Preprint of: Andreas Reinhardt, Delphine Christin, Salil S. Kanhere: *Predicting the Power Consumption of Electric Appliances through Time Series Pattern Matching*. In: Proceedings of the 5th ACM Workshop on Embedded Systems for Energy-Efficient Buildings (BuildSys), pp. 1-2, ACM Press, November 2013.

Predicting the Power Consumption of Electric Appliances through Time Series Pattern Matching

Andreas Reinhardt
The University of New South Wales
Sydney, Australia
andreasr@cse.unsw.edu.au

Delphine Christin
Technische Universität Darmstadt
Darmstadt, Germany
delphine.christin@cased.de

Salil S. Kanhere
The University of New South Wales
Sydney, Australia
salilk@cse.unsw.edu.au

ABSTRACT

We present a system to forecast the power consumption of electric household appliances. Accurate load prediction has numerous application domains, e.g., the facilitation of peak load prediction at a much higher resolution than permitted by state-of-the-art load profiles. Our solution is based on the identification and isolation of representative characteristic signatures from previously collected power consumption traces. Subsequently, time series pattern matching is applied to detect these signatures in real-time data, and emit predictions of an appliance's future consumption based thereupon. We evaluate the prediction accuracy of our approach with thousands of device-level power consumption traces and highlight the achievable prediction horizon.

Categories and Subject Descriptors

C.3 [**Special-purpose and Application-based Systems**]: Real-time and Embedded Systems; G.3 [**Probability and Statistics**]: Time Series Analysis; I.5.4 [**Pattern Recognition**]: Applications—*waveform analysis*

1. INTRODUCTION

Many electrical appliances show a recurring pattern of power consumption throughout their operational phases. We leverage this property and present a system that predicts an appliance's future power consumption based on its previous power consumption data. First, the continuous time series of past consumption data is divided into *segments*, i.e., time series representations of the appliance's power consumption during individual phases of its operation. Because each segment has a finite duration, all further extracted characteristics can subsequently relate to the start or end of the segment. In a second step, short characteristic power consumption snippets (*signatures*) are extracted from the segments. These signatures are later matched against real-time data in order to make predictions. All signatures are annotated by their distance to the end of the underlying segment and the corresponding power consumption behavior until this point.

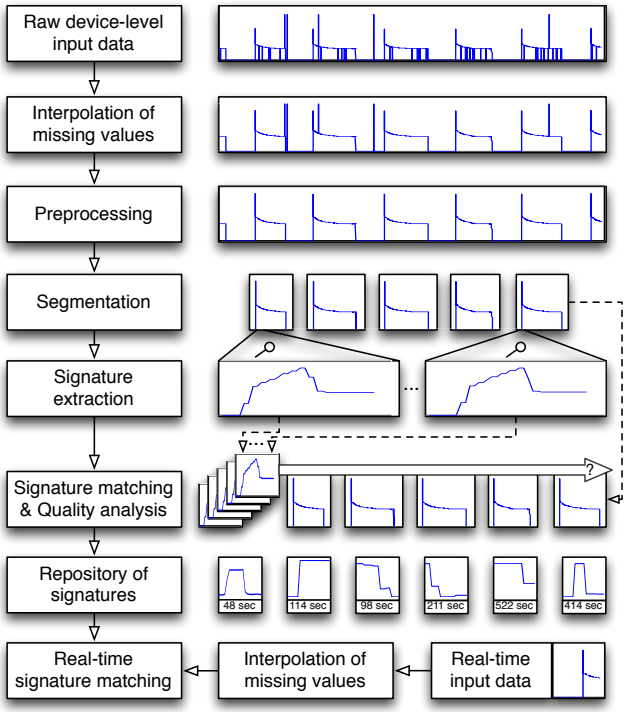

Figure 1: Flow of the data processing

Before relying on the extracted signatures to make predictions, we must verify if they can unambiguously indicate the remaining operational time and power consumption. Their extraction is thus followed by a subsequent quality analysis, which eliminates all signatures that have either resulted in incorrect predictions of deactivation time or power consumption until then. Only the remaining unambiguous signatures are eventually stored in a repository, based on which the future consumption behavior of real-time collected data can be forecast. The complete sequence of the steps taken during the system's initial training phase, including the required preprocessing of the input data, is visualized in Fig. 1 and described in the following section.

During the system's regular operation, incoming power consumption data is matched against the signature repository established in the previous step. Due to the real-time nature of the input data, however, preprocessing steps that require knowledge of future data points cannot be applied. As a result, only data interpolation is used to ensure that the temporal resolution of signatures and input traces match.

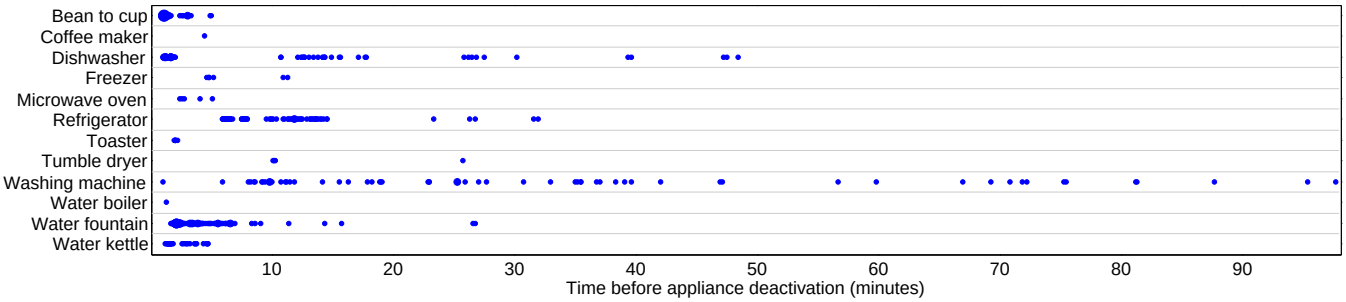

Figure 3: Achievable prediction horizon for selected device types (marker size indicates occurrence frequency)

2. DATA PROCESSING

In order to quantify the achievable accuracy of the predictions in the optimum case, we conduct our analysis with device-level consumption data. We source our input data from the Tracebase [1] and process it as shown in Fig. 1. First, all traces are interpolated to a sampling rate of one sample per second, i.e., 86,400 samples per day, in order to operate on an identical temporal resolution. A following preprocessing step removes very short power consumption bursts and eliminates outliers introduced by the measurement devices. Subsequently, individual activity segments are extracted from the traces; periods of standby power consumption are thus inherently disregarded. The final step of our data processing chain is the extraction of signatures from the segments. All our extracted signatures are comprised of an event, i.e., a notable change to the power consumption, surrounded by additional consumption samples. These additional readings are used to improve the reliability of a signature match. All extracted signatures are annotated by their distance to the end of their originating activity segment as well as the corresponding power consumption.

Before the extracted signatures can be used to reliably forecast appliance deactivations, they must be confirmed to be unambiguous. To this end, we have matched all extracted signatures against all activity segments for each appliance type. Only signatures that correctly indicate the appliance's deactivation time (±30 seconds) as well as making an accurate prediction of the power consumption are regarded further. As an exact match is virtually impossible to achieve due to measurement and quantization errors, discrepancies were allowed up to a threshold of ±5W.

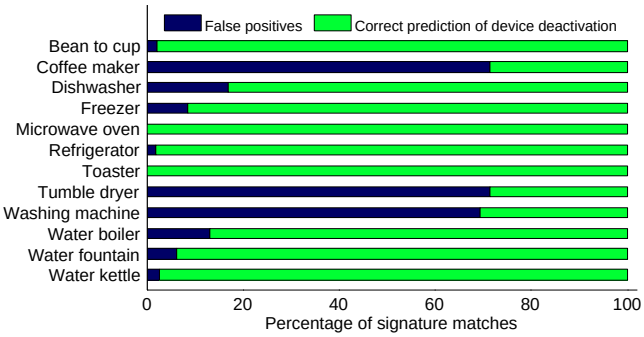

Figure 2: Percentage of correct predictions per type

3. PREDICTION ACCURACY & HORIZON

We consider twelve household appliance types in our evaluation. First, we assess the number of erroneous predictions, which can occur because signatures from all individual appliance types have been combined into a single repository. The results are visualized in Fig. 2 and show that many appliance types (e.g., refrigerators) feature a set of highly unique signatures that allow for a large number of correct predictions. The similarity of signatures of some appliance types, however, leads to incorrect predictions for these devices when using the combined signature set.

In a second evaluation, we show the achievable prediction horizon (cf. Fig. 3). For this evaluation, only the correct predictions from the previous experiment have been used and predictions less than one minute prior to the appliance's deactivation were omitted for the sake of visual clarity. The figure confirms that, based on the recurring power consumption patterns observed across different operating cycles, power predictions for extended periods of time, up to more than 90 minutes, can be made.

4. CONCLUSIONS

We have analyzed to which extent the power consumption of electric appliances can be predicted solely based on historical consumption data. Our approach is based on the extraction of signatures, i.e., potential indicators for future appliance behavior, from an appliance's past power consumption characteristics. The presence of these signatures in subsequently collected consumption data is then used to forecast the remaining duration of an appliance's operation. We have shown that predictions can be made more than 90 minutes in advance. The solution is been based on device-level measurements instead of using circuit-level or even meter-level data. We have deliberately chosen this approach due to the availability of a comprehensive data set which has allowed us to extract the signatures for each device individually. The resulting model can, however, also be applied at circuit or meter level, permitted that suitable means to extract signatures from the aggregated load exist.

5. REFERENCES

[1] 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 *Proceedings of the 2nd IFIP Conference on Sustainable Internet and ICT for Sustainability (SustainIT)*, pages 1–9, 2012.