Improving Device-level Electricity Consumption Breakdowns in Private Households Using ON/OFF Events

Christian Beckel, Wilhelm Kleiminger
Institute for Pervasive Computing
ETH Zurich
8092 Zurich, Switzerland
{beckel, kleiminger}@inf.ethz.ch

Thorsten Staake
Information Management
ETH Zurich
8092 Zurich, Switzerland
tstaake@ethz.ch

Silvia Santini
Wireless Sensor Networks Lab
TU Darmstadt
64283 Darmstadt, Germany
santinis@wsn.tu-darmstadt.de

ABSTRACT

Smart meters can measure the electricity consumption of a household at a fine temporal granularity. By adequately processing this aggregated data an estimation of the consumption of individual appliances can be retrieved and used to provide novel services, such as personalized recommendations on how to reduce the overall energy consumption of the household. In this paper, we build upon existing work in consumption data disaggregation and consider smart meter data along with additional information made available by networked sensors and household appliances. In particular, we investigate the use of ON/OFF events, which signal when appliances have been turned on or off, to improve the accuracy of a state-of-the art disaggregation algorithm that uses such events along with smart meter data to estimate the consumption of single appliances. Our results, obtained by applying the algorithm to a publicly available dataset, show that the accuracy of the algorithm quickly deteriorates as the number of available ON/OFF events decreases. We thus suggest possible countermeasures to cope with this limitation and to provide accurate electricity consumption breakdowns in private households.

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Miscellaneous

1. INTRODUCTION

The energy sector is undergoing a massive paradigm shift by reducing the dependency on fossil fuels towards an increasing share of renewable energy sources. To manage integration of these highly dynamic energy sources into the electricity network, smart meters are being installed in millions of private households worldwide [1]. A smart meter is a sensing device that can measure electric power consumption (from here on also referred to as *electricity consumption*) and can report the collected readings at given time intervals, e.g., every second, through a wireless or wired communication interface. The availability of smart meter data (possibly along with additional sensory information) enables the design and development of novel services and applications [10]. For instance, household inhabitants can be provided with fine-grained data about the contribution of individual appliances to the overall electricity consumption. The availability of this information can potentially motivate users to an overall more thrifty usage of electricity, e.g., by inducing them to purchase more efficient devices [9, 2, 13]. Also, knowledge about the use of appliances can help determining the occupancy state of a household and thus enable other services, like automatic heating control [7].

Recently, several approaches that derive the consumption of single devices by adequately processing the data collected by a smart meter have been proposed [4, 8, 11, 5]. However, because there are many possible combinations of appliances contributing to the electricity consumption at the same time, these centralized, single-sensor approaches usually achieve limited accuracy in real deployments [8, 11]. An alternative, fully distributed approach consists in instrumenting all household appliances in order to measure their individual electricity consumption. For instance, so-called smart power outlets can be used to measure the consumption signatures and report them to a central processing unit [4]. These solutions can clearly provide a very accurate consumption breakdown but they also result in high deployment and maintenance costs. A third possibility consists in combining knowledge of the total electricity consumption, retrieved through a smart meter, with additional sensory information like the sequences of *ON/OFF events* of each device [5, 6, 12].[1] Clearly, the lesser the amount of additional information that is needed to achieve a given disaggregation accuracy, the more attractive is its use in practical settings. Previous studies have indeed shown that the knowledge of the total electricity consumption along with the sequence of ON/OFF events of each appliance is sufficient to estimate the consumption breakdown accurately [5]. However, the accuracy of the estimation decreases significantly if only partial knowledge about the sequence of ON/OFF events is available.

In this paper, we focus on these latter hybrid approaches and provide two main contributions. First, we evaluate the performance of an existing state-of-the-art load disaggregation algorithm that relies on the use of ON/OFF events along with smart meter data. In particular, we quantify the deterioration of the algorithm's performance in terms of device identification accuracy as the number of collected ON/OFF events decreases. To this end, we use a publicly available dataset of electricity consumption data [8]. Building upon our evaluation study, we then propose a set of mechanisms that can contribute in making the considered load disaggregation algorithm more robust against the presence of appliances that do not provide ON/OFF events. To

[1] An ON/OFF event signals whether the device has been turned on or off.

this end, cooperation between different sensors and devices (e.g., smart meters, intelligent power outlets, light sensors) plays a crucial role in capturing additional sensory information, which is required to disaggregate the overall electricity consumption and attribute it accurately to its individual contributors.

2. RELATED WORK

Accurate device-level consumption breakdowns could easily be measured if appliances were able to autonomously measure and report their own electricity consumption. This would however require costly hardware enhancements of the devices and is thus an impractical solution, especially for cheap and old appliances. Alternatively, smart power outlets (such as the ones from Plugwise[2]) can measure devices' electricity consumption at the "socket level". Nonetheless, accordingly equipping each socket of a private household would also cause a significant cost and deployment effort.

To overcome these limitations, several authors focused on the concept of non-intrusive load monitoring (NILM). NILM strategies typically process the aggregated electricity consumption profile [14] to estimate the corresponding device-level consumption breakdown. For instance, Kolter et al [8] resort to a Factorial Hidden Markov Model (FHMM) to perform load disaggregation. The model is trained on consumption data from multiple households and when used on data from a test household it achieves an breakdown accuracy of 47.7%. This evaluation relies on a metric that captures accuracy every 10 seconds, thus avoiding that errors even out over time. Parson et al. focus on the three appliances that consume the highest amount electricity and use Hidden Markov Models (HMM) to guess their current state [11]. The corresponding experimental evaluation shows that restricting the set of appliances leads to a consumption breakdown accuracy of 83%. Both Kolter et al.'s and Parson et al.'s approaches rely solely on smart meter data. In the context of our work, we instead aim at using also information from other sensors – such as infrared or light sensors – and thus trade off a slightly higher system complexity with a corresponding increase in estimation accuracy.

Several authors have indeed already abandoned the idea of disaggregating electricity consumption using single-point measurements only [6, 5, 3, 12]. For instance, the ViridiScope system [6] leverages magnetic field sensors and light sensors to indirectly sense electricity consumption of appliances. Smart meter data is only used for the purpose of sensor calibration. Within ViridiScope, uninstrumented appliances whose consumption cannot be directly measured are defined as *ghost power* consumers [6]. While our approach to consumption breakdown relies on ON/OFF events of appliances as in [5], ViridiScope approximates electricity consumption through indirect sensing. Furthermore, in our work we explicitly model the influence of ghost power of the final estimation while ViridiScope assumes that the ghost power of non-instrumented appliances is constant and rather small. Jung and Savvides take into account knowledge about ON/OFF states of all appliances [5] to estimate device-level consumption breakdowns. To evaluate their approach the authors gathered three days of data from both a central electricity meter and from sensors mounted next to the appliance's switches that recorded ON/OFF events. On this dataset, the consumption breakdown algorithm achieves an accuracy of 90%. However, the algorithm assumes that complete knowledge about occurring ON/OFF events is available. This assumption is unfortunately hard to meet in real deployments. In the following section, we show that the performance of this disaggregation algorithm decreases significantly when ON/OFF events are recorded for only a subset of the available appliances, instead of for all of them. To cope with this problem, we suggest to explicitly include ghost power in the disaggregation model.

Other approaches also leverage ON/OFF events but use different techniques to detect them. The ElectriSense system [3], for example, can detect consumer electronics devices and fluorescent lighting leveraging the fact that these appliances use switch mode power supplies and thus generate measurable electromagnetic interference (EMI) during their operation. This approach, however, requires adequate (expensive) hardware to measure high frequency switching events within the household's electrical circuit. Instead, our disaggregation algorithm relies on low frequency data from an ordinary smart meter along with ON/OFF events measured directly at the appliance. An example of a low-cost sensor that enables such direct measurements is provided in [12]. Placed next to the appliances to be observed – such as refrigerators, lights, desktop computers, or televisions – this sensor can detect their state changes by measuring variations in electromagnetic fields.

3. PERFORMANCE ANALYSIS OF A LOAD DISAGGREGATION ALGORITHM

In this section we discuss the performance of a reference algorithm to perform load disaggregation using ON/OFF events. In particular, we quantify the deterioration in terms of device identification accuracy as the number of collected ON/OFF events decreases. In the following, we first elaborate on the opportunity of using ON/OFF events to perform load disaggregation and then describe in detail the reference algorithm considered for our performance evaluation. Before presenting the final results, we also describe the dataset that we have used for running our experiments.

3.1 Collection of ON/OFF events

The availability of smart meter data alone is often not sufficient to achieve high load disaggregation accuracies [8, 11] and the use of additional information, like the sequence of ON/OFF events, is thus often unavoidable. Capturing ON/OFF events is particularly interesting as it only requires a lightweight (i.e., cheap and easy to deploy) sensing infrastructure [6, 5]. However, instrumenting all or most devices within an household would still cause too high deployment and maintenance costs. Therefore, algorithms that can cope with partial information about ON/OFF events would significantly increase the attractiveness of approaches based on this load disaggregation technique.

Figure 1 exemplarily shows the sensing infrastructure that can be used to gather ON/OFF events. The system includes a smart meter that captures data at a frequency of 1Hz and a number of smart power outlets (e.g., from Plugwise) that

[2]www.plugwise.com

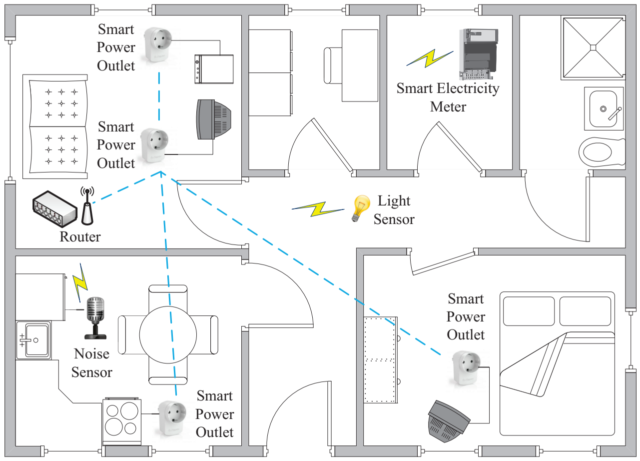

Figure 1: Data collection architecture

can measure and notify when a device is turned on or off. Although the specific smart power outlets used in our experiments are able to capture the actual consumption curve of a device, we only assume the availability of a sensor that is able to determine the operational state of the appliance. Additional sensors may also be used to indirectly measure ON/OFF events. For instance, a light sensor can measure the state of a lamp while a microphone or accelerometer can allow inferring the current state of a washing machine. However, in order to make the collection of ON/OFF events feasible in practical scenarios, it must be implemented as simple, cheap, and unobtrusively as possible. Considering the steadily growing trend towards embedding communication capabilities (like Ethernet or WiFi interfaces) in common appliances, we believe that it is reasonable to assume that (at least a subset of) devices within a household will be able to determine and communicate their operation status autonomously. However, a given number of devices is likely to remain unobservable, e.g., old or cheap appliances. Therefore, a given amount of the total power consumption will remain non-attributable to the actual devices causing it. Borrowing the definition reported in [6], we refer to this "non-attributable" power as *ghost power*. Unlike [6], however, we do not assume ghost power to be small and constant. Instead, we include ghost power as an explicit variable in our model so as to increase accuracy of the overall consumption breakdown.

3.2 Reference disaggregation algorithm

To the best of our knowledge, Jung and Savvides's algorithm [5] is the only currently available algorithm that combines smart meter data with ON/OFF events in a comprehensive way. It thus represents an obvious choice as a reference algorithm to analyze the performance of load disaggregation approaches based on ON/OFF events.

The disaggregation algorithm of Jung and Savvides [5] solves a linear optimization problem to estimate the contribution of each appliance to the overall power consumption. To this end, it maintains a trace of the total electricity consumption as well as a state vector of active appliances. States with the same set of active appliances are merged on the fly by averaging the total power consumption and increasing a state counter. Using this data, the algorithm computes the average consumption of each appliance by minimizing the mean square error between the sum of estimates of all active appliances and the total electricity consumption. To improve the estimation accuracy samples of the ON/OFF state vector that show fewer appliances in the ON state as well as samples that occur frequently are given higher weight in the estimation. Similarly, stationary loads are also given higher importance, as they can be estimated more accurately. The estimation procedure is performed over a specific time interval (e.g., one hour) and then restarted, whereas estimations from previous intervals are "remembered" for each successive iteration.

In order to perform the experimental study presented below, we implemented Jung and Savvides's algorithm in Matlab. For the sake of simplicity, our implementation does not include the above mentioned weighting for stationary loads and we consider a single time interval only. We believe this simplification does not affect the general validity of our conclusions about the robustness of the algorithm. Nonetheless, we plan to use the original version of Jung and Savvides's algorithm in our future work.

3.3 REDD dataset

In order to evaluate non-intrusive load monitoring algorithms using actual electricity consumption measurements, Kolter et al. released the REDD dataset [8], which is available at `http://redd.csail.mit.edu`. The initial release of the dataset (version 1.0) contains electricity consumption measurements from six households in the USA collected in April and May 2011. There are approximately 20 consecutive days of measurements available for each house. The REDD dataset provides data from the two main phases of each house at a granularity of one reading per second and measurements from 11 to 26 individual circuits – depending on the house – measured every 3–4 seconds. Some of the circuits contain a single appliance (e.g., a dishwasher) and thus qualify for device-level consumption breakdown. Other circuits contain multiple appliances (e.g., lights, kitchen outlets), which can then only treated as a group of devices by the consumption breakdown algorithm. The REDD dataset represents to date one of the largest and richest publicly available dataset of electricity consumption measurements.

3.4 Performance analysis

To evaluate the performance of Jung and Savvides's algorithm on the REDD dataset, we first extract ON/OFF events from the electricity consumption of each individual circuit. We then apply our simplified version of the disaggregation algorithm on the dataset to obtain the electricity consumption of each appliance.

As a second step, we analyze the estimation accuracy for each appliance using the circuit-level electricity consumption as ground truth information. Next, we investigate the influence of ghost power on the final estimation accuracy by selecting a single base appliance as a test device and progressively filtering out ON/OFF events of selected appliances. Clearly, this procedure does not affect the data relative to the aggregated electricity consumption. The choice of the appliances whose ON/OFF events are removed from the dataset is based on two strategies. The first strategy

Appliance	ON/OFF Events	Mean Power Consumption	Estimation Error
Oven	2	1,991W	12.9%
Fridge	29	193W	-0.9%
Dishwasher	2	552W	15%
KitchenOutlets2	11	56W	-61.8%
KitchenOutlets3	4	89W	12.6%
KitchenOutlets4	2	1,436W	29.6%
Lighting 1	5	90W	39.8%
Lighting 2	5	80W	4.7%
Washer/Dryer	4	1,897W	4.7 %
Microwave	20	1,239W	-4.3%
Bathroom	2	1,525W	21.2%
Total ON/OFF events: 86 Mean of relative errors: 19.8% Relative error weighted by contribution: 14.3%			

Table 1: Estimation of consumption breakdown obtained by analyzing 24 hours of the REDD dataset.

removes ON/OFF events of the appliance with the highest number of state transitions. The second strategy removes ON/OFF events of appliances that consume a large amount of electricity.

For the sake of simplicity and without loss of generality, we report results obtained from the REDD data relative to house 1 (which has a high number of individual circuits) and to April 24, 2011 (a day that exhibits many ON/OFF events). Table 1 shows the results obtained when applying our implementation of Jung and Savvides's disaggregation algorithm to the aggregated electricity consumption data and including all ON/OFF events. The first column lists appliances that provide ON/OFF events during the 24-hours long observation interval. The total electricity consumption is a result of the aggregation of the consumption of each of these devices. The second column shows that 86 state transitions occur during the whole time frame of 24 hours. The column *Mean Power Consumption* shows the average electricity consumption of each device during its ON phase, and the last column denotes the relative error of the estimated average consumption compared to its actual, measured value. The total relative error of 14.3% is obtained by comparing the aggregated values of the estimation and the actual mean electricity consumption of each appliance.

Figure 2 and Figure 3 show the resulting relative error depending on the number of appliances with missing ON/OFF events. Figure 2 illustrates the effect on the estimation error when following the first strategy by removing appliances with the higher number of state transitions. Similarly Figure 3 shows the relative estimation error obtained by removing appliances that consume a large amount of electricity. Both graphs illustrate that consumption from appliances with missing ON/OFF events is virtually spread over the rest of the appliances and thus significantly reduces accuracy of the estimation.

4. PROPOSED IMPROVEMENTS

The results presented in the previous section show that the accuracy of Jung and Savvides's disaggregation algorithm decreases considerably as the number of reported ON/OFF events decreases. However, robustness against missing ON/OFF events is required in order to make the algorithm able to run reliably also in houses that contain unobserved appliances. Unlike existing work [5, 6] we propose to explicitly model power consumption of uninstrumented appliances (i.e. ghost power). To this end, we integrate a *virtual ghost power consumer* into Jung and Savvides' disaggregation model. Acting as an "always-on" appliance it accounts for that unobservable part of the power consumption that would have been otherwise wrongly assigned to other appliances. Similar to the procedure explained in Section 3.4 we apply the extended disaggregation algorithm on the REDD dataset and successively filter out ON/OFF events of appliances. As Figure 4 shows (compared to Figure 2), our version of Jung and Savvides's disaggregation algorithm outperforms the original version in case uninstrumented appliances are present.

Building upon these results we believe that the accuracy can be further improved by estimating power consumption of non-instrumented appliances more accurately. To this end, we plan to integrate the following mechanism into Jung and Savvides's algorithm: (1) Performing approximation over time; (2) Using characteristics of the load curve; (3) Employing state information gathered from sensors or the devices themselves; (4) Sharing consumption patterns. We describe these mechanisms in detail below, while the experimental

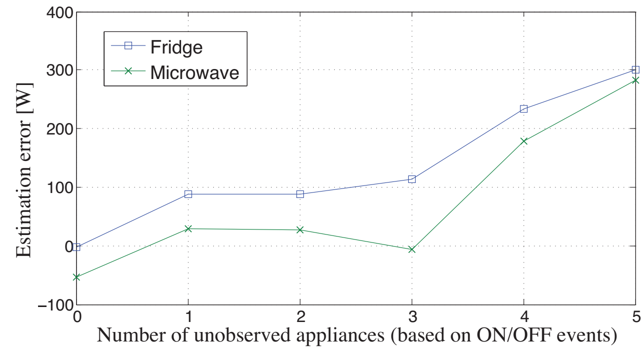

Figure 2: Estimation accuracy based on ghost power of appliances with a large number of ON/OFF events.

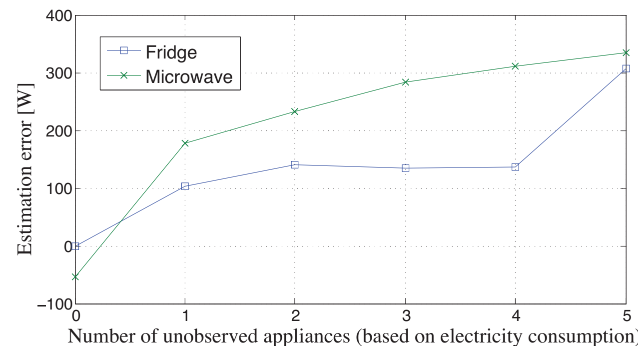

Figure 3: Estimation accuracy based on ghost power of appliances with high electricity consumption.

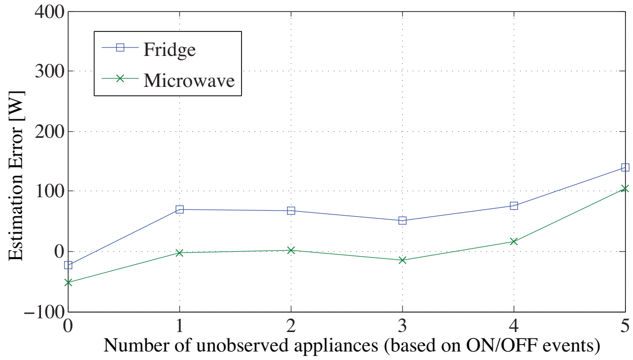

Figure 4: Improved estimation accuracy using a simple model to attribute power consumption of unobserved appliances.

evaluation of the actual gains achieved by such measures in terms of estimation accuracy is planned for future work.

Value of time. Over time it is possible to identify fractions of power consumption that are caused by non-instrumented appliances. This is done by observing specific (instrumented) appliances and discriminating ghost power from actual changes in the appliance's consumption pattern with a certain probability. In particular this works well for steady-state appliances, which exhibit a constant power consumption for a given time interval.

Appliances with a constant power consumption (such as lights) can be used as an indicator for ghost power, in particular when the total power consumption exhibits a large variability while the set of running appliances remains constant. Observing such features allows to estimate the proportion of non-instrumented appliances in a house. Such an estimate then supports calibration of the ghost power estimation itself, which is essential as we learned from our experiments performed on the REDD dataset.

Load curve characteristics. The consumption patterns of many appliances exhibit characteristics such as periodicity (e.g., cooling appliances), a certain change in power when being switched on or off (e.g., lights, kettle), or a particular shape of the load curve once the appliance is running. This information, which is currently not included in Jung and Savvides' algorithm, could contribute to identifying ghost power consumption as follows:

- *Periodicity*: Some appliances (e.g., cooling appliances) exhibit a periodic consumption pattern that can be detected in the load curve. We first derive edges from the power consumption. In case the temporal occurence of these edges correlates with ON/OFF events observed from a sensor that reports such events, we assume that the edges are caused by an appliance whose events are reported by this sensor. Otherwise, we assume that they are caused by an appliance that is not instrumented to report ON/OFF events. In this case some part of the power consumption between these edges can be classified as ghost power.
- *Change in power consumption*: Measuring the increase in power consumption of an appliance that provides ON/OFF events at the time it is being switched on gives information about the device's power consumption right at the beginning of its operation. Similarly, the decrease in power consumption when an appliance is being switched off provides an estimate about the device's power consumption right before being switched off. First, this effect can be employed to estimate ghost power that disturbs the consumption pattern of appliances with a steady power consumption (e.g., a lamp). Second, ghost power caused by appliances with a steady power consumption can be estimated by investigating switching events of appliances that are instrumented to report ON/OFF events.
- *Shape*: Over time, observing the shape of the load curve of a device that provides ON/OFF events reveals certain load curve characteristics. Therefore, changes in the consumption pattern that are not caused by observed appliances are possibly – but not necessarily – evidence for ghost power.

States. Many household appliances exhibit a consumption pattern that is based on the states the appliance is running through, while having a constant power consumption throughout each state.Figure 5 compares the electric consumption of a washing machine at different temperature settings. Here we can distinguish between two main states – the heating of the water and the spinning of the drum. In order to measure the total power consumed by the washing machine accurately, knowing the time it spends in each state provides valuable input to the algorithm. Determining state transitions requires more sophisticated sensors than ON/OFF sensors such as vibration sensors or even network-connected appliances. The choice of these sensors depends on the appliance. Employing knowledge about state transitions of instrumented devices can increase accuracy in the estimation of the power consumption of this appliance. Hence, it improves the accuracy of ghost power estimation and thus the overall accuracy of the consumption breakdown.

Cooperative consumption analysis. In order to differentiate between ghost power and the effects caused by instrumented appliances, comparing consumption patterns with more instrumented households potentially increases accuracy of less instrumented households, since devices might be identified based on their consumption pattern. Sharing consumption patterns of individual appliances comes at the expenses of an increased communication burden as well as potential privacy losses. The trade-off between these costs and the accuracy that is gained from cooperation must be thoroughly analyzed.

Integrating and evaluating these improvements on the REDD dataset highlights what proportion of appliances must be instrumented to provide an electricity breakdown with high accuracy in a real world setting. Thus it contributes to make load disaggregation applicable in a real world environment.

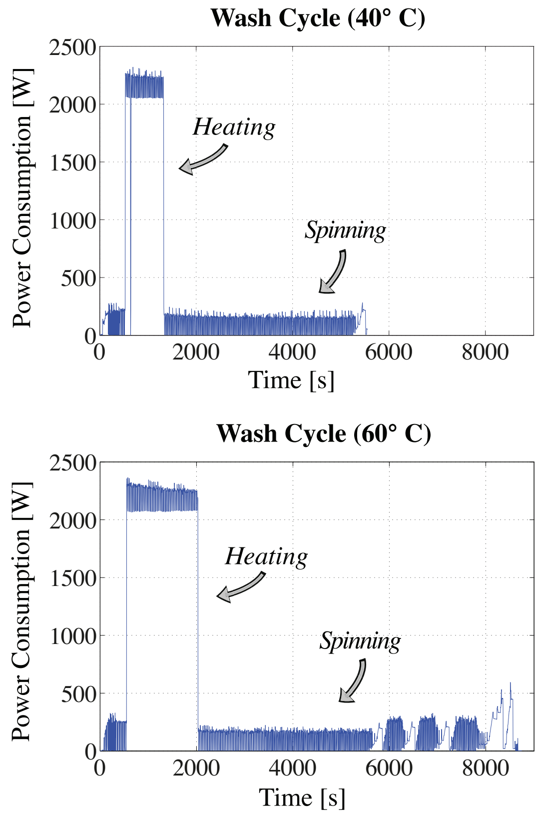

Figure 5: Washing mashine cycles at 40°C (top) and 60°C (bottom).

The information gained from analyzing the REDD dataset further provides the basis for our real world deployment, in which we collect electricity consumption information over a longer time period.

5. CONCLUSIONS

This paper presents the preliminary design of a novel approach to obtain electricity consumption breakdowns in residential settings. Our approach builds and improves upon Jung and Savvides's disaggregation algorithm, which relies on smart meter data and ON/OFF event reports to achieve a reliable consumption breakdown. We provide a quantitative analysis, based on the publicly available REDD dataset, to outline the limitations of Jung and Savvides's algorithm. In particular, we showed that the performance of the algorithm in terms of estimation accuracy quickly decreases as the number of missing ON/OFF events increases. Starting from this observation, we suggest a set of possible improvements, to perform disaggregation with only partial event knowledge.

Future work includes an extensive experimental evaluation of the proposed improvements and an analysis of what proportion – and type – of appliances needs to be instrumented in order to perform electricity consumption breakdown with high accuracy. In addition we investigate the scenario in which ON/OFF events are transmitted unreliably by extending the estimation algorithm to account for losses of events.

Acknowledgements

This work has been partially supported by the Hans L. Merkle Foundation and by the LOEWE research initiative of the State of Hesse, Germany, within the Priority Program Cocoon.

6. REFERENCES

[1] J. Burgess and M. Nye. Re-Materialising Energy Use Through Transparent Monitoring Systems. *Energy Policy*, 36(12):4454–4459, 2008.

[2] C. Fischer. Feedback on Household Electricity Consumption: a Tool for Saving Energy? *Energy Efficiency*, 1(1):79–104, 2008.

[3] S. Gupta, M. Reynolds, and S. Patel. ElectriSense: Single-Point Sensing Using EMI for Electrical Event Detection and Classification in the Home. In *Proceedings of the 12th International Conference on Ubiquitous Computing (UbiComp 2010), Copenhagen, Denmark*, 2010.

[4] X. Jiang, S. Dawson-Haggerty, P. Dutta, and D. Culler. Design and Implementation of a High-fidelity AC Metering Network. In *Proceedings of the 8th International Conference on Information Processing in Sensor Networks (IPSN 2009), San Francisco, CA, USA*, 2009.

[5] D. Jung and A. Savvides. Estimating Building Consumption Breakdowns using ON/OFF State Sensing and Incremental Sub-meter Deployment. In *Proceedings of the 8th Conference on Embedded Network Sensing (SenSys 2010), Zurich, Switzerland*, 2010.

[6] Y. Kim, T. Schmid, Z. Charbiwala, and M. Srivastava. ViridiScope: Design and Implementation of a Fine Grained Power Monitoring System for Homes. In *Proceedings of the 11th International Conference on Ubiquitous Computing (UbiComp 2009), Orlando, FL, USA*, 2009.

[7] W. Kleiminger, S. Santini, and M. Weiss. Opportunistic sensing for smart heating control in private households. In *Proceedings of the 2nd International Workshop on Networks of Cooperating Objects (CONET 2011), Chicago, IL, USA*, 2011.

[8] J. Z. Kolter and M. J. Johnson. REDD: A Public Data Set for Energy Disaggregation Research. In *Proceedings of the KDD Workshop on Data Mining Applications in Sustainability (SustKDD 2011), San Diego, CA, USA*, 2011.

[9] F. Mattern and C. Floerkemeier. From the Internet of Computers to the Internet of Things. In *Kai Sachs, Ilia Petrov, Pablo Guerrero (Eds.): From Active Data Management to Event-Based Systems and More*, volume 6462 of *LNCS*, pages 242–259. Springer, 2010.

[10] F. Mattern, T. Staake, and M. Weiss. ICT for Green – How Computers Can Help Us to Conserve Energy. In *Proceedings of the ACM e-Energy Conference*, 2010.

[11] O. Parson, S. Ghosh, M. Weal, and A. Rogers. Using Hidden Markov Models for Iterative Non-intrusive Appliance Monitoring. In *Machine Learning for Sustainability Workshop (MLSUST 2011), Sierra Nevada, Spain*, 2011.

[12] A. Rowe, M. Berges, and R. Rajkumar. Contactless Sensing of Appliance State Transitions Through

Variations in Electromagnetic Fields. In *Proceedings of the 2nd Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings (BuildSys 2010), Zurich, Switzerland*, 2010.

[13] M. Weiss, T. Graml, T. Staake, F. Mattern, and E. Fleisch. Handy Feedback: Connecting Smart Meters with Mobile Phones. In *Proceedings of the 9th International Conference on Mobile and Ubiquitous Multimedia (MUM 2009), Cambridge, UK*, 2009.

[14] M. Zeifman and K. Roth. Nonintrusive Appliance Load Monitoring: Review and Outlook. *IEEE Transactions on Consumer Electronics*, 57(1):76–84, 2011.