$$

where $\tilde{\varphi}$ is given by

$$\begin{aligned} \tilde{\varphi}(\xi, y, b)(w_+) &= \frac{\sum_{x \in \mathcal{X}, w \in \mathcal{W}} P_X(x) \mathbb{1}_{w_+} \{y + w - x\} b(y|w) \xi(w)}{\sum_{w \in \mathcal{W}} b(y|w) \xi(w)}. \end{aligned}$$

Proof. The proof is similar to the proof of Lemma 5. \square

For any $b \in \mathcal{B}$ and $\xi \in \mathcal{P}_W$, let us define the Bellman operator $\tilde{\mathcal{B}}_b : [\mathcal{P}_W \rightarrow \mathbb{R}] \rightarrow [\mathcal{P}_W \rightarrow \mathbb{R}]$ as follows:

$$[\tilde{\mathcal{B}}_b V](\xi) = I(b; \xi) + \sum_{y \in \mathcal{Y}, w \in \mathcal{W}} \xi(w) b(y|w) V(\tilde{\varphi}(\xi, y, b)).$$

Theorem 3. Under assumption (A), there is no loss of optimality in restricting attention to optimal policies of the form $q_t(y_t|w_t, \xi_t)$ in Problem A.

1) For the finite horizon case, we can identify the optimal policy $\mathbf{q}^* = (q_1^*, \dots, q_T^*)$ by iteratively defining value functions $\tilde{V}_t : \mathcal{P}_W \rightarrow \mathbb{R}$. For any $\xi \in \mathcal{P}_W$, $\tilde{V}_{T+1}(\xi) = 0$, and for $t = T, T-1, \dots, 1$,

$$\tilde{V}_t(\xi) = \min_{b \in \mathcal{B}} [\tilde{\mathcal{B}}_b \tilde{V}_{t+1}](\xi). \quad (30)$$

Let $f_t^\circ(\xi)$ denote the arg min of the right hand side of (30). Then, the optimal policy $\mathbf{q}^* = (q_1^*, \dots, q_T^*)$ is given by

$$q_t^*(y_t|w_t, \xi_t) = b_t(y_t|w_t), \text{ where } b_t = f_t^\circ(\xi_t).$$

Moreover, the optimal (finite horizon) leakage rate is given by $\tilde{V}_1(\xi_1)/T$, where $\xi_1(w) = \sum_{(x,s) \in \mathcal{D}(w)} P_X(x) P_{S_1}(s)$.

2) For the infinite horizon, the optimal charging policy is time-homogeneous and is given by the solution of the following fixed point equation:

$$\tilde{J} + \tilde{v}(\xi) = \min_{b \in \mathcal{B}} [\tilde{\mathcal{B}}_b \tilde{v}](\xi), \quad \forall \xi \in \mathcal{P}_S. \quad (31)$$

where $\tilde{J} \in \mathbb{R}$ is a constant and $\tilde{v} : \mathcal{P}_S \rightarrow \mathbb{R}$. Let $\mathbf{f}^\circ(\xi)$ denote the arg min of the right hand side of (31). Then, the time-homogeneous policy $\mathbf{q}^* = (q^*, q^*, \dots)$ given by

$$q^*(y_t|w_t, \xi_t) = b_t(y_t|w_t), \text{ where } b_t = f^\circ(\xi_t)$$

is optimal. Moreover, the optimal (infinite horizon) leakage rate is given by \tilde{J} . \square

Proof. Lemma 5 implies that $\{\xi_t\}_{t \geq 1}$ is a controlled Markov process with control action b_t . Lemma 4, part 2), implies that the per-step cost can be written as

$$\frac{1}{T} \mathbb{E} \left[\sum_{t=1}^T I(b; \xi) \right].$$

Thus, by standard results in Markov decision theorem [24], the optimal solution is given by the dynamic program described above. \square

B. Weak achievability

To simplify the analysis, we assume that we are free to choose the initial distribution of the state of the battery, which could be done by, for example, initially charging the battery to a random value according to that distribution. In principle, such an assumption could lead to a lower achievable leakage rate. For this reason, we call it *weak* achievability. In the next section, we will show achievability starting from an arbitrary initial distribution, which we call *strong* achievability.

Definition 1. Any $\theta \in \mathcal{P}_S$ and $\xi \in \mathcal{P}_W$ are said to *equivalent* to each other if they satisfy the transformation in Lemma 3.

Definition 2 (Constant-distribution policy). A time-homogeneous policy $\mathbf{f}^\circ = (f^\circ, f^\circ, \dots)$ is called a constant-distribution policy if for all $\xi \in \mathcal{P}_W$, $f^\circ(\xi)$ is a constant. If $f^\circ(\xi) = b^\circ$, then with a slight abuse of notation, we refer to $\mathbf{b}^\circ = (b^\circ, b^\circ, \dots)$ as a constant-distribution policy.

Recall that under a constant-distribution policy $b \in \mathcal{B}$, for any realization y^t of Y^t , θ_t and ξ_t are given as follows:

$$\begin{aligned} \theta_t(s) &= \mathbb{P}(S_t = s | Y^t = y^t, B^{t-1} = b^{t-1}) \\ \xi_t(w) &= \mathbb{P}(W_t = w | Y^t = y^t, B^{t-1} = b^{t-1}). \end{aligned}$$

1) *Invariant Policies*: We next impose an invariance property on the class of policies. Under this restriction the leakage rate expression will simplify substantially. Subsequently we will show that the optimal policy belong to this restricted class.

Definition 3 (Invariance Property). For a given distribution θ_1 of the initial battery state, a constant-distribution policy $b \in \mathcal{B}$ is called a *invariant* policy if for all t , $\theta_t = \theta_1$ and $\xi_t = \xi_1$, where ξ_1 is equivalent to θ_1 .

Remark 1. An immediate implication of the above definition is that under any invariant policy \mathbf{b} , the conditional distribution $\mathbb{P}^{\mathbf{b}}(X_t, S_t, Y_t | Y^{t-1})$ is the same as the joint distribution $\mathbb{P}^{\mathbf{b}}(X_1, S_1, Y_1)$. Marginalizing over (X, S) we get that $\{Y_t\}_{t \geq 1}$ is an i.i.d. sequence.

Lemma 6. *If the system starts with an initial distribution θ of the battery state, and ξ is equivalent to θ , then an invariant policy $\mathbf{b} = (b, b, \dots)$ corresponding to (θ, ξ) achieves a leakage rate*

$$L_T(\mathbf{b}) = I(W_1; Y_1) = I(b; \xi)$$

for any horizon T .

Proof. The proof is a simple corollary of the invariance property in Definition 3. Recall from the dynamic program of Theorem 3, that the performance of any policy $\mathbf{b} = (b_1, b_2, \dots)$ such that $b_t \in \mathcal{B}$, is given by

$$L_T(\mathbf{b}) = \frac{1}{T} \mathbb{E} \left[\sum_{t=1}^T I(b_t; \xi_t) \right].$$

Now, we start with an initial distribution $\xi_1 = \xi$ and follow the constant-distribution structured policy $\mathbf{b} = (b, b, \dots)$. Therefore, by Lemma 6, $\xi_t = \xi$ for all t . Hence, the leakage rate under policy \mathbf{b} is

$$L_T(\mathbf{b}) = I(b; \xi). \quad \square$$

Remark 2. Note that Lemma 6 can be easily derived independently of Theorem 3. From Proposition 1, we have that:

$$L_T(\mathbf{b}) = \frac{1}{T} \sum_{t=1}^T I^{\mathbf{b}}(S_t, X_t; Y_t | Y^{t-1}). \quad (32)$$

From Remark 1, we have that $I^{\mathbf{b}}(S_t, X_t; Y_t | Y^{t-1}) = I^{\mathbf{b}}(S_1, X_1; Y_1) = I^{\mathbf{b}}(W_1; Y_1)$, which immediately results in Lemma 6.

For invariant policies we can further express the leakage rate in the following fashion, which is useful in the proof of optimality.

Lemma 7. *For any invariant policy b ,*

$$I^{\mathbf{b}}(W_1; Y_1) = I^{\mathbf{b}}(W_1; X_1).$$

Proof. Consider the following sequence of simplifications:

$$\begin{aligned} I^{\mathbf{b}}(W_1; Y_1) &= H^{\mathbf{b}}(W_1) - H^{\mathbf{b}}(W_1 | Y_1) \\ &= H^{\mathbf{b}}(W_1) - H^{\mathbf{b}}(W_1 + Y_1 | Y_1) \\ &\stackrel{(a)}{=} H^{\mathbf{b}}(W_1) - H^{\mathbf{b}}(S_2 | Y_1) \\ &\stackrel{(b)}{=} H^{\mathbf{b}}(W_1) - H^{\mathbf{b}}(S_1) \\ &\stackrel{(c)}{=} H^{\mathbf{b}}(W_1) - H^{\mathbf{b}}(S_1 | X_1) \\ &\stackrel{(d)}{=} H^{\mathbf{b}}(W_1) - H^{\mathbf{b}}(W_1 | X_1) \\ &= I^{\mathbf{b}}(W_1; X_1). \end{aligned}$$

where (a) is due to the battery update equation (1); (b) is because b is an invariant ; (c) is because S_1 and X_1 are independent; and (d) is because $S_1 = W_1 + X_1$. \square

2) *Structured Policy*: We now introduce a class of policies that satisfy the invariance property in Def. 3. This will be then used in the proof of Theorem 2.

Definition 4 (Structured Policy). Given $\theta \in \mathcal{P}_S$ and $\xi \in \mathcal{P}_W$, a constant-distribution policy $\mathbf{b} = (b, b, \dots)$ is called a *structured policy* with respect to (θ, ξ) if:

$$b(y|w) = \begin{cases} P_X(y) \frac{\theta(y+w)}{\xi(w)}, & y \in \mathcal{X} \cap \mathcal{Y}_o(w) \\ 0, & \text{otherwise.} \end{cases}$$

Note that it is easy to verify that the distribution b defined above is a valid conditional probability distribution.

Lemma 8. *For any $\theta \in \mathcal{P}_S$ and $\xi \in \mathcal{P}_W$, the structured policy $\mathbf{b} = (b, b, \dots)$ given in Def. 4 is an invariant policy.*

Proof. Since (θ_t, ξ_t) are related according to Lemma 3, in order to check whether a policy is invariant it is sufficient to check that either $\theta_t = \theta_1$ for all t or $\xi_t = \xi_1$ for all t . Furthermore, to check if a *time-homogeneous* policy is an invariant policy, it is sufficient to check that either $\theta_2 = \theta_1$ or $\xi_2 = \xi_1$. We will prove that $\theta_2 = \theta_1$.

Let the initial distributions $(\theta_1, \xi_1) = (\theta, \xi)$ and the system variables be defined as usual. Now consider a realization s_2 of S_2 and y_1 of Y_1 . This means that $W_1 = s_2 - y_1$. Since Y_1 is chosen according to $\xi(\cdot | w_1)$, it must be that $y_1 \in \mathcal{X} \cap \mathcal{Y}_o(w_1)$. Therefore,

$$\begin{aligned} \mathbb{P}^{\mathbf{b}}(S_2 = s_2, Y_1 = y_1) &= \mathbb{P}^{\mathbf{b}}(S_2 = s_2, Y_1 = y_1, W_1 = s_2 - y_1) \\ &= \xi_1(s_2 - y_1) b(y_1 | s_2 - y_1) \\ &= P_X(y_1) \theta_1(s_2), \end{aligned} \quad (33)$$

where in the last equality we use the fact that $y_1 \in \mathcal{X} \cap \mathcal{Y}_o(s_2 - y_1)$. Note that if $y_1 \notin \mathcal{X} \cap \mathcal{Y}_o(s_2 - y_1)$, then $\mathbb{P}^{\mathbf{b}}(S_2 = s_2, Y_1 = y_1) = 0$. Marginalizing over s_2 , we get $\mathbb{P}^{\mathbf{b}}(Y_1 = y_1) = P_X(y_1)$.

Consequently, $\theta_2(s_2) = \mathbb{P}^{\mathbf{b}}(S_2 = s_2 | Y_1 = y_1) = \theta_1(s_2)$. Hence, \mathbf{b} is invariant as required. \square

Remark 3. As argued in Remark 1, under any invariant policy, $\{Y_t\}_{t \geq 1}$ is an i.i.d. sequence. As argued in the proof of Lemma 8, for a structured policy the marginal distribution of Y_t is P_X . Thus, an eavesdropper cannot statistically distinguish between $\{X_t\}_{t \geq 1}$ and $\{Y_t\}_{t \geq 1}$.

Proposition 4. Let θ^* , ξ^* , and b^* be as defined in Theorem 2. Then,

- 1) (θ^*, ξ^*) is an equivalent pair;
- 2) b^* is a structured policy with respect to (θ^*, ξ^*) .
- 3) If the system starts in the initial battery state θ^* and follows the constant-distribution policy $\mathbf{b}^* = (b^*, b^*, \dots)$, the leakage rate is given by J^* .

Thus, the performance J^* is achievable.

Proof. The proofs of parts 1) and 2) follows from the definitions. The proof of part 3) follows from Lemmas 6 and 7. \square

This completes the proof of the achievability of Theorem 2.

C. Binary Model (Revisited)

We revisit the binary model in Section II-B. Recall that the policy suggested in [8] is as follows:

$$q_t(y_t = 0|x_t, s_t) = q(y_t = 1|x_t, s_t) = 1/2, \quad x_t = s_t \quad (34)$$

and $q_t(y_t|x_t, s_t) = \mathbf{1}_{x_t}(y_t)$ otherwise. It can be easily verified that this policy is a structured policy in Def. 4, where the initial state is taken as $\theta(s_1 = 0) = \theta(s_1 = 1) = 1/2$. Thus using Lemma 6 and Lemma 7 it follows that the leakage rate equals $I(W_1; X_1) = 0.5$. This yields an analytical proof of the result in [8].

D. Strong achievability

Lemma 9. Assume that for any $x \in \mathcal{X}$, $P_X(x) > 0$. Let $(\theta^\circ, \xi^\circ)$ an equivalent pair and $\mathbf{b}^\circ = (b^\circ, b^\circ, \dots)$ be the corresponding structured policy.

Assume that $\theta^\circ \in \text{int}(\mathcal{P}_S)$ or equivalently, for any $w \in \mathcal{W}$ and $y \in \mathcal{X} \cap \mathcal{Y}_\circ(w)$, $b^\circ(y|w) > 0$. Suppose the system starts in the initial state (θ_1, ξ_1) and follows policy \mathbf{b}° . Then:

- 1) the process $\{\theta_t\}_{t \geq 1}$ converges weakly to θ° ;
- 2) the process $\{\xi_t\}_{t \geq 1}$ converges weakly to ξ° ;
- 3) for any continuous function $c: \mathcal{P}_W \rightarrow \mathbb{R}$,

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T E[c(\xi_t)] = c(\xi^\circ). \quad (35)$$

- 4) Consequently, the infinite horizon leakage rate under \mathbf{b}° is

$$L_\infty(\mathbf{b}^\circ) = I(b^\circ, \xi^\circ).$$

Proof. The proof of parts 1) and 2) is presented in Appendix D. From 2), $\lim_{t \rightarrow \infty} \mathbb{E}[c(\xi_t)] = c(\xi^\circ)$, which implies (35). Part 4) follows from part 3) by setting $c(\xi_t) = I(b^\circ, \xi_t)$. \square

Lemma 2 implies that θ^* defined in Theorem 2 lies in $\text{int}(\mathcal{P}_S)$. Then, by Lemma 9, the constant-distribution policy $\mathbf{b}^* = (b^*, b^*, \dots)$ (where b^* is given by Theorem 2), achieves the leakage rate $I(b^*, \xi^*)$. By Lemma 7, $I(b^*, \xi^*)$ is same as J^* defined in Theorem 2. Thus, J^* is achievable starting from any initial state (θ_1, ξ_1) .

E. Dynamic programming converse

We provide two converses. One is based on the dynamic program of Theorem 3, which is presented in this section; the other is based purely on information theoretic arguments, which is presented in the next section.

In the dynamic programming converse, we show that for J^* given in Theorem 2, $v^*(\xi) = H(\xi)$, and any $b \in \mathcal{B}$,

$$J^* + v^*(\xi) \leq [\tilde{\mathcal{B}}_b v^*](\xi), \quad \forall \xi \in \mathcal{P}_W, \quad (36)$$

Thus, J^* is a lower bound of the optimal leakage rate (see [27], [28]).

To prove (36), pick any $\xi \in \mathcal{P}_W$ and $b \in \mathcal{B}$. Suppose $W_1 \sim \xi$, $Y_1 \sim b(\cdot|W_1)$, $S_2 = Y_1 + W_1$, X_2 is independent of W_1 and $X_2 \sim P_X$ and $W_2 = S_2 - X_2$. Then,

$$\begin{aligned} [\tilde{\mathcal{B}}_b v^*](\xi) &= I(b; \xi) + \sum_{(w_1, y_1) \in \mathcal{W} \times \mathcal{Y}} \xi(w_1) b(y_1|w_1) v^*(\hat{\varphi}(\xi, y_1, b)) \\ &= I(W_1; Y_1) + H(W_2|Y_1) \end{aligned} \quad (37)$$

where the second equality is due to the definition of conditional entropy. Consequently,

$$\begin{aligned} [\tilde{\mathcal{B}}_b v^*](\xi) - v^*(\xi) &= H(W_2|Y_1) - H(W_1|Y_1) \\ &= H(W_2|Y_1) - H(W_1 + Y_1|Y_1) \\ &\stackrel{(a)}{=} H(S_2 - X_2|Y_1) - H(S_2|Y_1) \\ &\stackrel{(b)}{\geq} \min_{\theta_2 \in \mathcal{P}_S} [H(\tilde{S}_2 - X_2) - H(\tilde{S}_2)], \quad \tilde{S}_2 \sim \theta_2 \\ &= J^* \end{aligned} \quad (38)$$

where (a) uses $S_2 = Y_1 + W_1$ and $W_2 = S_2 - X_2$; (b) uses the fact that $H(A_1|B) - H(A_1 - A_2|B) \geq \min_{P_{A_1}} [H(A_1) - H(A_1 - A_2)]$ for any joint distribution on (A_1, A_2, B) .

The equality in (38) occurs when b is an invariant policy and θ_2 is same as θ^* defined in Theorem 2. For ξ that are not equivalent to θ^* , the inequality in (38) is strict.

We have shown that Eq. (36) is true. Consequently, J^* is a lower bound on the optimal leakage rate \tilde{J} .

F. Information theoretic converse

Consider the following inequalities: for any admissible policy $\mathbf{q} \in \mathcal{Q}_B$, we have

$$\begin{aligned} I(S_1, X^T; Y^T) &= \sum_{t=1}^T I(S_t, X_t; Y_t|Y^{t-1}) \\ &\stackrel{(a)}{\geq} \sum_{t=1}^T I(W_t; Y_t|Y^{t-1}) \end{aligned} \quad (39)$$

where (a) follows from the data processing inequality.

Now consider

$$\begin{aligned} I(W_t; Y_t|Y^{t-1}) &= H(W_t|Y^{t-1}) - H(W_t|Y^t) \\ &= H(W_t|Y^{t-1}) - H(W_t + Y_t|Y^t) \\ &\stackrel{(b)}{=} H(W_t|Y^{t-1}) - H(S_{t+1}|Y^t) \\ &\stackrel{(c)}{=} H(W_t|Y^{t-1}) - H(S_{t+1}|Y^t, X_{t+1}) \\ &= H(W_t|Y^{t-1}) - H(S_{t+1} - X_{t+1}|Y^t, X_{t+1}) \\ &\stackrel{(d)}{=} H(W_t|Y^{t-1}) - H(W_{t+1}|Y^t, X_{t+1}) \end{aligned} \quad (40)$$

where (b) follows from (1); (c) follows because of assumption (A); and (d) also follows from (1).

Substituting (40) in (39) (but expanding the last term as $H(W_T|Y^{T-1}) - H(W_T|Y^T)$), we get

$$\begin{aligned} I(S_1, X^T; Y^T) &\geq \sum_{t=1}^T [H(W_t|Y^{t-1}) - H(W_{t+1}|Y^t, X_{t+1})] \\ &= H(W_1) + \sum_{t=1}^{T-1} [-H(W_t|Y^{t-1}, X_t) + H(W_t|Y^{t-1})] \\ &\quad - H(W_T|Y^T) \\ &= H(W_1) + \sum_{t=2}^T I(W_t; X_t|Y^{t-1}) - H(W_T|Y^T). \quad (41) \end{aligned}$$

Now, we take the limit $T \rightarrow \infty$ to obtain a lower bound to the leakage rate:

$$\begin{aligned} L_\infty(\mathbf{q}) &= \limsup_{T \rightarrow \infty} \frac{1}{T} I(S_1, X^T; Y^T) \\ &\geq \limsup_{T \rightarrow \infty} \frac{1}{T} \left[H(W_1) + \sum_{t=2}^T I(W_t; X_t|Y^{t-1}) - H(W_T|Y^T) \right] \\ &\stackrel{(a)}{=} \lim_{T \rightarrow \infty} \frac{1}{T} \left[\sum_{t=2}^T I(W_t; X_t|Y^{t-1}) \right] \\ &\stackrel{(b)}{\geq} \min_{P_S \in \mathcal{P}_S} I(S - X; X) = J^* \end{aligned}$$

where (a) is because the entropy of any discrete random variable is bounded and (b) follows from the observation that every term in the summation is only a function of the posterior $P(S_t|Y^{t-1})$. Therefore, minimizing each term over a $P_S \in \mathcal{P}_S$ results in a lower bound. This shows that J^* is a lower bound to the minimum (infinite horizon) leakage rate.

V. CONCLUSIONS AND DISCUSSION

In this paper, we study a smart metering system that uses a rechargeable battery to partially obscure the user's power demand. Through a series of reductions, we show that the problem of finding the best battery charging strategy can be recast as a Markov decision process. Consequently, the optimal charging strategies and the minimum information leakage rate are given by the solution of an appropriate dynamic program.

For the case of i.i.d. demand, we provide an explicit characterization of the optimal battery policy and the leakage rate. In this special case it suffices to choose a memoryless strategy where the distribution of Y_t depends only on W_t . Our achievability results rely on restricting attention to a class of invariant policies. Under an invariant policy, the consumption $\{Y_t\}_{t \geq 1}$ is i.i.d. and the leakage rate is characterized by a single-letter mutual information expression. We then further restrict attention to what we call structured policies under which the marginal distribution of $\{Y_t\}_{t \geq 1}$ is P_X . Thus, under the structured policies, an eavesdropper cannot statistically distinguish between $\{X_t\}_{t \geq 1}$ and $\{Y_t\}_{t \geq 1}$. We provide two converses; one is based on the dynamic programming argument while the other is based on a purely information theoretic argument. It is worth highlighting that the weak achievability

and both converses extend to continuous alphabets under mild technical conditions; see [31] for details. It is only the strong achievability result that relies on the finiteness of the alphabets.

Extending of our MDP formulation to incorporate an additive cost, such as the price of consumption, is rather immediate. However, the approach presented in this work for explicitly characterizing the optimal leakage rate in the i.i.d. case may not immediately extend to such general cost functions. The study of such problems remains an interesting further direction.

APPENDIX A

PROOF OF PROPERTY 2

For any $\theta \in \text{int}(\mathcal{P}_S)$ and $\delta(s) : \mathcal{S} \rightarrow \mathbb{R}$ such that $\sum_{s \in \mathcal{S}} \delta(s) = 0$. Let $\theta_\alpha(s) := \theta(s) + \alpha\delta(s)$. Then for small enough α , $\theta_\alpha \in \mathcal{P}_S$. Given such a θ_α , let $\mathbb{P}_{W,X}(w, x) = \mathbb{P}_{W|X}(w|x)P_X(x) = \theta_\alpha(w+x)P_X(x)$. Then to show that $I(W; X)$ is strictly convex on \mathcal{P}_S we require $\frac{d^2 I(W; X)}{d\alpha^2} > 0$. From Property 1, $I(W; X) = H(W) - H(S)$. Therefore,

$$\begin{aligned} \frac{dI(W; X)}{d\alpha} &= \frac{d[-H(S) + H(W)]}{d\alpha} \\ &= \sum_{\tilde{s}} \delta(\tilde{s}) \ln \theta_\alpha(\tilde{s}) - \sum_{w \in \mathcal{W}, s \in \mathcal{S}} P_X(s-w) \delta(s) \ln P_W(w) \\ \frac{d^2 I(W; X)}{d\alpha^2} &= \sum_s \frac{\delta(s)^2}{\theta_\alpha(s)} - \sum_{w \in \mathcal{W}} \frac{(\sum_{\tilde{s} \in \mathcal{S}} P_X(\tilde{s}-w) \delta(\tilde{s}))^2}{P_W(w)}. \end{aligned}$$

Let $a_w(s) = \delta(s) \sqrt{\frac{P_X(s-w)}{\theta_\alpha(s)}}$ and $b_w(s) = \sqrt{\theta_\alpha(s) P_X(s-w)}$. Using the Cauchy-Schwarz inequality, we can show that

$$\begin{aligned} \frac{d^2 I(W; X)}{d\alpha^2} &= \sum_s \frac{\delta(s)^2}{\theta_\alpha(s)} - \sum_{w \in \mathcal{W}} \frac{(\sum_{\tilde{s} \in \mathcal{S}} a_w(\tilde{s}) b_w(\tilde{s}))^2}{P_W(w)} \\ &> \sum_s \frac{\delta(s)^2}{\theta_\alpha(s)} - \sum_{w \in \mathcal{W}} \frac{(\sum_{\tilde{s} \in \mathcal{S}} a_w(\tilde{s})^2) (\sum_{\tilde{s} \in \mathcal{S}} b_w(\tilde{s})^2)}{P_W(w)} \\ &= \sum_s \frac{\delta(s)^2}{\theta_\alpha(s)} - \sum_{w \in \mathcal{W}} \left(\sum_{\tilde{s} \in \mathcal{S}} a_w(\tilde{s})^2 \right) = 0. \end{aligned}$$

The strict inequality is because a and b cannot be linearly dependent. To see this, observe that $\frac{a(s)}{b(s)} = \frac{\delta(s)}{\theta(s) + \alpha\delta(s)}$ cannot be equal to a constant for all $s \in \mathcal{S}$ since δ must contain negative as well as positive elements.

APPENDIX B

PROOF OF PROPOSITION 1

The proof of Proposition 1 relies on the following intermediate results (which are proved later):

Lemma B.1. For any $\mathbf{q} \in Q_A$,

$$I^{\mathbf{q}}(S_1, X^T; Y^T) \geq \sum_{t=1}^T I^{\mathbf{q}}(X_t, S_t; Y_t|Y^{t-1})$$

with equality if and only if $\mathbf{q} \in Q_B$.

Lemma B.2. For any $\mathbf{q}_a \in Q_A$, there exists a $\mathbf{q}_b \in Q_B$, such that

$$\sum_{t=1}^T I^{\mathbf{q}_a}(X_t, S_t; Y_t|Y^{t-1}) = \sum_{t=1}^T I^{\mathbf{q}_b}(X_t, S_t; Y_t|Y^{t-1}).$$

Combining Lemmas B.1 and B.2, we get that for any $\mathbf{q}_a \in \mathcal{Q}_A$, there exists a $\mathbf{q}_b \in \mathcal{Q}_B$ such that

$$I^{\mathbf{q}_a}(S_1, X^T; Y^T) \geq I^{\mathbf{q}_b}(S_1, X^T; Y^T).$$

Therefore there is no loss of optimality in restricting attention to charging policies in \mathcal{Q}_B . Furthermore, Lemma B.1 shows that for any $q \in \mathcal{Q}_B$, $L_T(q)$ takes the additive form as given in the statement of the proposition.

Proof of Lemma B.1. For any $\mathbf{q} \in \mathcal{Q}_A$, we have

$$\begin{aligned} I^{\mathbf{q}}(S_1, X^n; Y^n) &\stackrel{(a)}{=} \sum_{t=1}^n I^{\mathbf{q}}(S_1, X^t; Y_t | Y^{t-1}) \\ &\stackrel{(b)}{=} \sum_{t=1}^n I^{\mathbf{q}}(X^t, S^t; Y_t | Y^{t-1}) \\ &\stackrel{(c)}{\geq} \sum_{t=1}^n I^{\mathbf{q}}(X_t, S_t; Y_t | Y^{t-1}) \end{aligned}$$

where (a) uses the chain rule of mutual information and the fact that $(Z^{t-2}, Y^{t-1}) \rightarrow X_{t-1} \rightarrow X_t$;⁴ (b) uses the fact that the battery process S^t is a deterministic function of S_1, X^t , and Y^t given by (1); and (c) uses the fact that removing terms lowers the mutual information. \square

Proof of Lemma B.2. For any $\mathbf{q}_a = (q_1^a, q_2^a, \dots, q_T^a) \in \mathcal{Q}_A$, construct a $\mathbf{q}_b = (q_1^b, q_2^b, \dots, q_T^b) \in \mathcal{Q}_B$ as follows: for any t and realization (x^t, s^t, y^t) of (X^t, S^t, Y^t) let

$$q_t^b(y_t | x_t, s_t, y^{t-1}) = \mathbb{P}_{Y_t | X_t, S_t, Y^{t-1}}^{\mathbf{q}_a}(y_t | x_t, s_t, y^{t-1}). \quad (42)$$

To prove the Lemma, we show that for any t ,

$$\mathbb{P}_{X_t, S_t, Y^t}^{\mathbf{q}_a} = \mathbb{P}_{X_t, S_t, Y^t}^{\mathbf{q}_b}. \quad (43)$$

By definition of \mathbf{q}_b given by (42), to prove (43), it is sufficient to show that

$$\mathbb{P}_{X_t, S_t, Y^{t-1}}^{\mathbf{q}_a} = \mathbb{P}_{X_t, S_t, Y^{t-1}}^{\mathbf{q}_b}. \quad (44)$$

We do so using induction.

For $t = 1$, $\mathbb{P}_{X_1, S_1}^{\mathbf{q}_a}(x, s) = P_{X_1}(x)P_{S_1}(s) = \mathbb{P}_{X_1, S_1}^{\mathbf{q}_b}(x, s)$. This forms the basis of induction. Now assume that (44) hold for t .

In the rest of the proof, for ease of notation, we denote $\mathbb{P}_{X_{t+1}, S_{t+1}, Y^t}^{\mathbf{q}_a}(x_{t+1}, s_{t+1}, y^t)$ simply by $\mathbb{P}^{\mathbf{q}_a}(x_{t+1}, s_{t+1}, y^t)$. For $t + 1$, we have

$$\begin{aligned} \mathbb{P}^{\mathbf{q}_a}(x_{t+1}, s_{t+1}, y^t) &= \sum_{(x_t, s_t) \in \mathcal{X} \times \mathcal{S}} \mathbb{P}^{\mathbf{q}_a}(x_{t+1}, x_t, s_{t+1}, s_t, y^t) \\ &= \sum_{(x_t, s_t) \in \mathcal{X} \times \mathcal{S}} Q(x_{t+1} | x_t) \mathbb{1}_{s_{t+1}} \{s_t - x_t + y_t\} q_a(y_t | x_t, s_t, y^{t-1}) \\ &\quad \times \mathbb{P}^{\mathbf{q}_a}(x_t, s_t, y^{t-1}) \\ &\stackrel{(a)}{=} \sum_{(x_t, s_t) \in \mathcal{X} \times \mathcal{S}} Q(x_{t+1} | x_t) \mathbb{1}_{s_{t+1}} \{s_t - x_t + y_t\} q_b(y_t | x_t, s_t, y^{t-1}) \\ &\quad \times \mathbb{P}^{\mathbf{q}_b}(x_t, s_t, y^{t-1}) \\ &= \mathbb{P}^{\mathbf{q}_b}(x_{t+1}, s_{t+1}, y^t) \end{aligned}$$

⁴The notation $A \rightarrow B \rightarrow C$ is used to indicate that A is conditionally independent of C given B .

where (a) uses (42) and the induction hypothesis. Thus, (44) holds for $t + 1$ and, by the principle of induction, holds for all t . Hence (43) holds and, therefore, $I^{\mathbf{q}_a}(X_t, S_t; Y_t | Y^{t-1}) = I^{\mathbf{q}_b}(X_t, S_t; Y_t | Y^{t-1})$. The statement in the Lemma follows by adding over t . \square

APPENDIX C PROOF OF PROPOSITION 3

To prove the result, we show the following:

Lemma C.1. *For any action $a \in A$, if $V : \mathcal{P}_{X,S} \rightarrow \mathbb{R}$ is concave, then $\mathcal{B}_a V$ is concave.*

The proof of Proposition 3 follows from backward induction. V_{T+1} is a constant and, therefore, also concave. Lemma C.1 implies that V_T, V_{T-1}, \dots, V_1 are concave.

Proof of Lemma C.1. The first term $I(a; \pi)$ of $[\mathcal{B}_a V](\pi)$ is a concave function of π . We show the same for the second term.

Note that if a function V is concave, then its perspective $g(u, t) := tV(u/t)$ is concave in the domain $\{(u, t) : u/t \in \text{Dom}(V), t > 0\}$. The second term in the definition of the Bellman operator (10)

$$\sum_{y \in Y} \left[\sum_{(x,s) \in \mathcal{X} \times \mathcal{S}} a(y|x, s) \pi(x, s) \right] V(\varphi(\pi, y, a))$$

has this form because the numerator of $\varphi(\pi, y, a)$ is linear in π and the denominator is $\sum_{x,s} a(y|x, s) \pi(x, s)$ (and corresponds to t in the definition of perspective). Thus, for each y , the summand is concave in π , and the sum of concave functions is concave. Hence, the second term of the Bellman operator is concave in π . Thus we conclude that concavity is preserved under \mathcal{B}_a . \square

APPENDIX D PROOF OF LEMMA 9

The proof of the convergence of $\{\xi_t\}_{t \geq 1}$ relies on a result on the convergence of partially observed Markov chains due to Kaijser [29] that we restate below.

Definition 5. A square matrix D is called *subrectangular* if for every pair of indices (i_1, j_1) and (i_2, j_2) such that $D_{i_1, j_1} \neq 0$ and $D_{i_2, j_2} \neq 0$, we have that $D_{i_2, j_1} \neq 0$ and $D_{i_1, j_2} \neq 0$.

Theorem 4 (Kaijser [29]). *Let $\{U_t\}_{t \geq 1}$, $U_t \in \mathcal{U}$, be a finite state Markov chain with transition matrix P^u . The initial state U_1 is distributed according to probability mass function P_{U_1} . Given a finite set \mathcal{Z} and an observation function $g: \mathcal{U} \rightarrow \mathcal{Z}$, define the following:*

- The process $\{Z_t\}_{t \geq 1}$, $Z_t \in \mathcal{Z}$, given by $Z_t = g(U_t)$.
- The process $\{\psi_t\}_{t \geq 1}$, $\psi_t \in \mathcal{P}_U$, given by $\psi_t(u) = \mathbb{P}(U_t = u | Z^t)$.
- A square matrix $M(z)$, $z \in \mathcal{Z}$, given by

$$[M(z)]_{i,j} = \begin{cases} P_{ij}^u & \text{if } g(j) = z \\ 0 & \text{otherwise} \end{cases} \quad i, j \in \mathcal{U}.$$

If there exists a finite sequence z_1^m such that $\prod_{t=1}^m M(z_t)$ is subrectangular, then $\{\psi_t\}_{t \geq 1}$ converges in distribution to a limit that is independent of the initial distribution P_{U_1} .

We will use the above theorem to prove that under policy b° , $\{\xi_t\}_{t \geq 1}$ converges to a limit. For that matter, let $\mathcal{U} = \mathcal{S} \times \mathcal{Y}$, $\mathcal{Z} = \mathcal{Y}$, $U_t = (S_t, Y_{t-1})$ and $g(S_t, Y_{t-1}) = Y_{t-1}$.

First, we show that $\{U_t\}_{t \geq 1}$ is a Markov chain. In particular, for any realization (s^{t+1}, y^t) of (S^{t+1}, Y^t) , we have that

$$\begin{aligned} & \mathbb{P}^{b^\circ}(U_{t+1} = (s_{t+1}, y_t) \mid U^t = (s^t, y^{t-1})) \\ &= \sum_{\tilde{x}_t \in \mathcal{X}} P(U_{t+1} = (s_{t+1}, y_t), X_t = \tilde{x}_t \mid U^t = (s^t, y^{t-1})) \\ &= \sum_{\tilde{x}_t \in \mathcal{X}} \mathbb{1}_{s_{t+1}} \{y_t + s_t - \tilde{x}_t\} b^*(y_t | s_t - \tilde{x}_t) P_X(\tilde{x}_t) \\ &= \mathbb{P}^{b^\circ}(U_{t+1} = (s_{t+1}, y_t) \mid U_t = (s_t, y_{t-1})). \end{aligned}$$

Next, let $m = 2m_s$ and consider

$$z^m = \underbrace{111 \dots 1}_{m_s \text{ times}} \underbrace{000 \dots 0}_{m_s \text{ times}}.$$

We will show that this z^m satisfies the subrectangularity condition of Theorem 4. The basic idea is the following. Consider any initial state $u_1 = (s, y)$ and any final state $u_m = (s', 0)$. We will show that

$$\mathbb{P}(S_{m_s} = m_s \mid U_1 = (s, y), Z^{m_s} = (111 \dots 1)) > 0, \quad (45)$$

and

$$\mathbb{P}(S_{2m_s} = s' \mid U_{m_s} = (s_m, 1), Z_{m_s+1}^{m_s} = (000 \dots 0)) > 0. \quad (46)$$

Eqs. (45) and (46) show that given the observation sequence z^m , for any initial state (s, y) there is a positive probability of observing any final state $(s', 0)$.⁵ Hence, the matrix $\prod_{t=1}^m M(z)$ is subrectangular. Consequently, by Theorem 4, the process $\{\psi_t\}_{t \geq 1}$ converges in distribution to a limit that is independent of the initial distribution P_{U_1} .

Now observe that $\theta_t(s) = \sum_{y \in \mathcal{Y}} \psi_t(s, y)$ and (θ_t, ξ_t) are related according to Lemma 3. Since $\{\psi_t\}_{t \geq 1}$ converges weakly independent of the initial condition, so do $\{\theta_t\}_{t \geq 1}$ and $\{\xi_t\}_{t \geq 1}$.

Let $\bar{\theta}$ and $\bar{\xi}$ denote the limit of $\{\theta_t\}_{t \geq 1}$ and $\{\xi_t\}_{t \geq 1}$. Suppose the initial condition is $(\theta^\circ, \xi^\circ)$. Since b° is an invariant policy, $(\theta_t, \xi_t) = (\theta^\circ, \xi^\circ)$ for all t . Therefore, the limits $(\bar{\theta}, \bar{\xi}) = (\theta^\circ, \xi^\circ)$.

Proof of Eq. (45). Given the initial state (s, y) , define $\bar{s} = m_s - s$, and consider the sequence

$$x^{m_s} = \underbrace{000 \dots 0}_{\bar{s} \text{ times}} \underbrace{111 \dots 1}_s.$$

Under this sequence of demands, consider the sequence of consumption $y^{m_s-1} = (11 \dots 1)$, which is feasible because the state of the battery increases by 1 for the first \bar{s} steps (at

⁵Note that given the observation sequence z^m , the final state must be of the form $(s', 0)$.

which time it reaches m_s) and then remains constant for the remaining s steps. Therefore,

$$\begin{aligned} & \mathbb{P}(S_{m_s} = m_s \mid U_1 = (s, y), \\ & \quad Y^{m_s-1} = (111 \dots 1), X^{m_s} = x^{m_s}) > 0. \end{aligned}$$

Since the sequence of demands x^m has a positive probability,

$$\begin{aligned} & \mathbb{P}(S_{m_s} = m_s, X^{m_s} = x^{m_s} \mid U_1 = (s, y), \\ & \quad Y^{m_s-1} = (111 \dots 1)) > 0. \end{aligned}$$

Therefore,

$$\mathbb{P}(S_{m_s} = m_s \mid U_1 = (s, y), Y^{m_s-1} = (111 \dots 1)) > 0$$

which completes the proof. \square

Proof of Eq. (46). The proof is similar to the Proof of (45). Given the final state $(s', 0)$, define $\bar{s}' = m_s - s'$ and consider the sequence

$$x_{m_s+1}^{2m_s} = \underbrace{111 \dots 1}_{\bar{s}' \text{ times}} \underbrace{000 \dots 0}_{s' \text{ times}}.$$

Under this sequence of demands and the sequence of consumption given by $y_{m_s-1}^{2m_s} = (00 \dots 0)$, the state of the battery decreases by 1 for the first \bar{s}' steps (at which time it reaches s') and then remains constant for the remaining s' steps. Since $x_{m_s+1}^{2m_s}$ has positive probability, we can complete the proof by following an argument similar to that in the proof of (45). \square

REFERENCES

- [1] A. Ipakchi and F. Albuyeh, "Grid of the future," in *IEEE Power Energy Mag.*, vol. 7, no. 2, pp. 52-62, Mar.-Apr. 2009.
- [2] A. Predunzi, "A neuron nets based procedure for identifying domestic appliances pattern-of-use from energy recordings at meter panel," in *Proc. IEEE Power Eng. Society Winter Meeting, New York*, Jan. 2002.
- [3] G. W. Hart, "Nonintrusive appliance load monitoring," *Proceedings IEEE*, vol. 80, no. 12, pp. 1870-1891, Dec. 1992.
- [4] H. Y. Lam, G. S. K. Fung, and W. K. Lee, "A novel method to construct taxonomy of electrical appliances based on load signatures," *IEEE Trans. Consumer Electronics*, vol. 53, no. 2, pp. 653-660, May 2007.
- [5] C. Efthymiou and G. Kalogridis, "Smart grid privacy via anonymization of smart metering data," in *Proc. IEEE Smart Grid Commun. Conf.*, Gaithersburg, Maryland, 2010.
- [6] G. Kalogridis, C. Efthymiou, S. Z. Denic, T. A. Lewis, and R. Cepeda, "Privacy for smart meters: towards undetectable appliance load signatures," in *Proc. IEEE Smart Grid Commun. Conf.*, Gaithersburg, Maryland, 2010.
- [7] I. Csiszar and J. Korner, "Information Theory: Coding Theorems for Discrete Memoryless Channels," New York: Academic, 1981.
- [8] 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 Sig. Proc. Prague, Czech Republic*, May 2011.
- [9] 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, pp. 1331-1341, 2012 Chicago
- [10] G. Giulio, D. Gunduz, and H. Vincent Poor, "Smart meter privacy with an energy harvesting device and instantaneous power constraints." *IEEE International Conference on Communications (ICC)*, 2015.
- [11] S. Han, T. Ufuk and G. Pappas, "Event-based information-theoretic privacy: A case study of smart meters", 2016 American Control Conference (ACC).
- [12] L. Bahl, J. Cocke, F. Jelinek, and J. Raviv, "Optimal Decoding of Linear Codes for minimizing symbol error rate", *IEEE Transactions on Information Theory*, vol. IT-20(2), pp. 284-287, March 1974.
- [13] L. Sankar, S.R. Rajagopalan, and H.V. Poor, "An Information-Theoretic Approach To Privacy," in *Proc. 48th Annual Allerton Conf. on Commun., Control, and Computing*, Monticello, IL, Sep. 2010, pp. 1220-1227.

- [14] D. Gunduz and J. Gomez-Vilardebo, "Smart meter privacy in the presence of an alternative energy source," in *2013 IEEE International Conference on Communications (ICC)*, June 2013.
- [15] S. Rajagopalan, L. Sankar, S. Mohajer, and V. Poor, "Smart meter privacy: A utility-privacy framework," in *2nd IEEE International Conference on Smart Grid Communications*, 2011.
- [16] S. Tatikonda and S. Mitter, "The capacity of channels with feedback," *IEEE Trans. on Inform. Theory*, vol. 55, no. 1, pp. 323-349, January 2009.
- [17] A. Goldsmith and P. Varaiya, "Capacity, mutual information, and coding for finite state Markov channels," *IEEE Trans. Inform. Theory*, vol. 42, pp. 868-886, May 1996.
- [18] H. Permuter, T. Weissman and B. Roy, "Capacity and zero-error capacity of the chemical channel with feedback." *IEEE International Symposium on IEEE*, 2007, Nice, France
- [19] R. Blahut, "An algorithm for computing the capacity of arbitrary discrete memoryless channels," *IEEE Trans. Inform. Theory*, pp. 14-20, 1972.
- [20] S. Arimoto, "Computation of channel capacity and rate-distortion functions," *IEEE Trans. Inform. Theory*, pp. 160-473, 1972
- [21] J. Yao and P. Venkatasubramanian, "On the Privacy of an In-Home Storage Mechanism," *52nd Allerton Conference on Communication, Computation and Control*, Monticello, IL, October 2013.
- [22] L. Wang, J.O. Woo and M. Madiman: "A lower bound on the Rényi entropy of convolutions in the integers," *Proc. IEEE Int. Symp. on Inform. Theory (ISIT)*, Honolulu, Hawaii, July 2014.
- [23] W. B. Powell, *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, Wiley, 2011.
- [24] D. P. Bertsekas, *Dynamic Programming and Optimal Control* Athena Publications, 2012.
- [25] G. Shani, J. Pineau, R. Kaplow, "A survey of point-based POMDP solvers," *Autonomous Agents and Multi-Agent Systems*, vol. 27, no. 1, pp. 1-51, Jul 2013.
- [26] R.G. Wood, T. Linder, S.Yüksel, "Optimal Zero Delay Coding of Markov Sources: Stationary and Finite Memory Codes," *arXiv:1606.09135*, June 2016.
- [27] A. Arapostathis, V. Borkar, E. Fernandez-Gaucherand, M. K. Gosh, and S. I. Marcus, "Discrete-time controlled Markov processes with average cost criterion: A survey," *SIAM J. Contr. Optimiz.*, vol. 31, no. 2, pp. 282-344, Mar. 1993.
- [28] M. Schal, "Average Optimality in Dynamic Programming with General State Space," *Mathematics of Operations Research* vol. 18 no. 1, pp. 163-172, Feb 1993.
- [29] T. Kaijser, "A limit theorem for partially observed Markov chains," *Ann. Probab.*, Vol.3, no. 4, pp. 667-696, 1975.
- [30] S. Li, A. Khisti, and A. Mahajan, "Structure of optimal privacy-preserving policies in smart-metered systems with a rechargeable battery," *Proceedings of the IEEE SPAWC*, pp. 1-5, Stockholm, Sweden, June 28-July 1, 2015.
- [31] S. Li, "Information Theoretic Privacy in Smart Metering Systems Using Rechargeable Batteries," *MASc thesis*, University of Toronto, Jan. 2016