Power Signature Obfuscation using Flexible Building Loads

Kyri Baker and Kaitlyn Garifi
College of Engineering and Applied Science
University of Colorado, Boulder
Email: {kyri.baker, kaitlyn.garifi}@colorado.edu

***Abstract*—This work provides a indication of the potential of using intelligent load shifting to increase the privacy of building occupants. Unwanted interception of smart meter communications and subsequent load disaggregation by third parties can potentially reveal private information about building occupants such as occupant demographics and occupancy/appliance schedules. Shifting loads and exploiting flexibility in appliance usage in a way that minimally impacts occupants can obfuscate revealing information in the aggregate power signature of a building. Unlike security measures such as encryption, the techniques proposed here use a building's inherent resources, removing the value from disaggregating the signal in the first place.**

I. Introduction

The emergence of "connected" devices in the power grid; that is, electronic devices which have wireless communication capabilities, is rapidly evolving. The network in which these devices communicate and through which information is readily available and easily accessible is sometimes deemed the "Internet of Things" (IoT) due to the interconnectedness and ability to monitor a variety of devices and appliances that is comparable to the internet. The wireless communication capabilities of many smart electricity meters, for example, allow for the ease of information exchange between buildings and utility companies on granular timescales [1]. This has enabled automated reading of electricity meters for billing, whereas previously the meters had to be read manually to charge consumers for energy usage. However, the increase in communication capabilities between buildings and the grid brings an increase in security and privacy concerns that is recognized in both industry, government, and even in the general public [2]–[4]. For example, if a building is performing a routine service of communicating their demand profile to the utility company and this data is intercepted, this interception could reveal private information regarding the occupancy of the building and behavior of the occupants. Because of the continuing advancement of load disaggregation technologies, these demand profiles can be classified into individual appliance categories, uncovering detailed behaviors about occupants such as what times the occupants are in the building and what appliances are in use [4]–[6].

Load disaggregation, also known as non-intrusive load monitoring (NILM), is a field of research that analyzes the aggregate power consumption of a building and disaggregates the signal into individual appliance signatures without requiring invasive and expensive monitoring of individual devices or circuit breakers [7], [8]. In contrast to monitoring each individual appliance, NILM allows for a less invasive, more cost effective solution to determining the usage patterns of building loads by simply using existing electricity meter technology. By using the information from the disaggregated load profiles, building owners can determine the breakdown of their energy consumption and determine how to use energy more efficiently [9]. Load disaggregation may also provide valuable diagnostics for appliances by monitoring changes in their typical power consumption and alerting a homeowner when something goes wrong with an appliance. While the benefits of disaggregating load certainly exist, many risks also exist. First and foremost, hackers can also use these services and technologies to their advantage; an interception of these load patterns can reveal private information about individuals, such as their typical occupancy patterns [5], and which appliances they are using at what times [6], [10].

Masking the energy patterns in households due to the emergence of smart meters and NILM technologies has been analyzed previously [11]–[14]. In [11]–[14], battery storage is strategically charged and discharged to balance out and flatten the load profile. However, using energy storage to mask recoverable information with NILM algorithms has the potential degrade the battery due to constant charging and discharging [15]. Instead, it may be beneficial to look at load obfuscation by using a variety of building loads. In this paper, we attempt to decrease the NILM accuracy as much as possible by shifting existing loads such as electric heating/cooling resources, and by considering new resources such as electric vehicles. We use the non-intrusive load monitoring toolkit (NILM-TK) [16] with homes from the REDD dataset [17] to test the impact of flexible loads on the rate of accuracy of appliance identification. Specifically, we examine two test cases: 1) when the home has an electric vehicle of varying energy capacity that can obfuscate the true appliance signatures; and, 2) when the electric furnace control is optimized, rather than controlled via typical hysteresis ("bang-bang") control. The results show an overall decrease in NILM accuracy, due to the load profile being flattened from the EV charging, and from the implementation of an unpredictable/atypical furnace power pattern.

II. Problem Definition

It was assumed in the following results that a supervised training period of 10 days was performed using the top 5-energy consuming appliances for each considered home. In

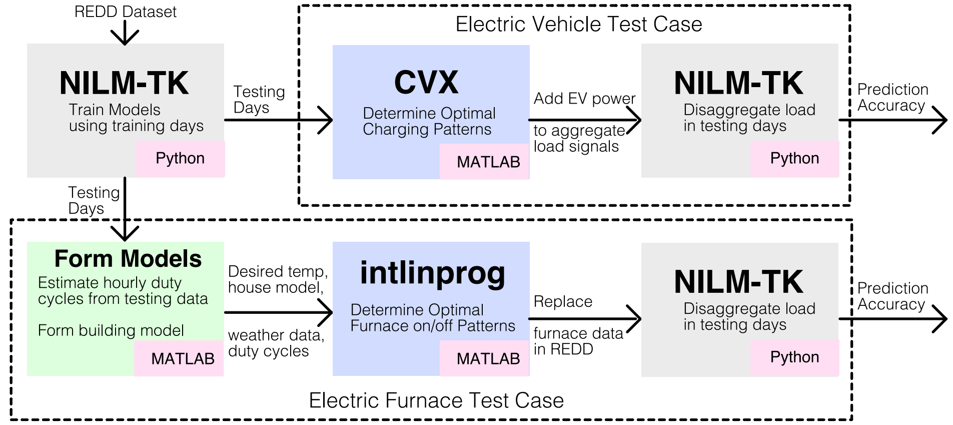
REDD Dataset
NILM-TK
Train Models using training days
Python
Testing Days
Electric Vehicle Test Case
CVX
Determine Optimal Charging Patterns
MATLAB
Add EV power to aggregate load signals
NILM-TK
Disaggregate load in testing days
Python
Prediction Accuracy
Testing Days
Form Models
Estimate hourly duty cycles from testing data
Form building model
MATLAB
Desired temp, house model, weather data, duty cycles
intlinprog
Determine Optimal Furnace on/off Patterns
MATLAB
Replace furnace data in REDD
NILM-TK
Disaggregate load in testing days
Python
Prediction Accuracy
Electric Furnace Test Case

Fig. 1. A flowchart of the overall training/testing procedure from start to finish for the two considered test cases, indicating what MATLAB/Python toolbox or function was used to solve the problem during particular steps.

reality, this is a conservative assumption, due to the fact that it is unlikely for a third-party to have this initial training set. Two popular algorithms for NILM were considered; the Factorial Hidden Markov Model (FHMM) and Combinatorial Optimization (CO) provided in NILM-TK. A five-minute timescale was used (i.e., the REDD data was downsampled to 5 minutes). In the following, the problem definitions for the two considered test cases are defined, with a graphical illustration of the training/testing process shown in Fig. 1.

A. Electric Vehicle Charging Obfuscation

Optimizing EV charging has many benefits [18], including the ability to obfuscate load signatures. In contrast to previous works which use energy storage for load obfuscation, we assume that the electric vehicle has additional constraints that stationary storage does not possess. Firstly, we assume the EV is unavailable during normal working hours (9:00 AM - 5:00 PM), and therefore is unable to contribute to load obfuscation during this time. Secondly, we assume the EV is 20% charged when the owner arrives at home at 5:00 PM every day, and that the owner requires the car to be charged at 100% capacity by 9:00 AM before the owner leaves for work. The optimization problem for EV scheduling, given these constraints, is given in (P1).

$$
\begin{aligned}
\text{(P1)}\ \min_{\mathbf{EV}_P}\ & \sum_{t=1}^{T-1} (P_{tot}^{(t+1)} - P_{tot}^{(t)})^2 \\
\text{subject to}\ & \\
& \underline{EV}_E \leq EV_E^{(t)} \leq \overline{EV}_E, \quad \forall t \\
& 0 \leq EV_P^{(t)} \leq \overline{EV}_P, \quad \forall t \\
& EV_E^{(t+1)} = EV_E^{(t)} + \Delta t EV_P^{(t)} \quad \forall t \\
& P_{tot}^{(t)} = EV_P^{(t)} + P_{agg}^{(t)}, \quad \forall t \\
& EV_E^{(t)} = 0.2 * \overline{EV}_E, \quad EV_P^{(t)} = 0, \quad t \in \tau \\
& EV_E^{(t)} = \overline{EV}_E, \quad t \in \tau'
\end{aligned}
$$

where $EV_E^{(t)}$ is the state of charge of the EV at time t, $EV_P^{(t)}$ is power consumption of the EV at time t, $\underline{EV}_E$ and $\overline{EV}_E$ are the lower and upper capacity limits of the EV, $\overline{EV}_P$ is the maximum charge rate, $P_{agg}^{(t)}$ is the aggregate power consumption of the other considered appliances at time t (uncontrollable input), Δt is the length of each timestep, τ includes time periods when the EV is unavailable (at work), τ' includes time periods when the user needs the EV to be fully charged, and T is the overall simulation time. The objective function aims to schedule EV charging in order to flatten the aggregate load profile.

B. Shifting Electric Furnace Usage

From the existing furnace data in the homes in the REDD dataset and the historical outdoor weather profile for that corresponding day/location, the building envelope (β_1), cooling/heating gain (β_2), and solar gain (β_3) were estimated and linked together in the following linear building model [19]:

$$
T_{in}^{(t+1)} = T_{in}^{(t)} + \beta_1 (T_{out}^{(t)} - T_{in}^{(t)}) - \beta_2 f_i^{(t)} + \beta_3 P_{rad}^{(t)},
$$

where $f_i^t \in \{0, 1\}$ is a binary variable indicating if the furnace is off (0) or on (1) at time t. State variable $T_{in}^{(t)}$ represents the indoor temperature at time t, $T_{out}^{(t)}$ is the outdoor temperature at time t, and $P_{rad}^{(t)}$ is the solar irradiance at time t.

The duty cycle of the furnace at hour h, $d_c(h) \in [0, 1]$, was estimated from the training data by counting the number of 5-minute periods the furnace was on divided by the total number of timeslots within an hour (12):

$$
d_c(h) = \frac{\sum_{k=1}^{12} \mathbb{1}_{[\underline{f_p}, +\inf]}(f_p(k))}{12}
$$

where $f_p(k)$ is the measured furnace power consumption at timestep k within hour h, $\underline{f_p}(k)$ is the minimum cutoff power measured where the furnace is considered "on" (to avoid labeling the furnace as on from measurement noise).

The objectives considered in this test case aim to flatten the aggregate load profile by shifting furnace usage, as in (P1), in addition to matching the original duty cycle of the furnace as close as possible per hour. This way, the same heating energy will be delivered to the home, aiming to keep the indoor air temperature within desirable bounds. The rated power of the furnace was inferred from the dataset and assumed to be the furnace's operating power when running the house model (a similar analysis could be performed during cooling seasons for air conditioning). A desired setpoint and deadband were estimated from the outdoor temperature and the measured duty cycle of the furnace. While the entire optimization problem for the electric furnace is not listed here due to space constraints, an overview of the procedure is seen in Fig. 1.

III. Test Cases

The two aforementioned test cases are discussed in this section, and results are shown indicating the benefit of load flexibility with respect to privacy preservation.

A. Electric Vehicle Charging

House 1 in the REDD dataset was used to demonstrate the privacy impacts with the addition of an electric vehicle. Fig. 2 shows the power consumption of the EV with respect to the aggregate energy from the top 5 energy consuming appliances in that house (dishwasher, fridge, lights, microwave, and sockets). The EV was assumed to have a state of charge of 20% when the owner comes home at 5:00 PM, and was required to be fully charged by the time the owner leaves at 9:00 AM. As seen in Fig. 2, the EV charging algorithm attempts to flatten the overall load profile, making it more difficult to recover unique load signatures. The root mean square error increases for nearly every appliance as compared to the case without an EV, as seen in Fig. 3. Further, as the available capacity of the EV increases, more flexibility is offered for load obfuscation, and the prediction accuracy decreases for both the FHMM and CO algorithms.

B. Electric Furnace Shifting

House 4 in the REDD dataset, which contains an electric furnace, was used in this test case. A baseline power consumption of 450 W was used for the furnace output, and the on/off decisions for the furnace were determined in the mixed-integer optimization problem, which aims to flatten load by shifting the furnace usage, while maintaining a desired duty cycle / heat energy output per hour. Fig. 4 shows how close to the original duty cycle the optimization problem was able to schedule the furnace for each hour of the considered test day, in order to preserve the indoor comfort requirements of the homeowner. The actual shifting of furnace power is shown in Fig. 5, where the new furnace power profile is plotted against the actual measured data. The furnace usage is not dramatically shifted throughout the day; however, by shifting the consumption slightly, the RMSE of the prediction algorithms still decreases overall, as tabulated in Table I. While the optimization of furnace usage indicates an overall decrease in accuracy, it can be seen from the table that the accuracy in identifying lights increased. This is because on/off thermal loads do not provide as much flexibility compared to the EV case, especially when constraints such as thermal comfort are considered. In addition, the indoor temperature, as modeled in the problem constraints by the aforementioned house model, cycles around the setpoint of 69° and stays within the given deadband of +/-2 degrees, as seen in Fig. 6. A more detailed model fitting strategy for the house model should be considered in the future; this is a main objective of future work on this topic.

TABLE I
RMSE increase during April 21 due to furnace usage shifting

% Increase in RMSE	FHMM	CO
Furnace	12.7%	37.0%
Lights	-29.6%	-22.7%
Sockets	24.7%	209.5%
Washer/Dryer	1.3%	18.1%

IV. Conclusion and Future Work

The results presented in this paper provide an initial indication of how leveraging a building's flexible loads can increase the overall resiliency of a building to unwanted load disaggregation. Having very flexible devices such as energy storage or plug-in electric vehicles can drastically decrease the NILM accuracy rate; whereas more constrained devices such as controlling an electric furnace can provide less flexibility but still improve the privacy of the homeowner. Currently, the formulation assumes perfect predictions of load, which is unrealistic. Future work will incorporate stochasticity into the problem, perhaps in a stochastic model predictive control approach. This will make the framework more robust to uncertainty (for example, the EV user coming home at a time other than 5:00 PM, or leaving at some point in the evening).

References

[1] S. Finster and I. Baumgart, "Privacy-aware smart metering: A survey," *IEEE Comms. Surveys Tutorials*, vol. 16, no. 3, pp. 1732–1745, 2014.

[2] C. Hawk and A. Kaushiva, "Cybersecurity and the smarter grid," *The Electricity Journal*, vol. 27, no. 8, pp. 84–95, 2014.

[3] B. McMillin and T. Roth, "Cyber-physical security and privacy in the electric smart grid," *Synthesis Lectures on Information Security, Privacy, and Trust*, vol. 9, no. 2, pp. 1–64, 2017.

[4] H. Souri, A. Dhraief, S. Tlili, K. Drira, and A. Belghith, "Smart metering privacy-preserving techniques in a nutshell," *Procedia Computer Science*, vol. 32, pp. 1087 – 1094, 2014.

[5] W. Kleiminger, C. Beckel, T. Staake, and S. Santini, "Occupancy detection from electricity consumption data," in *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*, ser. BuildSys'13. New York, NY, USA: ACM, 2013, pp. 10:1–10:8.

[6] A. Zoha, A. Gluhak, M. A. Imran, and S. Rajasegarar, "Non-intrusive load monitoring approaches for disaggregated energy sensing: A survey," *Sensors*, vol. 12, no. 12, pp. 16 838–16 866, 2012.

[7] W. Wichakool, Z. Remscrim, U. A. Orji, and S. B. Leeb, "Smart metering of variable power loads," *IEEE Transactions on Smart Grid*, vol. 6, no. 1, pp. 189–198, Jan 2015.

[8] H. H. Chang, K. L. Lian, Y. C. Su, and W. J. Lee, "Power-spectrum-based wavelet transform for nonintrusive demand monitoring and load identification," *IEEE Transactions on Industry Applications*, vol. 50, no. 3, pp. 2081–2089, May 2014.

[9] K. C. Armel, A. Gupta, G. Shrimali, and A. Albert, "Is disaggregation the holy grail of energy efficiency? the case of electricity," *Energy Policy*, vol. 52, pp. 213–234, 2012.

[10] A. Molina-Markham, P. Shenoy, K. Fu, E. Cecchet, and D. Irwin, "Private memoirs of a smart meter," in *Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*. New York, NY, USA: ACM, 2010, pp. 61–66.

[11] S. McLaughlin, P. McDaniel, and W. Aiello, "Protecting consumer privacy from electric load monitoring," in *The 18th ACM Conf. on Comp. and Comms. Security*. New York, NY, USA: ACM, 2011, pp. 87–98.

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

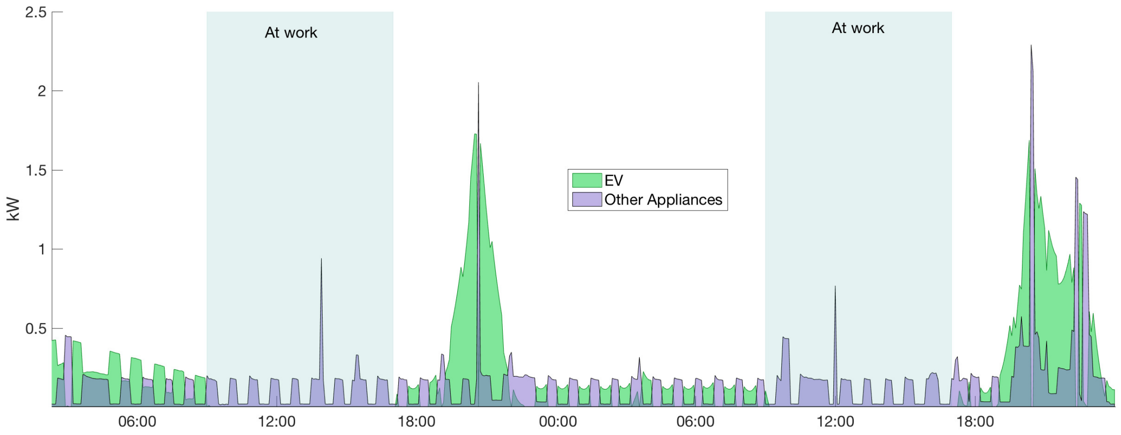

Fig. 2. EV and aggregate appliance power consumption for House 1 for two days during April 2011, with an available 60 kWh EV capacity. Using strategic electric vehicle charging can flatten the aggregate load profile, obfuscating individual appliance signatures. In contrast to stationary energy storage, additional constraints were imposed on the EV; during work hours (9:00 AM - 5:00 PM), the EV is unavailable to participate in load obfuscation. In addition, the EV comes home each day charged at 20% of its full capacity, and it is required that the EV is fully charged by morning (9:00 AM).

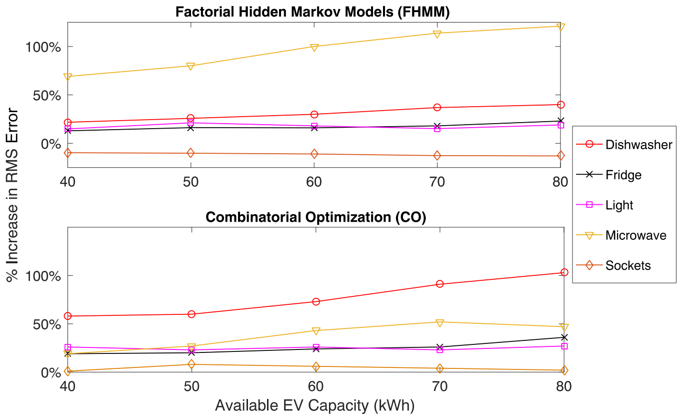

Fig. 3. Root mean square error (RMSE) for House 1 for a single day in April, 2011. Compared to the baseline RMSE from the home without load obfuscation, using strategic EV charging results in an overall increase in prediction error.

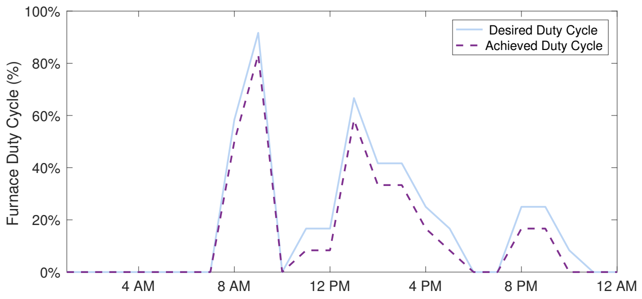

Fig. 4. To minimize the indoor comfort impact, the optimized duty cycle of the furnace for each hour attempts to match the original duty cycle.

[13] G. Kalogridis, C. Efthymiou, S. Z. Denic, T. A. Lewis, and R. Cepeda, “Privacy for smart meters: Towards undetectable appliance load signatures,” in *IEEE Int. Conf. on Smart Grid Comm.*, Oct 2010, pp. 232–237.

[14] J. X. Chin, T. T. D. Rubira, and G. Hug, “Privacy-protecting energy management unit through model-distribution predictive control,” *IEEE Transactions on Smart Grid*, vol. 8, no. 6, pp. 3084–3093, Nov 2017.

[15] P. Fortenbacher, J. L. Mathieu, and G. Andersson, “Modeling and optimal operation of distributed battery storage in low voltage grids,” *IEEE Trans. on Power Sys.*, vol. 32, no. 6, pp. 4340–4350, Nov 2017.

[16] N. Batra, J. Kelly, O. Parson, H. Dutta, W. Knottenbelt, A. Rogers, A. Singh, and M. Srivastava, “NILMTK: an open source toolkit for non-intrusive load monitoring,” in *Proceedings of the 5th international conference on Future energy systems*. ACM, 2014, pp. 265–276.

[17] Z. Kolter and M. Johnson, “REDD: A public data set for energy disaggregation research,” in *Proceedings of the SustKDD workshop on data mining applications in sustainability*, Aug 2011.

[18] A. J. B. Brush, J. Krumm, S. Gupta, and S. Patel, “Evhomeshifter: Evaluating intelligent techniques for using electrical vehicle batteries to shift when homes draw energy from the grid,” in *ACM International Joint Conference on Pervasive and Ubiquitous Computing*, ser. UbiComp ’15. New York, NY, USA: ACM, 2015, pp. 1077–1088.

[19] X. Jin, K. Baker, D. Christensen, and S. Isley, “Foresee: A user-centric home energy management system for energy efficiency and demand response,” *Applied Energy*, 2017.

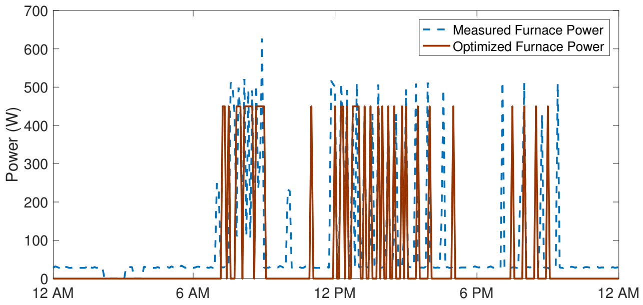

Fig. 5. Instead of assuming the furnace operates under hysteresis control, the on/off control of the furnace was determined in an optimization problem to attempt to flatten the aggregate load profile.

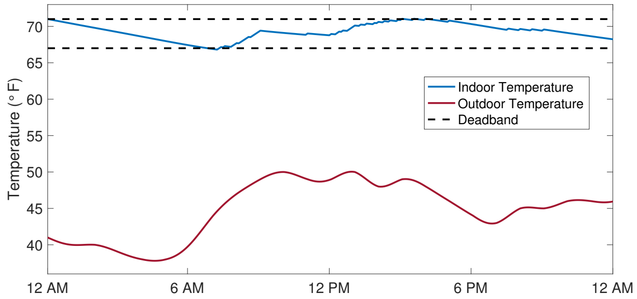

Fig. 6. The house model, estimated from the measured furnace data, is used as a constraint in the furnace optimization to ensure shifting the furnace usage does not result in undesirable indoor temperatures.