Smart Meter Privacy for Multiple Users in the Presence of an Alternative Energy Source

Jesús Gómez-Vilardebó and Deniz Gündüz

***Abstract*— Smart meters (SMs) measure and report users' energy consumption to the utility provider (UP) in almost real-time, providing a much more detailed depiction of the consumer's energy consumption compared to their analog counterparts. This increased rate of information flow to the UP, together with its many potential benefits, raise important concerns regarding user privacy. This paper investigates, from an information theoretic perspective, the privacy that can be achieved in a multiuser SM system in the presence of an alternative energy source (AES). To measure privacy, we use the mutual information rate between the users' real energy consumption profile and SM readings that are available to the UP. The objective is to characterize the privacy-power function, defined as the minimal information leakage rate that can be obtained with an average power-limited AES. We characterize the privacy-power function in a single letter form when the users' energy demands are assumed to be independent and identically distributed over time. Moreover, for binary and exponentially distributed energy demands, we provide an explicit characterization of the privacy-power function. For any discrete energy demands, we demonstrate that the privacy-power function can always be efficiently evaluated numerically. Finally, for continuous energy demands, we derive an explicit lower bound on the privacy-power function, which is tight for exponentially distributed loads.**

***Index Terms*— Smart meter, privacy, rate-distortion, information leakage.**

I. Introduction

WITH the adoption of smart meters (SMs) in energy distribution networks the utility providers (UPs) are able to monitor the grid more closely, and predict the changes in the demand more accurately. This, in turn, allows the UPs to increase the efficiency and the reliability of the grid by dynamically adjusting the energy generation and distribution, as well as the prices, thereby, also influencing the user demand. SMs also benefit the users by allowing them to monitor their own energy consumption profile in almost real time. Consumers can use this information to cut unnecessary consumption, or to reduce the cost by dynamically shifting consumption based on the prices dynamically set by the UPs.

SM deployment is spreading rapidly worldwide [1]. In Europe, the adoption of SMs has been mandated by a directive of the European Parliament [2], which requires 80% SM adoption in all European households by 2020 and 100% by 2022. However, the massive deployment of SMs at homes have also raised serious concerns regarding user privacy [3]. High resolution SM readings can allow anyone who has access to this data to infer valuable private information regarding user behaviour, including the type of electrical equipments used, the time, frequency and duration of usage [4], and even the TV channel that is being watched, as reported in [5]. The privacy of smart meter data is more critical for businesses, such as data centers, factories, etc., whose energy consumption behaviour can reveal important information about their business to competitors. As pointed out in [6], depending on the monitoring granularity different consumption patterns can be identified. With a granularity of hours or minutes, one can identify the user's presence, with a granularity of minutes or seconds one can infer the activities of appliances such as TV or refrigerator, and with a granularity of seconds one could detect bursts of power and identify the activity of appliances such as microwaves, coffee machines or toasters.

Several methods have been proposed in the literature to provide privacy to SM users while keeping the benefits of SMs for control and monitoring of the grid. In [7] user anonymization is proposed by the participation of a trusted third party. Bohli et al. [8] propose sending the aggregated energy consumption of a group of users and in [9] users protect their privacy by adding random noise to their SM readings before being forwarded to the UP. Similarly, [10] proposes quantization of SM readings.

In all of the above work, privacy is obtained by distorting/transforming the SM readings before being forwarded to the UP. However, energy is provided to the user by the UP, and in principle, the UP can easily track user's energy consumption by installing its own smart measurement devices at points where the user connects to the grid. It seems that no level of privacy can be achieved under such a strong assumption; however, users can conceal the patterns corresponding to individual devices and usage patterns by manipulating their energy consumption. This can be achieved either by filtering the energy consumption over time by means of a storage device such as an electric car battery [11]–[14], or by considering the availability of an alternative energy source (AES) [14], [15]. An AES can model a connection to a second energy grid, such as a microgrid, or a renewable energy source, such as a solar panel.

In our model, we assume that the users can satisfy part of their energy demand from the AES. While the UP can track the energy it provides to the users perfectly, it does not have access

Manuscript received April 4, 2014; revised July 17, 2014; accepted October 21, 2014. Date of publication October 27, 2014; date of current version December 17, 2014. This work was supported in part by the Catalan Government under Grant SGR2014-1567, and in part by the Spanish Government under Grant PCIN-2013-027 (ECROPS) and Grant TEC2013-44591-P (INTENSYV). The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Deepa Kundur.

J. Gómez-Vilardebó is with the Centre Tecnològic de Telecomunicacions de Catalunya, Castelldefels 08860, Spain (e-mail: jesus.gomez@cttc.cat).

D. Gündüz is with the Department of Electrical and Electronic Engineering, Imperial College London, London SW7 2AZ, U.K. (e-mail: d.gunduz@imperial.ac.uk).

Digital Object Identifier 10.1109/TIFS.2014.2365365

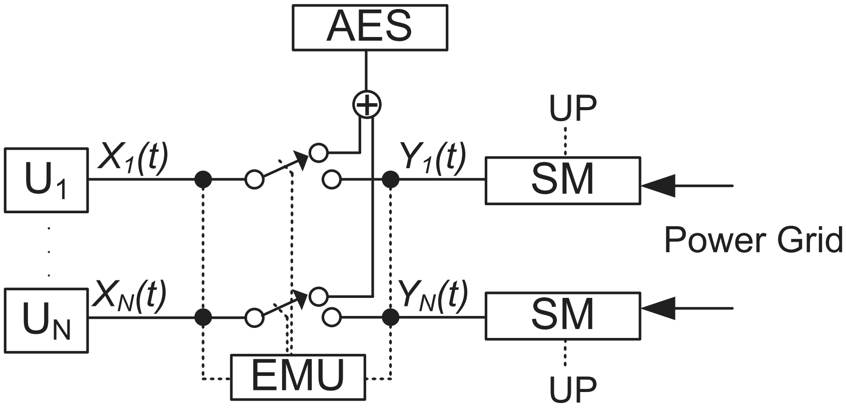

Fig. 1. Studied in this paper. The EMU receives the energy demand from multiple users, $U_1, \ldots, U_N$, and decides how much of the energy demand of each user should be provided from the AES. The remainder of the energy demands are satisfied from the grid, which are measured and reported by the SMs to the UP. The privacy is measured through the information leakage rate, which measures how much information the UP receives about the input load $[\mathbf{X}(1), \ldots, \mathbf{X}(n)]$ by observing the SM readings $[\mathbf{Y}(1), \ldots, \mathbf{Y}(n)]$.

to the instantaneous values of the amount of energy the user receives from the AES. Hence, a certain level of privacy will be achieved depending on the amount of power available from the AES. For instance, if the power that the AES can provide is sufficient enough to satisfy, at any time, all the energy demand of the appliances, the privacy problem can be resolved in a straightforward manner, as no power is requested from the power grid. However, in general, the AES will be limited in terms of the average power it can support, and as we show in this paper, how the user utilizes the energy provided by the AES is critical from the privacy perspective. We measure the privacy through the mutual information rate between the user's real energy consumption and the energy provided by the UP (the SM readings). Mutual information has previously been proposed as a measure of privacy in several works [16]–[18], and in particular, for SM systems in [10], [12], and [14].

In our previous work [15], [19] we have characterized the minimum information leakage rate in the case of a single user with an average and peak power constrained AES. We have shown that there is a very close connection with this problem and the rate-distortion problem in lossy source compression [20] albeit with significant differences. Here we generalize our results to multiple users. In this scenario (see Fig. 1), multiple users, each with its own independent energy demand, share a single AES. The reason for users to share an AES can be economical. AESs, such as solar panels, and efficient energy storage units are expensive facilities, and may be shared by multiple parties to reduce cost. There could be also energy efficiency reasons: consider a scenario in which multiple smart meters belong to the same user; for example, different buildings of the same company. In such a case, the most energy efficient solution requires the centralized management of the AES for all the components of the system.

We assume that there is one separate SM for each user, and the privacy is measured by the total information leaked to the UP about the users' energy consumption. A single energy management unit (EMU) receives users' instantaneous energy demands and decides how much energy to provide to each user from the AES, while satisfying the average power constraint. We first introduce the *privacy-power function* which characterizes the minimal information leakage rate to the UP for a given AES average power constraint. We then provide a single-letter information theoretic characterization of the privacy-power function for the multi-user scenario when the input loads are independent and identically distributed (i.i.d.) random variables. While the EMU can employ energy management policies with memory, our result shows that a memoryless energy management policy that randomly requests energy from the AES is optimal, significantly simplifying the implementation.

We consider both discrete and continuous input loads. For discrete input load distributions, we first show that the optimal output alphabet can be limited to the input alphabet without loss of optimality, which allows us to write the privacy-power function as the solution of a convex optimization problem with linear constraints. As a result, the privacy-power function with discrete input loads can be evaluated numerically in polynomial time. We also provide a closed-form expression for the privacy-power function when the input loads are independent and binary distributed. Using numerical optimization, we compare the optimal privacy-power function with two heuristic power allocation schemes. We consider a time-division heuristic scheme which, at each time instant, obtains the requested energy either from the grid or from the AES, but not from both simultaneously. We also consider an output load limiting heuristic scheme which limits the output load to a fixed maximum value in order to cover up any variation in the energy demand beyond this value. We numerically show that our optimal scheme provides significant privacy gains compared to these heuristic energy management policies.

While the numerical evaluation of the privacy-power function for general continuous input load distributions is elusive, we derive the Shannon lower bound (SLB) on the privacy-power function, and show that this lower bound is tight when users have independent exponentially distributed input loads. For the latter case, we also show that the optimal allocation of the energy generated by the AES among the users can be obtained by the *reverse waterfilling* algorithm [20]. The users with low average input load satisfy all their demand from the AES, while the users with higher average load receive the same amount of energy from the grid.

The rest of the paper is organized as follows. In Section II, we introduce the system model, and provide a single-letter information theoretic characterization of the privacy-power function when users have i.i.d. energy demands over time. Then we show that the privacy-power function for independent users can be solved by simply minimizing the sum of the individual privacy-power functions with a sum average power constraint. The derivation of the privacy-power function for discrete input loads and its particularization to binary input loads is addressed in Section III. Then in Section IV the privacy-power function for continuous input loads is studied and particularized to the exponential distribution. Numerical results are provided in Section V. Finally, conclusions are drawn in Section VI.

II. System Model

We consider the discrete time SM model depicted in Fig. 1. We have N users connected to the energy grid. The energy requested by user i at time instant t is denoted by $X_i(t) \in \mathcal{X}_i$, where $\mathcal{X}_i$ is the support set of the energy demand of user i. We consider the availability of an AES in the system. The AES can provide energy to the users at a maximum average power of P. The AES reduces the energy requested from the grid; but the primary use of the AES here is to create privacy against the UP and other third parties.

The energy flow in the system is managed by the EMU. The EMU receives, at time t, the energy demands of all the users, i.e., the vector $\mathbf{X}(t) = [X_1(t), \ldots, X_N(t)]$. Part of the energy demand of the users can be supported by the AES, while the remainder is provided directly from the energy grid. We denote by $Y_i(t) \in \mathcal{Y}_i$, the amount of energy user i gets from the grid at time t, or equivalently, the reading of SM i at time t. We define $\mathbf{Y}(t) = [Y_1(1), \ldots, Y_N(t)]$ as the aggregated SM readings available to the UP at time t. The energy demand of each user has to be satisfied fully at any time, that is, we do not allow outages or delaying/shifting the user demand. Moreover, we do not allow increasing privacy at the expense of wasting energy, i.e., we have $0 \leq Y_i(t) \leq X_i(t)$ for all t.

At the EMU, we consider energy management policies which, at each time instant t, decide on the amount of power that will be provided from the AES to each of the users based on the input loads up to time t, $\mathbf{X}^t = [\mathbf{X}(1), \ldots, \mathbf{X}(t)]$, and the output loads up to the previous time instant, $\mathbf{Y}^{t-1} = [\mathbf{Y}(1), \ldots, \mathbf{Y}(t-1)]$. We allow stochastic energy management policies, that is, the output load at time t, $\mathbf{Y}(t)$, can be a random function of $\mathbf{X}^t$ and $\mathbf{Y}^{t-1}$. We assume that, while the UP knows P, the average power generated by the AES, it does not have access to the instantaneous values of the energy users receive from the AES.

Definition 1: Denote the vector of input and output load alphabets for all the users as $\mathcal{X}^N = [\mathcal{X}_1, \ldots, \mathcal{X}_N]$ and $\mathcal{Y}^N = [\mathcal{Y}_1, \ldots, \mathcal{Y}_N]$, respectively. A length-n *energy management policy* is composed of, possibly stochastic, power allocation functions

$$f_t : \mathcal{X}^{N \times t} \times \mathcal{Y}^{N \times (t-1)} \to \mathcal{Y}^N, \tag{1}$$

for $t = 1, \ldots, n$, such that

$$\mathbf{Y}(t) = f_t(\mathbf{X}(1), \ldots, \mathbf{X}(t), \mathbf{Y}(1), \ldots, \mathbf{Y}(t-1)), \tag{2}$$

with $X_i(t) \geq Y_i(t) \geq 0$ for all $1 \leq i \leq N$ and $1 \leq t \leq n$.

We measure the privacy achieved by an $n-$length energy management policy with the *information leakage rate*. Assuming that the statistical behavior of the energy demand is known by the UP, its initial uncertainty about the real energy consumption can be measured by the entropy rate $\frac{1}{n}H(\mathbf{X}^n)$. This uncertainty is reduced to $\frac{1}{n}H(\mathbf{X}^n|\mathbf{Y}^n)$ once the UP observes the output load. Hence, the information leaked to the UP can be measured by the reduction in the uncertainty, or equivalently, by the mutual information rate between the input and the output loads $I_n \triangleq \frac{1}{n}I(\mathbf{X}^n; \mathbf{Y}^n)$. Notice that if we could provide all the energy required by the users from the AES, we could achieve perfect privacy, i.e., we would have $I_n = 0$ for all n, by letting $Y_i(t) = 0$ for all i and t. However, in general the AES will be limited in terms of the average power it can provide.

We are thus interested in characterizing the *achievable* level of privacy as a function of the average power P that is provided by the AES, given by

$$P_n = \mathbb{E}\left[\sum_{i=1}^{N}\frac{1}{n}\sum_{t=1}^{n}(X_i(t) - Y_i(t))\right], \tag{3}$$

where the expectation is take over the joint probability distribution of the input and output loads.

Definition 2: An information leakage rate - average power pair (I, P) is said to be *achievable* if there exists a sequence of energy management policies of duration n with $\lim_{n\to\infty} I_n \leq I$, and $\lim_{n\to\infty} P_n \leq P$.

Definition 3: The *privacy-power function*, $\mathcal{I}(P)$, is the infimum of the information leakage rates I such that (I, P) is achievable.

The privacy-power function characterizes the level of privacy that can be achieved by an average power limited AES. The goal of the EMU is to achieve the minimum information leakage rate by optimally allocating the limited energy from the AES over the users and time.

This model of an AES is appropriate for energy sources with their own large energy storage unit, which can provide energy reliably at a certain rate for a sufficiently long duration of time. A peak power constraint on the AES, in addition to the average power constraint, is also considered in [19]. On the other hand, in [14] we have explicitly considered the energy generation process at the AES, in which case the EMU is limited not only by the average power it can pull from the AES, but also the generated energy plus the energy available in the battery at each time instant. Such instantaneous constraints that vary over time depending on the energy management policy and the energy arrival process at the AES, render the analysis significantly harder as they prevent us from invoking information theoretic arguments that will be instrumental in obtaining the single-letter results in this work.

Our goal here is to give a mathematically tractable expression for the privacy-power function, and identify the optimal energy management policy that achieves it. In the rest of the paper, we consider i.i.d. input loads for simplicity, as this will allow us to obtain a single-letter expression for the privacy-power function. Note that in most real-life applications there is a significant correlation among energy demands over time. The i.i.d. assumption allows us to characterize the optimal privacy-preserving solutions, which will be instrumental in identifying solutions for more realistic energy consumption models. Moreover, the i.i.d input load model might be valid in scenarios where the energy consumption either does not have memory at any time scale, or can be modelled as i.i.d. over the time scale of interest. This could be the case, for example, when there is a huge number of applications in use at any time, e.g., in a data center, where the input load can be modelled as i.i.d. over time for different traffic/load states.

In the next theorem, we show that if the input load vectors $\mathbf{X}(t)$ are i.i.d. over time with $f_{\mathbf{X}}(\mathbf{x})$, we can characterize the function $\mathcal{I}(P)$ in a single-letter format. Note that the instantaneous energy demands of the users can be correlated with each other.

Theorem 1: The privacy-power function $\mathcal{I}(P)$ for an i.i.d. input load vector $\mathbf{X} = [X_1, \ldots, X_N]$ with distribution $f_{\mathbf{X}}(\mathbf{x})$ is given by

$$\mathcal{I}(P) = \inf_{\substack{f_{\mathbf{Y}|\mathbf{X}}(\mathbf{y}|\mathbf{x}):\mathbb{E}\left[\sum_{i=1}^{N}(X_i-Y_i)\right]\leq P,\\ 0\leq Y_i\leq X_i,\ i=1,..N}} I(\mathbf{X};\mathbf{Y}), \tag{4}$$

where $\mathbf{Y} = [Y_1, \ldots, Y_N]$ is the corresponding vector of SM readings.

Some basic properties of the privacy-power function $\mathcal{I}(P)$ are characterized in the following lemma. The proof follows from standard techniques based on time-sharing arguments [20].

Lemma 1: The privacy-power function $\mathcal{I}(P)$, given above, is a non-increasing convex function of P.

Next we prove Theorem 1.

Proof: We first prove the achievability. Given a conditional probability distribution $f_{\mathbf{Y}|\mathbf{X}}(\mathbf{y}|\mathbf{x})$ that satisfies (4), we generate each $\mathbf{Y}(t)$ independently using $f_{\mathbf{Y}|\mathbf{X}}(\mathbf{y}(t)|\mathbf{x}(t))$. The mutual information leakage rate is then given by $I(\mathbf{X};\mathbf{Y})$ whereas the average power constraint in (4) is trivially satisfied. For the converse, assume that there is an $n-$length energy management policy that satisfies the instantaneous and average constraints in (4). Let $H(\mathbf{X})$ denote the entropy of the random variable $\mathbf{X}$. The information leakage rate of the resulting output load vector will satisfy the following chain of inequalities:

$$\frac{1}{n}I(\mathbf{X}^n;\mathbf{Y}^n) = \frac{1}{n}\left[H(\mathbf{X}^n) - H(\mathbf{X}^n|\mathbf{Y}^n)\right], \tag{5a}$$

$$= \frac{1}{n}\sum_{t=1}^{n}\left[H(\mathbf{X}(t)) - H(\mathbf{X}(t)|\mathbf{X}^{t-1}\mathbf{Y}^n)\right], \tag{5b}$$

$$\geq \frac{1}{n}\sum_{t=1}^{n}\left[H(\mathbf{X}(t)) - H(\mathbf{X}(t)|\mathbf{Y}(t))\right], \tag{5c}$$

$$= \frac{1}{n}\sum_{t=1}^{n} I(\mathbf{X}(t);\mathbf{Y}(t)), \tag{5d}$$

$$\geq \frac{1}{n}\sum_{t=1}^{n}\mathcal{I}\left(\mathbb{E}\left[\sum_{i=1}^{N}X_i(t) - Y_i(t)\right]\right), \tag{5e}$$

$$\geq \mathcal{I}\left(\frac{1}{n}\sum_{t=1}^{n}\mathbb{E}\left[\sum_{i=1}^{N}X_i(t) - Y_i(t)\right]\right), \tag{5f}$$

$$\geq \mathcal{I}(P), \tag{5g}$$

where (5b) follows from the assumption that the input loads are i.i.d. over time, (5c) follows as conditioning reduces entropy; (5e) follows from the definition of the privacy-power function $\mathcal{I}(\cdot)$; (5f) follows from the convexity of function $\mathcal{I}(\cdot)$ stated in Lemma 1 and Jensen's inequality; and finally (5g) follows since the energy management policy has to satisfy the average power constraint and $\mathcal{I}(\cdot)$ is a non-increasing function of its argument. ■

Remark 1.1: The achievability part of the proof reveals that the optimal energy management policy is memoryless; that is, it can be achieved by simply looking at the instantaneous input load, and generating the output load randomly using the optimal conditional probability, which simplifies the operation of the EMU significantly. This results in a stochastic energy management policy rather than a deterministic one.

We note here that the same performance in Theorem 1 can also be achieved by a deterministic block-based energy management policy if the user knew all the future energy demands over a block of n time instants.

We also note the similarity between the privacy-power function in (4) and the classical rate-distortion function [20]. The characterization of the privacy-power function for a multi-user SM system is equivalent to the rate-distortion function for a vector source with a difference distortion measure

$$d(\mathbf{x},\mathbf{y}) = \begin{cases}\sum_{i=1}^{N} x_i - y_i, & \text{if } y_i \leq x_i,\ \forall i\\ \infty, & \text{otherwise.}\end{cases} \tag{6}$$

However, despite the similarity between the expressions of the rate-distortion and the privacy-power functions, their operational definitions are quite different. In the case of lossy source compression, there is an encoder and a decoder and the rate-distortion function characterizes the minimum number of bits per sample that the encoder should send to the decoder, such that the decoder can reconstruct the source sequence within the specified average distortion level. In lossy source compression, the encoder observes the whole block of n source samples, and maps them to an index from the compression codebook, which is agreed upon in advance.

There are major differences between the two problems. In the SM privacy problem, there is neither an agreed codebook nor a digital interface. Here $\mathbf{Y}^n$ is the direct output of the "encoder", rather than the reconstruction of the decoder based on the transmitted index. The EMU does not operate over blocks of input load realizations; instead, the output load is decided instantaneously based on the previous input and output loads. Similarly, in the SM privacy problem, there is no encoder or decoder either, although the EMU can be considered as an encoder and $\mathbf{Y}^n$ as the reconstruction of the input load $\mathbf{X}^n$. However, the "distortion" constraint between the input and output loads in the SM privacy problem stems from the constraint on the available power that the AES can generate, rather than the limited rate of encoding as in the rate - distortion problem.

Having clarified the distinctions between the privacy-power and rate-distortion functions, we also remark the differences between our formulation of the SM privacy problem and the *privacy-utility framework* studied in [10]. In our privacy model the SM readings are not tempered, and thus, the SMs report the exact amount of energy received from the grid. On the other hand, in [10], the SM readings are considered as the samples of an information source, which are compressed before being forwarded to the UP in order to hide their real values; and hence, privacy is achieved at the expense of distorting the SM measurements. The distortion constraint in [10] is explicit and measures the utility of the compressed SM samples.

If the users' input loads are independent from each other, but not necessarily identically distributed, the multi-user privacy-power function in (4) simplifies further. The following chain of inequalities lower bound the privacy-power function under this assumption:

$$\begin{aligned} I(\mathbf{X};\mathbf{Y}) &= \sum_{i=1}^{N} H(X_i) - H\left(X_i|X^{i-1}, Y^N\right), && \text{(7a)}\\ &\geq \sum_{i=1}^{N} H\left(X_i\right) - \sum_{i=1}^{N} H\left(X_i|Y_i\right), && \text{(7b)}\\ &= \sum_{i=1}^{N} I\left(X_i;Y_i\right), && \text{(7c)}\\ &\geq \sum_{i=1}^{N} \mathcal{I}_{X_i}\left(P_i\right), && \text{(7d)} \end{aligned}$$

where we have defined $P_i = \mathbb{E}[X_i - Y_i]$, and $\mathcal{I}_{X_i}(\cdot)$ denotes the privacy power function for a system with an input load distribution $f_{X_i}(x_i)$. We can achieve equality in (7b) with independent EMU policies for individual users, $f_{\mathbf{Y}|\mathbf{X}}(\mathbf{y}|\mathbf{x}) = \prod_i^N f_{Y_i|X_i}(y_i|x_i)$. Consequently, we can achieve equality in (7d) by using the single user optimal energy management policy for each of the input loads separately, while satisfying the total average power constraint, $\sum_{i=1}^{N} P_i \leq P$.

Following the above arguments, the problem of characterizing the optimal privacy-power function for a multi-user SM system is reduced to the following optimization problem

$$\mathcal{I}(P) = \inf_{\sum_{i=1}^{N} P_i \leq P} \sum_{i=1}^{N} \mathcal{I}_{X_i}\left(P_i\right). \qquad (8)$$

In the following sections, we use the information theoretic single-letter characterization of the privacy-power function in order to obtain either closed-form solutions or numerical algorithms that give us the optimal energy management policies in multi-user SM systems with certain input load distributions and an average power constraint on the AES.

III. Discrete Input Loads

In the previous section we have characterized the privacy-power function for i.i.d. input loads as an optimization problem in a single-letter format in (4). Now we will show that this problem can always be efficiently solved for any discrete input load distribution. In addition, for the particular case where all the users have binary input loads, we give a closed-form expression for the privacy-power function.

For discrete input and output alphabets, the characterization of the privacy-power function $\mathcal{I}(\mathcal{P})$ in (4) is a convex optimization problem since the mutual information is a convex function of the conditional probabilities, $f_{\mathbf{Y}|\mathbf{X}}(\mathbf{y}|\mathbf{x})$, for $\mathbf{y} \in \mathcal{Y}^N$, $\mathbf{x} \in \mathcal{X}^N$, and the constraints are linear. Then, (4) can be solved numerically, e.g., by the efficient Blahut-Arimoto (BA) algorithm [20]. However, while the input load alphabet, defined by the system based on the energy demand profiles of the users, can be discrete, the output load alphabet is not necessarily discrete, and the output load, in general ,can take any real value. The next theorem shows that for discrete input load alphabets, the output load alphabet can be constrained to the input alphabet without loss of optimally, i.e., $\mathcal{Y} = \mathcal{X}$, and consequently, for any given discrete input alphabet the privacy-power function can always be computed efficiently. This result is only valid for i.i.d. input loads, but does not require users' input loads to be independent from each other.

Theorem 2: Without loss of optimality, for discrete input load alphabets, the output load alphabet $\mathcal{Y}^N$ can be constrained to the input load alphabet, i.e., $\mathcal{Y}^N = \mathcal{X}^N$.

Proof: Let the discrete input load alphabets for each user be defined as a possibly infinite set

$$\mathcal{X}_i = \{x_{i,1}, \ldots, x_{i,m_i} : x_{i,j} < x_{i,j+1}\},$$

where $m_i = +\infty$ if the input alphabet is countably infinite.

Define $\mathcal{X}_i^C$ as the set of non-negative real numbers that are not in the input load alphabet for each user i.

For any vector $\mathbf{x} = [x_1, \ldots, x_N] \in \mathcal{X}^N$ define the set

$$\Omega(\mathbf{x}) \triangleq (x_1^-, x_1] \times \cdots \times (x_N^-, x_N]$$

where × denotes the Cartesian product and $x_i^- = \max\{x \in \{0, \mathcal{X}_i\} : x < x_i\}$. Now assume that the optimal privacy-power function in (4) is achieved by the conditional probability distribution $f_{\mathbf{Y}|\mathbf{X}}(\mathbf{y}|\mathbf{x})$, which might take positive values for some $y_i \in \mathcal{X}_i^C$. We define the following new conditional probability distribution:

$$f_{\hat{\mathbf{Y}}|\mathbf{X}}(\hat{\mathbf{y}}|\mathbf{x}) = \begin{cases} 0, & \text{if } \exists i : \hat{y}_i \in \mathcal{X}_i^C, \\ \int_{\Omega(\hat{\mathbf{y}})} f_{\mathbf{Y}|\mathbf{X}}(\mathbf{y}|\mathbf{x}) d\mathbf{y}, & \text{if } \hat{y}_i \in \mathcal{X}_i, \ \forall i. \end{cases}$$

The new conditional probability function does not allow any output value in $\mathcal{X}_i^C$ for any i, i.e., the output alphabet is limited to the input alphabet. Instead, any output vector $\mathbf{y} = [y_1, \ldots, y_N]$, which has a non-zero probability according to $f_{\mathbf{Y}|\mathbf{X}}(\mathbf{y}|\mathbf{x})$, is assigned to a new output vector $[\hat{y}_1, \ldots, \hat{y}_N]$ such that

$$\hat{y}_i = \min\{x \in \mathcal{X}_i : x \geq y_i\}. \qquad (9)$$

Notice that the energy management policy, $f_{\hat{\mathbf{Y}}|\mathbf{X}}(\hat{\mathbf{y}}|\mathbf{x})$, is still feasible since the output load, at any time instant, is still less than what is requested by the appliances, i.e., $\hat{y}_i \leq x_i, \ \forall i$. Moreover, with this new conditional distribution the power load demanded from the AES can only have a smaller average value compared to the original energy management policy, since the output load is not reduced for any input load value.

Thus, it only remains to show that the new conditional distribution leaks at most the same amount of information to the UP. Notice that the new output load $\hat{\mathbf{Y}}$ is a deterministic function of $\mathbf{Y}$ define in (9). Hence, from the information processing inequality, we have that $\mathbf{X} - \mathbf{Y} - \hat{\mathbf{Y}}$ form a Markov chain, and consequently, $I(\mathbf{X}, \mathbf{Y}) \geq I(\mathbf{X}, \hat{\mathbf{Y}})$, which completes the proof. ■

A. Binary Input Loads

The simplest discrete input load model we can consider is a binary input alphabet with independent Bernoulli input load distributions for all the users, i.e., $X_i \sim \mathbf{Ber}(p_i)$, where $p_i = p_{X_i}(L_i)$ and $\mathcal{X}_i = \{L_i, H_i\}$ for $i = 1, \dots, N$. Observe that the average power required by the $i-$th user is given by $P_{X_i} = L_i + \Delta_i(1 - p_i)$, where $\Delta_i = H_i - L_i$. This power consumption model corresponds to a scenario in which the users, at each time instant, require either a constant high power load level H_i, or a constant low power load level L_i, i.e., the standby power consumption level. When there is a power demand, the EMU fulfills this demand either obtaining the energy from the UP, or from the AES according to $p_{\mathbf{Y}|\mathbf{X}}$.

From Theorem 2, the optimal output distribution $\mathcal{Y}_i$ is also binary for all i. Hence, the power allocated from the AES to each user is a binary random variable over the set $\{0, \Delta_i\}$. Note that, since we require $Y_i \leq X_i$, we can only provide energy from the AES to user i if $X_i(t) = H_i$ and $Y_i(t) = L_i$, and consequently, $p_{X_iY_i}(L_i, H_i) = 0$ and $p_{X_iY_i}(L_i, L_i) = p_{X_i}(L_i) = p_i$. The energy obtained from the AES is then directly related to $p_{X_iY_i}(H_i, L_i)$ by $P_i = \Delta_i p_{X_iY_i}(H_i, L_i)$, and we can express the mutual information $I(X_i; Y_i)$ for the bivariate binary distribution

$$p_{X_iY_i} = \begin{bmatrix} p_i & 0 \\ \dfrac{P_i}{\Delta_i} & 1 - p_i - \dfrac{P_i}{\Delta_i} \end{bmatrix},$$

as a function of P_i as follows:

$$I_{\mathsf{B}_i}(P_i) = \frac{P_i}{\Delta_i} \log_2 \left(\frac{P_i}{\Delta_i} \right) - \left(p_i + \frac{P_i}{\Delta_i} \right) \log_2 \left(p_i + \frac{P_i}{\Delta_i} \right) - (1 - p_i) \log_2 (1 - p_i).$$

Observe that $I_{\mathsf{B}_i}(P_i)$ is a monotonically decreasing function of P_i, and $I_{\mathsf{B}_i}(\Delta_i(1 - p_i)) = 0$. Consequently, the privacy-power function for the binary model for a single user is given by

$$\mathcal{I}_{\mathsf{B}_i}(P_i) = (I_{\mathsf{B}_i}(P_i))^+, \tag{10}$$

where $(x)^+ = \max(x, 0)$.

By particularizing (8) with $\mathcal{I}_{\mathsf{X}_i}(P_i) = \mathcal{I}_{\mathsf{B}_i}(P_i)$ for all i, and solving the resultant problem, we find the optimal power allocation P_i^* as

$$P_i^* = \begin{cases} \Delta_i p_i \dfrac{1 - p_{\Delta_i}}{p_{\Delta_i}} & \text{if } p_i < p_{\Delta_i}, \\ \Delta_i(1 - p_i) & \text{otherwise,} \end{cases} \tag{11}$$

where $p_{\Delta_i}(\lambda) = 1 - e^{-\lambda \Delta_i}$, and λ is chosen such that $\sum_{i=1}^N P_i^* = P$. Note that p_{Δ_i} satisfies $0 \leq p_{\Delta_i} \leq 1$. Then, the privacy-power function for the multiple users with independent binary input load distributions is given by

$$\begin{aligned} \mathcal{I}_{\mathbf{B}}(P) &= \sum_{i=1}^N \mathcal{I}_{\mathsf{B}_i}(P_i^*), & (12) \\ &= \sum_{i=1}^N \left(H_{\mathsf{B}}(p_i) - \frac{p_i}{p_{\Delta_i}} H_{\mathsf{B}}(p_{\Delta_i}) \right)^+, & (13) \end{aligned}$$

where $H_{\mathsf{B}}(p)$ denotes the entropy of a $\mathbf{Ber}(p)$ distribution.

Each user can achieve full privacy $\mathcal{I}_{\mathsf{B}_i}(P_i^*) = 0$ by obtaining an average power of $P_{X_i} - L_i = \Delta_i(1 - p_i)$ from the AES, the remaining power L_i is obtained from the grid without incurring any lost of privacy. However, if the average power obtained from the AES is below $P_{X_i} - L_i$ then the energy obtained from the grid comes at the expense of a loss in privacy. Note that P_i^* and $\mathcal{I}_{\mathbf{B}}(P)$ depend on the input load parameters P_{X_i}, L_i, Δ_i, and p_i in a non-straightforward manner. We postpone the detailed analysis of this privacy-power function to Section V.

IV. Continuous Input Loads

For continuous input loads, the optimal output alphabet is also continuous. Consequently, efficient algorithms, such as the BA algorithm, do not yield the optimal solution to (4). In this case, we provide a lower bound on the privacy-power function by using the Shannon lower bound. We then show that this lower bound is achievable when the users have independent exponentially distributed input loads.

Using the SLB [20], for any input load distribution, we have

$$\mathcal{I}_{X_i}(P_i) \geq (\mathsf{h}(X_i) - \ln(P_i))^+ \text{ nats}, \tag{14}$$

where $\mathsf{h}(X)$ denotes the differential entropy of the continuous random variable X. Observe that,

$$\begin{aligned} I(X_i, Y_i) &= \mathsf{h}(X_i) - \mathsf{h}(X_i|Y_i), & \text{(15a)} \\ &= \mathsf{h}(X_i) - \mathsf{h}(X_i - Y_i|Y_i), & \text{(15b)} \\ &\geq \mathsf{h}(X_i) - \mathsf{h}(X_i - Y_i), & \text{(15c)} \\ &\geq \mathsf{h}(X_i) - \mathsf{h}(\mathsf{Exp}(\mathbb{E}[X_i - Y_i])), & \text{(15d)} \\ &= \mathsf{h}(X_i) - \ln(P_i), & \text{(15e)} \end{aligned}$$

where we have used $\mathsf{Exp}(\lambda)$ to denote an exponential random variable with mean λ. In the above chain of inequalities, (15c) follows as conditioning reduces entropy, and (15d) follows since exponential distribution maximizes the entropy among all nonnegative distributions with a given mean value [20].

Next, we present the necessary and sufficient conditions for any piecewise continuous input load distribution $f_X(x)$ to achieve the SLB, together with the conditional probability distribution $f_{Y|X}(y|x)$ achieving it. We denote by $u(x)$, the unit step function which assigns 0 for $x < 0$, and 1 for $x \geq 0$. The Dirac delta function is denoted by $\delta(x)$. We use $f'(x)$ to denote the first order derivative of $f(x)$ and $f(x_i^+) = \lim_{x \to x_i^+} f(x)$ and $f(x_i^-) = \lim_{x \to x_i^-} f(x)$ and $x \to x_i^+$ and $x \to x_i^-$ mean that $x \to x_i$ from the left and right, respectively. Finally, we define $\Delta_f(x_i) = f(x_i^+) - f(x_i^-)$.

Theorem 3: Suppose that the input load distribution $f_X(x)$ is continuous on $\mathcal{R}_+$ except for a countable number of jump discontinuities or non-differentiable points $\mathcal{X}_D = \{x_1, \ldots, x_D\}$. Then, the SLB (14) is achieved for all P satisfying $g_Y(y) \geq 0$, $\forall y \in \mathcal{R}_+$, where

$$g_Y(y) = g_{Y_C}(y) + g_{Y_D}(y) \tag{16}$$

is a mixture of a continuous and a discrete function specified as follows:

$$g_{Y_C}(y) = f_X(y) + \mathbb{E}[V] f'_X(y), \; y \in \mathcal{R}_+/\mathcal{X}_D,$$
$$g_{Y_D}(y) = \mathbb{E}[V] \sum_{i=0}^{D} \Delta_X(x_i)\delta(y - x_i), \; y \in \mathcal{X}_D.$$

For all P, at which the SLB is achieved, the output distribution is given by $f_Y(y) = g_Y(y)$ and the optimal conditional output load distribution reads $f_{Y|X}(y|x) = f_V(x-y)\frac{f_Y(y)}{f_X(x)}$ where $f_V(v) = \frac{1}{\mathbb{E}[V]} e^{-\frac{v}{\mathbb{E}[V]}} u(v)$.

Proof: To show this results, we need to find the conditional distribution $f_{Y|X}(y|x)$ that satisfies the SLB with equality [20]. We require the random variables $V = X - Y$ and Y to be independent, and V to be distributed according to an exponential distribution $V \sim \mathsf{Exp}(P)$ with mean P. We first obtain the output distribution $f_Y(y)$ from its Laplace transform $\mathcal{L}f_Y(s) = \mathcal{L}(f_Y(y))(s)$ as

$$\mathcal{L}f_Y(s) = \frac{\mathcal{L}f_X(s)}{\mathcal{L}f_V(s)},$$
$$= \mathcal{L}f_X(s)\left(1 + \mathbb{E}[V]s\right).$$

Then, it follows that $f_Y(y)$ is given by (16). The conditional distribution $f_{Y|X}(y|x)$ is obtained using the fact that $f_{X|Y}(x|y) = f_V(x-y)$. Finally, it can be shown that $\int_0^\infty f_Y(y)dy = 1$; and thus, the achievability is guaranteed by requiring $f_Y(y) \geq 0$, $\forall y \in \mathcal{R}^+$. ■

Remark 3.1: If the achievability condition in Theorem 3 is satisfied for a given $P_{\max}$, it is satisfied at any $P \leq P_{\max}$. Then it follows that, there is a unique critical average power level, P_0, such that $\mathcal{I}_X(P) = \mathsf{h}(X) - \ln(P)$ for all $P \leq P_0$ and $\mathcal{I}_X(P) > \mathsf{h}(X) - \ln(P)$ for all $P > P_0$.

To find a lowerbound on the privacy-power function in the case of multiple users with continuous input load distributions, we replace $\mathcal{I}_{X_i}(P_i)$ with $(\mathsf{h}(X_i) - \ln(P_i))^+$ in (8), and find the corresponding optimal power allocation P_i^* as

$$P_i^* = \begin{cases} \lambda, & \text{if } e^{\mathsf{h}(X_i)} > \lambda, \\ e^{\,\mathsf{h}(X_i)}, & \text{otherwise,} \end{cases} \tag{17}$$

where λ is chosen such that $\sum_{i=1}^N P_i^* = P$. Then the privacy-power function for multiple users can be lower-bounded by

$$\mathcal{I}_{\mathbf{X}}(P) \geq \sum_{i=1}^{N} (\mathsf{h}(X_i) - \ln(\lambda))^+ \text{ nats.} \tag{18}$$

A. Exponential Input Loads

For an exponential input load distribution with mean λ_i, i.e., $X_i \sim \mathsf{Exp}(\lambda_i)$, the SLB in (14) is achievable by using the conditional distribution [19]

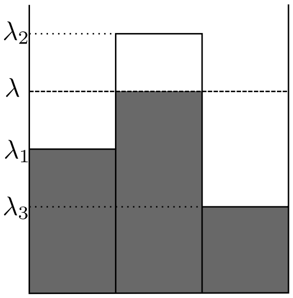

Fig. 2. The reverse waterfilling solution for the optimal power provided to each user from the AES.

$$f_{Y_i|X_i}(y|x) = \frac{\lambda_i}{P_i} e^{-\frac{(x-y)}{P_i}} e^{\frac{x}{\lambda_i}} f_{Y_i}(y),$$

where f_{Y_i} is a mixture of a continuous and a discrete distribution specified by

$$f_{Y_i}(y) = \left(1 - \frac{P_i}{\lambda_i}\right)\frac{1}{\lambda_i} e^{-\frac{y}{\lambda_i}} + \frac{P_i}{\lambda_i}\delta(y).$$

Then the privacy-power function for a single user with an exponential input load with mean λ_i can be explicitly characterized as follows:

$$\mathcal{I}_{\mathsf{E}_i}(P_i) = \begin{cases} \ln\left(\frac{\lambda_i}{P_i}\right), & \text{if } P_i \leq \lambda_i, \\ 0, & \text{otherwise.} \end{cases} \tag{19}$$

By particularizing (8) with $\mathcal{I}_{\mathsf{X}_i}(P_i) = \mathcal{I}_{\mathsf{E}_i}(P_i)$ for all i, and solving the resultant problem, we find the optimal AES power allocation among users, P_i^*, as the well-known reverse waterfilling solution $P_i^* = \lambda$, if $\lambda < \lambda_i$, and $P_i^* = \lambda_i$, if $\lambda \geq \lambda_i$, where λ is chosen such that $\sum_{i=1}^N P_i^* = P$.

The reverse waterfilling power allocation is illustrated in Fig. 2 for three users with independent exponentially distributed energy demands with means λ_1, λ_2, and λ_3, respectively. The optimal reverse water level is given by λ, where the heights of the shaded areas in the figure correspond to the average AES powers allocated to the different users. We observe that the optimal energy management policy satisfies all the energy demands of the users whose average input load is below λ, directly from the AES. Hence, no information is leaked to the UP about the energy consumption of these users; user 1 and user 3 in the figure. The rest of the users receive exactly the same amount of power λ from the AES, and the remainder of their energy demand is satisfied from the grid. Finally, the privacy-power function for multiple users with exponential input loads can be expressed as

$$\mathcal{I}_{\mathbf{E}}(P) = \sum_{i=1}^{N} \left(\ln\left(\frac{\lambda_i}{\lambda}\right)\right)^+. \tag{20}$$

V. Numerical Results

In this section we numerically analyze the privacy-power function in a SM system with various input load distributions and number of users, by explicitly evaluating the information theoretic optimal leakage rate expressions.

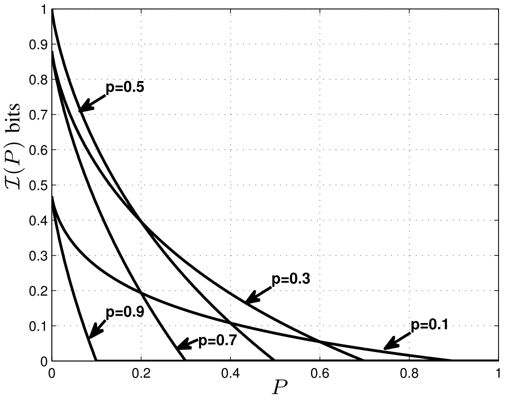

Fig. 3. Privacy-power function for a binary input-output system with different p values.

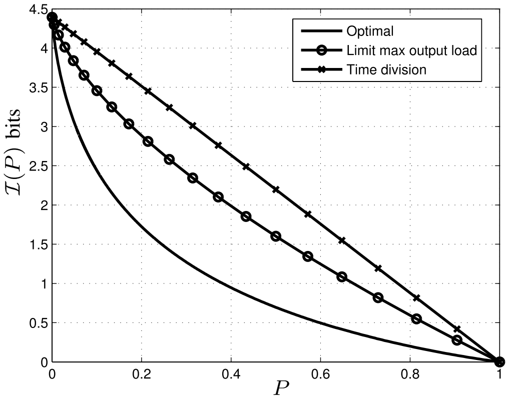

Fig. 4. Privacy-power function for a uniform input load, and different EMU policies.

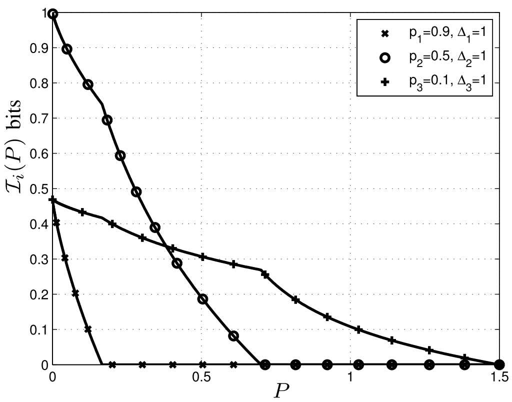

Fig. 5. Individual privacy achieved by three users $\mathcal{I}_{\mathsf{X}_i}(P)$, $i = \{1, 2, 3\}$, all with the same input load alphabet $\delta_i = 1$ $i = \{1, 2, 3\}$, but with different input load distributions as a function of the average AES power P.

A. Single User Scenario

In order to illustrate the behaviour of the privacy-power function for a simple binary input load system, we first consider a single user with an input load alphabet $\mathcal{X} = \mathcal{Y} = \{0, 1\}$, and $p_X(0) = p$. We plot the $\mathcal{I}(P)$ function for the binary input load in Fig. 3 for different p values. As expected, the required average power from the AES is maximum when the user wants perfect privacy, and it is zero when no privacy is required. We also observe that the privacy-power function is decreasing in P and convex. Another interesting observation from the figure is the fact that the $\mathcal{I}(P)$ curves for two different input load distributions, i.e., different p values, might intersect. This means that, to achieve the same level of privacy a lighter input load might require lower or higher average power than a heavier input load. Also note that the two different input load distributions, say $p = 0.1$ and $p = 0.9$, have the same level of privacy when there is no AES in the system; however, the input load with lower average energy demand, i.e., the one with $p = 0.9$, achieves perfect privacy with a much lower P value.

Next, we use the discrete uniform distribution to compare the privacy protection achieved by the information theoretical optimal policy derived here, with different heuristic policies. In this case, the input load has a uniform distribution $U(x)$ with input load alphabet $\mathcal{X} = \{0, c, 2c, \ldots, (N-1)c\}$, where $c = \frac{2}{N-1}$ is a constant used to impose a mean value of $\mathbb{E}[X] = 1$. Based on Theorem 2, the output load alphabet can be limited to $\mathcal{X}$ without loss of optimality. We set $N = 21$ and in Fig. 4 we plot the privacy-power function for the optimal strategy obtained by the BA algorithm together with the privacy-power functions of the following two heuristic strategies:

1) Time Division: In this policy, at each time instant, the EMU gets all the energy needed by the user, either from the AES or from the grid, but not from both simultaneously. Then, to satisfy the average power constraint at the AES, the EMU obtains energy from the AES with probability $\frac{P}{\mathbb{E}[X]}$. The information leaked to the UP, is thus given by

$$\begin{aligned} I(X;Y) &= H(X) - H(X|Y=0)\frac{P}{\mathbb{E}[X]} \\ &\quad - H(X|Y=x)\left(1 - \frac{P}{\mathbb{E}[X]}\right), \\ &= \left(1 - \frac{P}{\mathbb{E}[X]}\right)\log_2 N. \end{aligned}$$

2) Limit Maximum Output Load: In this policy, we use the AES to limit the maximum energy received from the grid. At each time instant, we get all the energy from the grid $X(t) = Y(t)$ if $X(t) < kc$, whereas if $X(t) \geq kc$ we get $Y(t) = kc$ from the grid and the remaining energy is taken from the AES. In this case, for each $k = 0, \ldots, N-1$, the average power requested from the AES is given by $P = (N-1-k)(N-k)\frac{c}{2N}$, and the information leaked to the UP is

$$\begin{aligned} I(X;Y) &= H(X) - \Pr(Y = kc)H(X|Y = kc), \\ &= \log_2 N - \frac{N-k}{2N}\log_2(N-k). \end{aligned}$$

In Fig. 4, we can observe that given an average power limited AES, the privacy achieved by both of these heuristics is significantly lower than that of the optimal EMU policy.

B. Multi-User Scenario

Next we consider a multi-user scenario with $N = 3$ users. We assume equal binary load levels $H_i = 1$ and $L_i = 0$, but

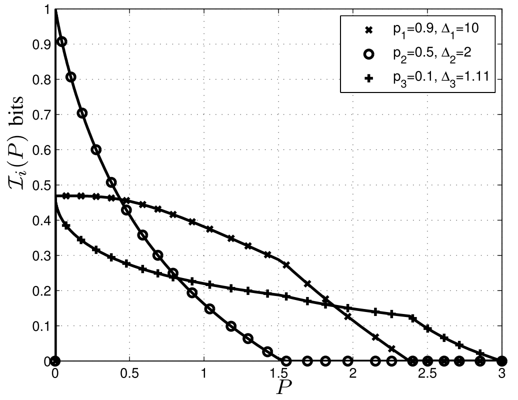

Fig. 6. Individual privacy achieved by three users $\mathcal{I}_{\mathsf{X}_i}(P)$, $i = \{1, 2, 3\}$, each with a different input load alphabet and input load distribution, as a function of the average AES power P.

different average energy demands with $p_1 = 0.9$, $p_2 = 0.5$, and $p_3 = 0.1$; thus we have $P_{X_1} = 0.1$, $P_{X_2} = 0.5$, $P_{X_3} = 0.9$. Fig. 5 illustrates the privacy for each user $\mathcal{I}_{\mathsf{B}_i}(P_i^*)$ as a function of the average power P available at the AES. Notice that, although users 1 and 3, in the absence of an AES, leak the same amount of information to the UP, since $H_{\mathsf{B}}(0.1) = H_{\mathsf{B}}(0.9)$, user 1 achieves perfect privacy much more rapidly since it has a lower average energy demand. Also note that, user 3 achieves perfect privacy for a much higher value of P, even compared to user 2, which leaks the highest amount of information when there is no AES, as it has the highest entropy.

Remember that, as opposed to the exponential input load scenario, in the binary case, the privacy-power function $\mathcal{I}_{\mathsf{B}_i}(P_i^*)$ for each user does not depend solely on the average power demand of the user, but on both of the parameters Δ_i and p_i. To illustrate this dependence, we consider a scenario again with $N = 3$ users, but with equal average power demands $P_{X_i} = \Delta_i(1 - p_i)$, while $L_i = 0$ for all i. We choose different parameters Δ_i and p_i for each user. Fig. 6 again shows the privacy of each user as a function of the average power P. Observe that the optimal power allocation quickly reduces the information leaked by user 2, and achieves perfect privacy for this user much before the other two, although this is the user leaking the most amount of information in the absence of an AES. The input power loads for users 1 and 3 have equal entropy, but with different behaviours; user 1 demands large amounts of energy but very rarely, while user 3 demands low amounts of energy very frequently. The optimal EMU policy seen by these users also differs significantly. While for user 1 the privacy-power function is a concave monotonically decreasing function, for user 3 the privacy-power function is monotonically decreasing but piecewise convex.

Next, we study the effect of the number of users on the privacy-power function. In Fig. 7, we depict the optimal information leakage rate with respect to the available average AES power for binary input loads with different number of users $N = \{1, 2, 3\}$. We can observe that with more than one user, we have different regimes of operation corresponding to the number of users that receive energy from the grid.

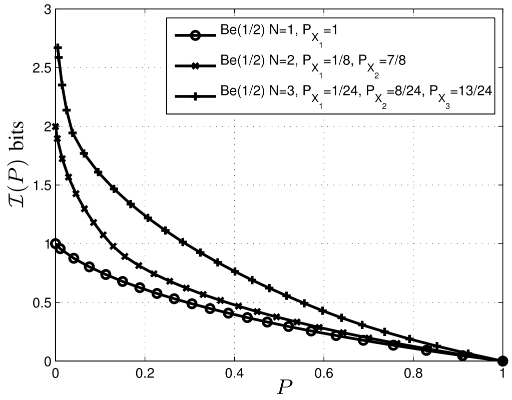

Fig. 7. $\mathcal{I}(P)$ with respect to the average AES power P for binary input loads with different number of users.

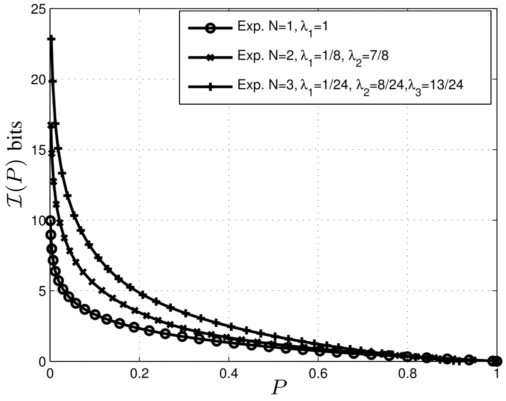

Fig. 8. $\mathcal{I}(P)$ with respect to the average AES power P for exponential input loads with different number of users.

Similarly, in Fig. 8 we consider the scenario with exponential input loads. In both models, regardless of the number of users in the system the total average power consumed by the users is fixed to P_X. In the figures we set $P_X = 1$. As expected, if the average power provided by the AES is equal to the total average power demanded by the users, perfect privacy can be achieved. Instead, as the average power of the AES goes to zero, all the information is revealed to the UP, and thus, the information leakage rate is equal to the sum of the entropies of all the input loads. In between these two extremes the privacy-power function exhibits a monotone decreasing convex behaviour, and the information leakage rate increases with the number of users in the system.

VI. Conclusions

We have introduced and studied the privacy-power function, $\mathcal{I}(P)$, which characterizes the achievable information theoretic privacy in a multi-user SM system in the presence of an AES. We have provided a single-letter information theoretic characterization for $\mathcal{I}(P)$, and showed that it can be evaluated numerically when the input loads are discrete. We have also provided explicit characterization of the privacy-power function for binary and exponential input load distributions. We have shown that the optimal allocation of the energy provided by the AES in the exponentially distributed input load scenario can be derived using the reverse waterfilling algorithm, which resembles the rate-distortion function for multiple Gaussian sources.

We believe that the proposed information theoretic framework for privacy in SM systems provides valuable tools to identify the fundamental challenges and limits for this critical problem, whose importance will only increase as SM adoption becomes more widespread. Many interesting research problems implore further studies, including time correlated input loads, systems with multiple EMUs, as well as cost and pricing issues considering dynamic pricing over time.

References

[1] P. Wunderlich, D. Veit, and S. Sarker, "Adoption of information systems in the electricity sector: The issue of smart metering," in *Proc. Amer. Conf. Inf. Syst.*, Seattle, WA, USA, Aug. 2012, paper 16.

[2] European Union "Directive 2009/72/EC of the European parliament and of the council of 13 July 2009 concerning common rules for the internal market in electricity and repealing directive 2003/54/EC," *Off. J. Eur. Union*, vol. 52, no. L211, pp. 55–93, Aug. 14, 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, May/Jun. 2009.

[4] A. Prudenzi, "A neuron nets based procedure for identifying domestic appliances pattern-of-use from energy recordings at meter panel," in *Proc. IEEE Power Eng. Soc. Winter Meeting*, New York, NY, USA, Jan. 2002, pp. 941–946.

[5] U. Greveler, B. Justus, and D. Loehr, "Multimedia content identification through smart meter power usage profiles," in *Proc. Comput., Privacy, Data Protection (CPDP)*, Brussels, Belgium, Jan. 2012, pp. 383–390.

[6] A. Molina-Markham, P. Shenoy, K. Fu, E. Cecchet, and D. Irwin, "Private memoirs of a smart meter," in *Proc. 2nd ACM Workshop Embedded Sens. Syst. Energy-Efficiency Building (BuildSys)*, 2010, pp. 61–66. [Online]. Available: http://doi.acm.org/10.1145/1878431.1878446

[7] C. Efthymiou and G. Kalogridis, "Smart grid privacy via anonymization of smart metering data," in *Proc. 1st IEEE Int. Conf. Smart Grid Commun.*, Gaithersburg, MD, USA, Oct. 2010, pp. 238–243.

[8] J.-M. Bohli, C. Sorge, and O. Ugus, "A privacy model for smart metering," in *Proc. IEEE ICC*, Cape Town, South Africa, May 2010, pp. 1–5.

[9] S. Wang *et al.*, "A randomized response model for privacy preserving smart metering," *IEEE Trans. Smart Grid*, vol. 3, no. 3, pp. 1317–1324, Sep. 2012.

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

[11] G. Kalogridis, C. Efthymiou, S. Z. Denic, T. A. Lewis, and R. Cepeda, "Privacy for smart meters: Towards undetectable appliance load signatures," in *Proc. 1st IEEE Int. Conf. Smart Grid Commun.*, Gaithersburg, MD, USA, Oct. 2010, pp. 232–237.

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

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

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

[15] D. Gunduz and J. Gomez-Vilardebo, "Smart meter privacy in the presence of an alternative energy source," in *Proc. IEEE Int. Conf. Commun.*, Budapest, Hungary, Jun. 2013, pp. 2027–2031.

[16] D. Agrawal and C. C. Aggarwal, "On the design and quantification of privacy preserving data mining algorithms," in *Proc. Symp. Principles Database Syst.*, Santa Barbara, CA, USA, May 2001, pp. 247–255.

[17] D. Rebollo-Monedero, J. Forné, and J. Domingo-Ferrer, "From t-closeness-like privacy to postrandomization via information theory," *IEEE Trans. Knowl. Data Eng.*, vol. 22, no. 11, pp. 1623–1636, Nov. 2010.

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

[19] J. Gomez-Vilardebo and D. Gunduz, "Privacy of smart meter systems with an alternative energy source," in *Proc. IEEE Int. Symp. Inf. Theory*, Istanbul, Turkey, Jul. 2013, pp. 2572–2576.

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

Jesús Gómez-Vilardebó received the M.Sc. and Ph.D. degrees in telecommunication engineering from the Universitat Politècnica de Catalunya, Barcelona, Spain, in 2003 and 2009, respectively, where he received the Ph.D. degree in signal theory and communications. He is currently with the Centre Tecnològic de Telecomunicacions de Catalunya, Castelldefels, Spain, as a Research Associate. His current research interests include information theory, stochastic signal processing, and their applications in wireless multiuser communications and information privacy.

Deniz Gündüz received the B.S. degree in electrical and electronics engineering from Middle East Technical University, Ankara, Turkey, in 2002, and the M.S. and Ph.D. degrees in electrical engineering from the New York University Polytechnic School of Engineering, New York, NY, USA, in 2004 and 2007, respectively. After his Ph.D. study, he served as a Post-Doctoral Research Associate with the Department of Electrical Engineering, Princeton University, Princeton, NJ, USA, and a Consulting Assistant Professor with the Department of Electrical Engineering, Stanford University, Stanford, CA, USA. Since 2012, he has been a Lecturer with the Department of Electrical and Electronic Engineering, Imperial College London, London, U.K. He was a Research Associate with the Centre Tecnològic de Telecomunicacions de Catalunya, Castelldefels, Spain. He was also a Visiting Researcher with Princeton University from 2009 to 2011.

He was a recipient of the Marie Curie Fellowship Award from the European Commission, and the Best Student Paper Award at the 2007 IEEE International Symposium on Information Theory. He is an Associate Editor of the IEEE Transactions on Communications, and served as a Guest Editor of the *EURASIP Journal on Wireless Communications and Networking*, Special Issue on Recent Advances in Optimization Techniques in Wireless Communication Networks.

Dr. Gunduz is serving as the Cochair of the IEEE Information Theory Society Student Committee. He served as the Cochair of the Network Theory Symposium at the 2013 and 2014 IEEE Global Conference on Signal and Information Processing, and was the Cochair of the 2012 IEEE European School of Information Theory. His research interests lie in the areas of communication theory and information theory with a special emphasis on joint source-channel coding, multiuser networks, energy-efficient communications, and information theoretic security and privacy.