

How to Auto-Configure Your Smart Home? High-Resolution Power Measurements to the Rescue

Frank Englert¹, Till Schmitt¹, Sebastian Kößler¹, Andreas Reinhardt²,
Ralf Steinmetz¹

¹ Multimedia Communications Lab
Technische Universität Darmstadt
Darmstadt, Germany
{frank.englert, till.schmitt,
sebastian.koessler, ralf.steinmetz}
@kom.tu-darmstadt.de

² School of Computer Science and Engineering
The University of New South Wales
Sydney, Australia
andreasr@cse.unsw.edu.au

ABSTRACT

Most current home automation systems are confined to a timer-based control of light and heating in order to improve the user's comfort. Additionally, these systems can be used to achieve energy savings, e.g., by turning the appliances off during the user's absence. The configuration of such systems, however, represents a major hindrance to their widespread deployment, as each connected appliance must be individually configured and assigned an operation schedule. The detection of active appliances as well as their current operating mode represents an enabling technology on the way to truly smart buildings. Once appliance identities are known, the devices can be deactivated to save energy or automatically controlled to increase the user's comfort.

In this paper, we propose an approach to have buildings informed about the presence and activity of electric appliances. It relies on distributed high-frequency measurements of electrical voltage and current and a feature extraction process that distills the collected data into distinct features. We utilize a supervised machine learning algorithm to classify readings into the underlying device type as well as its operation mode, which achieves an accuracy of up to 99.8%.

Categories and Subject Descriptors

B.m [Hardware]: Miscellaneous; H.4.m [Information Systems]: Miscellaneous

Keywords

Smart Home, Device Classification, High Accuracy, Operating Mode Identification, Harmonic Spectrum Analysis

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

e-Energy'13, May 21–24, 2013, Berkeley, California, USA.
Copyright 2013 ACM 978-1-4503-2052-8/13/05 ...\$15.00.

1. INTRODUCTION

The rise of home automation systems has greatly improved the comfort in modern homes. Current systems are, however, generally confined to controlling temperature [14] and lighting settings and they do not yet offer complete building automation. Besides increasing the user comfort, various further functionalities can be realized, e.g., detecting unexpected behavior to realize building security or deactivating appliances to achieve energy savings. In recent years, numerous researchers have addressed the latter issue. Example applications include to cut off standby loads [9], trim appliances to the most energy efficient setting [21], defer the usage of devices [13] until the grid-wide energy demand is low, or infer information about the user activities at home [1].

However, an information gap exists between the building automation system and the electric appliances, as the system is not aware of available devices, their location, and their state. In order to realize such functionalities without manual configuration of all appliances and their operation state, the building automation system needs to autonomously acquire this information. Only when detailed information about appliances and their mode of operation are known, can the system exert control over the electric appliances and actuate them to save energy and increase the user comfort.

In this paper, we present our solution to this problem, consisting of a hardware board for high resolution power measurements and a software framework for classifying the connected appliance and its state. The hardware board is carefully designed to achieve galvanic insulation between mains and the line-level output, as well as maintaining a low noise floor and a flat frequency response. In order to fulfill these properties, we have conducted an evaluation of different current transducers to select the best match for our use case. The resulting hardware board outputs line-level signals for both voltage and current. These signals can be simply converted to the digital domain by a computer system's audio interface or a dedicated analog-to-digital converter to process them further with our proposed software framework. This setup allows us to sample the power consumption of a connected device with up to 96 kHz sample rate.

Our software framework analyses the current and voltage waveform obtained from the connected appliance, and ex-




Figure 1: A schematic system overview: Our hardware board measures the voltage and the current of the load. The resulting signals are then converted to the digital domain and post processed by a machine learning toolkit.

tracts different features from the shape of the voltage and current consumption curves. Those features are used to build up two kinds of models, namely (1) a coarse-grained model to detect the appliance type attached as well as (2) a fine-grained model to differentiate between products of an appliance type and also to determine its mode of operation. Splitting the classification step into two stages allows to extend the list of supported devices dynamically. Our software framework is specifically tailored to cope with high data rates obtained from the hardware board. The system is designed for low latencies: Typically it takes at most one second for recording, feature extraction, and classification. Our approach requires no disaggregation algorithm because each device has its own metering unit.

This work closes the information gap between existing building automation systems and the electric appliances available in a household. By inferring type and state information, our system allows building automation systems to control electric appliances, e.g., in order to reduce their standby power consumptions. We make the following contributions:

- We present our hard- and software design for the device classification and operation mode determination system.
- We evaluate the classification accuracy values for the device type identification based on a comprehensive collection of traces.
- We assess the accuracy of the operation mode determination based on harmonic component fingerprints.

This paper is structured as follows. We discuss work related to the field of device classification and operation mode detection in Section 2, and subsequently introduce our general concept in Section 3. The hardware design for our power measurement device is detailed in Section 4, which also summarizes our comprehensive analysis of several current transducers in order to achieve a high spectral resolution and low noise levels. In the subsequent Section 5, we present our PISI software system, which extracts the features from the

collected traces and inserts them into a machine learning model. We evaluate the detection sensitivity and the overall classification accuracy in Section 6 where we also discuss the results. Finally, we conclude this paper in Section 7.

2. RELATED WORK

Numerous approaches exist to classify electric appliances based on their power consumption. These approaches can be divided into three different groups: Non-Intrusive approaches meter the current draw at a central position and then apply a disaggregation step to separate between different devices. In contrast Intrusive approaches provide a measuring unit for each household appliance. Finally integrated approaches require the household appliance to contain a metering and communication module. In this section we will describe these three different groups and discuss their advantages and disadvantages.

Centralized approaches meter a household’s whole electricity consumption in order to detect attached appliances. The first who introduced such a solution was Hart in 1992. His Non Intrusive Appliance Monitoring paper [7] described a system consisting of a power meter for real and reactive power measurement and an attached state machine to track steps in the power consumption. Since then many researchers picked up his idea and improved it by sampling with a higher frequency [16, 17], better features [20, 5] and better disaggregation algorithms [2, 10, 15]. The main advantage of such a system is the requirement to place only one metering unit at a central position next to the fuse box. But as Zeifman summarizes in [24], centralized approaches are known to suffer from problems in detecting low-power appliances and devices with variable power consumption. Also the centralized approaches only allow the monitoring of currently running appliances and not controlling them from remote.

Alternatively distributed smart meters could be used to measure the power consumption at appliance level. One of the first researchers to explore this field was Ito [8]. He built up a system for classifying devices connected to a special

wall outlet based on their power consumption. Many researchers improved his system by using better features [22, 10] or a wireless sensor network for metering devices [11, 4]. Kim even used different audio, light, temperature, and vibration sensors in combination with power meters to classify the currently running appliances [12]. As a main advantage, this solution does not rely on an error-prone load disaggregation method to split the power trace of a whole household into power traces for single devices. Furthermore, the distributed smart meters can control the attached appliances. But on the other hand each appliance requires its own metering unit. To the best of our knowledge, none of these systems can classify the state of an attached appliance.

There is a third field of related work namely the fact that an increasing number of business and household appliances will be equipped with networking capabilities in the future. If those appliances expose an interface to their functionality they become controllable from remote. A smartphone application to control home and office appliances was shown by Nichols [19]. Other researchers developed smart home servers to control such appliances from the Web [18]. Although those systems would make it possible to forward the internal state of appliances to a smart home system without the necessity of deploying additional measurement devices, we do not expect such functionalities to be present in low-end products in the coming years. As a result, this renders the integration of intelligent monitoring components into everyday appliances unlikely in the mid-term.

3. CONCEPT DESCRIPTION

In this section, we describe the general concept of our system, followed by the selected system architecture and a short description of its individual components.

In order to extract suitable features from voltage and current waveforms, the collection of these analog quantities is required, followed by their conversion to the digital domain and their appropriate processing. Hence, a hardware system with transducers for both physical parameters is needed, which needs to be capable of its interfacing to a data processing system. As a result, we have conducted both hard- and software design in order to address the problem at hand. We have especially designed our system to fulfill the following fundamentals:

- The system has to record power traces including associated current and voltage waveforms to enable the calculation of real, reactive, and apparent power, as well as the calculation of the phase shift.
- The collected power traces need to have a high resolution in both time and quantization. This also makes the usage of low-noise components necessary in order to minimize the impact of unwanted error signals.
- The provided current waveforms must be processed appropriately in order to enable machine learning tools to find and evaluate current and voltage waveform characteristics of appliances.
- The stored data has to be normalized with respect to the power value and grid frequency in order to cater for the repeatable nature of the conducted evaluations.

The resulting architecture of our system is shown in Figure 1. Its main components can be decomposed into three

tiers: a hardware board, the software framework and the machine learner. The hardware board is plugged between the wall outlet and the electric appliance to easily measure the power consumption. It features voltage and current sensors to analyze device-specific high frequency characteristics of its energy consumption. It outputs two analog signals, one which is proportional to the grid voltage, whilst the second signal is directly proportional to the device's current draw. Those signals are then sampled by an analog-to-digital converter and further processed by our software framework. This framework is responsible for extracting features from the power recordings and to forward the extracted features to the widely used Weka data mining toolkit [6]. In order to maintain workplace safety, the hardware board furthermore ensures a galvanic decoupling from the mains voltage.

4. POWER MEASUREMENT

In order to collect the power consumption of electric appliances at a high resolution, distributed measurement units with high sample rates are required. Commercially available metering platforms, like Plugwise¹ or Wattson² however, often output mean power consumption values instead of allowing for direct access to the collected samples. Additionally, the phase shift is commonly not reported by these units, such that real power, but not reactive power components, can be measured.

As a result, these existing solutions are not suitable for the task at hand, and in consequence, we have designed a specific circuit board that fulfills the following set of criteria.

- High sample rate: In order to collect all possibly relevant characteristics from an appliance's current consumption, a sample rate as high as possible is desired. This requirement is based on the observations made in related work (e.g., [16]). At the same time, higher sampling frequencies allow for a better spectral resolution of the frequency domain plot, and thus better means for the subsequent signal analysis.
- Galvanic decoupling: The most straightforward approach to collect voltage and current measurements relies on the use of resistors (shunts for current measurement and resistor dividers to sample the mains voltage). Both approaches, however, suffer from the drawback of galvanic coupling to the mains voltage, which can destroy expensive measurement equipment and is not touch-safe. As a result, our prerequisite to the data collection unit was the galvanically decoupled transmission of both voltage and current to the sampling unit.
- Linear transducers: Especially when higher frequency components are analyzed in the signal, non-linear signal distortions can falsify the signal. As a result, the transducers used to galvanically decouple the signals from the mains voltage need to behave as linear as possible and add a minimum amount of noise onto the signals.

¹Plugwise: Smart Wireless Solution for Energy Saving, Monitoring and Switching.

<http://www.plugwise.com>. Accessed: 10.01.2013

²Wattson: Energy Monitoring from Energeno.

<http://www.diykyoto.com/uk/>. Accessed: 10.01.2013

- Line-level output: In order to process the collected readings using small-signal processing components, e.g. studio recording equipment, the device needs to provide adjustable output voltages up to 2V RMS. In our case 775mV RMS for line-level was used.

4.1 Hardware Design

As a result of the commercial unavailability of galvanically decoupled high-resolution voltage and current measurement units in affordable price classes, we have developed our own printed circuit board that meets aforementioned design criteria. The resulting device is visualized in Figure 2, in which its four main functional sections are highlighted.




Figure 2: The PCB of the hardware board. Part 1 shows the capacitive power supply. Part 2 illustrates the small signal processing, Part 3 and 4 show the voltage and current measuring components.

4.1.1 Mains Voltage Components

The mains voltage is connected to the voltage and current transducers. In our practical realization, we have chosen a current transducer manufactured by LEM (a detailed discussion on its selection is given in Section 4.2). In order to enable our device to also collect voltage readings, an Avago Technologies HCPL-7840 optocoupler has been used, which caters for a galvanic decoupling of analog signals. While the LEM transducer does not require any power supply on the mains voltage side, the linear voltage optocoupler requires an input voltage of 5 volts DC to operate, which is not galvanically decoupled from mains. We have hence included a capacitive power supply on the mains side, which provides the linear optocoupler with its required operating voltage. The widely available switch mode power supplies have deliberately not been used in order to avoid switching noise on the captured signals.

4.1.2 Current and Voltage Sampling

The second part of our circuit board comprises the signal conditioning of the current signal. Instead of directly interfacing the transducer’s output voltage to the signal processing system, however, we have included an operational amplifier in the signal path. The operational amplifier is used as voltage follower to increase the output driving strength.

Similar to the current conditioning part, the output of the linear optocoupler is interfaced to an operational amplifier. This amplifier setup adjusts the fully differential output signal of the optocoupler to different line levels and increases the output driving strength. Both voltage and current signals at the secondary side are decoupled using capacitors.

4.2 Sensor Selection

Many different sensor technologies are readily available for current sensing. Emerald [3] and Ziegler [25] give a good overview of the performance characteristics, advantages and disadvantages for the major sensor technologies. For the hardware board, we require a current sensor with a current range from 5mA up to 15A, which corresponds to a power draw ranging from 1.15W up to 3.5kW given the European grid voltage of 230V. The sensor should be capable of measuring frequencies from 50Hz up to 50kHz with a flat frequency response. Ideally it should have linear I/V correlation. Because we are mainly interested in the characteristics of different devices a precise calibration for current and voltage was not a major requirement.

To evaluate the sensors we generated sinusoidal, rectangular, and SMPS-like (Switched Mode Power Supply) signals using a HAMEG HMF 2525 frequency generator. Those signals were amplified using a BEAK BAA 1000 high current amplifier with a maximum output power of 1250VA and a maximum frequency of 50kHz. To avoid measuring errors caused by distortions of the amplifier a reference current was metered using a shunt resistor. To simulate a load we used high power resistors wrapped in two CPU coolers. To keep the temperature and thus the resistance of the load network nearly constant we heated the whole circuit until its temperature reached a steady state.

For each sensor we measured the I/V linearity, the frequency response using a sweep, the noise floor and the correlation with the reference signal for rectangular and sinusoidal signals and repeated all of those measurements with different currents.

We evaluated a set of different technologies current transducers, namely LTS 15-NP (hall effect), ACS712 (hall effect), CAS 15-NP (fluxgate), and ACS1050 (current transformer) called CT using the stated criteria. For each criterion we ranked the sensors and finally we calculated the average rank of each sensor. According to the criteria the CT and the CAS performed best. Both have the same average rank. The CT has the lowest non-linear distortions and a flat frequency response. On the other hand the signal obtained from the CAS has the highest correlation with the reference signal and also the lowest noise floor. As the frequency response could also be flattened with a post processing step we decided to build up the hardware board with a CAS current sensor.

4.3 Analog-to-Digital Conversion

To digitize the data a Realtek ACL 880 Intel HDA compatible sound card was connected to our hardware board via line-in. The sound card has a maximum sample rate of 96kHz. With respect to the Nyquist Theorem this results in a maximum usable sample rate of 48kHz. But due to an internal analog low pass filter, frequencies above 32kHz are attenuated and the maximum feasible frequency ranges up to 40kHz. Above this frequency, the signals are too weak to process them further. This low pass behavior is important to

avoid aliasing effects if signal components above the Nyquist frequency are present in the input. The sound card has a sampling resolution of 20 Bit and an inbound voltage range of $1V_{RMS}$. To obtain the readings from the sound card we use the ALSA sound system in combination with the octave package *aarecord*.

We used a soundcard for recording the current and voltage signals because of the wide availability of such Intel HDA compatible sound capturing devices. Nearly every laptop and desktop computer is equipped with such a hardware component. This should allow a flexible change of the data processing component, e.g. different PCs.

5. SOFTWARE FRAMEWORK

As we have described in Section 3, a software framework is responsible for recording and post processing the digitized sample stream. Our software framework called PISI is written in GNU Octave and performs a number of processing steps, including a segmentation of the recorded data, the feature extraction, and the export to the Attribute Relations File Format (ARFF), supported by the Weka toolkit. Next, we describe each of these steps in detail.

The first step in data processing is the generation of raw sample recordings. Those raw recordings contain many periods of the alternating voltage and current signals. As the recording starts at a random point in time, the alternating voltage signal has an undetermined phasing. To avoid any influence of this indeterminism, the software framework performs a segmentation of the recorded data. A segment starts on a zero-crossing of the voltage signal from a negative to a positive value and has a fixed length of N periods. The segmentation process is graphically shown in Figure 3. Shorter values of N are faster to obtain, whereas larger values of N are expected to yield better classification results due to the smaller impact of unexpected signal variations.




Figure 3: The segmentation process of the PISI .

In the second stage, our software framework extracts features from the segments which we present in more detail in Section 5.1. The resulting feature vector is stored and can be exported as a ARFF file. This is an important task to interface the Weka machine learning toolkit [6, 23] which is used for the classification. Numerous feature generation functions exist to extract the feature vector. Those functions take a segment as input and produce a set of features as output. PISI was designed with the requirement in mind to add new feature generation functions on demand. Therefore PISI organizes all recordings in a project environment. The environment keeps track of historic recordings with their corresponding class, the available feature generation functions and a list of available classes.

Last but not least the software framework is able to live predict the type and state of an electric appliance, which is attached to the hardware board. In order to do so, the software framework records the current and voltage signals for several seconds. Then it extracts segments from this recording and generates the features for each segment. The last step is the invocation of a Weka machine learner with a pre-built context model to predict the type of electric appliance. Then PISI loads a device related model which contains all possible states of the given device. This machine learning model is used to classify the state of the currently connected appliance.

5.1 Feature Extraction

One of the most important issues in the machine learning process is the feature extraction. Hence, a modular and flexible feature extraction solution was chosen to enable post processing the data prior to their forwarding to the data mining component. In general, we differentiate between two classes of features, namely *waveform* and *classical* features. In the context of this paper we use active power $P[W]$, phase shift $\varphi[^\circ]$, and the current's crest factor C_i as the so called classical features. For continuous values the real power P_{cont} can be calculated from voltage and current waveforms over one or more periods T (cf. Equation 1). Its corresponding iterative calculation multiplies current and voltage samples over one or more periods N_P (cf. Equation 2).

$$P_{cont} = \frac{1}{T} \int_t^{t+T} v(t) \cdot i(t) dt \quad (1)$$

$$P = \frac{1}{N_P} \sum_{x=n}^{n+N_P} v_x \cdot i_x \quad (2)$$

Equation 3 gives the shift from a current to an according voltage signal in degrees. In this context the voltage is supplied from grid, and thus considered to have constant amplitude during the short sampling window. Hence, only the current is considered to vary.

$$\varphi = \varphi_v - \varphi_i \quad (3)$$

At this point we have to mention that the phase shift is calculated using the cross-correlation of v and i due to the diversity of current draw signals.

The crest factor measurement is the ratio from a waveforms amplitude peak to its RMS (Root Mean Square) value. For ideal sinusoidal signals it results in the factor of $\sqrt{2}$, which can be assumed for the grid voltage. On the other hand the current signal i varies and its crest factor C_i is specified according to Equation 4. The crest factor and its variance from $\sqrt{2}$ allow getting a first estimation of a signal's waveform characteristics.

$$C_i = \frac{i_{peak}}{i_{RMS}} \quad (4)$$

To acquire specific information about the appliance current draw characteristics, a harmonic analysis is commonly used. In our implementation a discrete Fourier transform (DFT) is applied to the digitized current signal. The result is the representation of the signal in the frequency domain. Then the current signal harmonics (H_I) are selected by integer multiples of the fundamental frequency up to the Nyquist frequency which is defined as the half sampling frequency f_s . Since the current is directly dependent on the

grid voltage, the first harmonic h_1 is allocated to the grid and thus the fundamental frequency f_{grid} . Hence the number of the highest harmonic h_N is defined by Equation 6.

$$H_I = [h_1, h_2, \dots, h_N] \quad (5)$$

$$\text{with: } N = \frac{f_S}{2 f_{grid}} \quad (6)$$

According to Parseval's theorem the energy in the time and frequency domain are equal. Because of this, the absolute amplitude of the harmonics depends on the power consumption of the metered appliance. To eliminate this coupling, we normalize the harmonics so that the first harmonic has a value of 0dB. This is leading to clearly separated features of signal strength (P) and signal waveform (H_I).

In theory, the even harmonics should be zero if the stimulating signal has a symmetric waveform and the stimulated system is linear time invariant (LTI). In practice, the even harmonics are not zero because the observed real-world systems were never totally linear. Therefore, we introduce two more features in addition to the harmonics for the evaluation of even and odd harmonic's influences (cf. Equations 7 and 8).

$$RMS_{h,odd} = \sqrt{\frac{2}{N} \sum_{x=1}^N h_{2x+1}^2} \quad (7)$$

$$RMS_{h,even} = \sqrt{\frac{2}{N} \sum_{x=1}^N h_{2x}^2} \quad (8)$$

Equation 9 shows the feature vector which is the base for this research in summary. The setup has a sampling frequency f_S of $96kHz$ which - together with the European grid frequency ($f_{grid} = 50Hz$) and Equation 6 - allows us to choose up to 960 harmonics to gain a wide range spectrum. Due to the low pass characteristics of the Metering Unit (c.f. Section 4.3), we have used only the first 800 harmonics in our implementation.

$$F = [P, \varphi, C_i, RMS_{h,odd}, RMS_{h,even}, h_1, h_2, \dots, h_{800}] \quad (9)$$

5.2 Instance Extraction

The feature vector introduced in Equation 9 holds all features of one segment and represents one class instance for the machine learner. A set of K instances which are stored together are named recording. Obviously it requires less space to save the features of one segment as vector F . But this reduction of the size causes no substantial information loss because the information is then represented by classical features and harmonics. According to Figure 3 this allows us to adopt the resolution of the feature by changing the number of periods per segment (M). This relation can be observed in Figure 4 which shows the odd harmonics spectrum of a single instance and the standard deviation borders (SD+, SD-) of many instances.

The spectrum of each instance is noisy in between the characteristic standard deviation borders. This is caused by aforementioned sampling resolution. We were able to identify this noise as non-deterministic by means of Figure 4b, which contains the same base data and standard deviation borders as in Figure 4, but shows a flattened spectrum in between these.

The spectrum was calculated as arithmetic average of odd harmonics from all instances of one recording, and effectively shows the noise cancellation as compared to a single spectrum. Obviously the flattened spectrum is visibly better suited for classification than a single harmonic spectrum of an instance. Therefore we have used the averaged spectrum in further waveform-related analysis.

We have also observed that there is a trade-off between harmonic features and features derived from the power consumption of one instance. Classical features of one instance do not show the huge variance which is visible for the harmonics as shown in Figure 4. Thus, their features can be extracted from short segments in a sufficiently good quality and being used for classification. On the other hand, averaged harmonic features could help to distinguish between different appliances of one type or with similar power consumptions with the drawback of an increased demand for data, time, and processing. To exploit both benefits during the evaluation process (cf. Section 6), we have decided to split the machine learning models into two parts: One coarse-grained model with power related features to separate between different appliances and fine-grained models to distinguish between devices of one appliance type with nearly the same power consumption using average values. Our software enables the user to build training and testing data sets using feature vectors defined in Equation 9 as input. One can extract new features from this base value F and define a set with a subset of features easily. It is also possible to merge instances with the above mentioned arithmetic mean calculation. All these combinations are stored as ARFF files and processed in the data mining framework Weka.


6. ACCURACY EVALUATION

We have evaluated the system in different settings, which are described as follows.


6.1 Evaluation Setup

Based on the observations in Section 4 we used the current transducer CAS 15-NP for measuring different electric appliances. We have interfaced different electric appliances, as shown in Table 1.

For each appliance, we have measured 30 segments, where each segment covers voltage and current readings of 10 periods. Hence the measurement of one segment needs 200ms at 50Hz European grid frequency and one recording consists of 30 instances where each holds a complete feature vector F . The measurements of the devices and their operating modes were repeated several times under different conditions. We used independent recordings to create the training and test set for the machine learning process. More precisely: Both sets were recorded on different days with the attempt of different conditions for the measured appliance (e.g. temperature and running time of device). Using this approach we built one machine learning model for the device classification to detect the appliance (Section 6.2.1). A second model was built to differentiate between products of one appliance type (Section 6.2.2). In a third step we analyzed specified operating modes of selected devices (Section 6.3). The quality of the machine learning models was optimized through the feature and the classifier selection.



(a) The noisy spectrum of a single instance.



(b) The mean spectrum of multiple instances.

Figure 4: The harmonic spectrum for a laptop device. Both plots are based on the same data.

6.2 Device Identification

For the evaluation, we selected different appliances which are typical for a household or an office environment. Those appliances are listed in Table 1. We took care to select diverse devices: Some of them are movable - others are stationary, several appliances have the same power consumption - others differ in a spread spectrum and a few devices are even of the same production series. We recorded all of those appliances according to the setup described in Section 6.1 and extracted the full set of features described in Section 5.1 from their voltage and current readings.

6.2.1 Coarse-Grained Classification

In the first step we tried to classify the appliance type only with the classical features as described in Section 5.1. We define this as coarse-grained classification. The first column of Table 1 lists the different appliances to be classified in this step. Some of them were recorded in distinctive major operating modes. This results in a total of 14 appliance types. Each set consists of four recordings with 30 instances each. We used the Random Forest classifier to build our machine learning model. With this setup we obtained an accuracy of 99.6% for the device types. With the additional features that are calculated through the RMS value of all even respectively all odd harmonics (see Equations 7 and 8) we can improve the result to 99.8%. This means that even if the device is in standby mode we can differentiate the device type. In general this is a hard task, because the standby power consumption is typically located in the range of 0.5-5 W. Since the classical feature values showed small variance, we were able to use single instances of recordings as classifier input. In our scenario an appliance type can hence be classified by a 200ms snapshot of its I/V characteristics.

6.2.2 Fine-Grained Classification

The classification of different series of devices within the same class using the coarse-grained classification was only partially satisfying. For example the six devices of the monitor type reached a maximum of 87% classification accuracy using the features described in the last section. This clearly shows that the coarse-grained classification would not scale

with respect to a broad training set. Again, the harmonic RMS features brought an accuracy improvement of 4%. This indicates that it might be useful to infer more information from the harmonic content. To do so, we introduced as second step, the so called fine-grained classification that considers also the device specific waveforms. This step was also motivated through the findings we made during the visualization of the current draw in the frequency domain (see Figure 5). Since all errors in the aforementioned classification appeared in between the products of the monitor type, we will use the fine-grained classification only for different monitors in standby and on status. Therefore the fine-grained classification is another escalation layer to separate between similar devices.

In a first step we processed a Greedy Stepwise feature ranking on a global set. This global set holds only one instance for each monitor which was calculated as the arithmetic average of all recordings. The aim of such a set is to provide a spectrum with product specific harmonics and thus to prevent the ranking of single instance's noise. We chose the top 20 ranked features, including 3 classical features (P , φ , $RMS_{h,even}$) and 17 harmonics: In total 3 even harmonics, 9 harmonics below 5kHz, and 7 harmonics up to 38.95kHz are present. These features build the new feature vector for the fine grained classification. As next step we created an independent training and test set with recordings from different days. To reduce the collected information and to concentrate on finding patterns we calculated the mean of 30 instances for each recording. Thus - and with a Random Forest classification model built from the train set - we gained a classification rate of 100%. This rate has to be interpreted as optimistic for a generalized case but it is valid for this specific scenario. It might decrease with the use of more products or with the use of test sets recorded under yet new conditions. To determine the influence of varying environmental conditions over the time, we investigated the long term behavior in Section 6.4.

6.3 Operating Mode Identification

In the last section, we have shown that it is possible to distinguish between different appliances. Furthermore the

question is, if it is even possible to detect the operating mode of an appliance by analyzing its power consumption characteristics. To do so, we have selected different devices with several operating modes. Those devices have either operating modes with different power consumptions in each mode or their power consumption is nearly constant in all operating modes. To detect the operating mode, we trained a machine learning model for each appliance. Due to this design decision, it is possible to use different feature sets for each machine learning model. This flexibility is especially useful if some features do not differ between operating modes.

To evaluate the models, we performed the classification of a test set for each model. Those results are shown in Table 2. The model of a smartphone charger showed a high accuracy. It was always possible to detect, whether the charger was running idle, loading an empty smartphone battery or trickle loading an attached smartphone. No incorrect classification occurred by the use of only waveform features. We repeated the same procedure to classify the number of clients wired to an Ethernet switch and also if those clients were transferring data or not. Our system can classify the number of wired ports accurately but it was impossible to detect whether there is a network transfer happening or not. This can be seen in Table 2: The state c0 shows the switch running idle, c1 shows the switch connected with one client, c2 with two clients, c2d1 shows the results with two clients and a pending network transfer. As one can see in the table, it is impossible, to separate between the state c2 and c2d1. Last but not least we built a model with several operating modes of a TFT monitor. As a TFT monitor consists of millions of different pixels and a dimmable back light it has a nearly infinitely large set of different operating modes. From all of these operating modes we have chosen seven different

Table 1: Description of the devices tested during the evaluation.

Appliance [status]	Product	Power/W
Monitor [on, stby]	Fujitsu Siemens 24"	[77, 1]
	Fujitsu Siemens 19"	[34, 1]
	Fujitsu Siemens 17"	[34, 2]
	3 x Dell 20"	[48, 1]
LCD TV [on, stby]	Samsung 40"	[178, 1]
Laptop	Lenovo T430s	28
	Lenovo T420s	22
	Macbook 2.1	22
	Dell	29
Freezer [on, idle]	Liebherr FKS 3602	[130, 13]
Lamp	110W CCFL	105
	10W LED	6
	60W resistive	63
	100W resistive	98
Charger [stby]	Samsung Phone	0.5
	Apple Laptop	0.5
	Lenovo Laptop	0.5
Switch	Netgear 8 Port	6
Fan [level 1-3]	Tevion	[32-45]
USB-Hub	LogiLink 4 Port	0.5

Table 2: Results of the operating mode evaluation.

a	b	c	d	e	f	g	h	i	j	k	l	m	n	o	classified as
60	0	0													a charger idle
0	60	0													b charger loading
0	0	60													c charger done
			30	0	0	0	0								d switch c0
			0	30	0	0	0								e switch c1
			0	0	22	8	0								f switch c2
			0	0	16	14	0								g switch c2d1
			0	0	0	0	30								h switch c3
								149	30	0	0	0	0	0	i monitor black
								0	26	0	0	0	0	154	j monitor red
								0	0	175	0	0	2	0	k monitor white
								0	0	0	170	0	10	0	l monitor picture 1
								0	0	0	6	172	0	0	m monitor picture 2
								0	0	40	0	0	139	0	n monitor website 1
								0	1	0	0	1	0	177	o monitor website 2
			100%			83%			81%						

screen contents. Those operating modes include screens of the colors black, red, white, two pictures and two web sites. The initial machine learning model with the classical feature set performed poorly. By looking at the distribution of the features we noticed that only the power varies for different screen outputs and there are no noticeable changes in the frequency domain. Therefore this model considers only the power feature. In spite of the fact that we are merely able to classify the operating mode by the power we achieved an accuracy of 81%. Most remarkably, colors with high contrast like black and white can be separated successfully. An opposite effect is shown regarding the website one, which is the eEnergy 2013 website, and the color white. In this case the classification may be mixed up because of the large white areas of the eEnergy 2013 website.

It appears from the measurements that different colors do not affect the power consumption. But differences in the contrast are significantly measurable. Because of this fact, the classification of different screen outputs is possible. The variation of the power consumption ranges from 0.1W to 3W for the different operating modes. To separate between different states of the TFT monitor, there must be a difference of 0.5W. Differences below 0.5W might be caused by environmental noise and thus should not be used as feature.

6.4 Long Term Stability

We measured the energy consumption of particular devices at different points in time. But due to changes in the environment (varying grid voltage, grid frequency, temperature, SMPS switching frequency), it might be the case that the recordings of devices vary over time in a way that causes the machine learner to fail. To analyze if variable environment conditions could cause problems, we measured the energy consumption of one monitor over the period of one week. The Figure 5 shows the spectral variance of the obtained recordings. The blue line shows the mean of a recording which was obtained by averaging 30 continuously recorded segments. The gray area indicates the standard deviation around the mean for each harmonic. The thickness of the band shows the noise level in this harmonic spectrum. This noise of single instances could either be caused by the environment or by the appliance itself. Characteristic for this device is the oscillation around the 325th harmonic which is caused by the switching power supply (SMPS). This peak was present in all recordings of this particular device. Therefore frequencies in this region are good features for recognizing the device.

As the Figure 5 shows, the instances contain time vari-




Figure 5: The variance of the spectrum of a monitor over a period of one week.

ant noise. The fine-grained device classification and also the operating mode classification have to deal with such distortions to avoid overfitting of the machine learning model to particular environment conditions. Two possible solutions exist to deal with those issues: One could either use a hand selected set of omnipresent features or one could repeat measurements over a period of time to face varying environmental conditions by averaging all measurements. Regardless of the time variant noise, the shape of the harmonics remains a characteristically feature for the classification process.

6.5 Discussion

As we have shown in Section 2, various related works perform a device classification based on the current draw of electrical appliances. But even recent works have diverse limitations: either they provide no contemporary information or the classification results are inaccurate for certain devices under certain circumstances. Moreover, to the best of our knowledge, no solution for operation mode estimation was presented yet.

Our work enhances the state of the art by carefully crafted methodologies for high accuracy device- and operating mode classification. Our two layered approach achieves a high precision of up to 99.8% together with low latencies. Furthermore, our evaluation shows, that both waveform and classical features carry valuable information with regard to the device classification.

These achieved improvements close the information gap between electric appliances and the smart home system, and thus enable the development of a new generation of smart home systems. These next generation smart home systems could improve the user comfort, reduce the energy consumption of the household or detect unexpected behavior and increase the safety. For example, one can build an energy saving module, which uses the provided information to cut off standby loads, trim appliances to the most energy efficient setting or defer devices until the time of use pricing is low.

In the next step, we will implement our approach on an embedded system. In this case, the amount of data to process is no serious problem: an adaptive sampling schema together with a decent Digital Signal Processor will easily

handle the feature extraction step. A classification of the attached appliance is only required, if turn-on or switching transients happen. This classification task could be handled by a centralized processing unit. The price of such a solution might be very low, if mass production cuts the costs per unit.

7. CONCLUSION

Enhancing current building automation systems by the capabilities to identify electric appliances is a major step on the way to truly smart buildings. We have thus presented an approach to determine both the type as well as the operation mode of an electric appliance based on measuring its voltage and current waveforms at high resolution. To this end, we have designed a hardware system to collect and condition the readings from the physical environment, as well as the PISI framework to extract relevant features from the raw data and add them to the model of a machine learner. PISI has been shown to classify the type and the state of connected electric appliances with a precision of up to 99.8%.

Our solution bridges the information gap between the building automation system and the electric appliances, and thus enables the development of the next generation home automation systems. Based on the appliance identification, novel features can be realized, e.g., saving energy by automated deactivation of appliances while their operation is not required by the user. Similarly, improvements to the user's comfort and safety as well as support for Ambient Assisted Living can be made when the presented functionality is available to building automation systems.

Acknowledgements

This work was supported by the BMWi Project IP-KOM-OeV. Sincere thanks are given to Christian Hatzfeld and Jan Lotichius from the research team of the Institute of Electromechanical Design for supporting us with equipment and helpful discussions for the sensor selection.

8. REFERENCES

- [1] C. Beckel, L. Sadamori, and S. Santini. Towards Automatic Classification of Private Households Using Electricity Consumption Data. In *4th Intl. Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings (BuildSys)*. ACM, 2012.
- [2] H. Chang, C. Lin, and J. Lee. Load Identification in Nonintrusive Load Monitoring Using Steady-State and Turn-On Transient Energy Algorithms. In *14th Intl. Conference on Computer Supported Cooperative Work in Design (CSCWD)*, pages 27–32. IEEE, 2010.
- [3] P. Emerald. Non Intrusive Hall Effect Current Sensing Techniques Provide Safe, Reliable Detection and Protection for Power Electronics. In *Intl. Appliance Technical Conference (IATC)*, pages 1–2. IEEE, 1998.
- [4] T. Ganu, D. Seetharam, V. Arya, R. Kunnath, J. Hazra, S. Husain, L. De Silva, and S. Kalyanaraman. nPlug: A Smart Plug for Alleviating Peak Loads. In *3rd Intl. Conference on Future Energy Systems: Where Energy, Computing and Communication Meet (e-Energy)*, pages 1–10. ACM, 2012.
- [5] S. Gupta, M. S. Reynolds, and S. N. Patel. ElectriSense: Single-Point Sensing Using EMI for

- Electrical Event Detection and Classification in the Home. In *12th Intl. Conference on Ubiquitous Computing (UbiComp)*, pages 139–148. ACM, 2010.
- [6] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten. The WEKA Data Mining Software: An Update. *SIGKDD Explorations*, 11(1):10–18, 2009.
- [7] G. Hart. Nonintrusive Appliance Load Monitoring. *Proceedings of the IEEE*, 80(12):1870–1891, 1992.
- [8] M. Ito, R. Uda, S. Ichimura, K. Tago, T. Hoshi, and Y. Matsushita. A Method of Appliance Detection Based on Features of Power Waveform. In *4th Intl. Conference on Symposium on Applications and the Internet (SAINT)*, pages 291–294. IEEE, 2004.
- [9] M. Jahn, M. Jentsch, C. Prause, F. Pramudianto, A. Al-Akkad, and R. Reiners. The Energy Aware Smart Home. In *5th Intl. Conference on Future Information Technology (FutureTech)*, pages 1–8. IEEE, 2010.
- [10] L. Jiang, S. Luo, and J. Li. An Approach of Household Power Appliance Monitoring Based on Machine Learning. In *5th Intl. Conference on Intelligent Computation Technology and Automation (ICICTA)*, pages 577–580. IEEE, 2012.
- [11] X. Jiang, S. Dawson-Haggerty, P. Dutta, and D. Culler. Design and Implementation of a High-Fidelity AC Metering Network. In *8th Intl. Conference on Information Processing in Sensor Networks (IPSN)*, pages 253–264. IEEE, 2009.
- [12] Y. Kim, T. Schmid, Z. M. Charbiwala, and M. B. Srivastava. ViridiScope: Design and Implementation of a Fine Grained Power Monitoring System for Homes. In *11th Intl. Conference on Ubiquitous Computing (UbiComp)*, pages 245–254. ACM, 2009.
- [13] D. Kirschen. Demand-Side View of Electricity Markets. *IEEE Transactions on Power Systems*, 18(2):520–527, 2003.
- [14] W. Kleiminger, S. Santini, and M. Weiss. Opportunistic Sensing for Smart Heating Control in Private Households. In *2nd Intl. Workshop on Networks of Cooperating Objects (CONET)*, 2011.
- [15] Kolter, Zico and Johnson, Matthew. REDD: A Public Data Set for Energy Disaggregation Research. In *1st Intl. Workshop on Data Mining Applications in Sustainability (SustKDD)*. ACM, 2011.
- [16] C. Laughman, K. Lee, R. Cox, S. Shaw, S. Leeb, L. Norford, and P. Armstrong. Power Signature Analysis. *IEEE Power and Energy Magazine*, 1(2):56–63, 2003.
- [17] K. Lee, S. Leeb, L. Norford, P. Armstrong, J. Holloway, and S. Shaw. Estimation of Variable-Speed-Drive Power Consumption from Harmonic Content. *IEEE Transactions on Energy Conversion*, 20(3):566–574, 2005.
- [18] T. Mantoro, M. A. Ayu, and E. E. Elnour. Web-Enabled Smart Home Using Wireless Node Infrastructure. In *9th Intl. Conference on Advances in Mobile Computing and Multimedia (MoMM)*, pages 72–79. ACM, 2011.
- [19] J. Nichols and B. Myers. Controlling Home and Office Appliances with Smart Phones. *Pervasive Computing*, 5(3):60–67, 2006.
- [20] S. Patel, T. Robertson, J. Kientz, M. Reynolds, and G. Abowd. At the Flick of a Switch: Detecting and Classifying Unique Electrical Events on the Residential Power Line. In *9th Intl. Conference on Ubiquitous Computing (UbiComp)*, pages 271–288. ACM, 2007.
- [21] J. Pierce, D. Schiano, and E. Paulos. Home, Habits, and Energy: Examining Domestic Interactions and Energy Consumption. In *28th Intl. Conference on Human Factors in Computing Systems (CHI)*, pages 1985–1994. ACM, 2010.
- [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. *Lamp*, 6:45, 2012.
- [23] I. Witten, E. Frank, and M. Hall. *Data Mining: Practical Machine Learning Tools and Techniques*. Morgan Kaufmann, 2011.
- [24] M. Zeifman and K. Roth. Nonintrusive Appliance Load Monitoring: Review and Outlook. *IEEE Transactions on Consumer Electronics*, 57(1):76–84, 2011.
- [25] S. Ziegler, R. Woodward, H. Iu, and L. Borle. Current Sensing Techniques: A Review. *IEEE Sensors Journal*, 9(4):354–376, 2009.