

Privacy-Aware Smart Metering

Progress and challenges



Giulio Giaconi, Deniz Gündüz, and H. Vincent Poor

The next-generation energy network, the so-called smart grid (SG), promises tremendous increases in efficiency, safety, and flexibility in managing the electricity grid as compared to the legacy energy network. This is needed today more than ever, as global energy consumption is growing at an unprecedented rate and renewable energy sources (RESs)

must be seamlessly integrated into the grid to assure a sustainable human development.

Smart meters (SMs) are among the crucial enablers of the SG concept. They supply accurate, high-frequency information about users' household energy consumption to a utility provider (UP), which is essential for time-of-use (ToU) pricing, rapid fault detection, and energy theft prevention, while also providing consumers with more flexibility and control over their consumption. However, highly accurate and granular SM

Digital Object Identifier 10.1109/MSP.2018.2841410
Date of publication: 13 November 2018

data also pose a threat to consumer privacy, as nonintrusive load monitoring (NILM) techniques enable a malicious attacker to infer many details of a user's private life.

This article focuses on privacy-enhancing energy management techniques that provide accurate energy consumption information to the grid operator without sacrificing consumer privacy. In particular, we focus on techniques that shape and modify actual user energy consumption by means of physical resources, such as rechargeable batteries (RBs), RESs, and demand shaping. A rigorous mathematical analysis of privacy is presented under various physical constraints on the available physical resources. Finally, open questions and challenges that need to be addressed to pave the way to the effective protection of users' privacy in future SGs are presented.

SMs for an SG

The current energy grid is one of the engineering marvels of the 20th century. However, it has become inadequate for satisfying the steadily growing global electricity demand of the 21st century. In fact, world energy consumption is predicted to increase 48% from 2012 to 2040 [1], driven by factors such as the growth of the global economy, the rise of the gross domestic product per person, the expansion of the planet's population, an increased penetration of electric vehicles, and a broader mobility revolution [2]. Other issues that need to be addressed are the effective integration of RESs and storage capabilities into the grid, the improvement of the grid's environmental sustainability, and the promotion of plug-in hybrid electric vehicles.

To address these challenges, SGs are being engineered. They are intended to substantially improve energy generation, transmission, distribution, consumption, and security, providing enhanced reliability and quality of the electricity supply, quicker detection of energy outages and theft, better matching of the energy supply with demand, and greater environmental sustainability by enabling an easier integration of distributed generation and storage capabilities. The smartness of an SG resides in its advanced metering infrastructure, which enables two-way communication between the utility and its customers and whose pivotal element in the distribution network is the SM, the device that monitors a user's electricity consumption in almost real time.

In contrast to legacy grids, in which billing data are gathered at the end of a use period, SMs send electricity consumption measurements automatically and at a much higher resolution. They enable two-way communication with the UP, the entity that sells energy to the customers, transmitting a great amount of detailed information. SMs collect and send bidirectional readings of active, reactive, and apparent power

and energy—i.e., so-called four-quadrant metering—that is purchased from the grid (or sold to it, if the user produces energy, e.g., by means of a photovoltaic panel). In the latter case, the user is referred to as a *prosumer*, i.e., at once a producer and consumer of electricity who can be financially rewarded for the energy sold to the grid.

SMs also keep track of historical consumption data over the previous days, weeks, and months and provide high-resolution consumption data analytics to the customers to enable them to monitor their energy consumption via an in-home display, web portal, or smartphone application in near real time. SMs also send alerts about voltage quality measurements, helping UPs fulfill their obligations toward customers concerning energy, power, and voltage quality, e.g., in accordance with the EN 50160 European standard. Examples of these measurements include the root mean square voltage variations, e.g., voltage dropout, sags and swells, and total harmonic distortion.

Data used for billing, such as the current ToU tariff, balance and debts, credit and prepayment modes, credit alerts, and topping up, are also sent to the UP. SMs can detect if tampering takes place and send relevant data about it, along with the security credentials for enabling the correct functioning of cryptographic protocols, e.g., hashing, digital signatures, and cyclic redundancy checks. Finally, SM firmware information and updates are also communicated.

The increased data resolution is crucial for enabling SG functionalities. Table 1 shows the smallest time resolution of some SMs currently in use, which is on the order of a few minutes. The European Union recommends a time resolution of at least 15 min to allow the new SG functionalities [3]. For example, the current SM specifications in the United Kingdom mandate that an SM send integrated energy readings every 30 min to the UP, while the data sent to a user's in-home display can have a resolution of up to 10 s [4]. It should be noted that, with the increased adoption of renewable energy generation by prosumers, the increased penetration of electric vehicles and energy storage technologies, and the diversification of the energy market, it is expected that SGs will become more volatile, requiring meter readings at a much higher rate in the near future.

SMs provide a wide range of benefits to all of the parties in an SG. Thanks to SMs, UPs can gain better knowledge of their customers' needs while reducing the cost of meter readings. SMs allow UPs to dynamically determine the electricity cost and produce more accurate bills, thus reducing customers' complaints and back-office rebilling. The implementation of ToU pricing can incentivize demand response and control customer behavior, while improved demand forecasts and load-shaping techniques can reduce peak electricity demands. Finally, energy theft can be detected more easily and quickly.

Distribution system operators (DSOs), i.e., the entities that operate the grid, benefit from SMs as well, by being able to better monitor and manage the grid. SMs allow DSOs to reduce operational costs and energy losses and improve grid efficiency and system design, distributed system state estimation, and volt/VAR control. Moreover, DSOs are able to better match distributed resources with the ongoing electricity demand and

Table 1. The time resolution of SMs currently in use.

SM Model	Time Resolution
ltron Centron [72]	1 min
REX2 [73]	5 min
Kamstrup Ominpower [74]	5 min
Enel Open Meter [75]	15 min

the grid's power delivery capability, thus reducing the need to build new power plants.

Consumers themselves take advantage of SMs to monitor their consumption in near real time, leading to better consumption awareness and energy usage management. Moreover, consumers receive accurate and timely billing services, with no more estimated bills, and benefit from ToU pricing by shifting nonurgent loads to off-peak price periods. Microgeneration and energy storage devices can be integrated more easily, and profits from selling the generated excess energy can be collected automatically. Failing or inefficient home appliances, unexpected activity or inactivity, and wasted energy are detected faster and more accurately; in addition, switching between UPs is made easier by requesting on-demand readings, which, in turn, increases the competition among UPs and reduces costs for consumers.

For the aforementioned reasons, the installation of SMs is proceeding rapidly and attracting massive investment globally. The SM market is expected to grow from an estimated US\$12.79 billion in 2017 to US\$19.98 billion by 2022, registering a compound annual growth rate of 9.34% [5]. Moreover, the global SM data analytics market, which includes demand response analytics and grid optimization tools, is expected to reach US\$4.6 billion by 2022 [6], while the global penetration of SMs is expected to climb from approximately 30% at the end of 2016 to 53% by the end of 2025 [7]. These figures show how timely and crucial the research in this field is, and they highlight the need to quickly resolve potential obstructions that can threaten the future benefits from this critical technology.

SM privacy risks

An SM's ability to monitor a user's electricity consumption in almost real time presents serious implications for consumer privacy. In fact, by employing NILM techniques, it is possible to identify the power signatures of specific appliances from aggregated household SM measurements. NILM approaches date back to the 1980s work of George Hart, who first proposed a prototype of an NILM device [9]. Since then, NILM methods have improved in different directions, e.g., by assuming either high- or low-frequency measurements, by considering known or learned signatures [10], and even by using off-the-shelf statistical methods without any a priori knowledge of household activities [11].

An example of a typical power consumption profile, along with some detected appliances, is illustrated in Figure 1. As shown in Figure 2, the UP, a third party that has access to SM data (by, for example, buying it from the UP), or a malicious eavesdropper may gain insights into users' activities and behaviors, and determine such information as a person's presence at home, religious beliefs, disabilities, illnesses, and even the TV channel being watched [12]–[14].

Apart from residential users, SM privacy is particularly critical for businesses, e.g., factories and data centers, as their power consumption profile may reveal sensitive information about the state of their businesses to their competitors. SM privacy has attracted significant public attention and continues

to be a topic of heated public and political debate. The issue even stopped the mandatory SM rollout plan in The Netherlands in 2009 after a court decided that the forced installation of SMs would violate consumers' right to privacy and be in breach of the European Convention of Human Rights [15]. Indeed, concerns about consumer privacy threaten the widespread adoption of SMs and can be a major roadblock for this multibillion-dollar industry.

It is worth pointing out that the privacy problem involving SMs is different from the SM data security problem [16]. In the latter, there is a sharp distinction between legitimate users and malicious attackers, whereas in the privacy problem in the SM context, any legitimate receiver of data can also be considered malicious. To benefit from the advantages provided by the SG, users need to share some information about their electricity consumption with the UP and DSO. However, by sharing accurate and high-frequency information about their energy consumption, consumers also expose their private lives and behavior to the UP, which is a fully legitimate user and, at the same time, a potential malevolent party. This renders traditional encryption techniques for data privacy ineffective in achieving privacy against the UP and calls for novel privacy measures and privacy-protection techniques.

Privacy-enabling techniques for SMs

There is a growing literature on SM privacy-preserving methods, which can be classified into two main families. The first, which we call the *SM data manipulation (SMDM)* approach, consists of techniques that process the SM data before reporting them to the UP; the practices in the second family, called *user demand shaping (UDS)*, aim at modifying the user's actual energy consumption. Considered within the first class are methods such as data obfuscation, data aggregation, data anonymization, and downsampling.

Data obfuscation, i.e., the perturbation of metering data by adding noise, is a classical privacy-protection method and has been adapted to SGs in [17] and [18]. In [19], differential




FIGURE 1. An example of a household electricity consumption profile with some appliances highlighted. (Figure courtesy of the Dataport database [8].)




FIGURE 2. Some questions an attacker may be able to answer by having access to SM data.

privacy, a well-established concept in the data-mining literature, is applied to SMs, where noise is added not only to the user's energy consumption, via the RB, but also to the energy used for charging the RB itself to provide differential privacy guarantees. Along these lines, the authors in [20] introduce an information-theoretic framework to study the tradeoff between the privacy obtained by altering the SM data and the utility of data for various SG functionalities. Note that the more noise added to the data, the higher the privacy but the less relevant and useful the data are for monitoring and controlling the grid. In [20], an additive distortion measure is considered to model the utility, which allows the characterization of the optimal privacy-versus-utility tradeoff in an information-theoretic single-letter form.

The data aggregation approach, proposed in [18], [21], and [22], considers sending the aggregate power measurements for a group of households so that the UP is prevented from distinguishing individual consumption patterns. The aggregation can be performed with or without the help of a trusted third party (TTP).

The data anonymization approach, on the other hand, mainly considers utilizing pseudonyms rather than the real identities of consumers [23], [24]. Another method, proposed in [25], reduces the SM sampling rate to a level that does not pose any privacy threat. However, the SMDM family suffers from the following shortcomings.

- Adding noise to the SM readings causes a mismatch between the reported values and the real energy con-

sumption, which prevents DSOs and UPs from accurately monitoring the grid state; rapidly reacting to outages, energy theft, or other problems; and producing accurate and timely billing services. These would significantly limit the SM benefits.

- DSOs, UPs, or, more generally, any eavesdropper can embed additional sensors right outside a household or a business (street-level measurements are already available to DSOs and UPs) to monitor the energy consumption without fully relying on SM readings.
- The anonymization and aggregation techniques that include the presence of a TTP only shift the problem of trust from one entity (the UP) to another (the TTP).

These issues are avoided by the UDS approaches, which directly modify the actual energy consumption profile of the user, called the *user load*, rather than modifying the data sent to the UP. In this family, the SM accurately reports the energy taken from the grid without any modification; however, this is not the energy that is actually consumed by the appliances. This is achieved by filtering the user's actual electricity consumption via a rechargeable energy storage device, i.e., an RB, or by exploiting an RES, which can be used to partially hide the consumer's energy consumption. Examples of RESs include solar panels and micro wind farms.

Another technique is to partially shift a consumer's demand. If we denote the energy received from the grid as the *grid load*, the idea is to physically differentiate the grid load from the user




FIGURE 3. An overview of the privacy-enabling approaches for SMs.

load. Note that the effect of using an RB or an RES can also be considered as adding noise to the household consumption, but the noise in this case corresponds to a physical variation in the energy received from the grid. Moreover, different from the approaches in the SMDM family, the SM measurements provided by the UDS methods are exact, and there is no issue of any data mismatch between the SM data and the effective user demand from the grid. Thus, when UDS methods are deployed, the utility of SMs for the SG is not diminished since the users' energy consumption is neither misreported nor distorted. As a result, while the privacy-versus-utility tradeoff is of particular concern for the SMDM techniques, with the UDS methods, SG utility is never diminished. Instead, other tradeoffs are considered, such as privacy versus cost or privacy versus wasted energy. Figure 3 shows an overview of the privacy-enabling approaches.

The focus of this article is on UDS techniques, which have been receiving growing attention from the research community in recent years. The physical resources these techniques rely on, such as RBs or RESs at consumer premises, are already becoming increasingly available, thanks to government incentives and the decreasing cost of solar panels and household and electric vehicle RBs. Moreover, shaping and filtering users' actual energy consumption by means of physical resources render any data misreporting or distortion unnecessary and thus, do not undermine the utility of the SG concept itself.

We present a signal processing perspective on SM privacy by treating the user load as a stochastic time series, which can be filtered and distorted by using an RB, an RES, and/or demand shaping/scheduling. The available energy generated by the RES can also be modeled as a random sequence, whose statistics depend on the energy source (e.g., solar or wind) and the specifications of the renewable energy generator. Additionally, the finite-capacity battery imposes instantaneous limitations on the available energy. We also note that such physical resources can be used as well for cost minimization purposes by the users, e.g., by acquiring and storing energy over low-cost periods and utilizing the stored energy in the RB and the energy generated by an RES over peak-cost periods. Accordingly, we also study the tradeoff between privacy and cost, as well as the minimization of the wasted renewable energy. Next, we describe and summarize the progress made in recent years toward quantifying SM privacy leakage in a rigorous manner, report the most significant results, and highlight a number of future research directions.

Current household batteries, typical energy demands, and renewable energy generation

Table 2 lists the storage capacity and peak power of some of the currently available RBs for residential use. As can be seen, the capacities are in the range of a few kilowatthours. For example, the peak power that batteries with a 4-kWh capacity

Table 2. The specifications of some currently available residential batteries.

Household RB	Capacity (kWh)	RB Charging Peak Power (kW)	RB Discharging Peak Power (kW)
Sunverge SIS-6848 [76]	7.7, 11.6, 15.5, 19.4	6.4	6
SonnenBatterie eco [77]	4–16	3–8	3–8
Tesla Powerwall 2 [78]	13.5	5	5
LG RESU 48V [79]	2.9, 5.9, 8.8	3, 4.2, 5	3, 4.2, 5
Panasonic battery system UJ-SK56A [80]	5.3	2	2
Powervault G200-LI2/4/6KWH [81]	2, 4, 6	0.8, 1.2	0.7, 1.4
Orison Panel [82]	2.2	1.8	1.8
SimpliPhi PHI 3.4–48V [83]	3.4	1.5	1.5




FIGURE 4. Examples of the user load, grid load, and target load over the course of a day when a piecewise target profile is considered [33]. The price periods are highlighted by arrows of different colors. Note that the target load assumes a different constant value for each price period.

can output sustainably is on the order of 1–2 kW. However, as the typical electricity consumption is very spiky (e.g., see Figure 1), current batteries cannot fully hide the consumption spikes, because of the charging/discharging peak power constraints. For example, while a 4-kWh battery can hide a constant consumption of 2 kW over 2 h, it cannot fully conceal spikes in the user load of more than 2 kW. An example of this effect can be noticed in the simulation of Figure 4.

The typical household average power consumption also lies within the range of a few kilowatts, as shown in Table 3, where the distribution of the average user power consumption values over different years obtained from various databases is reported, with various time resolutions. Analyzing the Dataport database [8], we observe that, independent of the period considered, the average user energy demand is fewer than 2 kWh 80–90%

of the time. Current batteries charged at full capacity would then be able to satisfy the demand continuously for only a few hours. However, completely covering the consumption over a few hours may come at the expense of revealing the energy consumption fully at future time periods. In fact, once the RB is discharged, it needs to be charged again before being able to hide the user consumption; hence, the use of the RB introduces memory into the system, as decisions taken at a certain time have an impact on the privacy performance at later times.

We should also remark that residential electricity consumption is forecast to increase significantly in the coming years [71], emphasizing the need to intelligently exploit limited-capacity storage devices to hide energy consumption behavior. We also would like to emphasize that the privacy leakage is caused mostly by these spikes, which are typically more informative (e.g., the oven, microwave, and heater) compared to more regular consumption (e.g., the refrigerator). Moreover, because of electricity price variations, users may prefer charging/discharging the battery during certain time periods, which limits the available energy that can be used for privacy. Finally, it is expected that the increasingly wider adoption of electric vehicles and the mass production and adoption of energy-hungry smart devices will inevitably increase the typical household electricity consumption, further limiting RBs' capability to fully hide the user load.

Table 4 shows the average power generated by typical residential solar panels, which are the most common residential RESs. The location, technology, inclination, and size of the solar panel affect the generated power, as shown in Table 5 for one of the databases considered, where *kWp* denotes the kilowatt peak, i.e., the output power achieved by a panel under full solar radiation. As expected, around 50% of the time, i.e., at night, no energy is generated at all, while there are differences in the distribution of the average values for the two databases considered, due to the different locations. Comparing these

Table 3. The distribution of the average household power consumption (resolution refers to the measurement frequency).

Source	Location	Resolution	Time Frame	Number of Houses	[0,0.5] kW	(0.5,1] kW	(1,2] kW	(2,3] kW	(3,4] kW	(4,+∞) kW
Dataport [8]	Texas, United States	60 min	1 January–31 May 2016	512	38	30	20	7	3	2
			1 January–31 December 2015	703	36	26	20	9	5	4
			1 January–31 December 2014	720	39	25	20	8	4	4
			1 January–31 December 2013	419	35	25	21	9	5	5
			1 January–31 December 2012	182	31	26	24	10	5	5
Intertek [26]	United Kingdom	2 min	1 May 2010–31 July 2011	251	18	24	47	11	0	0
Dred [27]	The Netherlands	1 s	5 July 2015–5 December 2015	1	98	1.8	0.4	0	0	0
Uci [28]	France	1 min	16 December 2006–26 November 2010	1	47	9	28	8	4	2

The values in each column indicate the percentage of time the average consumption falls into the corresponding interval.

Table 4. The distribution of the average power generated by residential photovoltaic systems.

Source	Location	Resolution	Time Frame	Number of Houses	0 kW	(0,0.5] kW	(0.5,1] kW	(1,2] kW	(2,3] kW	(3,4] kW	(4,+∞) kW
Dataport [8]	Texas, United States	60 min	1 January 2012–31 May 2016	351	49	17	7	9	7	6	5
Microgen [29]	United Kingdom	30 min	1 January–31 December 2015	100	51.7	36.4	9.8	2	0.1	0	0

The values in each column indicate the percentage of time the average generation falls into the corresponding interval.

Table 5. The specifications of the solar panels studied in the Microgen [29] database.

Solar Panel Area (m ²)					Solar Panel Cell Type		Nominal Installed Capacity (kWp)			
(0,15]	(15,20]	(20,25]	(25,30]	(30,+∞)	Monocrystalline	Polycrystalline	(0,2]	(2,3]	(3,4]	(4,+∞)
5	35	44	15	1	93	7	4	36	59	1

The values in each column indicate the percentage of solar panels that satisfy the corresponding property.

values with those in Table 2, we note that the battery capacities are large enough to store many hours of average solar energy generated by the solar panels most of the time.

A signal processing perspective on SM privacy

A generic discrete-time SM system model is depicted in Figure 5. In this model, each time slot is normalized to unit time; therefore, the power and energy values within a time slot are used interchangeably. $X_t \in \mathcal{X}$ denotes the total power demanded by the appliances in the household in time slot t , i.e., the user load, where $\mathcal{X}$ is the user load alphabet, i.e., the set of values that X_t can assume. The sequence $\{X_t\}$ represents the user's private information that needs to be protected. $Y_t \in \mathcal{Y}$ is the power received from the grid in time slot t , i.e., the grid load, which is measured and reported to the UP by the SM, while $\mathcal{Y}$ denotes the grid load alphabet. We assume that the user load and grid load power values remain constant within a time slot. In practice, this can be considered as a discrete-time linear approximation of a continuous-load profile. This approximation can be made as accurate as desired by reducing the time slot duration.

In current systems, where no energy manipulation is employed, $Y_t = X_t, \forall t$; that is, the actual energy consumption of the appliances is reported to the UP by the SM. Instead, we will assume that an RB and an RES are available to the user to physically distort the energy consumption, so that what the user receives from the grid, Y_t , does not reveal too much

information about the energy used by the appliances, X_t . We remark here that the time slots in our model correspond to time instants when the electricity is actually requested by the user and drawn from the grid, rather than the typically longer sampling interval used for sending SM measurements to the UP. In fact, we assume that the SM measures and records the output power values at each time slot. This is because our aim is to protect consumers' privacy not only from the UP but also from the DSO or any other attacker that may deploy a sensor on the consumer's power line, recording the electricity consumption in almost real time.




FIGURE 5. The system model. X_t , Y_t , E_t , and B_t denote the consumer's power demand, i.e., the user load; the SM readings, i.e., the grid load; the power produced by the RES; and the battery state of charge at time t , respectively. The dashed line represents the meter readings being reported to the UP. EMU: energy management unit.

The amount of energy stored in the RB at time t is $B_t \in [0, B_{\max}]$, where $B_{\max}$ denotes the maximum battery capacity. $X_t - Y_t$ is the power taken from the RB, and the battery charging and discharging processes are often constrained by the so-called charging and discharging power constraints $\hat{P}_c$ and $\hat{P}_d$, respectively, i.e., $-\hat{P}_c \leq X_t - Y_t \leq \hat{P}_d, \forall t$. There is also typically a constraint on the average energy that can be retrieved from an RB, imposed by an average power constraint $\bar{P}$, i.e., $\mathbb{E}[\frac{1}{n} \sum_{t=1}^n (X_t - Y_t)] \leq \bar{P}$. Losses in the battery charging and discharging processes may also be taken into account to model a more realistic energy management system. The renewable energy generated at time t by the RES is represented by $E_t \in \mathcal{E}$, where $\mathcal{E} = [0, E_{\max}]$. RBs and RESs are expensive facilities, and installation and operation costs can be reduced if they are shared by multiple users, e.g., users within the same neighborhood or block of apartments. Moreover, sharing these resources allows the centralized management of the energy system, which also leads to a more efficient use of the available resources. The renewable energy can be stored in the RB or used immediately so that a user can

- increase privacy by not reporting the actual power consumption to the UP
- decrease electricity costs by purchasing and storing electricity from the grid when it is cheaper and using it to satisfy future demand—or even selling it back to the UP when the price increases
- increase energy efficiency by reducing the waste of generated renewable energy when it is not needed and, when it is not profitable, to sell it to the UP.

The random processes X and E are often modeled as Markov processes or as sequences of independent and identically distributed (i.i.d.) random variables. Although the UP typically does not know the instantaneous outcomes of these processes, it may well know their statistics. In some cases, the UP may know the realizations of the renewable energy process E , for example, if it has access to additional information from sensors deployed near the household that measure different parameters, e.g., the solar or the wind power intensity, and if it knows the specifications of the user's renewable energy generator, e.g., the model and size of the solar panel.

Given these definitions, the equation expressing the evolution of the energy level in the battery is

$$B_{t+1} = \min \{B_t + E_t - (X_t - Y_t), B_{\max}\}. \quad (1)$$

Sometimes, the user load does not need to be satisfied immediately in its entirety. In fact, it can be further classified into demand that must be met immediately (e.g., lighting or cooking) and demand that can be satisfied at a later time, the so-called elastic demand (e.g., charging an electric vehicle or running a dishwasher or washing machine). For the latter demand, the user's only concern is that a certain task needs to be finished by a certain deadline (e.g., the electric car must be fully charged by 8 a.m.), and it does not matter exactly when

the consumption takes place. This flexibility allows the consumer to employ demand response to increase privacy and lower the energy cost.

The electricity unit cost at time t , denoted by C_t , can be modeled as a random variable or in accordance with a specific ToU tariff. The cost incurred by a user to purchase Y_t units of power over a time interval τ_t at price C_t is thus given by $\tau_t Y_t C_t$. When the presence of an RES is considered, the prosumer may be able to sell part of the energy generated to the grid to further improve privacy and minimize the energy cost. If this occurs, the net metering approach is typically considered, i.e., the utilities purchase consumer-generated electricity at the current retail electricity rate. The battery wear and tear due to charging and discharging the RB can also be taken into account and modeled as an additional cost [30].

The energy management policy

The energy management unit (EMU) is the intelligence of the system, located at the user's premises, where the SM privacy-preservation and cost-optimization algorithms are physically implemented. The energy management policy (EMP), implemented by the EMU, determines at any time t the amount of energy that should be drawn from the grid and the RB, given the previous values of the user load X^t , renewable energy E^t , level of energy in the battery B^t , and grid load Y^{t-1} , i.e.,

$$f_t: X^t \times \mathcal{E} \times \mathcal{B} \times \mathcal{Y}^{t-1} \rightarrow \mathcal{Y}, \quad \forall t, \quad (2)$$

where $f \in \mathcal{F}$, and $\mathcal{F}$ denotes the set of feasible policies, i.e., policies that produce grid load values that satisfy the RB and RES constraints at any time as well as the battery update equation in (1). The optimal policy is chosen to minimize the long-term information leakage about a consumer's electricity consumption, possibly along with other criteria, such as the minimization of electricity cost or wasted energy. The EMP prevents outages, and typically it is not allowed to draw more energy from the grid to be wasted simply for the sake of increased privacy.

The policy f_t in (2) corresponds to an online EMP, i.e., one in which the action taken by the EMU at any time slot depends only on the information available causally right up to that time. Alternatively, in an offline optimization framework, the policy takes actions based also on future information about the system state, i.e., the user load and RES energy generation, in a noncausal fashion. In the SM privacy literature, both offline and online SM privacy-preserving algorithms have been considered. Online algorithms are more realistic and relevant for real-world applications; however, offline algorithms may lead to interesting intuition or bounds on the performance. Moreover, noncausal knowledge of the electricity price process is a realistic assumption in today's energy networks; and even the noncausal knowledge of power consumption may be valid for certain appliances, such as refrigerators, boilers, heaters, and electric vehicles, whose energy consumption can be accurately predicted over certain finite time frames.

A heuristic privacy measure: Variations in the grid load profile

As in many other problems involving privacy, a wide consensus over the best privacy measures for SMs has not yet been reached. A number of privacy measures have been proposed in the literature, each with its own benefits and limitations. Although it is clear that privacy is achieved when the UP cannot infer a user's behavior on the basis of SM measurements, it is challenging to define a corresponding mathematical measure that is independent of the particular detection technique employed by the attacker.

Grid load variance as a privacy measure

One can argue that privacy in SMs can be ensured by opportunistically charging and discharging the RB so that the grid load is always constant. In fact, the differences in consecutive load measurements $Y_t - Y_{t-1}$ are indicative of the appliances' switch-on/off events, the so-called features, and are typically exploited by the existing NILM algorithms. Ideally, a completely flat grid load profile would not reveal any feature and would only leak a user's long-term average power consumption. However, this would require a very large battery capacity and/or a powerful RES. Alternatively, the level of privacy can be measured by what we might call the *distance* of the grid load from a completely flat target load profile, based on the intuition that the smaller the distance, the higher the level of privacy achieved [31]. Accordingly, privacy can be defined as the grid load variance around a prefixed target load profile W , i.e.,

$$\mathcal{V}_n \triangleq \frac{1}{n} \sum_{t=1}^n \mathbb{E}[(Y_t - W)^2], \quad (3)$$

where the expectation is over X_t and Y_t , and typically $W = \mathbb{E}[X]$.

Another important concern for consumers is their energy cost. With the integration of unreliable RESs into the grid, it is expected that the unit cost of energy from different UPs will fluctuate over time. RBs for residential use provide flexibility to consumers, as they can buy and store energy during low-cost periods to be used during peak-price periods. The impact of RBs in reducing the cost of energy to consumers has been extensively studied in the literature [32]. Note, however, that the operation of the EMU to minimize the energy cost does not necessarily align with the goal of minimizing privacy leakage. Therefore, it is essential to jointly optimize the electricity cost and user privacy. If the cost of energy and battery wear and tear are considered, the overall optimization problem becomes

$$\min \frac{1}{n} \sum_{t=1}^n \mathbb{E}[C_t Y_t + \mathbb{1}_B(t) C_B + \alpha (Y_t - W)^2], \quad (4)$$

where $\mathbb{1}_B(t) = 1$ if the battery is charging/discharging at time t and equals 0; otherwise, C_B is the battery operating cost due to the battery deterioration caused by charging and discharging the RB, and α strikes the tradeoff between privacy and cost. The expectation in (4) is over the probability distributions of all of the involved random variables, i.e., X_t , Y_t , and C_t .

If $W_t = \mathbb{E}[X]$, $\forall t$, the EMU tries to achieve a flat grid load profile around the average user energy consumption with as few deviations as possible. This scenario is illustrated in Figure 6, where the straight blue line is the fixed target consumption profile W_t and the red line indicates the achieved grid load profile Y_t . For i.i.d. X and C processes, an online EMP can be obtained using Lyapunov optimization [30]. The online control algorithm can be formulated as a Lyapunov function with a perturbed weight, and the drift-plus-penalty framework is adopted, which is typically used for stabilizing a queuing network, by minimizing the so-called drift while at the same time minimizing a penalty function. Here, the penalty is represented by the optimization target, while the Lyapunov drift is defined as the difference of the level of energy in the RB at successive time instants. The authors in [30] show that this approach leads to a mixed-integer nonlinear program, which they solve by decomposing it into multiple cases and finding a closed-form solution to each of them.

This problem can also be studied in an offline framework by assuming that the future user demand profile can be accurately estimated for a certain time horizon and that the energy cost is known in advance. When privacy and cost of energy are jointly optimized over a certain time horizon, one can characterize the points on the Pareto boundary of the convex region formed by all of the cost and privacy leakage pairs by solving the following convex optimization problem [31]:

$$\min_{Y_t \geq 0} \sum_{t=1}^n [(1 - \alpha) Y_t C_t + \alpha (Y_t - W)^2]. \quad (5)$$

It is shown in [31] that the optimal offline solution has a water-filling interpretation. However, unlike the classical water-filling algorithm, which appears as the solution of the power allocation problem across parallel Gaussian channels under a total power constraint, here the water level is not constant but changes across time because of the instantaneous power constraints.




FIGURE 6. An example of the user load, grid load, and constant target load profiles, where the distance $Y_t - W_t$ is highlighted. The aim of the algorithms presented in the “Grid Load Variance as a Privacy Measure” section is to minimize the average squared distance.

A completely flat consumption profile may not be feasible or even desirable—e.g., if the cost varies greatly during the system operation because of ToU tariffs. Thus, it is reasonable to assume that a user requests more energy during off-peak price periods as compared to peak price periods and hence allows a piecewise constant target load [33]. An example of this strategy is shown in Figure 4, where it is applied to real power consumption data from the UK-DALE data set [34]. The optimization problem (5) becomes

$$\min_{Y, W} \frac{1}{N} \sum_{i=1}^M \left[\sum_{t=t_c^{(i-1)}}^{t_c^{(i)}} (1-\alpha) Y_t C^{(i)} + \alpha (Y_t - W^{(i)})^2 \right], \quad (6)$$

where $C^{(i)}$ and $W^{(i)}$ are the cost of the energy purchased from the UP and the target profile during the i th price period,




FIGURE 7. The ToU tariff and timing conventions used for a piecewise target profile [33]. The time instants at which the price of energy changes are represented by $t_c^{(i)}$, for $i = 1, \dots, M$, and $t_c^{(0)} = 0$.




FIGURE 8. The privacy-versus-cost tradeoff when using a Powervault G200-LI-4KWH battery for the strategies in [31] and [33].

respectively, where $1 \leq i \leq M$; M is the total number of price periods during time T ; and the i th price period spans from time slot $t_c^{(i-1)}$ to $t_c^{(i)}$. Figure 7 depicts the timing convention considered in this scenario. Energy can be sold to the UP to further improve the privacy-versus-cost tradeoff, as assumed in [33]. Considering a piecewise target profile improves the overall privacy-versus-cost tradeoff compared to a constant target profile, as shown in Figure 8 for a Powervault G200-LI-4KWH RB when using power consumption data from [34].

A possible extension of the latter work is to consider the multiuser scenario, where, in principle, each user can fix its own target profile. As long as the target profile does not depend on the user's energy consumption profile, the UP does not receive much information about the consumer's activities. On the other hand, the UP can implicitly incentivize users to choose different target profiles by setting different ToU prices for different consumers. Since consumers will tend to buy more energy when it is cheaper, each of the users in the neighborhood will shift the load to a different time slot, also balancing the total load on the grid.

Markov decision process formulation

In the online optimization framework, where the user load and the energy generated by the RES can be modeled as Markov processes (or as i.i.d. sequences as a special case), the SM privacy problem can be cast as a Markov decision process (MDP). An MDP is a discrete-time state-transition system that is formally characterized by the following:

- a state space
- an action space, which includes the possible actions that can be taken by the decision maker at each state
- the transition probabilities from the current state to the next state, which describe the dynamics of the system
- the reward (or inversely the cost) process, which indicates the reward received (or cost incurred) by the decision maker by taking a particular action in a particular state.

The goal of an MDP is to find the optimal policy that minimizes the average (or discounted) cost either by a specified time in the future, i.e., by considering the so-called finite-horizon setting, or over an indefinite time period, by considering the infinite-horizon setting. To solve the corresponding MDP, the Bellman optimality equations should be formulated [35], which can be solved to obtain the optimal policy at each state and time instant. The problem can be solved numerically for the finite-horizon setting, while the value iteration algorithm can be employed to obtain the optimal stationary policy in the infinite-horizon scenario.

In the SM problem, the state at any time t is typically represented by a combination of the current level of energy in the battery B_t , user demand X_t , and renewable energy E_t . The action, performed by the EMU, is represented by the current grid load and the energy used from the RB and RES. State transitions are modeled by the battery update equation, which is typically assumed to be deterministic, and by transitions in the user demand and renewable generation states, which typically do not depend on users' actions. The cost function is the privacy

loss that is experienced when moving from one level of energy in the battery to another by following a certain action.

However, to consider privacy as the cost function in an MDP, it is necessary to formulate the privacy leakage in an additive form across time, so that the total loss of privacy over multiple time slots is given by the summation of the privacy leakage at different time slots. This may be challenging, depending on the privacy measure employed. For example, measuring privacy via the squared distance of the grid load from a constant target profile has a straightforward additive formulation, while the same does not hold when privacy is measured by the mutual information (MI) between the user and grid load sequences. This is because the MI takes into account the dependence between the realization of the user load at time t , X_t and the current, past, and future realizations of $Y, Y_1, \dots, Y_t, \dots$.

When the state and action spaces are continuous, it is necessary to discretize them to solve the problem numerically. The accuracy of the numerical solution can be improved by decreasing the discretization step size, though at the expense of significantly higher computational complexity. When the dimensions of the state and action spaces render numerical evaluation of the optimal policy unfeasible, one can resort to suboptimal solutions that are easier to optimize and compute numerically yet may provide near-optimal performance or interesting intuition. Also, when the information-theoretic privacy measures are used, it may be possible to simplify the infinite-horizon optimization problem and write it in a single-letter form. We will provide further insight into this next.

The SM problem is cast as an MDP in [36], where the loss of privacy is measured by the fluctuations of the grid load around a constant target load, and the joint optimization of privacy and cost is studied. The optimal privacy-preserving policies are characterized by minimizing the expected total cost. Denote by u_t the action at time t . To solve the MDP, the transition probabilities $p(X_t | X_{t-1})$ and $p(B_t | B_{t-1}, u_t)$ need to be known; however, this is normally not the case, as the user load and the energy storage usage are typically nonstationary. The authors in [36] overcome this issue by adopting the Q-learning algorithm [37], which is an iterative algorithm used for characterizing the expected cost for each state–action pair by alternating the exploitation and exploration phases. The corresponding offline optimization policy is also characterized in [36] to be considered as a benchmark for the online algorithm. The paper characterizes the privacy–cost tradeoff curves and also evaluates the performance of the proposed algorithm by means of the empirical mutual information.

Temporal and spatial similarities in the grid load as a privacy measure

Variations in the grid load profile can be captured by considering the power traces of single appliances and computing differences in power consumption both in the time domain, i.e., the consumption deviation over time of a specific appliance, and in the space domain, i.e., the consumption profiles of different appliances. As these variations are computed over a certain time horizon, when an online algorithm is considered,

future user electricity consumption is estimated by forecasting the future electricity prices and running Monte Carlo simulations. The optimal decision at any time is characterized by considering both the current inputs and the forecasts through a rolling online stochastic optimization process.

Load shifting, i.e., the scheduling of the user's flexible electricity demand in accordance with privacy as well as cost concerns, can also be considered. Load shifting is analyzed in [38] and [39], where privacy, cost of energy, and battery wear and tear are jointly optimized and an online algorithm is formulated. The objective is to minimize the sum of the current and expected electricity and charging/discharging costs together with the weighted power profile differences measured through the similarity parameters for an entire day. In [39], the effectiveness of three similarity measures are examined separately and jointly, considering only four typical appliances—an oven, a clothes dryer, a dishwasher, and an electric vehicle—for the sake of simplicity.

Heuristic algorithms

While the grid load can be flattened by minimizing its variation around a constant consumption target, several works in the literature propose heuristic battery charging and discharging algorithms that keep the grid load variations limited. An intuitive approach is to try to keep the grid load equal to its most recent value by discharging (or charging) the RB when the current user load is larger (or smaller) than the previous one. This approach, called the *best-effort (BE) algorithm* in [40], tends to eliminate the higher-frequency components of the user load while still revealing the lower-frequency components.

In [40], the similarity between the two probability distributions of the user and grid loads is quantified via the empirical relative entropy, i.e., the Kullback–Leibler (KL) divergence [41]. In the same work, the authors also consider cluster classification, whereby data are clustered according to power levels and cross-correlation and regression procedures, according to which the grid load is shifted in time at the point of maximum cross-correlation with the user load, and regression methods are then used to compare the two aligned signals.

The study in [42] considers a slightly more sophisticated approach, called the *nonintrusive load-leveling (NILL) algorithm*, in which more than one grid load target value, namely, a steady-state target and low and high recovery state targets, are allowed, and where the EMU tries to maintain the grid load at one of these values across time. If the steady-state load cannot be maintained, the EMU switches to a high (or low) recovery state in case of persistent light (or heavy) user demand. When one of the recovery states is reached, the target load is adapted accordingly to permit the battery to charge or discharge, similar to the empirical strategies outlined in [43]. The value of the steady-state target load can be updated whenever a recovery state is reached, to reduce the occurrences of recovery states, which is achieved by using an exponential weighted moving average of the demand. To assess their proposed approach, the authors in [42] count the number of features, i.e., the number of times a device is recognized as being on or off based on the grid load as compared to the user load.

As also pointed out in [44], these heuristic algorithms suffer from precise load change recovery attacks that can identify peaks of user demand. We note that the NILL algorithm essentially quantizes the input load to three values with the help of the RB. This idea is generalized in [44] by considering an arbitrary number of quantization levels. Since quantization is a many-to-few mapping, converting the grid load to a step function is inherently a nonlinear and irreversible process, which can be used to provide privacy by maximizing the quantization error under battery limitations. More specifically, the grid load is forced to be a multiple of a quantity β , i.e., $Y_t = h_t \beta$, where h_t is an integer value and β is the largest value that satisfies the battery's maximum capacity and power constraints. At any time slot, given the user load, the grid load is chosen between the two levels adjacent to the user load, namely, $\lceil \frac{X_t}{\beta} \rceil$ and $\lfloor \frac{X_t}{\beta} \rfloor$, where $\lceil \cdot \rceil$ and $\lfloor \cdot \rfloor$ denote the ceiling and floor functions, respectively.

The study in [44] proposed three stepping algorithms that have different quantization levels: 1) the lazy stepping algorithm, which tries to maintain the external load constant for as long as possible; 2) the lazy charging algorithm, which keeps charging (or discharging) the battery until it is full (or empty); and 3) the random charging algorithm, which chooses its actions at random. While the simulation results show that these algorithms outperform the BE and NILL algorithms, with the lazy stepping algorithm typically performing the best, it is hard to make general claims because of the heuristic nature of these algorithms. In fact, these approaches do not provide theoretical guarantees on the level of privacy achieved. Thus, they are not able to make any general claim about the strength of the proposed privacy-preserving approaches and their absolute performance. This is an important limitation, as consumers would like to know the level of privacy they can achieve, even if it is in statistical terms. Also, because such heuristics are often based on deterministic schemes, they are prone to be easily reverse-engineered.

Theoretical guarantees on SM privacy

One of the challenges in SM privacy is to provide theoretical assurances and fundamental limits on the information leaked by an SM system, independently of any assumption on the capability of an attacker or of the particular NILM algorithm employed. This is essential in privacy research, as privacy-preserving techniques may perform extremely well against some NILM algorithms and very poorly against others. Moreover, the privacy assurances should not be based on the complexity limitations of a potential attacker, as techniques that are currently thought to be not feasible may become available to attackers in the future if computational capabilities improve or if new methods are developed.

Last but not least, establishing a coherent mathematical framework would allow us to rigorously compare various SM scenarios and the use of different physical resources, e.g., RBs of various capacities, RESs of various kinds, and so forth. Accordingly, signal processing and information-theoretic tools have been employed in the literature to provide theoretical privacy assurances. We will overview various statistical measures for privacy, particularly conditional entropy [43], Fisher infor-

mation (FI) [45], and type II error probability for detecting user activity [46].

In this statistical framework, it is commonly assumed that the statistics of the user load and the RES are stationary over the period of interest and known to the EMU. This assumption is reasonable, especially if the period of stationarity is sufficiently long for the EMU to observe and learn these statistics [47]–[49]. On the other hand, an online learning theoretic framework can also be considered to account for the convergence time of the learning algorithm. Alternatively, most of the works in the literature that carry out a theoretical analysis also propose suboptimal policies that can be applied on real power traces, thus allowing the reader to gain an idea of the practical application and performance of these theoretically motivated techniques. We take a worst-case approach and assume that the statistics governing the involved random processes are known by the attacker. Note that this can only empower the attacker and strengthen the stated privacy guarantees.

The significance of single-letter expressions

It is expected that a meaningful privacy measure should consider the leakage of a user's information over a certain time period of reasonable length, because of the memory effects introduced by the RB and the RES. The energy consumption over a short period of time can be easily covered by satisfying all of the demand from the RB or the RES over this period, but this may come at the expense of fully revealing the energy consumption at future time periods. Therefore, the information-theoretic analysis typically considers an average information rate measured over a given finite time period and often studies its infinite-horizon asymptotics as well.

However, increasing the time horizon also increases the problem complexity, and one of the challenges of the information-theoretic analysis is to obtain a so-called single-letter expression for the optimal solution, which would significantly reduce the problem complexity, particularly when the involved random variables are defined over finite alphabets. Unfortunately, to date, closed-form or single-letter expressions for the information leaked in an SM system have been characterized only for specific settings under various simplifications, e.g., considering an i.i.d. or Markov user load or RES generation.

MI as a privacy measure

The entropy of a random variable X , $H(X)$ is a measure of the uncertainty of its realization. The MI between random variables X and Y , $I(X; Y)$, measures the amount of information shared between the two random variables [41]. The MI can also be considered as a measure of dependence between the random variables X and Y , and it is equal to zero if and only if they are independent. Rewriting the MI as $I(X; Y) = H(X) - H(X|Y)$, where $H(X|Y)$ is the conditional entropy, we can also interpret MI as the average reduction in the uncertainty of X from the knowledge of Y . Therefore, we can measure the privacy leakage about the input load sequence X^n through the SM readings Y^n by the MI between the two sequences $I(X^n; Y^n)$. This will measure the

reduction in the uncertainty of the UP about the real energy consumption of the appliances X^n after receiving the SM measurements Y^n . For an SM system with only an RB (but no RES) and a given EMP f in (2) running over n time slots, the average information leakage rate $\mathcal{I}_f^n(B_{\max}, \hat{P}_d)$ is defined as

$$\mathcal{I}_f^n(B_{\max}, \hat{P}_d) \triangleq \frac{1}{n} I(X^n; Y^n) = \frac{1}{n} [H(X^n) - H(X^n | Y^n)], \quad (7)$$

where $0 \leq X_t - Y_t \leq \hat{P}_d$. The parameters $B_{\max}$ and $\hat{P}_d$ emphasize the dependence of the EMP and, therefore, of the achievable information leakage rate on the battery capacity and the discharging peak power constraint. The optimal EMP and the corresponding minimum information leakage rate are obtained by minimizing (7) over all of the feasible policies $f \in \mathcal{F}$ to obtain $\mathcal{I}^n(B_{\max}, \hat{P}_d)$.

Privacy with an RES

Alternatively, one can also consider the SM system of Figure 5 with an RES but no RB. Assume that the renewable energy that can be used over the operation period is constrained by an average and a peak power constraint. We do not allow selling the generated renewable energy to the UP, as our goal is to understand the impact of the RES on providing privacy to the user. The minimum information leakage rate achieved under these assumptions and for an i.i.d. user load can be characterized by the so-called privacy-power function $\mathcal{I}(\hat{P}, \hat{P}_d)$ and formulated in the following single-letter form:

$$\mathcal{I}(\hat{P}, \hat{P}_d) = \inf_{p_{Y|X} \in \mathcal{P}} I(X; Y), \quad (8)$$

where $\mathcal{P} \triangleq \{p_{Y|X} : y \in \mathcal{Y}, \mathbb{E}[(X - Y)] \leq \hat{P}, 0 \leq X - Y \leq \hat{P}_d\}$. This formulation is presented in [50] for a discrete user load alphabet (i.e., X can assume only values that are multiples of a fixed quantum) and in [51] for a continuous user load alphabet (i.e., X can assume any real value within the limits specified by the peak power constraints of the appliances). The optimal EMP that minimizes (8) is stochastic and memoryless; that is, the optimal grid load at each time slot is generated randomly via the optimal conditional probability that minimizes (8) by only considering the current user load. Another interesting observation is that (8) is in a form similar to the well-known rate-distortion function in information theory, which characterizes the minimum compression rate R of data, in bits per sample, that is required for the receiver to reconstruct the source sequence within a specified average distortion level D [41]. Formally, the rate-distortion function $R(D)$ for an i.i.d. source $X \in \mathcal{X}$ with distribution p_X , reconstruction alphabet $\hat{\mathcal{X}}$, and distortion function $d(\hat{x}, x)$, where the distortion between sequences X^n and $\hat{X}^n$ is given by $\frac{1}{n} \sum_{i=1}^n d(x_i, \hat{x}_i)$, characterizes the minimum rate with which an average distortion of D is achievable. The compression rate specifies the size of the codebook 2^{nR} required to compress the source sequence of length n , X^n . Shannon showed that the rate-distortion function can be obtained in the following single-letter form:

$$R(D) = \min_{p_{\hat{X}|X}: \sum_{(\hat{x}, x)} p_X p_{\hat{X}|X} d(x, \hat{x}) \leq D} I(\hat{X}; X). \quad (9)$$

The analogy between (8) and (9) becomes clear considering the following distortion measure:

$$d(x, y) = \begin{cases} x - y, & \text{if } 0 \leq x - y \leq \hat{P}_d, \\ \infty, & \text{otherwise,} \end{cases} \quad (10)$$

and such an analogy enables the use of tools from rate-distortion theory to evaluate the privacy-power function for an SM system. However, it is important to highlight that, despite the functional similarity, there are major conceptual differences between the two problems: 1) in the SM privacy problem, Y^n is the direct output of the encoder rather than the reconstruction at the decoder side, and 2) unlike the lossy source encoder, the EMU does not operate over blocks of user load realizations; instead, it operates symbol by symbol, acting instantaneously after receiving the appliance load at each time slot.

For discrete user load alphabets, the grid load alphabet can be constrained to the user load alphabet without loss of optimality [52], and, since MI is a convex function of the conditional probability $p_{Y|X} \in \mathcal{P}$, the privacy-power function can be written as a convex optimization problem with linear constraints. Algorithms such as the Blahut–Arimoto (BA) algorithm can be used to numerically compute the optimal conditional distribution [41]. For continuous user load distributions, the Shannon lower bound is derived in [52], which is a computable lower bound on the rate-distortion function widely used in the literature and is shown to be tight for exponential user load distributions.

These results can be generalized to a multiuser scenario in which N users, each equipped with a single SM, share the same RES [52]. This scenario is represented in Figure 9, where the objective is to minimize the total privacy loss of N consumers (or devices) considered jointly, rather than minimizing the privacy loss for each of them separately. This requires the EMU to allocate the shared RES among all of the users in the most effective manner. The average information leakage rate can still be written as in (7), by replacing X_t and Y_t with $\mathbf{X}_t = [X_{1,t}, \dots, X_{N,t}]$ and $\mathbf{Y}_t = [Y_{1,t}, \dots, Y_{N,t}]$, where the bold-face characters denote the vectors representing the N power measurements. The privacy-power function has the same expression as in (8), and, for the case of independent but not




FIGURE 9. A single EMU and RES are shared among N users, each equipped with an SM. The EMU decides how much energy each user can retrieve from the RES and from the grid.

necessarily identically distributed user loads, the optimization problem (ignoring the peak power constraint) can be cast as

$$\mathcal{I}(\bar{P}) = \inf_{\sum_{i=1}^N P_i \leq \bar{P}} \sum_{i=1}^N \mathcal{I}_{X_i}(P_i), \quad (11)$$

where $\mathcal{I}_{X_i}(\cdot)$ denotes the privacy-power function for the i th user having user load distribution $p_{X_i}(x_i)$. For continuous and exponential user loads, the optimal allocation of the energy generated by an RES can be obtained by the reverse water-filling algorithm, according to which energy from the RES is used only to satisfy the users with a low average load, while users with higher average load need to request energy from the grid as well.

Privacy with an RB

We can also consider the presence of just an RB in the system, which is thus charged only via the grid (no RES is available to the EMU). Including an RB complicates the problem significantly, and the level of energy in the battery B_t plays an important role when designing a feasible EMP.

This problem can be solved by putting it in the form of an MDP and finding a suitable additive formulation for the privacy cost function [53]. The optimization problem is formulated as

$$L^* \triangleq \min_f \frac{1}{n} I(B_1, X^n; Y^n), \quad (12)$$

where f can be any feasible policy, as specified in (2) (without including the renewable energy process). The approach in [53] casts (12) in an additive formulation by noting that there is no loss of optimality in restricting the focus to charging strategies f' that decide on the grid load, based only on the current values of the user load X_t and level of energy in the battery B_t and on the past values of the grid load Y^{t-1} . That is, the general strategy f in (2) is specified as $f'_t: \mathcal{X} \times \mathcal{B} \times \mathcal{Y}^{t-1} \rightarrow \mathcal{Y}$, $\forall t$, because of the following inequality:

$$\frac{1}{n} I(X^n, B_1; Y^n) \geq \frac{1}{n} \sum_{t=1}^n I(X_t, B_t; Y_t | Y^{t-1}). \quad (13)$$

The conditional distributions in (13) grow exponentially with time because of the term Y^{t-1} , and the problem becomes computationally infeasible very quickly. To overcome this issue, the knowledge of Y^{t-1} is summarized into a belief state, defined as $p(X_t, B_t | Y^{t-1})$, which can be computed recursively and interpreted as the belief that the UP has about (X_t, B_t) at time t , given its past observations Y^{t-1} . This way, the Bellman equations can be formulated, and the optimal policy can be identified numerically (with a discretization of the belief state).

For an i.i.d. user load, the single-letter characterization of the minimum information leakage rate is given by [53] as

$$J^* \triangleq \min_{\theta \in \mathcal{P}_B} I(B - X; X), \quad (14)$$

where θ is the probability distribution over B , given the past output and actions, i.e., $\theta_t \triangleq p(b_t | y^{t-1}, a^{t-1})$, and the action a_t is defined as the transition probability from the current

belief, user load, and level of energy in the battery to the current grid load. This result is obtained by considering a belief on $W_t \triangleq B_t - X_t$, rather than (B_t, X_t) , and by further restricting to policies of the type $f'_t: \mathcal{W} \times \mathcal{Y}^{t-1} \rightarrow \mathcal{Y}$, $\forall t$. Since (14) is convex in θ , the optimal θ^* may be obtained by using the BA algorithm. The optimal grid load turns out to be i.i.d. and indistinguishable from the demand, while the optimal policy is memoryless, and the distribution of Y_t depends only on W_t . Such a characterization is provided in [54] for a binary i.i.d. user load, while the authors extend it to an i.i.d. user load of generic alphabet size in [53], [55], and [56].

Another approach is to model the level of energy in the RB as a trapdoor channel [57]. In a trapdoor channel, a certain number of red or blue balls are within the channel, and a new ball of either color is inserted into it as the channel input at each time step. After the new ball is inserted, one of the balls present in the channel is randomly selected and removed from the channel. In an SM setting, the finite-capacity RB can be viewed as a trapdoor channel, whereby inserting or extracting a ball from the channel represents charging or discharging the RB, respectively. An upper bound on the information leakage rate is characterized in [58] through this model by minimizing the information leakage rate over the set of stable output balls, i.e., the set of feasible output sequences Y^n that can be extracted from the channel, given a certain initial state and an input sequence X^n and by taking inspiration from codebook construction strategies in [59]. Such an upper bound is characterized in [58] as

$$\frac{1}{n} I(X^n; Y^n) \leq \frac{1}{\lfloor (B_{\max} + 1)/X_{\max} \rfloor}, \quad (15)$$

where $X_{\max}$ is the largest value X can assume. It is also shown in [58] that the average user energy consumption determines the level of achievable privacy.

Apart from only maximizing privacy, it is of interest to also minimize the cost. Different from privacy, the cost of energy has an immediate additive formulation and can be easily incorporated into the MDP construction. Considering the random price vector $C^t = (C_1, \dots, C_t)$, where C_t denotes the unit cost of energy at time slot t , privacy can be defined in the long time horizon as

$$\mathcal{P} \triangleq \lim_{t \rightarrow \infty} \frac{H(X^t | Y^t, C^t)}{t}. \quad (16)$$

This formulation is presented in [43], where the corresponding MDP is constructed and two suboptimal algorithms are proposed. The first is a greedy algorithm, which maximizes at any time the current instantaneous reward, while the second is a battery-centering approach that is aimed at keeping the battery at a medium level of charge so that the EMU is less constrained by the battery or the demand in determining the grid load. In the latter approach, if the grid load depends not on the current user load or the battery level but only on the current electricity price, the system is said to be in a hidden state, while it is said to be in a revealing state otherwise. The latter strategy is analyzed for an i.i.d. user load by considering the system as a recurrent Markov chain and adopting random walk theory.

Privacy with both an RES and an RB

When both an RES and an RB are present, the information-theoretic privacy analysis becomes more challenging. As an initial step, we can consider infinite and zero battery capacities, which represent, respectively, lower and upper bounds on the privacy leakage achievable for a practical SM system with a finite-capacity battery [60], [61]. When $B_{\max} = \infty$, the problem can be shown to be equivalent to the average and peak power-constrained scenario, and, interestingly, the privacy performance does not deteriorate, even if the UP knows the exact amount of renewable energy generated. This shows that keeping the renewable energy generation process private is more critical when the RB has a limited capacity.

Two different EMPs are shown to achieve the lower bound in [61]. In the BE policy, at any time slot, the optimal EMP derived from (11) is employed independently of the level of energy in the RB if there is sufficient energy in the RB, while otherwise all of the energy request is satisfied from the grid. The latter approach leads to full leakage of user consumption, but it can be shown that these events are rare enough that the information leakage rate does not increase. In the alternative store-and-hide policy, an initial storage phase is employed, during which all of the user's energy requests are satisfied from the grid while all of the generated renewable energy is stored in the battery. In the following hiding phase, the EMU deploys the optimal policy designed under average and peak power constraints.

On the other extreme, when $B_{\max} = 0$, the renewable energy that can be used at any time slot is limited by the amount of energy generated within that period. As expected, assuming that the UP has knowledge of the renewable energy process significantly degrades the privacy performance for this scenario. Figure 10 compares the minimum information leakage rate with respect to the renewable energy generation rate p_e for $|\mathcal{X}| = |\mathcal{E}| = |\mathcal{Y}| = 5$ when $B_{\max} = \{0, 1, 2, \infty\}$. In this figure, the curves for a finite battery capacity of $B_{\max} = 1$ and $B_{\max} = 2$ are obtained numerically by considering a suboptimal EMP [61].

The presence of a finite-capacity battery increases the problem complexity dramatically because of the memory effects induced by the finite battery, and single-letter expressions are still lacking for this scenario. A possible approach to find a theoretical solution to this problem is by extending the MDP formulation, as investigated in [62].

Detection error probability as a privacy measure

So far, we have considered approaches that try to hide the complete user energy demand from the UP. However, rather than hiding the entire energy consumption profile, in some cases it may be more meaningful to keep specific user activities private, such as whether there is anyone at home, whether the alarm has been activated, or whether someone is eating microwaved food. To keep the answer to such details private, the goal of the EMU is to maximize the attacker's probability of making errors when attempting to discern them.

Let the consumer's behavior that needs to be kept private belong to a set of M possible activities. Thus, we can treat the

attacker's decision and the user's action as an M -ary hypothesis test, i.e., $H \in \mathcal{H} = \{h_0, h_1, \dots, h_{M-1}\}$. When $M = 2$, the hypothesis test is said to be binary, and, by convention, the hypothesis h_0 , called the *null hypothesis*, represents the absence of some factor or condition, while the hypothesis h_1 , called the *alternative hypothesis*, is the complementary condition. For example, answering the question, "Is somebody at home?" corresponds to a binary hypothesis test, where h_0 is the hypothesis "somebody is not at home" and h_1 is the hypothesis "somebody is at home." It is reasonable to assume that the input load will have different statistics under these two hypotheses; accordingly, we assume that under hypothesis h_0 (h_1), the energy demand at time slot t is i.i.d. with $p_{X|h_0}$ ($p_{X|h_1}$). Based on the SM readings, the attacker aims at determining the best decision rule $\hat{H}(\cdot)$, i.e., the optimal map between the SM readings and the underlying hypothesis. In other words, the space of all possible SM readings $\mathcal{Y}^n$ is partitioned into two disjoint decision regions $\mathcal{A}_0$ and $\mathcal{A}_1$, defined as follows:

$$\mathcal{A}_0 \triangleq \{y^n | \hat{H}(y^n) = h_0\}, \quad (17)$$

$$\mathcal{A}_1 \triangleq \{y^n | \hat{H}(y^n) = h_1\}, \quad (18)$$

which correspond to the subsets of the SM readings for which the UP decides for one of the two hypotheses. The attacker's binary hypothesis test can incur two types of errors:

- type 1 error, in which a decision h_1 is made when h_0 is the true hypothesis (a false positive or false alarm); the type 1 error probability is $p_I = p_{Y^n|h_1}(\mathcal{A}_0)$
- type 2 error, in which a decision h_0 is made when h_1 is the true hypothesis (a false negative or miss); the type 2 error probability is $p_{II} = p_{Y^n|h_0}(\mathcal{A}_1)$.

The Neyman–Pearson test minimizes the type 2 error probability for a fixed maximum type 1 error probability and makes decisions by thresholding the likelihood ratio




FIGURE 10. The minimum information leakage rate with respect to the renewable energy generation rate p_e with $\mathcal{X} = \mathcal{E} = \mathcal{Y} = \{0, 1, 2, 3, 4\}$. The leakage for $B_{\max} = \infty$ has been found by setting $\hat{P} = 4$ [61].

$(p_{Y^n|h_0}(y^n|h_0))/(p_{Y^n|h_1}(y^n|h_1))$. Consider the worst case of an all-powerful attacker that has perfect knowledge of the EMP employed in the asymptotic regime $n \rightarrow \infty$ and denote by $p_{\Pi}^{\min}$ the minimal type 2 probability of error subject to a constraint on the type 1 error probability. Assuming that a memoryless EMP is employed by the EMU—i.e., the grid load at any time slot t depends only on the input load at the same time slot—then the attacker runs a Neyman–Pearson detection test on the grid load. We note that the memoryless EMP assumption is not without loss of optimality. However, it is justified on the grounds that characterizing the more general optimal policy with memory seems to be significantly more challenging and is unlikely to lend itself to a single-letter expression. The Chernoff–Stein lemma [41] links the minimal type 2 error probability $p_{\Pi}^{\min}$ to the KL divergence $D(\cdot \parallel \cdot)$ between the grid load distributions conditioned on the two hypotheses in the limit of the number of observations going to infinity:

$$\lim_{n \rightarrow \infty} -\frac{\log p_{\Pi}^{\min}}{n} = D(p_{Y|h_0} \parallel p_{Y|h_1}), \quad (19)$$

where the KL divergence between two probability distribution functions on X , p_X and q_X , is defined as [41]

$$D(p_X \parallel q_X) \triangleq \sum_{x \in \mathcal{X}} p_X(x) \log \frac{p_X(x)}{q_X(x)}. \quad (20)$$

Not surprisingly, to maximize privacy, the goal of the EMU is to find the optimal grid load distribution, which, given the user load X and the true hypothesis H , minimizes the KL divergence in (19) or, equivalently, minimizes the asymptotic exponential decay rate of $p_{\Pi}^{\min}$. However, the EMU is constrained by the available resources in making the two input load distributions produce similar grid load distributions. In particular, we impose a constraint on the average RES it can use. Thus, the objective is to solve the following minimization problem:

$$\min_{p_{Y|H} \in \mathcal{P}_{Y|H}} D(p_{Y|h_0} \parallel p_{Y|h_1}), \quad (21)$$

where $\mathcal{P}_{Y|H}$ is the set of feasible EMPs, i.e., those that satisfy the average RES generation rate $\bar{P}$, so that $(1/n)\mathbb{E}[\sum_{i=1}^n X_i - Y_i | h_j] \leq \bar{P}$, with $j = 0, 1$. This setting is studied in [63], where the asymptotic single-letter expressions of two privacy-preserving EMPs in the worst-case scenario are considered, i.e., when the probability of a type 1 error is close to 1. The first policy is a memoryless hypothesis-aware policy that decides on Y_t based only on the current X_t and H , while the second policy is unaware of the correct hypothesis H but takes into account all of the previous realizations of X and Y .

It is noteworthy that, even if the hypothesis-unaware policy with memory does not have access to the current hypothesis, it performs at least as well as the memoryless hypothesis-aware policy. This is because the hypothesis-unaware policy is able to learn the hypothesis with negligible error probability after observing the energy demand process

for a sufficiently long period. Additionally, the energy supply alphabet can be constrained to the energy demand alphabet without loss of optimality, which greatly simplifies the numerical solution to the problem.

FI as a privacy measure

FI is another statistical measure that can be employed as a measure of SM privacy [45]. Let some sample data x be drawn according to a distribution depending on an underlying parameter. Then, FI is a measure of the amount of information that x contains about the parameter. In the SM setting, Y^n is the sample data available to the attacker, while X^n is the parameter underlying the sample data that is to be estimated by the UP. Let $\hat{X}^n$ denote the estimate of the UP. The FI can be generalized to the multivariate case by the FI matrix, defined as

$$\begin{aligned} \mathcal{FI}(x^n) = \\ \int_{y^n \in \mathcal{Y}^n} p(y^n | x^n) \left[\frac{\partial \log(p(y^n | x^n))}{\partial x^n} \right] \left[\frac{\partial \log(p(y^n | x^n))}{\partial x^n} \right]^T dy^n. \end{aligned} \quad (22)$$

Assuming an unbiased estimator at the attacker, i.e., the difference between the estimator's expected value and the true average value of the parameter being estimated is zero, the variance of the estimation error can be bounded via the Cramér–Rao bound as follows:

$$\mathbb{E}[\|x^n - \hat{x}^n(y^n)\|_2^2] \geq \text{Tr}(\mathcal{FI}(x^n)^{-1}), \quad (23)$$

where $\|x^n - \hat{x}^n(y^n)\|_2^2$ denotes the squared Euclidean norm and $\text{Tr}(A)$ represents the trace of the matrix A . To maximize the privacy, it is then necessary to maximize the trace of the inverse of the FI matrix. In [45], two SM settings with RBs are studied, specifically when the battery charging policy is independent of the user load and when it is dependent noncausally on the entire user load sequence. For both cases, single-letter expressions are obtained for the maximum privacy. Moreover, the case of biased estimators, wear and tear on the batteries, and peak power charging and discharging constraints are also briefly analyzed in [45].

Empirical MI as a privacy measure

Approaches aimed at determining theoretical privacy limits provide important insights and intuition for the optimal EMP to limit privacy leakage. However, they are often difficult to optimize or even evaluate numerically, and the relatively simplified formulations obtained in various special cases rely on restrictive assumptions, e.g., i.i.d. user load and infinite RB capacity. An alternative is to follow a suboptimal or heuristic EMP. Although such a policy does not provide theoretical privacy guarantees, one can evaluate the corresponding privacy leakage numerically using empirical MI.

One way to compute the empirical MI is by simulating a discrete time system for a large enough time interval and sampling

the resulting X^n and Y^n sequences [64]. The MI between two observed sequences x^n and y^n can be approximated as

$$I(X; Y) \approx -\frac{1}{n} \log p(y^n) - \frac{1}{n} \log p(x^n) + \frac{1}{n} \log p(x^n, y^n), \quad (24)$$

where $p(y^n)$, $p(x^n)$ and $p(x^n, y^n)$ are calculated recursively through a sum-product computation. When using this method, the RB is modeled as a finite state machine (FSM), and the level of energy in the RB evolves in time through a Markov chain with transition probabilities depending on the specific policy implemented. An example of an FSM is illustrated in Figure 11, where all the processes are considered to be binary and Bernoulli distributed, and the parameters are $q_x = \Pr\{X = 1\}$, $p_e = \Pr\{E = 1\}$ and p_v , the latter being the probability of using energy from the battery, provided there is available energy. The support space for the parameters is discretized, and the optimal combination of parameters is found, which minimizes the empirical MI.

This approach is followed in [65], where only a binary RB is present and for an i.i.d. Bernoulli-distributed user demand, and in [66] where an RES is also considered. The latter work also analyzes the wasted energy and characterizes the privacy-versus-energy efficiency tradeoff for the binary scenario and the equiprobable user load and renewable energy generation processes. For larger battery capacities and for an equiprobable user load, the authors note that there is a symmetry and complementarity in the optimal transition probabilities in the FSM model, which simplifies the numerical analysis. This model is also employed in [60] and [61] by considering an RES and designing a suboptimal policy, which, at each time instant, decides among using all of the available energy, half of it, or no energy at all, according to a probability chosen to minimize the overall information leakage.

Another technique for approximating MI is to assume X and Y to be i.i.d. over a time interval and approximate the MI via the relative frequency of events (X_t, Y_t) during the same time window. In [67], this approach is enriched by additive smoothing, i.e., avoiding zero probability estimates by adding a positive scalar, and it is employed together with a model-dis-

tribution predictive controller, such that, at each time slot t , the EMU chooses its actions for a prediction horizon of length T , i.e., up to time $t + T$. Privacy and cost are jointly optimized by considering noncausal knowledge of the renewable energy generation process, user load, and energy prices, while the EMU's actions, i.e., the energy that is requested from the grid and the battery, are forecast over the prediction horizon. The user and grid load processes are assumed to be i.i.d. within a time window $N \gg T$, which also includes the prediction horizon T , and the finite alphabets $\mathcal{X}$ and $\mathcal{Y}$ are considered. As $N \gg T$, first-order Taylor approximation of the logarithm function is used, and the corresponding mixed-integer quadratic program is formulated, which is of manageable size and can be solved recursively whenever new SM readings are available.

Results show that considering a relatively small prediction horizon T prevents the EMU from fully utilizing the RB capacity, as the user load that is considered by the algorithm is generally smaller than the RB capacity. Allowing a longer prediction interval dramatically improves the performance in terms of both privacy and cost, at the expense of a much higher computational complexity. The work also shows that by increasing the alphabet sizes of $\mathcal{X}$ and $\mathcal{Y}$, better privacy performance can be achieved.

Empirical MI normalized by the empirical entropy of the user load is considered in [68], where an RB is used to minimize the energy cost subject to privacy constraints. Here, two cost tariffs are considered, a low price and a high price, and a dynamic programming approach is developed to maximize the energy stored in the battery at the end of the low-price period and minimize it at the end of the high-price period. At every time slot, the optimal probability distribution of the grid load is computed, which is forced to be independent of the user load distribution.

Concluding remarks and future challenges

Privacy and the so-called right to be let alone are considered to be an individual's inalienable, fundamental rights, which are safeguarded in many national constitutions worldwide. In Europe, the General Data Protection Regulation (GDPR) [69],




FIGURE 11. An example of RB evolution modeled as a finite state machine, with $\mathcal{B} = \{0, 1, \dots, B_{\max}\}$ and $\mathcal{X} = \mathcal{E} = \mathcal{Y} = \{0, 1\}$. The (x, e, v, y) represent, for every time t , the values of the user load, the renewable energy produced, the energy taken out of the battery by the EMU, and the grid load, respectively [61].

which went into effect on 25 May 2018, sets even more stringent requirements for every technology or device that collects and processes customer data, including SMs. For these reasons, addressing the SM privacy problem is crucial for the adoption of the SG concept. In fact, considering consumers' growing privacy concerns over SMs and many other emerging technologies [70], a critical growth in the adoption of SMs and other SG technologies will take place only when consumers are given full control of their privacy and feel they have clear and honest information on how their data are being used. Only then can consumer resistance be overcome and users' trust be assured, thus paving the way to a more fertile and fair ground for new products and increased innovation in this domain.

UPs and their partners, including governments, may be too keen on collecting users' data indiscriminately and not well incentivized to develop privacy-enhancing technologies. Therefore, legislators, public commissions, consumer advocacy groups, and researchers have important roles to play in tackling the SM privacy problem and preventing SM data from being gathered haphazardly and sold to third parties without explicit user consent or even passed to government intelligence agencies for mass surveillance. The GDPR is a good example of the initiatives that are needed.

However, given that such a legal framework is still lacking and not yet fully developed globally, it becomes imperative to push forward the concept of privacy by design, according to which privacy should be designed in to new products and services rather than considered only after user complaints and regulatory impositions. This is because more options are available during the design stage as compared to the completion stage, when the product has to be modified following a privacy incident or a user complaint. Achieving privacy by design is the ultimate goal of the techniques analyzed in this article.

In this article, we have focused exclusively on techniques that adopt physical resources, such as RESs and RBs, to provide privacy to users. The main motivation and benefits of these techniques is that they do not undermine the value of the SG concept. Each of the outlined techniques has its unique advantages and disadvantages and focuses on a particular aspect of privacy. However, despite the considerable efforts put into developing SM privacy-preserving techniques, the full extent of the SM privacy problem is far from completely understood, and a unified and coherent vision for SM privacy (just as in many other domains) is still elusive.

In the context of SM privacy, UDS-based methods manipulate a physical quantity, energy, to ensure privacy for users. This entails that physical constraints, such as those related to an RB or an RES, play a crucial role in finding the optimal privacy-preserving strategy. We expect that the techniques developed for enhancing SM privacy can prove useful in other privacy-sensitive settings in which physical quantities are involved, such as gas and water meters and location privacy.

Research challenges

Various challenges must be addressed before privacy by design can become a reality in SM systems. First, a generic privacy

measure or a combination of different measures must be determined and adopted to formally quantify loss of privacy, in the same way a user's electricity bill is computed. Such a measure should be device independent and should enable the comparison of various privacy-preserving strategies. It is also necessary to understand the implications of the various privacy measures on the grid load. From this point of view, theoretical measures may be preferable because of their abstract and fundamental nature, i.e., they are independent of any assumptions about the attacker's algorithms. However, their relevance in real-world scenarios must be assessed further, and, if necessary, valid suboptimal privacy measures or algorithms should be put forward and standardized as a proxy for more rigorous privacy assurances.

Another important goal is to give consumers as much flexibility as possible in setting their desired level of privacy, trading off privacy with the cost of electricity or other services. It is also essential to allow consumers the possibility of setting different privacy requirements for different devices, as users may consider the information about the usage of a certain device as more sensitive compared to others. This may happen because certain devices are naturally more correlated to the user's activities or presence at home, such as the use of a teakettle, a microwave, or an oven, or because a user may decide to hide the usage of a certain appliance for personal reasons.

In the near future, a wider use of electric vehicles will also bring additional complications to the SM privacy problem, as mobility patterns may be inferred by analyzing the charging and discharging events. This problem can be tackled by load shifting, which is expected to play an important role in jointly optimizing electricity cost and privacy. Load shifting, as well as other privacy-preserving techniques introduced here, will be more accurate and relevant thanks to the development of reliable prediction techniques for future electricity consumption, e.g., by using machine-learning techniques. The proliferation of various energy-hungry smart devices will complicate the problem further and overburden RBs even more.

Finally, the use of shared physical resources should also be investigated in more depth, as cities are becoming more and more densely populated and consumers may want to team up to install storage devices or energy generators that are still rather costly. In cities, solar panels or mini wind turbines may be installed on the roofs of blocks of apartments, and RBs may be put in communal areas; these resources can be used jointly by all of the users in a building. Such resource-sharing models make the privacy problem even more complicated and challenging and might call for a game-theoretic formulation of the problem.

Overall, we hope that presenting this overview of the SM privacy problem and current solutions will further encourage research and development in this area, so that remaining open issues will be solved and SMs' full potential will be realized.

Acknowledgments

Giulio Giaconì gratefully acknowledges the Engineering and Physical Sciences Research Council (EPSRC) of the United Kingdom for funding his Ph.D. studies (award reference #1507704). This work was also supported in part by the

EPSRC through the project COPES (#173605884), and by the U.S. National Science Foundation under grants CNS-1702808 and ECCS-1549881.

Authors

Giulio Giaconi (g.giaconi@imperial.ac.uk) received his B.Sc. and M.Sc. degrees (honors) in communications engineering from Sapienza University of Rome, Italy, in 2011 and 2013, respectively, and his Ph.D. degree from the Department of Electrical and Electronic Engineering, Imperial College London, in 2018. He is currently a research scientist with BT Applied Research, Security Futures Practice. In 2013, he was a visiting student with Imperial College London, working on indoor localization via visible light communications. His current research interests include cybersecurity, data privacy, information and communication theory, signal processing, and machine learning. In 2014, he received the Excellent Graduate Student Award from Sapienza University of Rome.

Deniz Gündüz (d.gunduz@imperial.ac.uk) received his B.S. degree from Middle East Technical University and his M.S. and Ph.D. degrees from the New York University Polytechnic School of Engineering in 2004 and 2007, respectively. He is a reader in information theory and communications in the Department of Electrical and Electronic Engineering, Imperial College London. He is an editor of *IEEE Transactions on Communications* and *IEEE Transactions on Green Communications and Networking*. He was the recipient of the IEEE Communications Society Communication Theory Technical Committee Early Achievement Award in 2017; a starting grant of the European Research Council in 2016; the IEEE Communications Society Best Young Researcher Award for the Europe, Middle East, and Africa Region in 2014; and the Best Paper Award at the 2016 IEEE Wireless Communications and Networking Conference.

H. Vincent Poor (poor@princeton.edu) is the Michael Henry Strater University Professor of Electrical Engineering at Princeton University, New Jersey. His interests include information theory and signal processing, with applications in wireless networks, energy systems, and related fields. He is an IEEE Fellow, a member of the National Academy of Engineering and National Academy of Sciences, and a foreign member of the Chinese Academy of Sciences and the Royal Society. He received the IEEE Signal Processing Society Technical Achievement and Society Awards in 2007 and 2011, respectively. Recent recognition of his work includes the 2017 IEEE Alexander Graham Bell Medal and a D.Sc. *honoris causa* from Syracuse University, also in 2017.

References

- [1] U.S. Energy Information Administration. (2016). International energy outlook 2016. U.S. Energy Inform. Admin. Washington, D.C. Tech. Rep. DOE/EIA-0484(2016). [Online]. Available: [https://www.eia.gov/outlooks/ieo/pdf/0484\(2016\).pdf](https://www.eia.gov/outlooks/ieo/pdf/0484(2016).pdf)
- [2] BP Global. (2016). BP energy outlook, 2016 edition. BP p.l.c. London. [Online]. Available: <https://biomasspower.gov.in/document/Reports/bp-energy-outlook-2016.pdf>
- [3] ERGEG. (2011). Final guidelines of good practice on regulatory aspects of smart metering for electricity and gas. European Regulators Group for Electricity and Gas. Brussels, Belgium. Tech. Rep. E10-RMF-29-05. [Online]. Available: http://www.smartgrids-cre.fr/media/documents/ERGEG_Guidelines_of_good_practice.pdf

- [4] Department of Energy and Climate Change. (2014). Smart metering equipment technical specifications: Version 2. Dpt. Energy and Climate Change, London. [Online]. Available: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/381535/SMIP_E2E_SMETS2.pdf
- [5] Markets and Markets, "Smart meters market by type (electric, water, and gas), application (commercial, residential, and industrial), technology (automatic meter reading and advanced metering infrastructure), and by region: Global forecasts to 2022," Rep. 2477, 2017.
- [6] "Smart grid data analytics market—Global industry analysis, size, share, growth, trends and forecast 2015–2022," Transparency Market Res., Albany, NY, Tech. Rep. TMRGI 3966, 2015.
- [7] Navigant Research. (2016). Market data: Smart meters. Navigant Research, Washington, D.C. [Online]. Available: <https://www.navigantresearch.com/reports/market-data-smart-meters>
- [8] Pecan Street Inc. Dataport. [Online]. Available: <https://dataport.cloud/>
- [9] G. Hart, "Prototype nonintrusive appliance load monitor," MIT Energy Lab. Tech. Rep. and Electric Power Res. Inst., Cambridge, MA, Tech. Rep. RP 2568-2, 1985.
- [10] G. Hart, "Nonintrusive appliance load monitoring," *Proc. IEEE*, vol. 80, no. 12, pp. 1870–1891, 1992.
- [11] A. Molina-Markham, P. Shenoy, K. Fu, E. Cecchet, and D. Irwin, "Private memoirs of a smart meter," in *Proc. ACM Workshop Embedded Sensing Systems Energy Efficiency Building*, Zurich, 2010, pp. 61–66.
- [12] A. Prudenzi, "A neuron nets based procedure for identifying domestic appliances pattern-of-use from energy recordings at meter panel," in *Proc. IEEE Power Engineering Society Winter Meeting*, New York, 2002, pp. 941–946.
- [13] E. Quinn, "Privacy and the new energy infrastructure," *Social Sci. Res. Network*, 2009. doi: 10.2139/ssrn.1370731.
- [14] I. Rouf, H. Mustafa, M. Xu, W. Xu, R. Miller, and M. Gruteser, "Neighborhood watch: Security and privacy analysis of automatic meter reading systems," in *Proc. ACM Conf. Computer and Communications Security*, Raleigh, NC, 2012, pp. 462–473.
- [15] C. Cuijpers and B.-J. Koops, "Smart metering and privacy in Europe: Lessons from the Dutch case," in *European Data Protection: Coming of Age*, S. Gutwirth, R. Leenes, P. de Hert, and Y. Pouillet, Eds. Dordrecht, The Netherlands: Springer-Verlag, 2012, pp. 269–293.
- [16] J. Loeb, "Smart meters: What would it take to stop the national rollout juggernaut?" *IET Eng. Technol.*, vol. 12, no. 5, pp. 32–34, 2017.
- [17] Y. Kim, E. Ngai, and M. Srivastava, "Cooperative state estimation for preserving privacy of user behaviors in smart grid," in *Proc. IEEE Int. Conf. Smart Grid Communications*. Brussels, Belgium, 2011, pp. 178–183.
- [18] J.-M. Bohl, C. Sorge, and O. Ugu, "A privacy model for smart metering," in *Proc. IEEE Int. Conf. Communications*. Cape Town, South Africa, 2010, pp. 1–5.
- [19] M. Backes and S. Meiser, "Differentially private smart metering with battery recharging," in *Proc. Int. Workshop Data Privacy Management and Autonomous Spontaneous Security*. Egham, U.K., 2014, pp. 194–212.
- [20] L. Sankar, S. Rajagopalan, S. Mohajer, and H. V. Poor, "Smart meter privacy: A theoretical framework," *IEEE Trans. Smart Grid*, vol. 4, no. 2, pp. 837–846, 2013.
- [21] F. D. Garcia and B. Jacobs, "Privacy-friendly energy-metering via homomorphic encryption," in *Proc. Int. Conf. Security and Trust Management*. Athens, Greece, 2010, pp. 226–238.
- [22] F. Li, B. Luo, and P. Liu, "Secure and privacy-preserving information aggregation for smart grids," *Int. J. Security Netw.*, vol. 6, no. 1, pp. 28–39, 2011.
- [23] R. Petric, "A privacy-preserving concept for smart grids," in *Proc. Sicherheit in vernetzten Systemen:18. DFN Workshop*, 2010, pp. B1–B14.
- [24] C. Efthymiou and G. Kalogridis, "Smart grid privacy via anonymization of smart metering data," in *Proc. IEEE Int. Conf. Smart Grid Communications*. Gaithersburg, MD, 2010, pp. 238–243.
- [25] A. Cárdenas, S. Amin, and G. A. Schwartz, "Privacy-aware sampling for residential demand response programs," in *Proc. ACM Int. Conf. High Confidence Networked Systems*, Beijing, China, 2012.
- [26] J.-P. Zimmermann, M. Evans, J. Griggs, N. King, L. Harding, P. Roberts, and C. Evans, "Household electricity survey: A study of domestic electrical product usage," Intertek Testing and Certification Ltd., Milton Keynes, U.K., Tech. Rep. R66141, 2012.
- [27] A. S. N. U. Nambi, A. R. Lua, and R. V. Prasad, "LocED: Location-aware energy disaggregation framework," in *Proc. ACM Int. Conf. Embedded Systems Energy-Efficient Built Environments*, 2015, pp. 45–54.
- [28] D. Dheeru and E. Karra Taniskidou. (2017). UCI machine learning repository. [Online]. Available: <http://archive.ics.uci.edu/ml>
- [29] University of Sheffield, Sheffield Solar, Microgen database. [Online]. Available: <https://microgen-database.sheffield.ac.uk/>
- [30] L. Yang, X. Chen, J. Zhang, and H. V. Poor, "Cost-effective and privacy-preserving energy management for smart meters," *IEEE Trans. Smart Grid*, vol. 6, no. 1, pp. 486–495, 2015.

- [31] O. Tan, J. Gómez-Vilardebó, and D. Gündüz, "Privacy-cost trade-offs in demand-side management with storage," *IEEE Trans. Inf. Forensics Security*, vol. 12, no. 6, pp. 1458–1469, 2017.
- [32] A. Lever, D. Sanders, N. Lehmann, M. Ravishankar, M. Ashcroft, G. Strbac, M. Aunedi, F. Teng, and D. Pudjanto. (2016). Can storage help reduce the cost of a future UK electricity system? Carbon Trust and Imperial Coll. London. [Online]. Available: <https://www.carbontrust.com/media/672486/energy-storage-report.pdf>
- [33] G. Giacon, D. Gündüz, and H. V. Poor, "Optimal demand-side management for joint privacy-cost optimization with energy storage," in *Proc. IEEE Int. Conf. Smart Grid Communications*, Dresden, Germany, 2017, pp. 265–270.
- [34] J. Kelly and W. Knottenbelt, "The UK-DALE dataset, domestic appliance-level electricity demand and whole-house demand from five UK homes," *Scientific Data*, vol. 2, Mar. 2015. doi: 10.1038/sdata.2015.7.
- [35] D. P. Bertsekas, *Dynamic Programming and Optimal Control*, vol. 2, 3rd ed. Belmont, MA: Athena Scientific, 2007.
- [36] Y. Sun, L. Lampe, and V. W. S. Wong, "Smart meter privacy: Exploiting the potential of household energy storage units," *IEEE Internet Things J.*, vol. 5, no. 1, pp. 69–78, 2018.
- [37] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press, 1998.
- [38] J. Wu, J. Liu, X. S. Hu, and Y. Shi, "Privacy protection via appliance scheduling in smart homes," in *Proc. IEEE/ACM Int. Conf. Computer-Aided Design*, Austin, TX, 2016, pp. 1–6.
- [39] Z. Chen and L. Wu, "Residential appliance DR energy management with electric privacy protection by online stochastic optimization," *IEEE Trans. Smart Grid*, vol. 4, no. 4, pp. 1861–1869, 2013.
- [40] G. Kalogridis, C. Efthymiou, S. Denic, T. Lewis, and R. Cepeda, "Privacy for smart meters: Towards undetectable appliance load signatures," in *Proc. IEEE Int. Conf. Smart Grid Communications*, Gaithersburg, MD, 2010, pp. 232–237.
- [41] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York: Wiley-Interscience, 1991.
- [42] S. McLaughlin, P. McDaniel, and W. Aiello, "Protecting consumer privacy from electric load monitoring," in *Proc. ACM Conf. Computer and Communications Security*, Chicago, 2011, pp. 87–98.
- [43] J. Yao and P. Venkatasubramanian, "The privacy analysis of battery control mechanisms in demand response: Revealing state approach and rate distortion bounds," *IEEE Trans. Smart Grid*, vol. 6, no. 5, pp. 2417–2425, 2015.
- [44] 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. Computer and Communications Security*, Raleigh, NC, 2012, pp. 415–427.
- [45] F. Farokhi and H. Sandberg, "Fisher information as a measure of privacy: Preserving privacy of households with smart meters using batteries," *IEEE Trans. Smart Grid*, 2017. doi: 10.1109/TSG.2017.2667702.
- [46] Z. Li and T. J. Oechtering, "Privacy on hypothesis testing in smart grids," in *Proc. IEEE Information Theory Workshop*, Jerusalem, Israel, 2015, pp. 337–341.
- [47] K. Qian, C. Zhou, M. Allan, and Y. Yuan, "Modeling of load demand due to EV battery charging in distribution systems," *IEEE Trans. Power Syst.*, vol. 26, no. 2, pp. 802–810, 2011.
- [48] P. A. Leicester, C. I. Goodier, and P. N. Rowley, "Probabilistic analysis of solar photovoltaic self-consumption using bayesian network models," *IET Renew. Power Gen.*, vol. 10, no. 4, pp. 448–455, 2016.
- [49] W. Labeeuw and G. Deconinck, "Residential electrical load model based on mixture model clustering and Markov models," *IEEE Trans. Ind. Informat.*, vol. 9, no. 3, pp. 1561–1569, 2013.
- [50] D. Gündüz and J. Gómez-Vilardebó, "Smart meter privacy in the presence of an alternative energy source," in *Proc. IEEE Int. Conf. Communications*, Budapest, Hungary, 2013, pp. 2027–2031.
- [51] J. Gómez-Vilardebó and D. Gündüz, "Privacy of smart meter systems with an alternative energy source," in *Proc. IEEE Int. Symp. Information Theory*, Istanbul, Turkey, 2013, pp. 2572–2576.
- [52] J. Gómez-Vilardebó and D. Gündüz, "Smart meter privacy for multiple users in the presence of an alternative energy source," *IEEE Trans. Inf. Forensics Security*, vol. 10, no. 1, pp. 132–141, 2015.
- [53] S. Li, A. Khisti, and A. Mahajan, "Information-theoretic privacy for smart metering systems with a rechargeable battery," *IEEE Trans. Inf. Theory*, vol. 64, no. 5, pp. 3679–3695, 2018.
- [54] S. Li, A. Khisti, and A. Mahajan, "Structure of optimal privacy-preserving policies in smart-metered systems with a rechargeable battery," in *Proc. IEEE Int. Workshop Signal Processing Advances Wireless Communication*, 2015, pp. 375–379.
- [55] S. Li, A. Khisti, and A. Mahajan, "Privacy preserving rechargeable battery policies for smart metering systems," in *Proc. Int. Zurich Seminar Communications*, Zurich, 2016, pp. 121–124.
- [56] S. Li, A. Khisti, and A. Mahajan, "Privacy-optimal strategies for smart metering systems with a rechargeable battery," in *Proc. American Control Conf.*, Boston, 2016, pp. 2080–2085.
- [57] H. Permuter, P. Cuff, B. Van Roy, and T. Weissman, "Capacity of the trapdoor channel with feedback," *IEEE Trans. Inf. Theory*, vol. 54, no. 7, pp. 3150–3165, 2008.
- [58] M. Arrieta and I. Esnaola, "Smart meter privacy via the trapdoor channel," in *Proc. IEEE Int. Conf. Smart Grid Communications*, Dresden, Germany, 2017, pp. 277–282.
- [59] R. Ahlswede and A. Kaspi, "Optimal coding strategies for certain permuting channels," *IEEE Trans. Inf. Theory*, vol. 33, no. 3, pp. 310–314, 1987.
- [60] G. Giacon, D. Gündüz, and H. V. Poor, "Smart meter privacy with an energy harvesting device and instantaneous power constraints," in *Proc. IEEE Int. Conf. Communications*, London, 2015, pp. 7216–7221.
- [61] G. Giacon, D. Gündüz, and H. V. Poor, "Smart meter privacy with renewable energy and an energy storage device," *IEEE Trans. Inf. Forensics Security*, vol. 13, no. 1, pp. 129–142, 2018.
- [62] G. Giacon and D. Gündüz, "Smart meter privacy with renewable energy and a finite capacity battery," in *Proc. IEEE Int. Workshop Signal Processing Advances Wireless Communication*, Edinburgh, U.K., 2016, pp. 1–5.
- [63] Z. Li, T. J. Oechtering, and D. Gündüz, "Smart meter privacy based on adversarial hypothesis testing," in *Proc. IEEE Int. Symp. Information Theory*, Aachen, Germany, 2017, pp. 774–778.
- [64] D.-M. Arnold, H.-A. Loeliger, P. Vontobel, A. Kavcic, and W. Zeng, "Simulation-based computation of information rates for channels with memory," *IEEE Trans. Inf. Theory*, vol. 52, no. 8, pp. 3498–3508, 2006.
- [65] D. Varodayan and A. Khisti, "Smart meter privacy using a rechargeable battery: Minimizing the rate of information leakage," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Prague, Czech Republic, 2011, pp. 1932–1935.
- [66] 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, 2013.
- [67] J. X. Chin, T. T. D. Rubira, and G. Hug, "Privacy-protecting energy management unit through model-distribution predictive control," *IEEE Trans. Smart Grid*, vol. 8, no. 6, pp. 3084–3093, 2017.
- [68] J. Koo, X. Lin, and S. Bagchi, "Privatus: Wallet-friendly privacy protection for smart meters," in *Proc. European Symp. Research Computer Security*, Pisa, Italy, 2012, pp. 343–360.
- [69] European Parliament and Council of the European Union, "Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation)," *Official J. European Union*, pp. L 119/1–L 119/88, May 2016.
- [70] "Consumer intelligence series: Protect.me (an in-depth look at what consumers want, what worries them, and how companies can earn their trust—and their business)," PwC, White Paper, Sept. 2017.
- [71] European Environment Agency. (2015, Sept. 4). Total electricity consumption: Outlook from IEA. [Online]. Available: <https://www.eea.europa.eu/data-and-maps/indicators/total-electricity-consumption-outlook-from-iea/total-electricity-consumption-outlook-from-1>
- [72] Itron. Itron Centron. [Online]. Available: <https://www.itron.com/na/technology/product-services-catalog/products/0/7/5/centron>
- [73] Honeywell. REX2 meter. [Online]. Available: https://www.elstersolutions.com/en/product-details-na/826/en/REX2_meter
- [74] Kamstrup. Omnipower electricity meters. [Online]. Available: <https://www.kamstrup.com/en-us/products-and-solutions/smart-grid/electricity-meters>
- [75] Enel SpA. (2016, June 27). Enel presents Enel Open Meter, the new electronic meter. [Online]. Available: https://www.enel.com/content/dam/enel-com/pressrelease/porting_pressrelease/1666038-1_PDF-1.pdf
- [76] Sunverge Energy Inc. Maximize the value of your solar power. [Online]. Available: <http://www.sunverge.com/energy-management/>
- [77] Sonnen. The sonnenBatterie. [Online]. Available: <https://sonnen-batterie.com/en-us/sonnenbatterie>
- [78] Tesla. Powerwall. [Online]. Available: <https://www.tesla.com/powerwall>
- [79] LG Chem. ESS battery. [Online]. Available: <http://www.lgchem.com/global/ess/ess/product-detail-PDEC0001>
- [80] Panasonic. LJ-SK56A residential storage battery system. [Online]. Available: <http://www.panasonic.com/au/consumer/energy-solutions/residential-storage-battery-system/lj-sk56a.html>
- [81] Powervault. [Online]. Available: <http://www.powervault.co.uk/technical/technical-specifications/>
- [82] Orison. Meet Orison: The first all-in-one, self-installable home battery system. [Online]. Available: <http://orison.energy/>
- [83] SimpliPhi Power. Phi 3.5 battery. [Online]. Available: <http://simpliphipower.com/product/phi3-4-smart-tech-battery/>