Influence of Data Granularity on Nonintrusive Appliance Load Monitoring

Günther Eibl
Josef Ressel Center for User-Centric Smart Grid Privacy, Security and Control
Salzburg University of Applied Sciences
Urstein Sued 1, Puch/Salzburg, Austria
guenther.eibl@en-trust.at

Dominik Engel
Josef Ressel Center for User-Centric Smart Grid Privacy, Security and Control
Salzburg University of Applied Sciences
Urstein Sued 1, Puch/Salzburg, Austria
dominik.engel@en-trust.at

ABSTRACT

Decreasing time resolution is the simplest possible privacy enhancing technique for energy consumption data. However, its impact on privacy analyses of load signals has never been studied systematically. Non-intrusive appliance load monitoring algorithms (NIALM) have originally been designed for energy disaggregation for subsequent energy feedback. However, the information on appliance use may also be misused for the extraction of personal information. In this work, the effect of decreasing the time resolution in the usual first step, namely edge detection, is studied. It is shown that event values can be estimated rather reliably, but the detection rate of events significantly decreases with increasing measurement time interval.

Categories and Subject Descriptors

I.5 [**Pattern Recognition**]: Design Methodology—*Pattern analysis*

General Terms

Privacy

Keywords

Privacy enhancement; smart metering; data representation; load disaggregation; edge detection

1. INTRODUCTION

There is a lot of public concern and discussions on the privacy impact of smart metering. However, the discussion is led without knowing the extent of personal information that can be read out of smart meter load profiles. Even more so, there is nearly a complete lack of knowledge about how the amount of personal information relates to the measured time interval. For example, in many European countries, it is planned, that people can opt-in for delivering their load data in 15 minute time intervals. To our knowledge, no one has tried to assess the amount of personal information that can be extracted on 15 minute time interval load profiles.

Note that the decrease in time resolution can be viewed as the most straightforward and simplest privacy enhancing technology (PET), cf. [3]. The goal of this work is making a first step towards the study of its actual impact. This work is a first step, because we focus on determining appliances. The main reasoning behind this approach is that activities of persons in the house trigger appliances that sum up to the total load. The activities themselves are already personal information of which some general habits could be deduced. However, such an analysis of general habits is out of scope of this work.

Information on the appliances are usually extracted from the load profiles by means of so-called 'non-intrusive appliance load monitoring analysis" (NIALM). There is a lot of literature on NIALM algorithms ([5, 15, 2, 1, 14, 8, 6, 13]). The goal of these algorithms is the disaggregation of the total load into the individual appliances loads, e.g., for sake of providing energy feedback to the end-user. From the privacy viewpoint, such NIALM analyses can be seen as a first step of attacking methods, which aim at the unauthorized extraction of personal information.

There are only a few papers treating the technical details of privacy implications of smart metering. In [9], load data were recorded with parallel video data which were processed into activity logs. A NIALM analysis was done yielding the input for subsequent behavior-extraction routines. Extracted behaviors include, e.g., presence, sleep cycles or meal times. In [12] the load profile is divided into so-called power segments using a density based clustering technique. These power segments are described by features such as start time, average power and duration. It is illustrated how such power events could be used for answering several privacy questions. In [4], it is shown that under ideal conditions load curves can be used to identify the currently viewed TV-program.

In this work, the impact of reducing the time granularity on the first part of typical low-frequency NIALM algorithms, namely edge detection ([5, 9, 1, 2, 8, 13]) is studied. In Section 2.1, event detection is described as part of low frequency NIALM analyses. In Section 2.2 the investigated edge detection methods are reviewed. After describing the experimental setup in Section 3, the performance of different edge detection methods is compared in Section 4.1. The core Section 4.2 of this work describes the effect of the time reso-

IHMMSEC 2014 Salzburg, Austria

http://dx.doi.org/10.1145/2600918.2600920.

Time	1s	3s	15s	20s	1min
Paper	[12, 11, 2, 1]	[8, 6]	[9]	[14]	[13]

Table 1: Time Granularities of low-frequency NIALM-studies

lution on the detection of events. Finally, Section 5 contains conclusion and outlook.

2. EVENT DETECTION METHODS

Event detection methods are the typical first analysis step in low frequency NIALM algorithms. Decreasing the performance of event detection is a countermeasure against a possible NIALM-privacy-attack and increase privacy. After discussing why event detection is such a useful first step of NIALM analysis, the event detection methods that are investigated in the experimental part are described.

2.1 Event Detection as Part of NIALM

NIALM approaches are divided broadly into two kinds of methods: high frequency methods look at the waveform of appliances or study transients or higher order harmonics. While high-frequency methods usually need a sampling in the range of kHz, low-frequency methods typically analyze load profiles which are sampled using time intervals in the order of seconds (see Table 1).

Since in this work the time granularity is decreased for privacy purposes, the focus is laid on low-frequency instead of high frequency NIALM methods. Supervised low frequency methods usually consist of several blocks: edge detection, cluster analysis and finding pairs of on-and off clusters for the determination of the duration of an appliance. Edges are sharp increases or decreases of the load signal due to turning on or off an appliance. More generally, edges arise due to the change from one state to another state of an appliance when modeled as a finite state machines (FSM). NIALM algorithms commonly use edges instead of the absolute values for two reasons: First, using absolute values in the presence of unknown appliances, these unknown appliances could be described as a combination of other known appliances. Second, there are adverse cases where a small change in the measured power would result in a big change in the configuration of used appliances, which is an implausible result [5]. Since edge detection is a common first step of a NIALM algorithm, if a decrease of time resolution is able to negatively influence edge detection, the subsequent part of the NIALM algorithm is expected to suffer significantly as well. Considering a possible abuse of NIALM algorithms the diminished disaggregation ability is beneficial from the privacy perspective. For sake of completeness it is noted that the use of edges is common but not mandatory, e.g., in [12] shape features are used instead of edges.

2.2 Investigated Event Detection Methods

In this section event detection methods used in the experimental part are reviewed. A main assumption of this work is the modeling of appliances as finite state machines (FSMs) having different power values for different states. In this work an event $e = (t_e, \Delta P_e)$ is a transition between two such states which is represented by its onset time t_e and the difference between the two power levels of the states ΔP_e. Many appliances have only two states and can simply only be turned on or off. Correspondingly, events for which the signal increases ($\Delta P > 0$) are called on-events because they should typically arise from turning on such an onoff-appliance. Analogously, events for which the signal decreases are called off-events.

The most straightforward method detects an edge, if the backward difference $\Delta P_i = P_i - P_{i-1}$ between consecutive points exceeds a threshold. Each detected edge is considered to be an event $e = (t_i, \Delta P_i)$. This method can be classified as one that focuses on the transition between two levels of a signal [10]. If the transition needs several time intervals, this method divides the transition between two levels in several edges having smaller values than the transition which is usually an unwanted behavior.

The drawback of the backward difference method can be accommodated by merging of subsequent occurring edges stemming from backward differences into a single event [1]. The value of the event is the sum of the individual edge values which can be both positive and negative. The time where the event occurs is defined as the onset time, i.e., the time of the first edge contributing to the event.

Another method proposed in [5] is called 'transient passing edge detection." As its names suggests it is a method focusing on the power levels of the two transition states instead of the transition itself. A transition is defined as being not steady. In the first step the method finds the steady subsequences of the signal. This is done using a sliding window approach where a point is considered part of a steady subsequence, if the range of it and the next $n - 1$ does not exceed a given threshold. The whole signal is thus divided into consecutive steady parts st and transitions tr. For the description of the event, all subsequences $(st_i, tr_{i+1}, st_{i+2})$ are considered. The onset-time t_{e_i} for the description of the event is the last time point of the first steady part st_i. The transition value ΔP_i is the difference between the median of the values of the first steady part st_i and the median of the values of the consecutive steady part st_{i+1}. Taking the median value over the whole steady part leads to a greater robustness in the determination of the event value ΔP_i.

3. EXPERIMENTAL SETUP

The experiments were done using a so-called low frequency dataset of the publicly available REDD-dataset [7]. This dataset consists of measurements of the apparent power for 6 different houses. Measurements are available for the main circuits `mains1` and `mains2`, and for subcircuits like for example kitchen outlets and measurements of individual appliances.

Although the decrease of the time granularity seems straightforward (integrating over the period), it is in fact not. There are several possibilities. First, considering a time interval, different statistics could be computed for this interval. The most straightforward statistic is the average load value which should be enough for most practical solutions such as normal billing or time-of-use billing. However, for some reasons, e.g., pricing based on the maximum load or for control reasons, the maximum load needed during the time interval, could be another useful number. Other statistics as for example the standard deviation of the load values, are also possible but will not be considered further. Finally, there is still the possibility of simple sampling, i.e., taking the load value at the specific point in time.

In the subsequent experiments, three variants are considered: (i) taking the average load in a time interval, (ii) taking the maximum load in a time interval and (iii) sampling at time points.

In order to account for noise, for all methods, events e with a value ΔP smaller than a threshold of 20 Watt are discarded. The same threshold was used for the detection of the stable parts of transient edge detection. The minimal required number of steady points n in transient passing was set to 3 which has good detection properties at reasonable stability.

4. RESULTS

In this section, different edge detection methods are compared with respect to their ability to detect events in smart meter load profiles. Then the effect of the decrease of time resolution on the events found is described.

4.1 Event detection

Since the results are based on the events found, the performance of the event detection methods is assessed for the highest available time resolution of 3 seconds first. A value of 20W was used as threshold for the removal of events occurring due to noise. If the threshold is set too low, additional edges can occur which tends to happen for high-power devices. For low-power devices such as lighting, a noise threshold that is in turn too high can lead to a loss of events. Therefore, the tradeoff between noise removal and the detection of events from low-power devices has to be considered.

The form of the load consumption of appliances can be quite complex. As an example, the load consumed for a full run of the dishwasher is shown in Figure 1. Since the

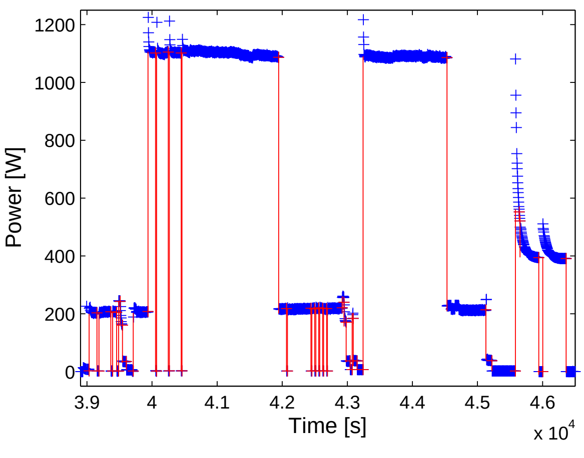

Figure 1: Dishwasher events, marked as "+", detected using transient detection at the highest time resolution

dishwasher's load profile has such a rich structure with long and short on-durations at different power levels and power levels that are decreasing, it was chosen for demonstration of effects of different edge detection settings and of the change of time granularity. Simpler devices for heating are usually purely ohmic and show high power values. These are the appliances whose load profiles have the highest similarity to a rectangular profile.

As expected, the simple backward difference yields more, but disturbing, events and can therefore not be recommended (compare Figures 1 and 2).

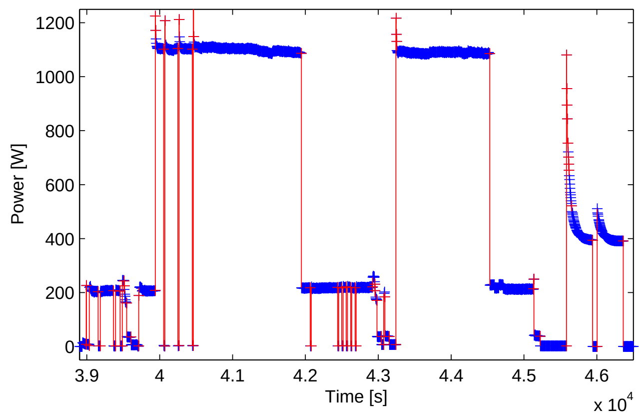

Figure 2: Dishwasher events, marked as "+", detected using backward differences at the highest time resolution. Too many events are detected (compare with Figure 1).

Generally, in terms of detecting appliances, both transient passing and edge merging give good and very similar results. There are also only tiny differences due to the use of the different variants of decreasing the time resolution. The correctness of the edges found was visually verified for all appliances. Additionally, the edge values of all appliances are shown in Figure 3. It can be seen that for all appliances rather distinct edge values can be found. The expected strong similarity of the absolute values of the on-events and the off-events leads to the symmetric look of Figure 3. More importantly, this figure suggests that some appliances such as washerDryer3 should be easily distinguishable from others. Other appliances such as kitchen outlets 2 and 4 are expected to be hardly distinguishable from others. For another class of appliances such as the dishwasher only some levels are distinguishable from the events of other appliances. The fact that the result of edge detection enables to formulate such an expected behavior shows the value of edge detection for a possibly privacy invading analysis of load profiles.

4.2 Effect of Decrease of Time Resolution

In this section, the influence of time granularity Δt on the events found above is studied. First, transient passing using the averaging statistic is studied. As can be seen in Figure 4 with increasing the time interval fewer edges are detected. Especially short-lived states cannot be detected anymore. The edges that are still detected have surprisingly stable heights ΔP.

Another remarkable point is that already with a time interval of 5 minutes, nearly the whole finer structure cannot be seen any more. These results can also be seen for the `mains` signals which was calculated as the sum of the

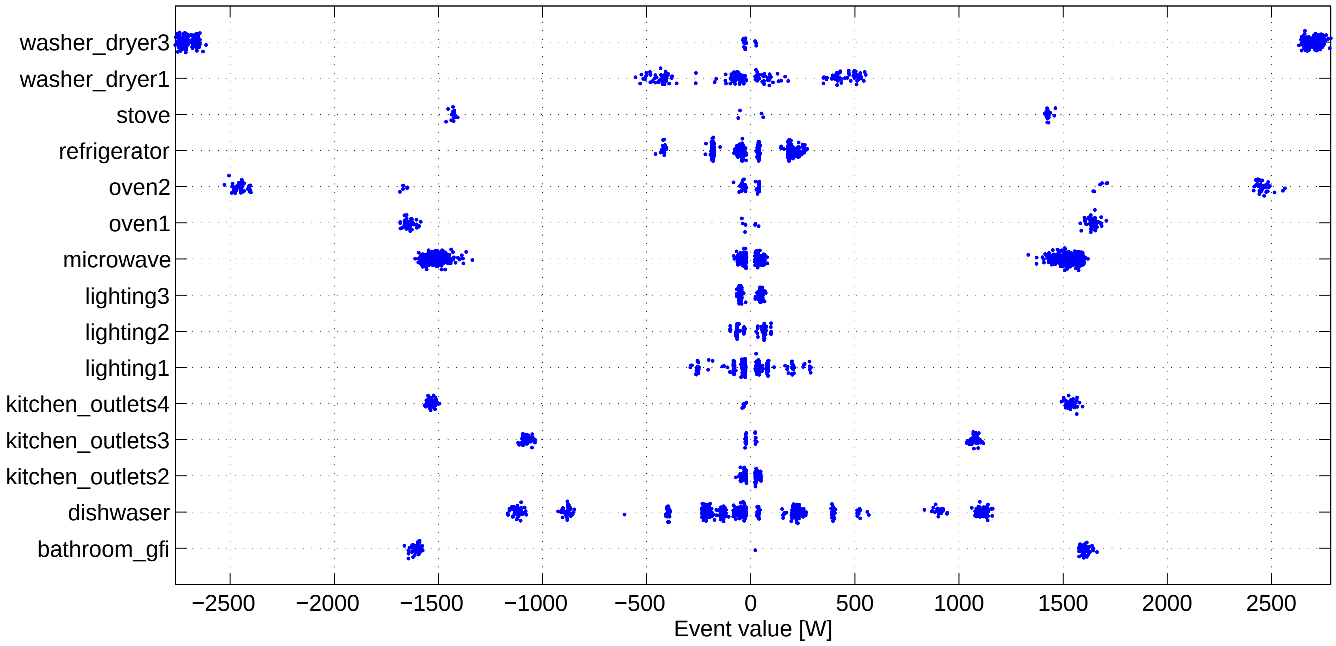

Figure 3: All events found at the highest time resolution, detected using transient detection. The symmetry of the figure stems from the strong similarity of the absolute values of on-events and their corresponding off-events.

`mains1` and the `mains2` signals. Using a 5 minute interval mostly privacy-irrelevant refrigerator events remain.

Possible effects on the decrease in privacy due to the decrease in time resolution can already be estimated. Since the edge heights are rather stable it seems reasonable that the edges of different appliances can still be distinguished at higher time intervals. However, the detection rate of appliances is diminished. In summary, the effect of a decrease in time resolution means that single events cannot be detected reliably. However, for the identification of habits, the detection of each single event is not necessary.

Comparing the different edge detection and time decrease variants, the following behavior could be seen: For high time resolution, edge merging and transient passing lead to nearly identical results, however, for lower time resolutions transient passing seems to better preserve the edge values. The results of both transient passing and edge merging are quite insensitive to the kind of statistic. Although still leading worse results, it should be noted that the performance of the backward difference method is better with taking the max statistic or with sampling than with taking the average statistic where extensive smearing of edge values occurs.

5. CONCLUSION AND OUTLOOK

The impact of decreasing the time resolution on privacy analysis of load signals obtained from smart metering to date has not been studied systematically. Based on the reasoning that knowledge about appliance use can be used as a first step in a privacy attack, the influence of the time interval on edge detection methods has been studied.

Three edge detection methods were investigated: the transient passing method [5], merging of backward differences and simple backward differences. Based on experiments with the REDD-data [7] the simple backward difference cannot be recommended as an edge detection tool in this setting leading to too many edges.

The decrease of the measurement time interval as a privacy enhancing operation has the effect that edge detection still works in the sense that edge heights can be detected in a stable manner. Privacy is enhanced in a way that not every edge is detected. The longer the time interval the fewer edges can be detected. Already with 5 minute intervals, for most of the appliances, the number of detected edges is significantly decreased. A potential privacy consequence would state that not every single event but rather regular habits can be detected.

This work constitutes the first, descriptive assessment of the effect of a decrease of data granularity on smart meter privacy focusing on the detection of appliance use. Next logical steps include the development of quantifiable performance indicators, e.g., based on the result of subsequent pattern recognition algorithms. Using these performance indicators the difference of the effect on different appliances should be described and visualized in a way that is also understandable for non-experts. Furthermore, when appropriate datasets are available, personal information such as activities or habits should be considered in addition to appliance usage.

6. ACKNOWLEDGMENTS

The financial support by the Austrian Federal Ministry of Economy, Family and Youth and the Austrian National Foundation for Research, Technology and Development is gratefully acknowledged. Funding by the Austrian Research Promotion Agency (FFG) under Bridge Project 832082 is gratefully acknowledged.

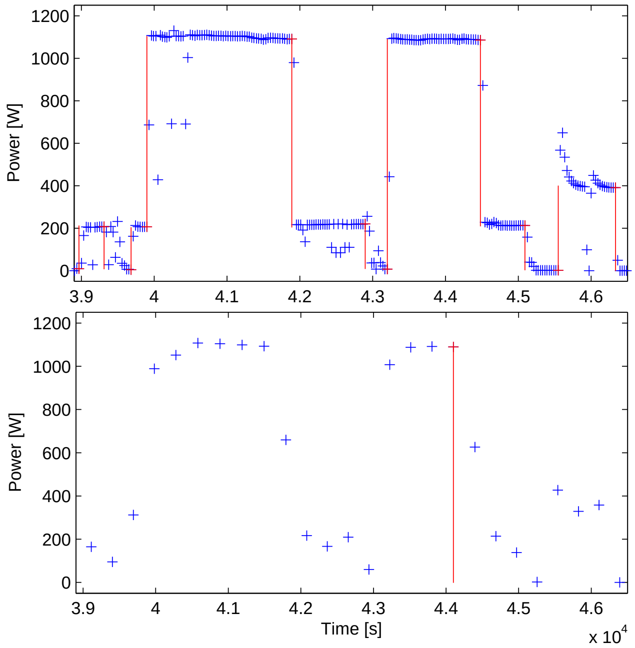

Figure 4: Dishwasher events, marked as "+", detected for Δt =30s (top) and 5 minutes (bottom).

7. REFERENCES

[1] M. Baranski and J. Voss. Genetic algorithm for pattern detection in nialm systems. In *IEEE International Conference on Systems, Man and Cybernetics*, 2004.

[2] D. C. Bergman, D. Jin, J. Juen, N. Tanaka, C. Gunter, and A. Wright. Distributed non-intrusive load monitoring. In *Proceedings of the IEEE/PES Conference on Innovative Smart Grid Technologies (ISGT 2011), Anaheim, CA, USA*, 2011.

[3] D. Engel. Wavelet-based load profile representation for smart meter privacy. In *Proc. IEEE PES Innovative Smart Grid Technologies (ISGT'13)*, pages 1–6, Washington, D.C., USA, Feb. 2013.

[4] U. Greveler, B. Justus, and D. Löhr. Multimedia content identification through smart meter power usage profiles. In *Proceedings of the 2012 International Conference on Information and Knowledge Engineering (IKE'12)*, Las Vegas, USA, 2012.

[5] G. Hart. Nonintrusive appliance load monitoring. *Proceedings of the IEEE*, 80(12):1870–1891, Dec. 1992.

[6] H. Kim, M. Marwah, M. Arlitt, G. Lyon, and J. Han. Unsupervised disaggregation of low frequency power measurements. In *The 11th SIAM International Conference on Data Mining*, pages 747–758, 2011.

[7] J. Kolter and M. Johnson. Redd: A public data set for energy disaggregation research. In *Workshop on Data Mining Applications in Sustainability (SIGKDD)*, pages 1–6, 2011.

[8] J. Z. Kolter and T. Jaakkola. Approximate inference in additive factorial hmms with application to energy disaggregation. *Journal of Machine Learning Research - Proceedings Track*, 22:1472–1482, 2012.

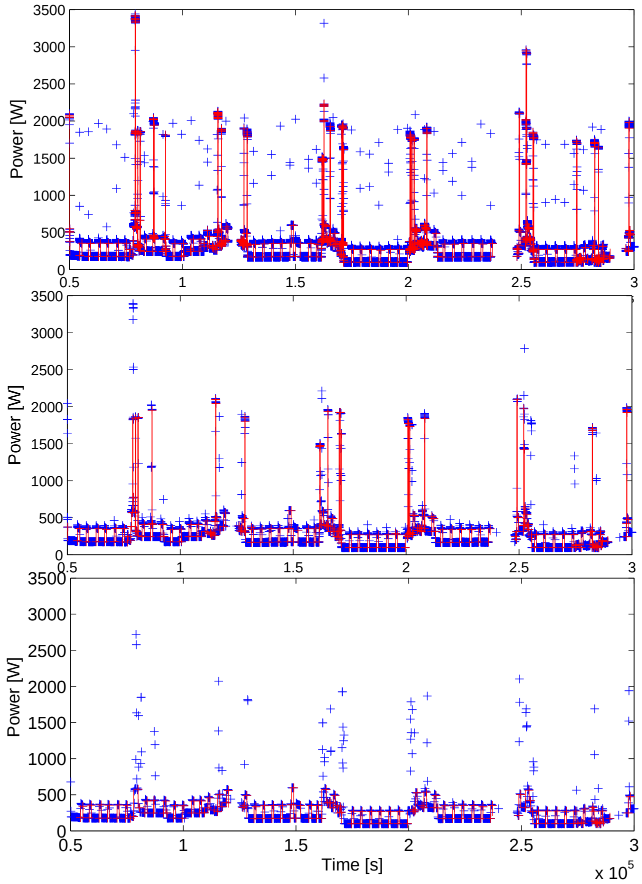

Figure 5: Mains events, marked as "+", detected with Δt=3s (top), 60s (middle) and 5min (bottom).

[9] M. Lisovich, D. Mulligan, and S. Wicker. Inferring personal information from demand-response systems. *IEEE Security & Privacy*, 8(1):11–20, 2010.

[10] M. A. Little and N. S. Jones. Generalized methods and solvers for noise removal from piecewise constant signals, background theory. In *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science*, volume 467, pages 3088–3114, 2011.

[11] A. Marchiori, D. Hakkarinen, Q. Han, and L. Earle. Circuit-level load monitoring for household energy management. *IEEE Pervasive Comp.*, 10:40–48, 2011.

[12] 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 Building*, BuildSys '10, pages 61–66, New York, NY, USA, 2010.

[13] O. Parson, S. Ghosh, M. Weal, and A. Rogers. Non-intrusive load monitoring using prior models of general appliance types. In *Twenty-Sixth Conference on Artificial Intelligence (AAAI-12).*, 2012.

[14] E. Vogiatzis, G. Kalogridis, and S. Z. Denic. Real-time and low cost energy disaggregation of coarse meter data. In *4th IEEE PES Innovative Smart Grid Technologies Europe (ISGT Europe)*, 2013.

[15] M. Zeifman and K. Roth. Nonintrusive appliance load monitoring: Review and outlook. *IEEE Transactions on Consumer Electronics*, 57:76–84, 2011.