When Privacy Protection Meets Non-Intrusive Load Monitoring: Trade-off Analysis and Privacy Schemes via Residential Energy Storage

Sangyoung Park

Smart Mobility Systems
Technische Universität Berlin
Einstein Center Digital Future
Berlin, Germany
Email: sangyoung.park@tu-berlin

Andrea Cominola

Smart Water Networks
Technische Universität Berlin
Einstein Center Digital Future
Berlin, Germany
Email: andrea.cominola@tu-berlin.de

***Abstract*—Non-Intrusive Load Monitoring (NILM) algorithms are actively being researched to disaggregate the electricity usage of a whole household into the contribution of individual appliances. While understanding the usage patterns of individual appliances can be beneficial for flattening the peak demand, reducing the electricity bill, and improving the energy usage efficiency, NILM algorithms raise privacy concerns. Residential energy storage could be used to relieve such concerns by modifying the monitored electricity profile. However, residential energy storage systems are yet costly, and hence assessing the financial overhead of privacy protection techniques is important. In this paper, we provide motivational examples and early results on how much residential energy storage would be required to fool a state-of-the-art NILM algorithm. Our preliminary results on the trade-off between NILM accuracy and privacy protection indicate that some intuitive approaches that require a significant amount of battery capacity are not necessarily the most effective in reducing the disaggregation accuracy.**

***Keywords*–Privacy; Non-Intrusive Load Monitoring; smart meter; energy storage.**

I. Introduction

Smart meters are becoming an essential component in smart grids because they allow high-resolution, real-time monitoring of electricity generation/consumption and communication of the pricing information [1]. Smart meter information is key for utility companies to provide information for the design and implementation of demand side management strategies and better cope with fluctuating electricity demands [2]. This results in more effective matching of the electricity demand and generation. In addition, smart meter information is expected to also bring benefits to electricity consumers, as it provides them with more transparent billing information and enables a better control over their electricity usage [3].

Motivated by the need to provide demand management programs with better understanding, models, and forecasts of electricity demand with high spatial and temporal resolution, there have been recent rapid advances in the development of Non-Intrusive Load Monitoring (NILM) algorithms. NILM algorithms allow automatic decomposition of the aggregated electric load measured by a smart meter at the household level into the electricity usage patterns of individual appliances (e.g., fridge, air conditioning system, etc.) [4]. One digital AC monitor, i.e., a smart meter, measuring the single-phase power into a household instead of individual sensors for each appliance, suffices to provide a low cost and non-intrusive solution. The information on electricity demand obtained via NILM algorithms can be used to support consumption-based feedback programs and customer segmentation for demand management, flatten the peak demand, identify faulty appliances, and provide hints on the effectiveness of demand side management programs [5].

However, NILM algorithms inherently generate privacy concerns. Besides detailed information on electricity usage behaviors, sensitive private information such as how many people are present in a home at a given time, absence of a resident, or even gender and age information could be potentially estimated using NILM algorithms [6][7]. What is even more of a concern is that, from the user side, it would be impossible to detect whether or not a NILM algorithm is being used by the utility provider or a third-party unless explicitly communicated. Resolving the existing trade-off between the expected benefits of coupled smart metering-NILM systems and the privacy challenges that they include is thus key to facilitate the development of privacy-aware smart meter deployments.

A potentially effective way of alleviating the privacy concerns is to physically modify the electricity profile seen by the smart meters [8]. This can be achieved by the use of a residential energy storage. Existing commercial products, such as Tesla Powerwall [9] or Encharge from Enphase [10], are primarily designed for the purpose of compensating the fluctuating power generation of the rooftop solar arrays. However, their marketability is still being carefully evaluated as it is highly dependent on the battery purchase cost, depreciation cost due to aging as well as the local electricity cost. Hence, using a residential energy storage for privacy protection could harm the return-on-investment unless it is properly sized and utilized.

In this paper, we investigate the cost-effectiveness of the potential usage of a residential energy storage for privacy protection against NILM algorithms. We first perturb the electricity signal seen by the single-point smart meter of a sample household to simulate different privacy protection scenarios. These scenarios include the superimposition of simple Gaussian noise on the smart meter signal, or a flattened signal due to the usage of residential energy storage (i.e., batteries) of different size and cost. We then investigate the trade-off between privacy protection and NILM capabilities by analyzing how the performance of a state-of-the-art NILM algorithm changes in scenarios determined by increasing privacy protection. The contributions of this paper can be summarized as follows.

- We analyze the accuracy and effectiveness of the NILM algorithm over modified household consumption profiles using residential energy storage (i.e., batteries).
- We analyze the trade-off between different privacy protection schemes and the battery size needed to implement them.
- We provide recommendations on how to effectively use residential energy storage to protect the privacy of electricity users against NILM algorithms.

The ultimate goal of this paper is not to condemn the detailed investigation of household energy demand at the end use level, which can be very important to provide demand management programs, or also coordinated water-energy conservation/peak shifting programs. Rather, we aim to foster the conversation about the privacy implications of NILM and to set the basis for a wider discussion about the conflicting trade-off of demand management programs in terms of technical feasibility, overall economic and environmental benefits and social acceptance.

The rest of the paper is organized as follows. Section II summarizes the related works on the usage of an energy storage in the residential sector and NILM. Section III describes the system setup and experimental settings for the NILM algorithm used in this study. Section V analyzes some early experimental results. Section VI concludes the paper.

II. Related Works

The development of smart energy grids and the uptake of smart metering devices have fostered many recent research efforts in two potentially conflicting directions, i.e., (i) electricity demand modeling at the end use level, and (ii) privacy protection.

On the one hand, the increased availability of smart meter data at the household level and recorded with sub-daily sampling resolution has fostered the development of NILM algorithms for electricity demand modelling at the end use level, following the 1992 seminal study by Hart [11]. NILM algorithms thus aim at estimating the electricity consumption (or consumption pattern) of each appliance contributing to the aggregate electricity profile recorded by a single-point smart meter, installed at the household level. Given the aggregate household electricity consumption Y_t caused by N appliances and recorded by the single-point smart meter at time t, NILM algorithms estimate the non-metered consumption y_t^i of each individual appliance i, where

$$Y_t = \sum_{i=1}^{N} y_t^i + e_t \tag{1}$$

and e_t is the measurement noise. This process is non-intrusive, as the installation of multiple sensors at the appliance level is avoided. In the last two decades, several NILM algorithms have been proposed in the literature: *optimization-based* methods (e.g., integer or sparse optimization [12]); *pattern recognition* methods, which model the temporal structure of electricity signals (e.g., methods based on Markov Models [13]); *hybrid* methods [14]; and *unsupervised* algorithms [15]. More recent approaches primarily exploited deep neural networks (e.g., [4]) and explored the potential of transfer learning to generalize and transfer NILM algorithms among different domains [16]. The above methods have been extensively tested and cross-compared on benchmark data sets [17][18] and with widely adopted performance metrics [19].

On the other hand, the state-of-the-art literature includes approaches for *privacy protection* that use a residential energy storage to modify the usage pattern of appliances and human activities [8][20]. These studies often target flattening out, i.e., water-filling, the household electricity profile for privacy protection. However, this usually results in an extensive use of the storage requiring a large capacity and leads to accelerated aging, thus rendering the solution impractical. An alternative approach to privacy protection consists of randomizing the household electricity profile by adding, for instance, Gaussian noise [21]. However, the lack of a precise definition of *privacy* has limited so far the possibility to come up with a cost effective and general solution. Most previous works made use of the concept of signal “flatness” to quantify the privacy level of a modified electricity signal, and formulate it by proxy indicators, including the sum of Root Mean Square (RMS) error, the mutual information [22], or the entropy [23] between the original and the modified profile. However, such metrics are only weakly linked to the people’s perception of privacy as it is hard to grasp the feeling of how much private information can be extracted by just flattening the electricity profile.

In this paper, we aim neither at introducing a comprehensive definition of privacy in the residential electricity sector, nor at proposing new NILM algorithms. Rather, here, we explore the influence of different privacy schemes on the accuracy of state-of-the-art NILM algorithms. The different privacy schemes are generated by perturbation of the household electricity signal as caused by operation of a residential energy storage device. Thus, we measure privacy protection as the ability of a residential energy storage to alter the household electricity signal seen by the smart meter and hamper accurate identification of end use electricity consumption via NILM algorithms. Such an approach entails that trade-offs between the privacy scheme defined by the usage (and thus size and cost) of the residential energy storage and the accuracy may emerge. The ultimate goal of this paper is, thus, to analyze this trade-off and come up with recommendations to foster the identification of cost-effective solutions for privacy protection in the residential electricity sector, as well as contributing to the overall discussion on the benefits and challenges of the advanced metering and analytics tools characterizing the ongoing digitalization of the utility and house sectors.

III. Material and Methods

A. Residential Electricity System

Figure 1 shows the system setup of the sample household considered in this study. A smart meter monitors the whole household energy consumption and communicates the value to the utility provider. The utility provider makes use of the smart meter information for demand side management. The utility provider could run NILM algorithms with the consent of the residents or there could be a third-party acting on its own initiative, potentially malicious, running NILM algorithms. In the household, there are a number of electricity appliances, such as washing machines, TVs, ovens, heat pumps, HVAC, dryers, etc. Optionally, the house includes a battery storage at home connected to the low-voltage AC network through a bidirectional DC/AC converter. The storage resembles products like Powerwall from Tesla (13.5 kWh) [9] or Encharge from Enphase (3.5 kWh and 10.5 kWh) [10]. There could be multiple purposes of installing the storage such as leveraging the electricity price difference over time, stabilizing the local low distribution grid by participating as a primary control reserve, etc. In this work, we exploit the presence of a battery storage for its additional function of privacy protection, as it can be used to modify the electricity signal seen by the smart meter, thus hiding the actual electricity usage of individual appliances.

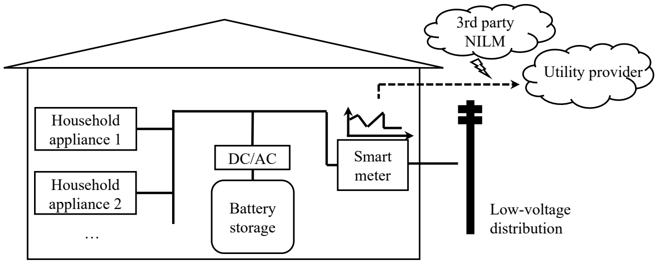

Figure 1. Electricity system setup for the sample household.

B. Sparse NILM algorithm

In this paper, we rely on state-of-the-art NILM algorithms to explore the influence of different privacy schemes via residential energy storage on NILM capabilities.

Specifically, we adapt the open source NILM algorithm proposed by [13], a computationally efficient algorithm for online real-time NILM based on a super-state Hidden Markov Model (HMM) and a modified Viterbi algorithm that efficiently manages the sparsity in HMM matrices (the original code used in [13] is available on GitHub at [24]). In short, the algorithm follows the following sequential process. First, the data sub-metered at the level of individual appliance are analyzed to create probability mass functions for each appliance and identify its load states. Second, a super-state HMM is created by combining the individual load states identified in the previous step. This builds the actual NILM model. Finally, the super-state HMM model is combined with a sparse Viterbi algorithm to perform NILM and disaggregate the observed aggregate smart meter load into the estimated contribution of each appliance. The Viterbi algorithm is a dynamic programming algorithm that is usually used to estimate the most likely sequence of hidden states associated with a measured output of a process modelled with HMM. The modified Viterbi algorithm proposed in [13] exploits the high rate of zero-probability terms in the HMM matrices, thus avoiding unnecessary computations and efficiently returning a solution. The resulting sparsity-based NILM algorithm has been demonstrated to outperform other state-of-the-art NILM algorithms. Moreover, it can handle disaggregation with data sampled with low-frequency (e.g., 1 min), while many other algorithms require higher frequencies, and is scalable, i.e., can handle a large number of super-states [13].

We assess the NILM performance of the sparsity-based NILM algorithm under different privacy-protection scenarios according to two performance metrics, i.e., (i) the *Finite-State F-score* (FS-fscore) and (ii) the *Mean Absolute Percentage Error* (MAPE). The FS-fscore was first introduced in [19] as an alternative to conventional F-score to account for inaccuracies in non-binary classification and is formulated as follows:

$$FS_i = \frac{2 \times PC_i \times RC_i}{PC_i + RC_i} \tag{2}$$

where RC_i and PC_i are the *recall* and *precision*, respectively, formulated for each appliance i in order to take into account the classification inaccuracy $RC_i = \frac{TP_i - inacc_i}{TP_i + FN_i}$ and $PC_i = \frac{TP_i - inacc_i}{TP_i + FP_i}$. The $inacc$ term can be defined, for each appliance i, as:

$$inacc = \sum_{t=1}^{H} \frac{\left|\hat{x}_t^i - x_t^i\right|}{K^i} \tag{3}$$

where $\hat{x}_t^i$ and x_t^i are the estimated and observed states of appliance i at time t, H is the considered time horizon, K^i is the number of states of appliance i, TP_i, FP_i, and FN_i are the number of true positive, false positive, and false negative events, respectively. Overall, the FS-fscore metric evaluates how good a NILM algorithm is in classifying the operating states of the considered appliances. MAPE is formulated as follows, for appliance i:

$$MAPE = \frac{1}{H} \sum_{t=1}^{H} \left| \frac{y_t^i - \hat{y}_t^i}{y_t^i} \right| \tag{4}$$

where the notation is consistent with the variables previously defined in this paper. After calculating the FS-fscore and MAPE performance metrics for each appliance, we compute their average value across appliances to assess how affected the NILM performance is on average, for each privacy protection scheme.

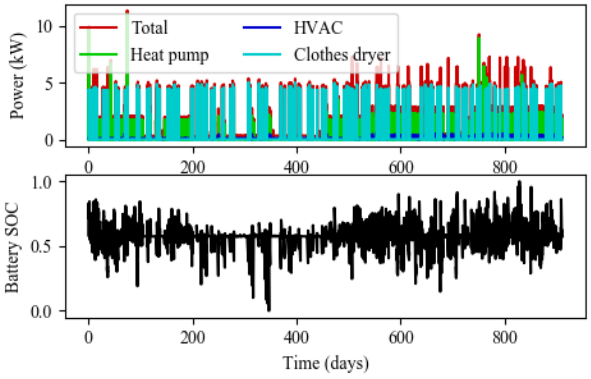

Figure 2. (Upper) Residential electricity profile example with usage of individual appliances shown [17], (Lower) Battery State-Of-Charge (SOC) for day-wise water-filling of a battery capacity of 15.04 kWh.

C. AMPds Dataset

To simulate the electricity consumption of the different appliances included in the sample household of Figure 1 and train the sparse NILM algorithm, we used the appliance-level data included in the AMPds dataset [17]. The AMPds dataset contains electricity, water, and natural gas measurements at one minute intervals for 2 years of monitoring. The data was collected at a house in Canada, where a family of three live. In this study, we considered the original version of the dataset, which includes data for 19 appliances monitored over a period of 1 year [25]. Electricity was measured with 1 Hz frequency, and 21 loads including a furnace, a fridge, a heat pump, a clothes dryer, TVs, a wall oven, etc., which makes it suitable for developing and evaluating NILM algorithms. In the trade-off analysis presented in this work, we considered an increasing number of appliances in the AMPds - from 3 to 10 - selected based on their ascending total electricity usage over the 1-year time series.

IV. Privacy Protection Schemes

One way to protect privacy is to completely flatten out the usage profile like water-filling strategy. However, water-filling strategy could be prohibitive due to the high usage of battery. For example, to completely filter out the whole house energy consumption reported in the AMPds dataset shown in Figure 2 (top) over a year, a battery of usable capacity 15.04 kWh has to be used. Assuming a battery system cost of 700 USD/kWh including the pack and the electronic equipment, this would result in more than 10,000 USD for purchasing the battery system. Assuming a lifetime of 10 years, which is equivalent to the warranty period of Tesla Powerwall, this would roughly equate to 1,000 USD/year solely for protecting privacy. There are other factors that add to cost, such as sub-optimal exploitation of Time-Of-Use (TOU) tariffs and accelerated aging due to repeated cycle charge/discharge. Such factors, however, will be left as future works, and we focus on the required capacity in this work. Next, we present various methods to modify the original electricity consumption profile to protect the privacy against the Sparse NILM algorithm presented in Section III-B.

M1: Add Gaussian noise to the whole duration of the profile. The method is intended to make the profile hard to analyze regardless of the time of use and appliance usage patterns. We control the intensity of the noise to see the impact. Larger noise leads to reduced accuracy of the NILM algorithms. However, it also involves more usage of the battery capacity. The upper graph in Figure 3 shows such an example. We have introduced a Gaussian noise with a standard deviation of 3 A. The mean of the noise over the long term is of course zero, but temporary increase and decrease in cumulative energy from the battery is inevitable, as can be seen from Figure 3. We observe that a naive implementation of M1 requires an unnecessarily large battery capacity, so we periodically reset the SOC of the battery to a pre-defined level. The usage pattern can still be modified while efficiently making use of the battery capacity.

M2: Add Gaussian noise only when a particular appliance of interest is used. For M2, we add noise only to time slots where a particular appliance is being used. We test this scheme in order to find out whether we can selectively hide the usage profile of an appliance, which will be using less battery capacity compared with M1.

M3: Water-filling for a particular appliance of interest. For M3, we see whether hiding the usage pattern by flattening out the appliance profile while it is being used is effective in reducing the accuracy of the NILM algorithm. Figure 4 shows how it modifies the profile. In the modified profile, the load looks rectangular rather than preserving all the "shapes" of the consumption. Because the load is averaged, the peak becomes lower than the original profile.

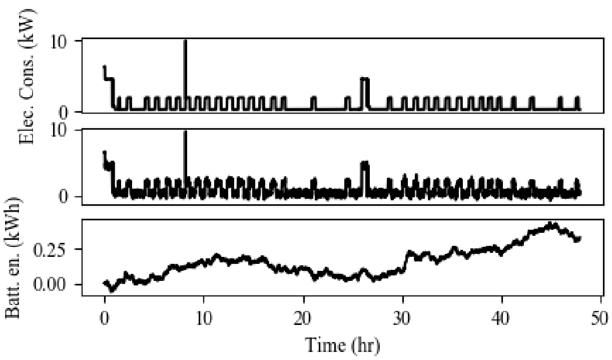

Figure 3. Original (upper) and modified (middle) whole household electricity consumption using privacy protection scheme M1, and the corresponding battery energy level (lower) over two days.

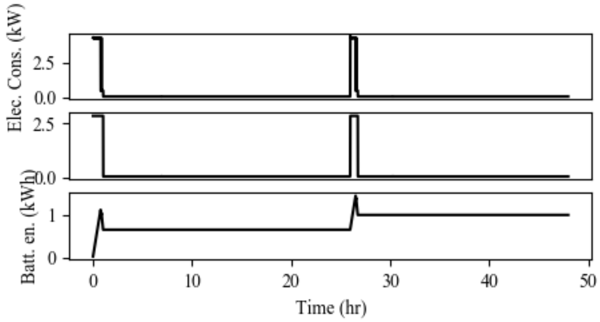

Figure 4. Original (upper) and modified (middle) profile of the clothes dryer using privacy protection scheme M3, and corresponding battery profile (lower) over two days.

M4: Spread-out the electricity consumption of a particular application. What can be seen from Figure 4 is that the appliance, which draws significant current, is still very visible to the human eye in the overall profile. Hence, we also attempt to modify the peak value with the support of a battery in varying degrees. Therefore, we explore M4, which reduces the size of the peak and spreads out the electricity consumption evenly among all time slots.

M5: Erase an appliance's consumption. This is similar to M4, but completely erases an appliance's usage profile from the whole household energy consumption by charging the battery storage. An equivalent amount of energy is discharged over the whole period of time just like M4.

M6: Day-wise water-filling for the whole electricity consumption profile. Finally, M6 performs a water-filling technique, which results in a flat profile seen by the smart meter. However, doing this for a whole year, i.e., flat profile for a year, is prohibitive because of the seasonal variations in electricity consumption, which mandates the use of an excessively large capacity battery, i.e., tens of kWh of battery capacity over the course of two years just for introducing the Gaussian noise of 6 A standard deviation. Therefore, we investigate day-wise water-filling, which means the smart meter will see a flat profile within a day, but varying electricity profile among days.

We evaluate the reduction in disaggregation accuracy and the required battery capacity, and hence the cost, for each method to provide an insight into developing cost-effective techniques.

V. Experimental Results

In this section, we present the FS-fscore and MAPE obtained for NILM under the different privacy protection schemes detailed in the previous section. The overall evaluation procedure is shown in Figure 5. For the sake of analysis, we have preprocessed the AMPds dataset and prepared 8 sets of the Whole Household Energy (WHE) consumption profile. Each set reconstructs the WHE profile using the top 3, 4, $\cdots$, 10 appliances, sorted by increasing contribution to the whole household energy consumption. Generally speaking, the disaggregation accuracy is expected to decrease when there are more appliances to disaggregate. Hence, we varied the number of appliances to assess the effectiveness of the NILM algorithm in all cases. We build the NILM models from the preprocessed WHE profiles. Then, the profile modification algorithms from Section IV are applied to simulate different privacy protection schemes. The modified profiles serve as inputs to the sparse Viterbi algorithm, which estimates the most likely sequence of hidden states (i.e., appliance states). Finally, the NILM performance is assessed by FS-fscore and the MAPE metrics.

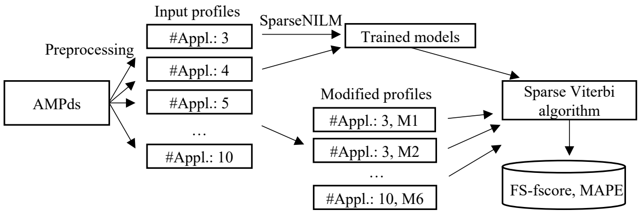

Figure 5. Overall evaluation procedure of NILM under different privacy protection schemes.

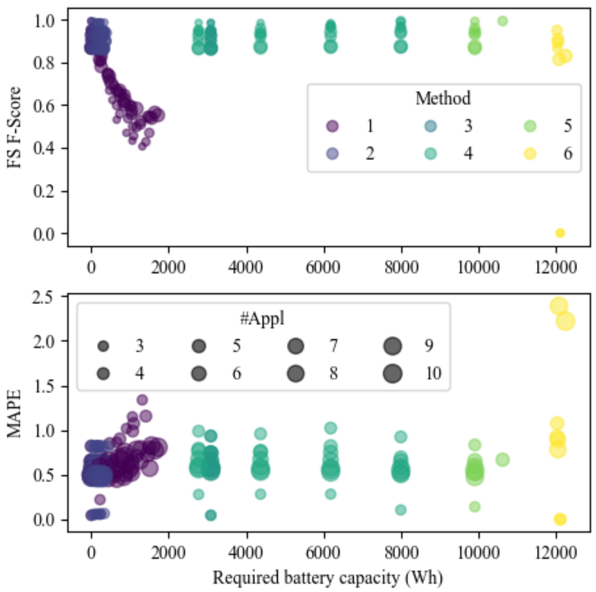

Figure 6. Mean FS-fscore (top) and MAPE (bottom) of the sparsity-based NILM algorithm across all appliances, privacy protection schemes, and their required battery capacity.

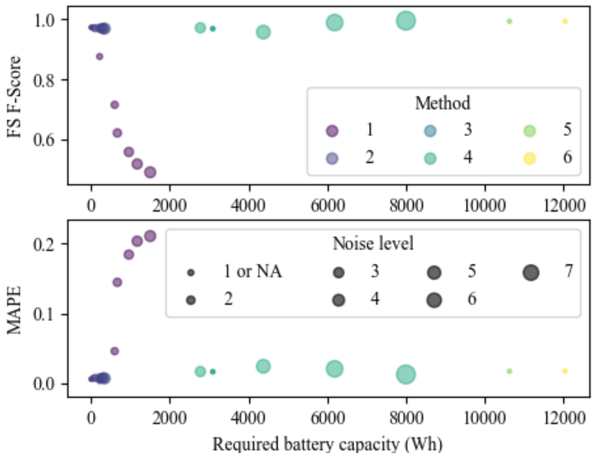

Figure 7. FS-fscore (top) and MAPE (bottom) of the cloths dryer (appliance of interest) vs required battery capacity of the profile modification methods. The number of appliances is 5 for all the results. Marker color varies for each profile modification method (M1-M6). Marker size is proportional to the noise level of each method (the higher, the more noisy).

Figure 6 shows the FS-fscore and MAPE averaged across all appliances versus the required battery capacity of the modification algorithm. The required battery capacity varies from 0 kWh to 12 kWh depending on the six modification methods. M6 requires the most amount of battery capacity as it aims at a completely flat profile. The sizes of markers correspond to the number of appliances and the colors of markers correspond to the privacy protection scheme. It can be seen that, overall, a larger number of appliances leads to slightly less accurate disaggregation within each method. The privacy protection schemes M1 and M2 require the least amount of battery capacity as Gaussian noise is relatively small and averages out over a long period of time, as can be seen from Figure 2. The most noticeable point is that both FS-fscore and MAPE do not dramatically change for methods M3 to M6.This might be due to the fact that the trained NILM algorithm can still provide accurate estimates for the status of most of those appliances that operate in a very limited range, or that fall within the range seen by the smart meter, and thus the average performance is still good in terms of FS-fscore. However, the majority of MAPE values are between 0.5 and 1.2, indicating that the different privacy protection algorithms can successfully hamper the identification of detailed appliance profiles. Overall, the two most cost-effective methods to protect privacy against the considered sparsity-based NILM algorithm appear to be, in this specific sample case, M1 and M2. The more battery capacity is utilized for M1 and M2, the worse FS-fscore and MAPE become. The FS-fscore drops almost to 0.4 when a battery capacity of around 2 kWh is used for obfuscating the original profile. Other methods require much more battery capacity, but they do not necessarily result in a larger drop in NILM accuracy, as far as the FS-fscore is concerned.

Figure 7 shows the FS-fscore and MAPE for a particular appliance of interest, in this case, the clothes dryer, instead of the mean of the metrics across all appliances. For the sake of illustration, only the results obtained for the NILM experiments including 5 appliances are reported. The marker size corresponds to the magnitude of profile modifications, in each privacy protection scheme. For M1 and M2, it is the standard deviation of the Gaussian noise. For M4, it is the amount of change in the mean of the appliance. For those methods not requiring different profile modification magnitude, the marker sizes are depicted as the smallest. In this graph, it is clearly visible that the disaggregation accuracy is lower for larger Gaussian noise with M1. Surprisingly, M5 and M6, which completely hide an appliance's profile, do not have a profound impact on the metrics of this appliance. This might be due to the dominance of the clothes dryer profile over other appliances. When Gaussian noise is imposed on the aggregate signal, this specific profile becomes harder to identify, but, in the other cases, a guess from the sparsity-based NILM algorithm for this specific appliance is still accurate. Another noticeable thing is that M2, which is dedicated to modifying the appliance's record, has nearly no impact on the metrics, while M1 degrades the performance of the sparsity-based NILM algorithm the most. Such an observation motivates further systematic research on the topic, including a sensitivity analysis on how these results change for different NILM algorithms and with different performance metrics.

VI. Conclusion and Future Work

In this paper, we investigated the effectiveness of a number of heuristic algorithms making use of a residential battery storage in preserving privacy against a NILM algorithm. Most of the prior works are based on the water-filling technique, i.e., flattening out the consumption profile as much as possible. However, they require a significant amount of battery storage while the amount of privacy protection is hard to quantify. Therefore, we focus on the accuracy metrics of a NILM algorithm and quantify the effectiveness of the profile modification methods. Our preliminary results indicate that some intuitive methods do not necessarily lead to a significant reduction in the disaggregation performance of the NILM algorithm, even though they require significant battery capacity. This points towards further systematic research tailored to providing protection against NILM algorithms while minimizing the battery cost overhead.

Also, the cost analysis in this paper is restricted to battery capacity. A holistic cost analysis incorporating the electricity bills under flexible pricing policy, the interplay with the rooftop solar arrays (when available), and the impacts of different battery usage schemes on battery aging rates should be performed. Finally, future research should look at better linking the battery capacity to the specific appliance (or appliances) that a customer would like to hide for privacy protection, if some are more privacy-sensitive than others.

Acknowledgment

The authors acknowledge the support of the Einstein Center Digital Future (Berlin, Germany).

References

[1] F. Tounquet and C. Alaton, "Benchmarking smart metering deployment in the EU-28," European Commission, Final Report, 2020.

[2] A. Mohsenian-Rad, V. W. S. Wong, J. Jatskevich, R. Schober, and A. Leon-Garcia, "Autonomous demand-side management based on game-theoretic energy consumption scheduling for the future smart grid," *IEEE Transactions on Smart Grid*, vol. 1, no. 3, pp. 320–331, 2010.

[3] S. S. S. R. Depuru, L. Wang, V. Devabhaktuni, and N. Gudi, "Smart meters for power grid — challenges, issues, advantages and status," in *IEEE/PES Power Systems Conference and Exposition*, 2011, pp. 1–7.

[4] J. Kelly and W. Knottenbelt, "Neural nilm: Deep neural networks applied to energy disaggregation," in *ACM International Conference on Embedded Systems for Energy-Efficient Built Environments*, 2015.

[5] H. Cao, C. Beckel, and T. Staake, "Are domestic load profiles stable over time? an attempt to identify target households for demand side management campaigns," in *IEEE Industrial Electronics Society Conference (IECON)*, 2013, pp. 4733–4738.

[6] A. Molina-Markham, P. Shenoy, K. Fu, E. Cecchet, and D. Irwin, "Private memoirs of a smart meter," in *ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building (BuildSys)*, 2010.

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

[8] L. Yang, X. Chen, J. Zhang, and H. V. Poor, "Cost-effective and privacy-preserving energy management for smart meters," *IEEE Transactions on Smart Grid*, vol. 6, no. 1, pp. 486–495, 2015.

[9] "Powerwall — Tesla," 2020, URL: https://www.tesla.com/powerwall [accessed: 2020-07-09].

[10] "Install Solar Backup Battery - Ensemble^{TM} Technology — Enphase," 2020, URL: https://enphase.com/en-us/ensemble-technology-enphase-installers [accessed: 2020-07-09].

[11] G. W. Hart, "Nonintrusive appliance load monitoring," *Proceedings of the IEEE*, vol. 80, no. 12, pp. 1870–1891, 1992.

[12] D. Piga, A. Cominola, M. Giuliani, A. Castelletti, and A. E. Rizzoli, "Sparse optimization for automated energy end use disaggregation," *IEEE Transactions on Control Systems Technology*, vol. 24, no. 3, pp. 1044–1051, 2015.

[13] S. Makonin, F. Popowich, I. V. Bajić, B. Gill, and L. Bartram, "Exploiting HMM sparsity to perform online real-time nonintrusive load monitoring," *IEEE Transactions on Smart Grid*, vol. 7, no. 6, pp. 2575–2585, 2016.

[14] A. Cominola, M. Giuliani, D. Piga, A. Castelletti, and A. E. Rizzoli, "A hybrid signature-based iterative disaggregation algorithm for non-intrusive load monitoring," *Applied Energy*, vol. 185, pp. 331–344, 2017.

[15] R. Bonfigli, S. Squartini, M. Fagiani, and F. Piazza, "Unsupervised algorithms for non-intrusive load monitoring: An up-to-date overview," in *2015 IEEE 15th International Conference on Environment and Electrical Engineering (EEEIC)*. IEEE, 2015, pp. 1175–1180.

[16] M. D'Incecco, S. Squartini, and M. Zhong, "Transfer learning for non-intrusive load monitoring," *IEEE Transactions on Smart Grid*, vol. 11, no. 2, pp. 1419–1429, 2019.

[17] S. Makonin, B. Ellert, I. V. Bajic, and F. Popowich, "Electricity, water, and natural gas consumption of a residential house in Canada from 2012 to 2014," *Scientific Data*, vol. 3, no. 160037, pp. 1–12, 2016.

[18] J. Z. Kolter and M. J. Johnson, "Redd: A public data set for energy disaggregation research," in *Workshop on data mining applications in sustainability (SIGKDD), San Diego, CA*, vol. 25, no. Citeseer, 2011, pp. 59–62.

[19] S. Makonin and F. Popowich, "Nonintrusive load monitoring (NILM) performance evaluation," *Energy Efficiency*, vol. 8, no. 4, pp. 809–814, 2015.

[20] G. Kalogridis, Z. Fan, and S. Basutkar, "Affordable privacy for home smart meters," in *IEEE International Symposium on Parallel and Distributed Processing with Applications Workshops*, 2011, pp. 77–84.

[21] E. Liu and P. Cheng, "Achieving privacy protection using distributed load scheduling: A randomized approach," *IEEE Transactions on Smart Grid*, vol. 8, no. 5, pp. 2460–2473, 2017.

[22] A. Pröbstl, S. Park, S. Steinhorst, and S. Chakraborty, "Cost/privacy co-optimization in smart energy grids," in *Design, Automation Test in Europe Conference Exhibition (DATE)*, 2019, pp. 872–877.

[23] S. McLaughlin, P. McDaniel, and W. Aiello, "Protecting consumer privacy from electric load monitoring," in *Conference on Computer and Communications Security (CCS)*, 2011, pp. 87–98.

[24] S. Makonin, *The super-state hidden Markov model disaggregator that uses a sparse Viterbi algorithm for decoding.*, 2018 (accessed August 17, 2020). [Online]. Available: https://github.com/smakonin/SparseNILM

[25] S. Makonin, F. Popowich, L. Bartram, B. Gill, and I. V. Bajić, "Ampds: A public dataset for load disaggregation and eco-feedback research," in *2013 IEEE Electrical Power & Energy Conference*. IEEE, 2013, pp. 1–6.