Andreas Reinhardt, Frank Englert, Delphine Christin: *Enhancing User Privacy by Preprocessing Distributed Smart Meter Data*. In: Proceedings of the 3rdIFIP Conference on Sustainable Internet and ICT for Sustainability (SustainIT), pp. 1-7, October 2013.

Enhancing User Privacy by Preprocessing Distributed Smart Meter Data

Andreas Reinhardt*, Frank Englert†, Delphine Christin‡
*School of Computer Science and Engineering, The University of New South Wales, Australia
†Multimedia Communications Lab, Technische Universität Darmstadt, Germany
‡Secure Mobile Networking Lab, Technische Universität Darmstadt, Germany
*andreasr@cse.unsw.edu.au †fenglert@kom.tu-darmstadt.de ‡delphine.christin@cased.de

Abstract

The increasing presence of renewable sources requires power grid operators to continuously monitor electricity generation and demand in order to maintain the grid's stability. To this end, smart meters have been deployed to collect real-time information about the current grid load and forward it to the utility in a timely manner. High resolution smart meter data can however reveal the nature of appliances and their mode of operation with high accuracy, and thus endanger user privacy. In this paper, we investigate the impact on user privacy when the consumption data collected by distributed smart metering devices are preprocessed prior to their usage. We therefore assess the impact on the successful classification of appliances when sensor readings are (1) quantized, (2) down-sampled at a lower sampling rate, and (3) averaged by means of an FIR filter. Our evaluation shows that a combination of these preprocessing steps can provide a balanced trade-off that is in the interests of both users (privacy protection) and utilities (near real-time information).

1. Introduction

The volatile nature of renewal sources requires electric utilities to constantly maintain up-to-date knowledge about generation and load in order to avert the risk of power outages. As a result, many countries have already deployed smart meters widely, or are currently in the process of doing so [1]. While this is of benefit to the utilities, the transmission of precise information about the current activity in people's households is often perceived as a threat to user privacy. This concern is underpinned by research results that have shown that information about the current user activities and even the television content can be inferred based solely on smart meter data (e.g., [2], [3]). So while users may be reluctant to provide high-resolution data to their utilities because of the possible privacy implications, utilities require this consumption data at a high temporal resolution in order to adapt their generation to the changing demand.

The field of privacy-aware data processing has received significant attention in orthogonal domains like participatory sensing ([4], [5]). Due to the different nature of the data collect by smart meters, e.g., the absence of location information, the applicability of such mechanisms is however very limited. Hence, we investigate to which extent preprocessing of the collected power readings can eliminate possibilities to infer appliance types solely based on their consumption data. To this end, we apply different mechanisms to obfuscate the data and analyze to which degree appliance types can still be identified after this preprocessing step. More specifically, we investigate how quantization, down-sampling, and averaging succeed in eliminating characteristic signatures from the data.

Instead of analyzing data that aggregates a complete household's consumption, we herein focus on distributed smart metering. In this scenario, individual metering devices are installed between each appliance's mains plug and the wall outlet. The reasons for selecting this application scenario are twofold. Firstly, existing approaches to infer device activity from smart meter data have shown that the disaggregation of loads performs significantly better when less appliances are connected at the same time [6]. A more efficient privacy protection is thus needed when individual appliances are being monitored. Secondly, very few household-wide meter data sets are available, and these only cover a small number of households (e.g., REDD [7]). In contrast, the Tracebase repository used in this paper contains more than 1,500 appliance power consumption traces, and thus allows for a better

generalization of our results. First, we provide an overview of related work from the domains of data privacy and smart metering in Sec. 2. Subsequently, we describe our designed software framework and the preprocessing steps in more detail in Sec. 3. Our evaluation settings are explained in Sec. 4, followed by the discussion of our evaluation results in Sec. 5. Finally, we conclude this paper in Sec. 6.

2. Related Work

The rise of smart meters has led to the availability of an unprecedented resolution of power consumption readings. To date, two major applications have emerged that rely on these data, namely supporting the utilities to match supply and demand as well as the creation of smart buildings. While smart building functionalities can be realized when accurate measurements are available (cf. [8], [9], [10]), the same methods can be applied by third parties (e.g., the utility or external attackers) to infer the current situation in a building. The CEN-CENELEC-ETSI Smart Grid Coordination Group outlines the information security requirements to the smart grid in [11]. Although proposing a separation of personal information and actual power consumption data, countermeasures to prevent inferring user activities from their meter data are not specifically regarded in the document. In order to protect users from such intrusions into their privacy, several solutions have thus been presented in related work. Cryptographic means to ensure a secure transport of data between end users and utilities are presented in [12], [13], [14]. Although these approaches ensure that third party attackers cannot retrieve the data, they still allow the utility to gain access to the readings in unaltered form. This limitation is addressed in [15] and [16], where solutions are presented that aggregate data collected by multiple meters before relaying it to the utility. While the users are protected against attacks by utilities in this case, they need to trust and cooperate with other households owners.

Approaches that operate locally have also been presented. Efthymiou and Kalogridis have shown that by transmitting smart meter data in an anonymized manner, utilities may be able to infer household activities, but are unable to link them to the actual households [17]. The addition of noise to the measurements can also be applied in smart grids in order to obfuscate user behavior [18], although it has not yet been analyzed in the domain of smart electricity grids. Also operating on a local basis, the privacy-aware data preprocessing step presented in [19] shows that privacy can be increased by applying filters that eliminate certain characteristics from the power meter data, but its efficacy is not analyzed in the domain of smart metering. Finally, instead of manipulating the collected readings, external storage components have been discussed as options to alter a building's consumption and thus eliminate characteristic features from the data. The use of batteries to smooth the load curve has been presented in [20], [21], but the limitations of state-of-the-art battery technology, e.g., decreasing capacities and high financial cost, render this technology inapplicable for many scenarios.

3. Concept and Software Framework

The primary objective of this paper is to evaluate the extent of privacy protection that can be achieved by preprocessing the data collected by distributed smart meters. In order to analyze the efficacy of this preprocessing, we first quantify the privacy threat resulting from the unprocessed transmission of power consumption data. Subsequently, results based on preprocessed data are compared to this baseline in order to draw reliable conclusions on the degree of additional protection attained by preprocessing. To establish the baseline detection accuracy, we have thus designed a system that is able to detect the type of an appliance based on its electric power consumption data. The system extracts specific characteristics that uniquely represent each appliance type based on its power consumption behavior, and memorize them in the form of a machine learning model. When the system is supplied with a power consumption trace collected from another device, it extracts the characteristic features from the trace, compares them to the knowledge stored in its model, and returns the device type with most similar characteristics. The objective of this paper, namely obfuscating device-specific characteristics in the power consumption data, should thus lead to larger number of false identifications. Hence, we use the fraction of appliances that can no longer be correctly identified as a measure of the efficacy of our data preprocessing.

3.1. Overall System Architecture

Our overall system is composed of distributed metering units that connect between the wall outlet and an electric appliance, as well as a server on which the data analysis is performed. This architecture is visualized in

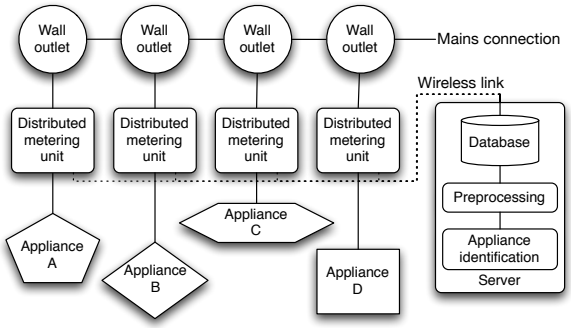

Figure 1. Overall system architecture

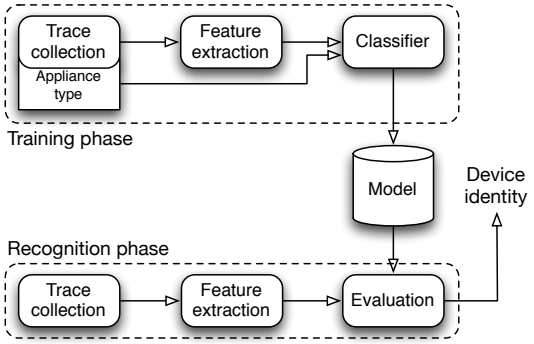

Figure 2. Appliance classification architecture

Fig. 1. Continuous lines indicate mains connections, whereas dashed lines reflect the wireless data transfer between the meters and the server. All metering units return their data once per second to the server, which records the power consumption traces in its database for their subsequent classification. Before feeding the resulting stream of power consumption data into the application identification component, we apply different preprocessing steps to the data, as outlined in Sec. 3.2. Subsequently, the appliance identification step extracts representative features from the data stream and uses a machine learning component to facilitate the classification of incoming data. Feature extraction and machine learning are detailed in Sec. 3.3.

3.2.Data Preprocessing

For our evaluation, we have selected three mechanisms to alter the data prior to their use for appliance identification. We explain them as follows and visualize their impact on an excerpt from a dishwasher's operation cycle, which is depicted in Fig. 3a.

Quantization. Value quantization is realized by rounding the actual power consumption values to a multiple of a pre-defined quantization factor q. Because the quantization step is stateless and requires no historical data, no delay is introduced by the introduction of this preprocessing step. The application of quantization to the dishwasher's consumption data is shown in Fig. 3b for $q = 80$ watts. It can be seen that quantization eliminates the slight slope on top of the power-intensive heating periods while the general shape is maintained.

Down-Sampling. This second preprocessing option reduces the temporal frequency at which measurements are made available by returning a sample of the actual power consumption only every w seconds. For all further samples until the next sampling point, the previously transmitted value is repeated instead. Like the quantization step, down-sampling does not introduce a delay, and Fig. 3c shows the output of the down-sampling step for $w = 150$ seconds.

Averaging. The third preprocessing alternative is the combination of an averaging of the input data over a time window of w seconds and its down-sampling according to the previous paragraph. The averaging is thus equivalent to an FIR filter with rectangular filter function. In contrast to the previously described preprocessing steps, averaging introduces a time lag of w seconds, and may thus only be applicable in scenarios where this can be tolerated by the utility. The output of our averaging preprocessor for $w = 150$ seconds is shown in Fig. 3d, from which the combination of averaging and down-sampling manifests itself in the form of steps on the steep edges of the power consumption curve.

3.3.Classification and Features

The actual classification is based on our previous appliance classification framework [22], which we briefly revisit as follows. Based on the overall process flow shown in Fig. 2, power consumption traces of 24 hours duration are first collected from electric appliances. In a subsequent step, characteristic features are extracted from each of the traces and stored in the form of a feature vector that is annotated by the actual appliance type. Similar to [9] and [23], our system utilizes more than 500 different features from different domains in order to describe the characteristic properties of the power consumption traces. We regard features from both the temporal and frequency domain in order to incorporate both the sudden changes

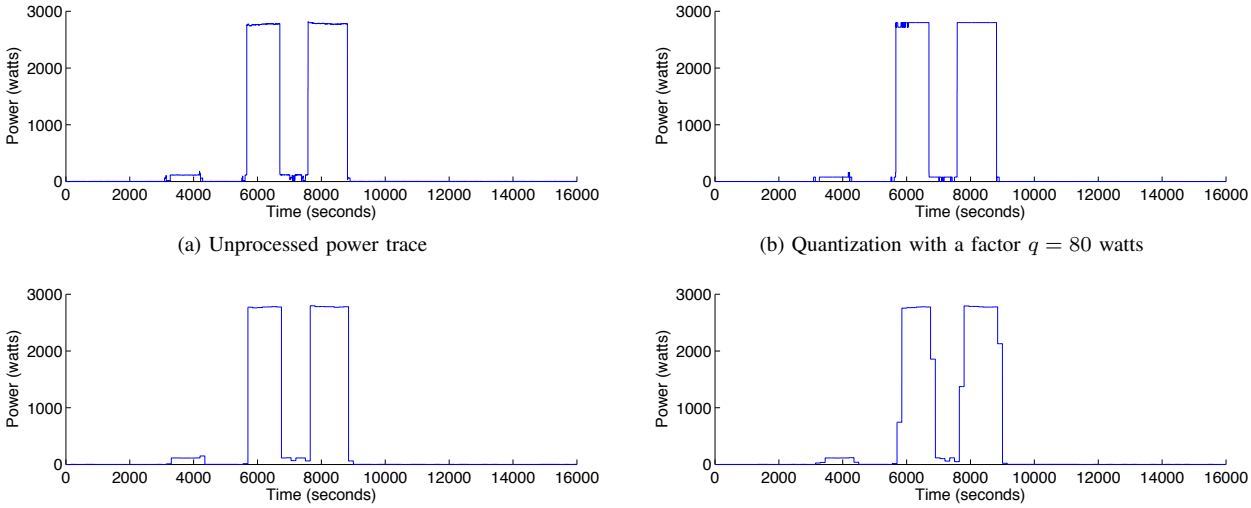

(a) Unprocessed power trace

(b) Quantization with a factor $q = 80$ watts

(c) Down-sampling with a window size of $w = 150$ seconds

(d) Averaged down-sampling with a window size of $w = 150$ seconds

Figure 3. Visualization of a dishwasher's consumption data before and after applying preprocessing steps

encountered on appliance activation as well as periodicities throughout the day into our classification model. We combine many different aspects of an appliance's power consumption (e.g., peak values and shape) to reach a high number of correct classification results. Each of the resulting annotated feature vectors is subsequently forwarded to the machine learning component, in which a classifier constructs its model based on the data.

Our previous results have shown that classification accuracy values in excess of 90% could be achieved when all of the presented features were being used for the appliance classification [22]. In other words, a very large fraction of the input data (composed of more than a thousand appliance traces) could be correctly classified solely based on their power consumption data throughout a day. In our evaluations we have demonstrated that maximum and average power consumption values are the most important features for the classification of appliances. Based on this observation, we have specifically chosen to preprocess the power consumption data in a way that alters the consumption characteristics and analyze its impact on the classification accuracy.

While our previous work has thus effectively promoted *anti-privacy* by identifying the types of electric appliances, we address the opposite target in this paper, namely how data preprocessing can render our appliance identification system ineffective.

4. Evaluation Setup

Our evaluation is based on the software system presented in Sec. 3. We have installed the server components on a dedicated machine that maintains the database, the preprocessing modules, and the appliance identification engine. For the construction of the classification model, we have used the Weka data mining toolkit [24]. Based on the comparison of different classifiers in our previous work, we have chosen to use the Random Forest classifier for the machine learning step, as it has been shown to result in both a high classification accuracy for the task at hand and a fast execution time.

The data for the classification has been taken from our Tracebase project [22]. The Tracebase already features more than 1,200 diurnal power consumption traces of more than 30 household appliance types. Furthermore, we have collected more than 300 additional traces in order to base our evaluation on an even larger corpus of data. On average, the power consumption traces have been collected at a high granularity of one sample per second and with a value resolution of one watt. We list the appliance types, the number of different instances, and the total number of traces used in our evaluation in Table 1.

In order to put the achieved device classification results into perspective, we compare them to the baseline, in which no preprocessing steps are applied (i.e.,

Table 1. Power traces used in our evaluation

Device type	# appliances	# traces
Alarm clock	1	5
Bean-to-cup coffee maker	1	43
Bread cutter	1	12
Coffee maker	5	77
Cooking stove	1	16
Desktop computer	9	126
Dishwasher	3	65
Ethernet switch	3	11
External USB hard disk drive	4	29
Freezer	1	9
HDTV media center	1	5
HiFi stereo amplifier	3	88
Internet router	1	20
Iron	1	3
Lamp	6	77
Laptop computer	6	50
Microwave oven	5	51
Monitor (CRT)	2	14
Monitor (TFT)	14	178
Playstation 3 console	2	12
Powered USB hub	1	10
Printer	1	6
Projector	1	8
Refrigerator	8	189
Solar-thermal system	1	8
Subwoofer	2	28
Television set	10	138
Toaster	4	21
Tumble dryer	2	9
Vacuum cleaner	1	1
Video projector	1	19
Washing machine	7	50
Water fountain	1	56
Water kettle	8	115
Xmas lights	1	6
Total	119	1,555

the parameters are chosen as q=1 watt, w=1 second, no averaging). Subsequently, we conduct a comprehensive analysis of the classification accuracy when varying the parameter values for q and w. We regard down-sampling window sizes of w=1...400 seconds and analyze quantization factors between q=1...180 watts. Additionally, we consider the case when data are averaged before the down-sampling and quantization steps are applied. For this case, we have also used the down-sampling window size parameter w as the averaging filter's window size.

5. Evaluation

We have conducted several evaluations in order to quantify the improvements to user privacy protection offered by the presented preprocessing steps. After determining the bounds for the classification success rates, we thus present the results of our comprehensive analysis of the parameter space and quantify the error that is added to the data.

5.1. Baseline Values

In order to put the evaluation results into perspective, we have first evaluated the baseline detection accuracy for the input data as listed in Table 1. In this case, the application identification component has returned an achievable accuracy value of 90.5%, i.e., nine out of ten devices could be correctly identified solely based on their power consumption. Likewise, the worst classification result is equal to the random selection of an appliance class, and can thus be calculated as $1/\#appliances$. For the given input set of 35 appliance types, the minimum accuracy thus equals 2.9%.

5.2. Quantization and Down-sampling

In this first evaluation step, we analyze the impacts of quantization and down-sampling only. As both preprocessing steps influence the classification results, we conduct a two-dimensional analysis for multiple combinations of down-sampling and quantization factors within the parameter ranges specified above. We show the results in Fig. 4a. Note that for a quantization factor of q=1 watt and a down-sampling window w=1 second, the classification accuracy is equal to the 90.4% reached in the case without preprocessing, as outlined above. While both approaches show a significant degradation of the classification accuracy already for small values of q and w, their behavior for larger values differs. More precisely, quantization achieves a degradation of the detection accuracy from 90.5% to 58.1% when a factor q=180 watt is chosen. In contrast, the impact of down-sampling in the temporal domain does not manage to reduce the classification accuracy below 73.8%, even for a window size of w=400 seconds. The extremal combination of both quantization and down-sampling with w=400 seconds and q=180 watts leads to an overall classification accuracy of 37.7%.

5.3. Quantization, Down-sampling & Averaging

While the down-sampling in the previous evaluation has taken a sample from the actual data every w seconds, we investigate the impact of its combination with averaging next. Again, we conduct a two-dimensional analysis, in which we vary the quantization factor q as well as the averaging and down-sampling window size w. Both preprocessing steps share the same window size, so effectively the average

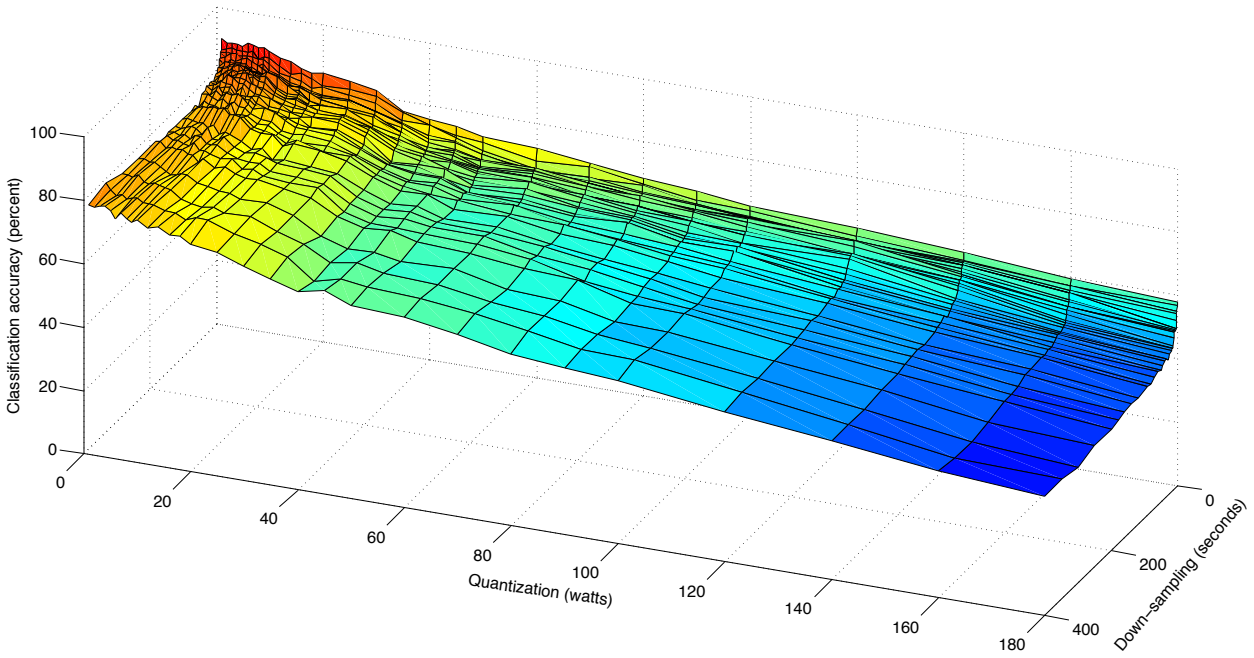

(a) Classification accuracy without averaging filter

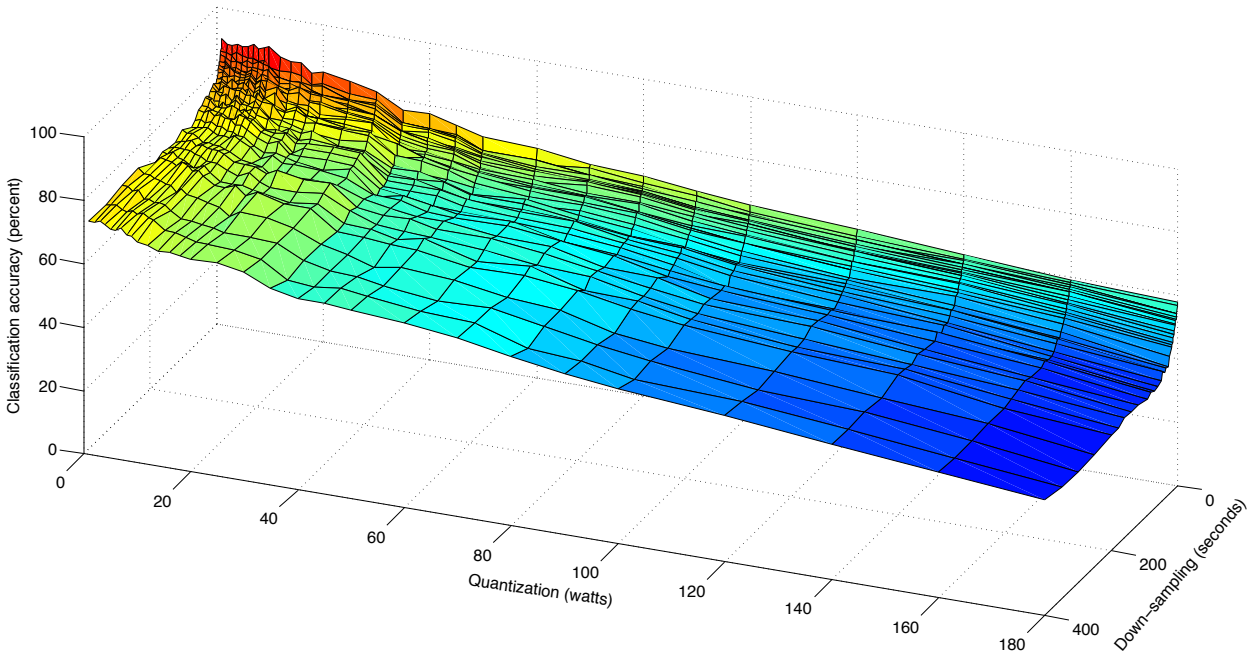

(b) Classification accuracy when an additional averaging filter is applied

Figure 4. Classification accuracy values for variations of quantization factor, down-sampling window, and averaging method

value of the last w seconds is reported for a duration of w seconds. The fact that the averaging component integrates over the actual power consumption in the time domain leads to the expectation of less privacy-compromising details to be contained in the trace. The results for this evaluation are visualized in Fig. 4b. Despite the averaging, however, the figure shows a very similar behavior to the previous analysis without averaging. In fact, the quantization step again achieves to degrade the detection accuracy to 58.1% when a factor q=180 watt is chosen. When averaging is used, however, down-sampling does not even manage achieve the degradations of the previous experiment, but always stays above 78.9%. Again, the using the maximal values of w=400 seconds and q=180 watts leads to an accuracy of slightly above 37.5%.

5.4.Quantization and Down-Sampling Errors

By applying any of the presented preprocessing steps, the signal is altered from its original form (cf. Fig. 3). As the transmitted smart meter readings are possibly used for the capacity planning of utilities, but might significantly differ from the actual readings due to the preprocessing, we complete our evaluation with an analysis of the error introduced by our preprocessing steps. We therefore determine the RMS error P_{RMS} between the original and the preprocessed power consumption traces. The results for three devices of different operating power ranges are shown in Fig. 5. In essence, they indicate that the quantized values show a comparably small difference to the original sequence, whereas down-sampling leads to more significant discrepancies and might hence be less favorable for electric utilities.

5.5.Discussion

The results in Fig. 4 show that increasing the parameters for both down-sampling and quantization (i.e., w and q) lead to increased levels of privacy protection. Their impact however differs, as down-sampling with a window size w=400 seconds only succeeds in reducing the classification accuracy by approximately 16 percentage points. In contrast, value quantization has been shown to have a significantly better performance in terms of privacy protection and succeeded in degrading the classification accuracy by up to 53 percentage points for large values of q. Averaging the consumption traces did not have any measurable impact on the privacy protection, and even led to a slightly worse privacy protection. In other words, the privacy-preserving effect of down-sampling is effectively reduced by prepending it with an averaging filter.

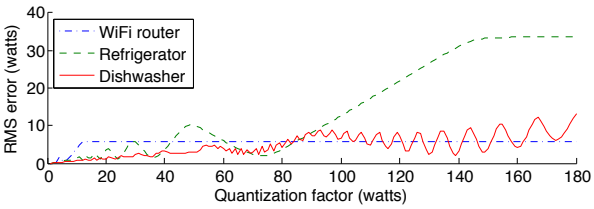

(a) RMS error for different quantization factors

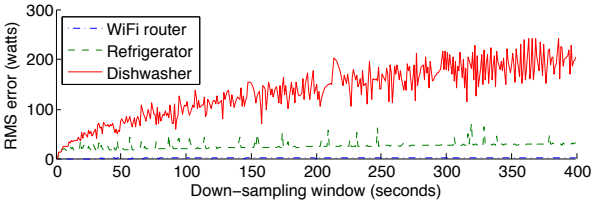

(b) RMS error for different down-sampling window sizes

Figure 5. RMS error analysis

6. Summary and Conclusions

In this paper, we have analyzed how preprocessing distributed smart metering data can decrease the possibility of attributing a power consumption trace to the underlying electric appliance. To this end, we have studied the impact of quantization, down-sampling, and averaging on more than 1,500 daily power consumption traces. More than 500 characteristic features were extracted from the traces, which allowed for the correct classification of appliances at 90.4% accuracy when no preprocessing was applied. After the application of the presented preprocessing steps, however, the detection accuracy experienced a measurable degradation. For example, quantization to q=40 watts and down-sampling to w=40 seconds already lead to a situation in which 30% less appliances could be detected properly, whereas the average error was still comparably small. In terms of privacy protection, value quantization has been shown to lead to better results than down-sampling or averaging alone. While keeping the error between the actual and the recorded data within definable bounds, its application to distributed smart metering data can strongly enhance user privacy. As a general result, our comprehensive study enables

application designers to carefully choose the required trade-off between timeliness, intentional inaccuracy, and privacy protection.

References

[1] R. Hierzinger, M. Albu, H. van Elburg, A. J. Scott, A. Łazicki, L. Penttinen, F. Puente, and H. Sæle, "European Smart Metering Landscape Report," Online: http://www.smartregions.net/, 2012.

[2] A. Molina-Markham, P. Shenoy, K. Fu, E. Cecchet, and D. Irwin, "Private Memoirs of a Smart Meter," in *Proc. BuildSys*, 2010, pp. 61–66.

[3] U. Greveler, B. Justus, and D. Loehr, "Multimedia Content Identification Through Smart Meter Power Usage Profiles," in *Proc. CPDP*, 2012.

[4] C. Cornelius, A. Kapadia, D. Kotz, D. Peebles, M. Shin, and N. Triandopoulos, "AnonySense: Privacy-aware People-centric Sensing," in *Proc. MobiSys*, 2008, pp. 211–224.

[5] D. Christin, A. Reinhardt, S. S. Kanhere, and M. Hollick, "A Survey on Privacy in Mobile Participatory Sensing Applications," *Journal of Systems and Software*, vol. 84, no. 11, pp. 1928–1946, 2011.

[6] A. Marchiori, D. Hakkarinen, Q. Han, and L. Earle, "Circuit-Level Load Monitoring for Household Energy Management," *IEEE Pervasive Computing*, vol. 10, no. 1, pp. 40–48, 2011.

[7] J. Z. Kolter and M. J. Johnson, "REDD: A Public Data Set for Energy Disaggregation Research," in *Proc. SustKDD*, 2011.

[8] M. Berges, E. Goldman, H. S. Matthews, and L. Soibelman, "Enhancing Electricity Audits in Residential Buildings with Nonintrusive Load Monitoring," *Journal of Industrial Ecology*, vol. 5, no. 14, 2008.

[9] J. Liang, S. K. K. Ng, G. Kendall, and J. W. M. Cheng, "Load Signature Study — Part I: Basic Concept, Structure, and Methodology," *IEEE Transactions on Power Delivery*, vol. 25, no. 2, 2010.

[10] M. Kazandjieva, O. Gnawali, B. Heller, P. Levis, and C. Kozyrakis, "Identifying Energy Waste through Dense Power Sensing and Utilization Monitoring," Stanford University, Tech. Rep. CSTR 2010-03, 2010.

[11] CEN-CENELEC-ETSI Smart Grid Coordination Group, "Smart Grid Information Security," Online: http://ec.europa.eu/energy/gas_electricity/smartgrids/doc/xpert_group1_security.pdf, 2012.

[12] A. R. Metke and R. L. Ekl, "Security Technology for Smart Grid Networks," *IEEE Transactions on Smart Grid*, vol. 1, no. 1, pp. 99–107, 2010.

[13] F. D. Garcia and B. Jacobs, "Privacy-Friendly Energy-Metering via Homomorphic Encryption," in *Proc. STM*, 2010, pp. 226–238.

[14] A. Rial and G. Danezis, "Privacy-Preserving Smart Metering," in *Proc. WPES*, 2011, pp. 49–60.

[15] L. Sankar, S. Kar, R. Tandon, and H. V. Poor, "Competetive Privacy in the Smart Grid: An Information-theoretic Approach," in *Proc. SmartGridComm*, 2011.

[16] K. Kursawe, G. Danezis, and M. Kohlweiss, "Privacy-Friendly Aggregation for the Smart-Grid," in *Proc. PETS*, 2011, pp. 175–191.

[17] C. Efthymiou and G. Kalogridis, "Smart Grid Privacy via Anonymization of Smart Metering Data," in *Proc. SmartGridComm*, 2010, pp. 238–243.

[18] V. Rastogi and S. Nath, "Differentially Private Aggregation of Distributed Time-Series with Transformation and Encryption," in *Proc. SIGMOD*, 2010, pp. 735–746.

[19] S. R. Rajagopalan, L. Sankar, S. Mohr, and H. V. Poor, "Smart Meter Privacy: A Utility-Privacy Tradeoff Framework," in *Proc. SmartGridComm*, 2011, pp. 190–195.

[20] G. Kalogridis, C. Efthymiou, S. Denic, T. Lewis, and R. Cepeda, "Privacy for Smart Meters: Towards Undetectable Appliance Load Signatures," in *Proc. SmartGridComm*, 2010, pp. 232–237.

[21] W. Yang, N. Li, Y. Qi, W. Qardaji, S. McLaughlin, and P. McDaniel, "Minimizing Private Data Disclosures in the Smart Grid," in *Proc. CCS*, 2012, pp. 415–427.

[22] A. Reinhardt, P. Baumann, D. Burgstahler, M. Hollick, H. Chonov, M. Werner, and R. Steinmetz, "On the Accuracy of Appliance Identification Based on Distributed Load Metering Data," in *Proc. SustainIT*, 2012, pp. 1–9.

[23] C. Laughman, K. Lee, R. Cox, S. Shaw, S. Leeb, L. Norford, and P. Armstrong, "Power Signature Analysis," *IEEE Power and Energy Magazine*, vol. 1, no. 2, 2003.

[24] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten, "The WEKA Data Mining Software: An Update," *SIGKDD Exploration Newsletter*, vol. 11, no. 1, 2009.