

Low-Overhead Power Trace Obfuscation for Smart Meter Privacy

Daniele Jahier Pagliari, Sara Vinco, Enrico Macii, Massimo Poncino
name.surname@polito.it
Politecnico di Torino
Turin, Italy

ABSTRACT

Smart meters communicate to the utility provider fine-grain information about a user’s energy consumption, which could be used to infer the user’s habits and pose thus a critical privacy risk. State-of-the-art solutions try to obfuscate the readings of a meter either by using a large re-chargeable battery to filter the trace or by adding random noise to alter it. Both solutions, however, have significant drawbacks: large batteries are prohibitively expensive, whereas digitally added noise implies that the user entrusts the utility provider to protect his/her privacy.

This work proposes a hybrid approach in which zero-average noise is inserted in the power trace by means of a small energy storage device (battery or supercapacitor); the distinguishing feature of our approach is that this obfuscating device is indistinguishable from any other load and therefore it complicates by construction the load disaggregation task performed by the provider or by a malicious third party. Simulation results show that our device can achieve comparable or superior privacy enhancement as that of a solution based on a large battery and therefore with smaller cost.

CCS CONCEPTS

• **Hardware** → **Energy generation and storage**; **Energy metering**; • **Security and privacy** → *Domain-specific security and privacy architectures*.

KEYWORDS

Smart Meter Privacy, Energy Storage, Load Disaggregation

1 INTRODUCTION

Smart meters are advanced devices measuring energy consumption with much more detail than a conventional meter: they can transmit to the utility information for monitoring or billing purposes and even communicate with a number of smart appliances to apply energy saving policies. The frequency of the readings of a smart meter can be at the second scale, thus raising possible privacy issues concerning user’s behavior [4, 7, 10, 17]. There is a rich literature of algorithms extracting detailed information on domestic appliances usage by analyzing and disaggregating household power traces; this operation, usually termed Non-Intrusive Load Monitoring (NILM), does not require the per-appliance installation of smart plugs, thus being totally transparent to the user [8].

Methods to hijack NILM algorithms for enhancing user privacy fall into two main categories. A first strategy uses an *energy storage device* (ESD) to filter out variations in the power trace; at the extreme, this allows to generate a completely flat profile equal to the average power consumption of the household [6, 11]. However,

regardless of the strategy adopted, this solution needs at least a 1 kWh ESD and a small inverter, thus implying a cost easily exceeding 1,000\$ [11]. Consequently, it is practical only in case of smart residential grids that already include an ESD.

A second solution is to *inject an additive digital noise* with zero average into the power trace [3, 4, 23]. This solution has a much lower cost than using an ESD. However, noise addition is performed within the meter: the utility company de facto knows its distribution, and could be able to filter it out. Thus, this method guarantees data privacy against third parties monitoring the traffic from meter to cloud, but not against the utility company itself.

In this work, we propose a hybrid approach in which *a small ESD (battery or supercapacitor)* is used to insert noise in the power trace. The ESD and its control circuitry (called noise generator, NG, hereafter) work as a special appliance that, at different times, either contributes to the energy consumption or provides energy. The proposed NG has several benefits with respect to existing strategies:

- We apply noise randomly, not necessarily in correspondence of switching events. As an effect, our method requires a smaller ESD with respect to filtering-based schemes, with benefits in total cost, volume and safety. Moreover, it does not require complex control policies to monitor the total power demand by other appliances.
- Our solution adds noise “as a load” to the actual power trace seen by the meter, thus not exposing the original power profile to the utility company, as done by techniques adding digital noise in the meter. Moreover, this paradigm makes the solution independent of the type of meter and of the NILM algorithm.

Experimental results show that our method obtains privacy levels comparable with previous approaches but using an ESD with smaller energy and current driving capability. This allows to use a supercapacitor instead of a battery, with the benefit of practically nullifying depreciation costs due to battery aging.

2 BACKGROUND AND MOTIVATION

2.1 Background on NILM Algorithms

NILM aims at estimating the power demand of individual appliances from an aggregate power trace gathered by a single meter [8].

NILM algorithms are based on detecting the appliance operations status (e.g., ON and OFF) from power measurements, and can be classified as event-based or state-based. *Event-based approaches* focus on state transitions generated by appliances: they detect changes in the aggregate load to identify the beginning/end of an event and the corresponding state change [2]. *State-based approaches* represent appliances as state machines, whose transitions are associated with a probability distribution, based on their usage patterns [9].

A key problem in NILM is the extraction of *appliance signatures*, which allow to identify appliance operations from the aggregated load. As an example, Figure 1 shows an aggregated load, where appliance signatures (colored lines) allowed to identify the activation of three appliances (colored boxes). Signatures are derived from load profiles and from appliance-specific information through a learning phase. In the subsequent inference phase, appliance states and power consumption are estimated from meter readings. Some NILM techniques use traces relative to individual appliances during training (*supervised NILM*), while others are agnostic of the appliances present in the aggregated profile (*unsupervised NILM*). For an exhaustive survey on NILM, the reader is referred to [5].

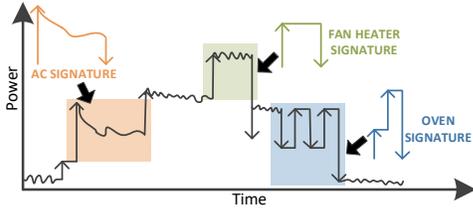


Figure 1: Example of appliance signatures identification in an aggregated load profile. Arrows represent state changes.

2.2 Related Work

Protection of smart meter data is typically achieved by either *filtering* the power consumption trace through an ESD, or *obfuscating* the trace by adding noise.

2.2.1 Filtering the Power Trace. Several works protect smart meter data privacy using a re-chargeable battery to filter the load profile, thus masking appliance features [11, 18, 23]. Ideally, a battery could be used to make the load profile perfectly flat over a billing period, with a value equivalent to the average power consumption [11]. However, this is generically not practical, and filtering policies are made less aggressive by activating battery charges/discharges only in correspondence of changes in the load profile, which are typically associated to an appliance activation [18].

This solution has cost as its main drawback; in order to achieve a good level of masking, and given typical household load fluctuations, battery packs with a significant energy are required. Previous works propose values ranging from 0.5kWh to 12kWh [11], corresponding to bulky battery packs and a significant cost for the user, e.g., around 1,200\$ for a typical 1.2kWh (12V, 100Ah) deep-cycle Li-Ion battery [22]. Importantly, this cost does not include the depreciation of the battery, which will become unusable and should be replaced after a number of charge/discharge cycles [11]. The use of batteries to “filter” the load profile is therefore practical only if (i) the house hosts a smart residential grid with an ESD and renewable resources [6], or (ii) its size can be reduced significantly.

2.2.2 Obfuscating the Power Trace. Obfuscation techniques operate blindly with respect to appliances and transition events, and try to add spurious signals to the power trace. A simple yet effective approach is presented in [3], where a zero-average random number in an interval of $[-X, +X]$ is added digitally to the meter reading, with X tuned by the user. The zero average guarantees that the utility provider can obtain a good approximation of the real total consumption by summing all received data.

The drawback of this approach is that the noise addition is implemented inside the meters; as such, the characteristics of the noise signal are decided by the utility provider, which may be able to remove the inserted noise upon receiving the data, for example by modeling its distribution in the NILM algorithm, or even by simply removing the exact value added by the meter, if noise is generated based on a deterministic pseudo-random sequence. So, even assuming that users are given the possibility of selecting some parameters of the noise insertion (e.g. X), this approach still implies *full trust* of the user on the energy provider effectively adding the noise as agreed, in a truly random way. In that, this approach is not substantially different from the simple encryption of meter data, paired with a data usage agreement between user and provider (e.g., the provider stating that user data will not be used for NILM applications). Conversely, our approach is based on directly generating a “scrambled” power trace whose alteration is controlled *only by the user* by randomly charging/discharging an ESD.

3 POWER TRACE OBFUSCATION METHOD

Figure 2 describes the basic concept of the proposed method. A custom device (*noise generator*, NG) is attached to the AC-bus downstream of the meter, as a regular load. The NG consists of an ESD (battery or supercapacitor), a bi-directional DC/AC converter (to allow flow of current in and out of the ESD), and a control circuitry, which decides whether the NG behaves as a load (consuming power) or as a generator (inserting power into the AC bus).

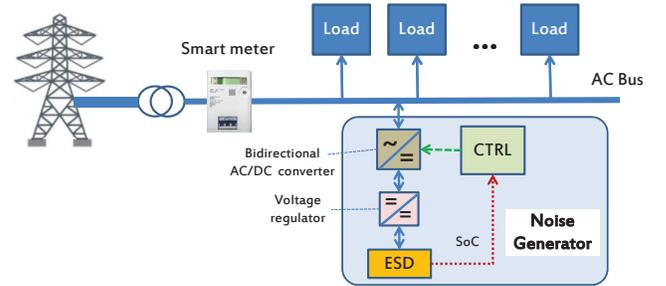


Figure 2: Conceptual scheme of the proposed method.

In this system, the “noise” is a *result of the randomized policy implemented by the controller (CTRL)*: regardless of the events on the AC bus, the NG adds or draws power according to some distribution which guarantees that the sum of the added and drawn power over a billing period is approximately zero. A perfect zero balance obfuscation is not required, as the maximum allowed error for active energy reading for residential consumers is $\pm 1\text{-}2\%$ [3, 21]. Note that our NG is agnostic of the load profile. Unlike other filtering approaches, we do not try to smooth the power trace in correspondence of the activation of an appliance: the controller operates autonomously by generating a random power demand or injection at random times.

3.1 System Operations

The NG operation is governed by the controller, which monitors and affects the operation of the ESD and of the bidirectional converter through status and control signals. At a given time point t , the controller generates a pair of random values (P_t, T_t) : P_t is the power

value to be added or drawn, and T_t is the time for which this power value is applied. Therefore, the controller repeatedly generates a request for an energy amount of $E_t = P_t \cdot T_t$, as conceptually shown in Figure 3. P_t is added to the total power demand requested by all other appliances in the household $P_{load,t}$, so that the meter “sees” a total power $P_{meter,t} = P_{load,t} + P_t$.

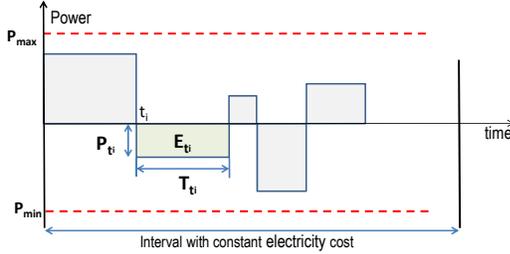


Figure 3: Basic operations of the Noise Generator.

3.1.1 Generation of P_t . The controller first generates a P_t power value according to a user-defined distribution. The extremes P_{min} and P_{max} of this distribution i.e., the maximum provided or absorbed power, are determined based on the operating limits of the ESD. The only requirement on the distribution is that $P_{min} < 0$ and $P_{max} > 0$ to guarantee that power can be both drawn and generated. However, there is no need that the power distribution has zero average or even that it is symmetric, since what should be approximately zero sum in every billing window is the *energy balance*, i.e. the sum of products between P_t and T_t .

P_{min} and P_{max} are derived as follows. First, the maximum charge/discharge current of the ESD $I'_{max,c}$, $I'_{max,d}$ are obtained from the datasheet and are converted to power values $P'_{max,c}$, $P'_{max,d}$ by multiplying them by the nominal voltage. Even if this does not hold in general, in this calculation we consider the nominal voltage to be constant, since the output of the ESD will be stabilized by a DC/DC converter before being connected to the AC bus.

These ESD power limits are then converted to output AC power limits according to the combined efficiency η of the bi-directional AC/DC converter and of the voltage regulators:

$$P_{min} = -\frac{P'_{max,c}}{\eta}, P_{max} = \eta \cdot P'_{max,d}$$

In general, the efficiency of a converter is a function of its input and output voltage and other operating conditions. By using a stabilized voltage at the input of the DC/AC converter, however, we can safely assume its efficiency as constant. For what concerns the regulator, instead, more accurate models can be used, as the ESD voltage can be variable (especially in case of a supercapacitor).

3.1.2 Generation of T_t . Once P_t is determined, the controller proceeds to the generation of a random time interval T_t . In general, T_t can be generated according to a different random distribution with respect to P_t , hence increasing the NG configurability and the effectiveness of load obfuscation.

The extremes T_{min} and T_{max} of the distribution of T_t determine a trade-off between the aging of the ESD and the effectiveness of the obfuscation for high-frequency events. Their values are constrained by several factors. Firstly, as the generated T s determine the length of each noise “pulse” inserted by the NG, their bounds

affect what *kind of load variation “events”* can be masked by our proposed method, with shorter pulses being able to alter higher frequency variations. Clearly, if the characteristics of the smart meter are known, T_{min} should not be set to a value smaller than the sampling frequency of the meter. T s are also constrained by the *characteristics of the ESD*: the selected pulse durations affect the depth-of-discharge (DoD) of the ESD, which for some devices like batteries can significantly impact their aging. Finally, the most stringent constraint comes from the observation that T should not be too large to *prevent over-charging and over-discharging of the ESD*. To this end, the controller must keep track of the State-of-Charge (SoC) of the ESD at all times (feedback from the ESD to CTRL in Figure 2). The SoC is then combined with the current $I_{ESD,t}$ absorbed/provided by the device in correspondence of a generated power value P_t , to determine the time of full-charge/discharge. This explains why P_t must be generated first: its value is necessary to determine the range of feasible T s.

$I_{ESD,t}$ is simply determined from P_t as $I_{ESD,t} = P_t / V_{ESD}$, where P_t is the power absorbed/provided by the ESD, and is related to P_t through the converter and regulator efficiency, as explained above. Then, the time of full-charge $T_{max,c}$ (when $P_t < 0$) and of full-discharge $T_{max,d}$ (when $P_t > 0$) are:

$$T_{max,c} = \frac{E(1 - SoC)}{|I_{ESD,t}|}, T_{max,d} = \frac{E \cdot SoC}{|I_{ESD,t}|}$$

where E denotes the nominal energy (capacity) of the ESD (e.g., in Ampere-hours for a battery and in Farad-Volt for a supercapacitor). Notice that there can be conditions (e.g., ESD almost completely charged or discharged) when $T_{max,c}$ or $T_{max,d}$ are smaller than the user-imposed T_{min} , i.e., full charge/discharge is reached with a pulse length smaller than the minimum allowed. In these cases, the controller simply discards the candidate P_t and generates a new one, until a feasible value is obtained.

3.1.3 Enforcing Zero Energy Balance. One important requirement of this solution is that the NG does not alter the total energy cost; the controller must then guarantee a zero-mean energy balance in each billing window, which typically lasts 4-12 hours for a residential plan. The lengths of billing windows are known to the users, as subscribers of a contract, and can thus can be configured in the system before deploying it.

To enforce a balance between the total energy absorbed and produced by the NG, we use a “lazy” compensation at the end of each window. The controller keeps track of:

- the energy balance from the start time t_0 of the current billing window $E_t = \sum_{\tau=t_0}^t P_\tau T_\tau$
- the *maximum energy* that can be provided by the ESD before the end of the window, computed as the product of the minimum (if $E_t < 0$) or maximum (if $E_t > 0$) power and the time remaining before the end of the window t_f , i.e.:

$$E_{off} = \begin{cases} P_{max}(t_f - t), & \text{if } E_t < 0 \\ P_{min}(t_f - t), & \text{if } E_t > 0 \end{cases}$$

While tracking E_t , when E_t and E_{off} become of the same magnitude, the controller overrides the random generation and absorbs/provides P_{min}/P_{max} for the rest of the billing window.

While this introduces some determinism in the operations, if the P_t distribution has zero-mean as well, the condition $E_t \approx E_{off}$ will occur only towards the end of period and the $t_f - t$ will thus be short (a few seconds) compared to the length of a billing window. Therefore, the obfuscation capabilities of the NG will not be degraded significantly. The study of more advanced balancing policies will be subject of future work.

Notice that E_{off} is guaranteed to be available in the ESD because, as a result of the zero average charge/discharge pattern, the ESD will tend to have a SoC equal to its initial value plus/minus the energy imbalance accumulated during the last billing period.

3.2 Design Issues

3.2.1 Choice of ESD. The proposed method is not restricted to a specific type of ESD (battery or supercapacitor). This is because, as shown later, an effective Noise Generator can be designed including an ESD with small energy capacity; previous methods that required capacity > 1 kWh would require hundreds of large supercaps.

Table 4 summarizes the pros and cons of the two types of ESD. In general, the cheapest solution depends on many parameters (mainly on $T_{t,min}$ and $T_{t,max}$), and on the relative cost of supercap and battery (replacing the batteries many times still costs less than a single supercap). A more quantitative analysis is carried on in Section 4.4.

	Cost	Energy Efficiency	Aging
Supercap	• Higher non-recurrent cost (100Wh cost >500 \$)	• Output voltage highly variable: need a DC/DC regulator at the output • Efficiency of the regulator can be low especially when voltage is low	• Virtually unlimited, $>10^6$ full cycles • Can assume NO aging cost
Battery	• Relatively low-cost (100Wh cost <50 \$)	• Output voltage relatively stable: does not need a DC/DC regulator at the output • Efficiency of the regulator can be kept high by appropriately choosing the voltage of the battery pack	• Cycle life is limited (500-1000) equivalent full cycles • A replacement cost should be included

Figure 4: Figures of merit for the two ESD options.

3.2.2 Bidirectional AC/DC Converter. The design of bidirectional AC/DC converters presents several challenges, e.g., the need of power factor correction (PFC), low distortion currents, high-quality DC output voltage [15, 16]. Most implementations in literature target large DC voltages typical of a big battery pack (e.g., for electric vehicles). In our case, the low currents and voltage levels of the ESD actually simplify the design of the device. In this work, as the design of such bidirectional AC/DC converter is out of the scope of our research, we consider this device as an off-the-shelf component. Notice that the design and cost issues relative to this component are totally ignored in previous works [11, 18, 23]. As we will show in Section 4, however, its cost can be non-negligible in the overall assessment.

3.2.3 Controller. The task of the controller can be implemented either using an embedded microcontroller [1] or entirely in HW using an embedded FPGA [16]. Whatever the implementation, the controller will consume some power, that should be taken into account during the generation of P_t . In practice, however, modern microcontrollers and FPGAs consume fractions of a Watt [19, 20]. Thus, the discussion relative to the generation of P_t is still valid, given the different power magnitudes between the ESD and the controller.

4 SIMULATION RESULTS

4.1 Models

For the privacy assessment of the proposed solution, modeling the non-idealities of the ESD and of the converters (in particular their inefficiency in transferring energy) has limited importance. As a matter of fact, these inefficiencies will simply imply that a given current drawn from (or inserted into) the AC bus will result into (i) a larger current into/from the ESD and (ii) some losses in the conversion process. For this reason, when evaluating the obfuscation effectiveness of our method, we do not use specific models for the ESD and the converters, and we use a single efficiency factor η , which lumps all the inefficiencies into a single quantity, as discussed in Section 3. These inefficiencies have rather impact on the aging of the ESD, as will be discussed in Section 4.4

4.2 Benchmarks and Metrics

4.2.1 Reference NILM and Adopted Dataset. Our reference NILM algorithm is the NeuralNILM approach proposed in [13], which is based on neural network architectures and proved to achieve state-of-the-art accuracy scores. NeuralNILM proposed two main architectures: the *denoising autoencoders approach* tries to recover a “clean” trace of power demand of an appliance from the aggregate power demand, while the *rectangles approach* identifies start time, end time and average power demand of each appliance activation. To allow a fair comparison, we select the same dataset used by NeuralNILM, i.e. UK-DALE [14], which is open-access and contains power consumption traces from five houses, sampled every 6 seconds. In our experiments, NeuralNILM models are trained in the same way as [13]. Disaggregation metrics are evaluated on four appliances (washing machine, microwave, dish washer and fridge) and on the UK-DALE “house 2” power traces.

4.2.2 Evaluation Metrics. We use the following metrics to estimate the privacy generated by our power trace obfuscation approach, and its impact on NILM algorithms:

- *Mean Absolute Error* (MAE) between the estimated power consumption and the actual power consumption;
- *F1-Score*: a common metric of classification accuracy, calculated as the harmonic average between (i) the proportion of positive results that are true positives (*precision*) and (ii) the proportion of true positives that are correctly identified (*recall*) [5];
- *Relative Error in Total Energy* (RETE), i.e., between the total predicted energy and the total actual energy.
- *Pearson’s correlation*, used to express privacy as the correlation between the obfuscated consumption profile and the real consumption profile [3]. A correlation value close to 0 has a high privacy level, while 1 has a low privacy level;
- *Relative entropy*, estimates how much information the adversary acquires from the observed (obfuscated) consumption profile [12]. This metric requires the estimation of the Probability Mass Function (PMF) of power traces, discretized into a set of histogram bins. To compare our method with [12], we computed PMFs using 200 bins.

The first three metrics assess the quality of a NILM algorithm, whereas the last two are NILM-independent privacy measures.

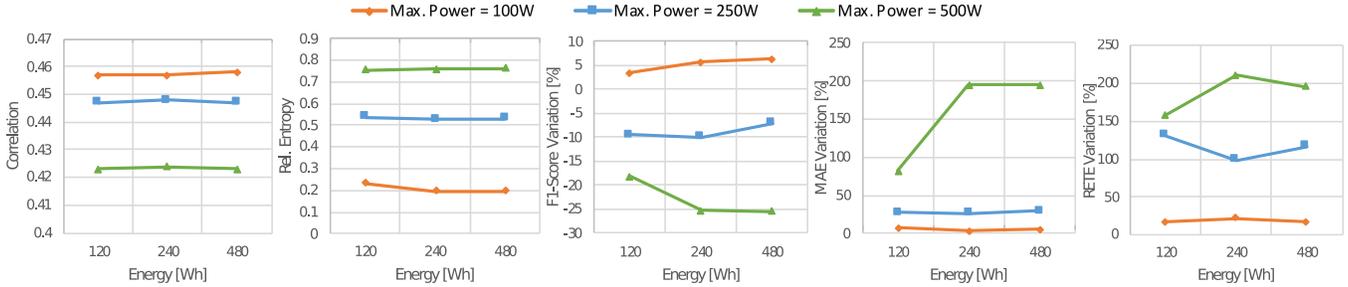


Figure 5: Impact of ESD parameters on different privacy and NILM “quality” metrics.

4.3 Parameter Space Exploration

In this first set of experiments, we explore the impact of two of the main ESD parameters on the effectiveness of the power trace obfuscation. Herein, we report generic results, valid for both batteries and supercapacitors since the goal is only to obtain an estimate of the ESD parameters required to obtain sufficient privacy.

We fix T_{min} and T_{max} of the noise generation algorithm to 5s and 100s respectively. These values are selected reasonably, but they are not explored. In fact, we assume that the NILM algorithm used to disaggregate power traces is unknown to the user. Thus, tuning these parameters to “disturb” a specific NILM method would be a form of over-fitting. In contrast, the minimum pulse width is selected to be comparable with the sampling period of the meter, whereas the maximum is chosen to be large enough to partially mask some relevant appliance events (e.g., the activation of the microwave, or of the centrifuge in a washing machine). We draw both P_t and T_t values from uniform distributions.

We then simulate different noise insertions on the UK-DALE data, exploring the impact of the minimum and maximum power values (P_{min} and P_{max}) added by the NG on the AC bus, and of the energy E of the ESD. We assume $P_{min} = -P_{max}$, and we report the results for three different values of energy (120Wh, 240Wh, 480Wh) and maximum power (100W, 250W and 500W).

4.3.1 Comparison with state-of-the-art privacy mechanisms. The leftmost two plots of Figure 5 report the results of this parameter exploration on NILM-independent privacy metrics, i.e. correlation and relative entropy. As expected, when the maximum power absorbed/provided by the ESD (i.e. the “height” of the power pulses) increases, the correlation among noisy and noiseless power traces decreases, being a measure of *similarity*. In contrast, the relative entropy increases, being a measure of the *difference* between traces. Interestingly, both metrics remain approximately constant for different ESD capacities at a given power level. This is explained by the fact that our method only inserts short power pulses of small absolute power. Indeed, even a pulse of maximum length ($T_{max} = 100$ s) and height ($P_{max} = 500$ W) discharges the smallest energy ESD (120Wh) of about 12%. Thus, the benefit of having a larger energy capacity is only being able to apply a longer sequence of power pulses with the same sign (i.e., charging or discharging the ESD multiple times), and has a limited effect on privacy.

These results can be used to compare our method with previous privacy mechanisms for smart meters. In terms of correlation, our method outperforms the one in [3], based on inserting noise digitally in the meter readings. Indeed, they reported a correlation of

0.489, which is larger than the one obtained with our method, even for the smallest power value of 100W. In terms of relative entropy, when considering an ESD with maximum output power of 250W, we obtain comparable results to those reported in [12], i.e., relative entropy ≈ 0.5 for the same power level (configuration “B1” in their paper). However, since the algorithm in [12] is based on filtering peaks in the power load, they must inevitably use a battery with large energy capacity (500Wh). In contrast, as shown by Figure 5, our noise generation approach obtains the same level of privacy with a 4x smaller energy (120Wh).

4.3.2 Effect on NILM disaggregation metrics. The rightmost three plots in Figure 5 report the disaggregation quality metrics obtained with NeuralNILM [13] for the considered power and energy combinations. The plots report the *percentage variation* of different scores with respect to the baseline obtained by NeuralNILM on the original UK-DALE traces. The scores reported are the average over the two considered NILM methods (*autoencoders* and *rectangles*) and over all considered appliances.

As expected, all metrics worsen when the power injected/absorbed by the ESD increases (the F1-score reduces, while errors increase). Interestingly, with an ESD power of 100W, the F1-score *improves* with respect to the baseline. While we do not have a certain explanation for why this happens, the small MAE variations (2-5% larger than the baseline) also show that this power level is not sufficient to significantly affect an advanced NILM algorithm. Conversely, when maximum power is increased to 250W, the average F1-score worsens of 7-10%, while the MAE increases of 27-30%, and the RETE explodes to +100-120%. Even larger variations are obtained with a maximum power of 500W. As for NILM-independent metrics, also in this case there is not a clear relation between ESD energy and scores. In general, the NG appears to obtain similar obfuscation results regardless of energy, except for the case of $P_{max} = 500$ W for the MAE metric.

In general, the choice of the appropriate ESD parameters depends on the desired level of privacy. However, given these results, a peak power of 250W and an energy of 120Wh can be considered sufficient, since they induce significant worsening in all considered metrics, even for an advanced disaggregation approach such as NeuralNILM. Unfortunately, a comparison with state-of-the-art methods for smart meter privacy is not possible in this case, as none of the previous works evaluated the effect of their methods on the performance of a real NILM algorithm.

4.4 Cost Analysis

In this section we analyze the overall operation cost for a specific design point of our NG, considering two possible choices of the ESD: a re-chargeable Li-Ion battery and a super-capacitor.

Based on the analysis of Section 4.3 we choose a maximum power of 250W as it guarantees a good privacy level. Concerning energy (capacity), the analysis of Section 4.3 would lead us choose the smallest point (120 Wh) as privacy metrics are basically independent of the ESD energy. However, different energies translate into a different impact of currents to/from the ESD, and this affects aging in case of a battery; therefore, we analyze all three energy values for the battery: B1 (120Wh), B2 (240Wh), and B3 (480Wh), while for the supercapacitor (SC) we stick to the 120Wh case.

The total privacy protection cost C_{priv} over a time T is obtained as $C_{priv}^T = C_{hw} + C_{loss}^T + C_{depr}^T$, where C_{hw} is the cost of the NG hardware components (ESD + converters + controller), C_{loss}^T is the energy cost due to battery and converter inefficiencies, and C_{depr}^T is the depreciation cost of the ESD in an interval $[0, T]$. For a battery, the latter depends on the number of full charge/discharge cycles completed from 0 to T and can be computed as the *effective number of cycles* $N_{cyc,eff}^T = \int_0^T |I(t)|dt/2Q_{nom}$, where Q_{nom} is the nominal charge in Ah of the ESD. The depreciation cost will therefore be $C_{depr}^T = C_{ESD} \cdot N_{cyc,eff}^T / N_{cyc,max}$, where $N_{cyc,max}$ is the maximum expected number of cycles (typically given in datasheets). Table 1 report the cost comparison among the battery and supercapacitor ESD configurations. We assume a cost of the electronics of 20\$, identical in all cases. Considering an overall *average* efficiency of the power conversion chain of 95%, we obtain 55kWh/year of power losses, which assuming 0.05\$/kWh yield a C_{loss}^T of about 3\$/year.

Table 1: Cost comparison ($N_{cyc,max} = 500$).

ESD	$N_{cyc,eff}^{1yr}$	C_{ESD}	C_{depr}^{1yr}	C_{depr}^{3yrs}	C_{priv}^{1yr}	C_{priv}^{3yrs}
B1	4960	25\$	248\$	744\$	296\$	795\$
B2	2483	50\$	248\$	744\$	321\$	824\$
B3	1241	100\$	248\$	744\$	371\$	873\$
SC	-	400\$	0\$	0\$	423\$	429\$

For the ESD cost, we assume battery packs built using 5Ah 18650 Li-Ion batteries (≈ 4 \$ each) and arranged in series (s)/parallel (p) to reach a nominal voltage of 24V. This is obtained with 7s for B1 (≈ 25 \$), 7s2p for B2 (≈ 50 \$), 7s3p for B3 (≈ 75 \$). For the supercap, we consider 3000F, 2.7V components ($\approx 20 - 30$ \$ each), arranged in a 9s2p pack, for a total of ≈ 400 \$).

We can notice how, as a result of the frequent charges/discharges on the battery, the number of effective cycles is quite large ($\approx 5,000$ in one year operation for the 120Wh case). Assuming a typical value of $N_{cyc,max} = 500$ cycles (the one used in the table), this implies replacing the whole pack 10 times a year. With these values, the large initial investment of for the SC becomes convenient after about 18 months of operations.

5 CONCLUSIONS

We have proposed a novel method to protect users privacy against load disaggregation from smart meter data. Thanks to an in-house noise-generating device, we solve the main privacy issue of previous

methods based on noise generation, while requiring an ESD with smaller energy compared to standard load-filtering approaches. This renders a supercapacitor-based implementation affordable, thus practically eliminating recurring costs. In future work, we plan to evaluate the costs of our approach using more advanced ESD models, including secondary effects of batteries (e.g. rated capacity) variable regulator/converter efficiencies, etc.

REFERENCES

- [1] M. P. Akter et al. Modified model predictive control of a bidirectional AC/DC converter based on lyapunov function for energy storage systems. *IEEE Trans. on Industrial Electronics*, 63(2):704--715, Feb 2016.
- [2] K. D. Anderson et al. Event detection for non intrusive load monitoring. In *Proc. of IEEE Industrial Electronics Society*, pp. 3312--3317, 2012.
- [3] P. Barbosa et al. Lightweight privacy for smart metering data by adding noise. In *Proc. of the ACM SIGAPP*, pp. 531--538, 2014.
- [4] J. Bohli, C. Sorge, and O. Ugus. A privacy model for smart metering. In *Proc. of IEEE ICC Workshops*, pp. 1--5, 2010.
- [5] A. Faustine et al. A survey on non-intrusive load monitoring methodologies and techniques for energy disaggregation problem. arxiv.org/abs/1703.00785, 2017.
- [6] G. Giacon, D. Gandaz, and H. V. Poor. Smart meter privacy with renewable energy and an energy storage device. *IEEE Trans. on Information Forensics and Security*, 13(1):129--142, 2018.
- [7] G. W. Hart. Residential energy monitoring and computerized surveillance via utility power flows. *IEEE Technology and Society Magazine*, 8(2):12--16, 1989.
- [8] G. W. Hart. Nonintrusive appliance load monitoring. *Proc. of the IEEE*, 80(12):1870--1891, 1992.
- [9] K. He et al. Non-intrusive load disaggregation using graph signal processing. *IEEE Trans. on Smart Grid*, 9(3):1739--1747, 2018.
- [10] A. Humayed et al. Cyber-physical systems security - a survey. *IEEE Internet of Things Journal*, 4(6):1802--1831, 2017.
- [11] G. Kalogridis et al. Privacy for smart meters: Towards undetectable appliance load signatures. In *Proc. of IEEE SmartGridComm*, pp. 232--237, 2010.
- [12] G. Kalogridis, Z. Fan, and S. Basutkar. Affordable privacy for home smart meters. In *Proc. of IEEE ISPA Workshops*, pp. 77--84, 2011.
- [13] J. Kelly and W. J. Knottenbelt. Neural NILM: deep neural networks applied to energy disaggregation. *Proc. of ACM BuildSys*, pages 55--64, 2015.
- [14] J. Kelly and W. J. Knottenbelt. The UK-DALE dataset. *Scientific Data*, 2, 2015.
- [15] H. Kim et al. High-efficiency isolated bidirectional AC/DC converter for a dc distribution system. *IEEE Trans. on Power Electronics*, 28(4):1642--1654, 2013.
- [16] Y. Liao. A novel reduced switching loss bidirectional AC/DC converter pwm strategy with feedforward control for grid-tied microgrid systems. *IEEE Trans. on Power Electronics*, 29(3):1500--1513, 2014.
- [17] P. McDaniel and S. McLaughlin. Security and privacy challenges in the smart grid. *IEEE Security Privacy*, 7(3):75--77, 2009.
- [18] S. McLaughlin, P. McDaniel, and W. Aiello. Protecting consumer privacy from electric load monitoring. In *Proc. of the ACM CCS*, pp. 87--98, 2011.
- [19] D. Jahier Pagliari et al. All-digital embedded meters for on-line power estimation. In *Proc. DATE*, pp. 737-742, 2018
- [20] Y. Chen et al. Battery-aware Design Exploration of Scheduling Policies for Multi-sensor Devices In *Proc. ACM GLSVLSI*, pp. 201-206, 2018
- [21] [Online] www.measurement.gov.au/Publications/PAREquirements/Documents/MI%2520M%25206-1.doc
- [22] [Online] www.lithiumion-batteries.com/products/product/12v-100ah-lithium-ion-battery.php
- [23] S. Wang et al. A randomized response model for privacy preserving smart metering. *IEEE Trans. on Smart Grid*, 3(3):1317--1324, 2012.