

# A Self-Calibrating Approach to Whole-Home Contactless Power Consumption Sensing

Md Tanvir Islam Aumi<sup>1</sup>, Sidhant Gupta<sup>1</sup>, Cameron Pickett<sup>1</sup>, Matt Reynolds<sup>1,2</sup>, Shwetak Patel<sup>1,2</sup>

<sup>1</sup>Computer Science & Engineering, <sup>2</sup>Electrical Engineering  
UbiComp Lab, DUB Group, University of Washington, Seattle, WA 98195  
{tanvir, sidhant, pickcj93, matt.reynolds, shwetak}@uw.edu

## ABSTRACT

In this paper, we present a significant improvement over past work on non-contact end-user deployable sensor for real time whole home power consumption. The technique allows users to place a single device consisting of magnetic pickups on the outside of a power or breaker panel to infer whole home power consumption without the need for professional installation of current transformers (CTs). The new approach does not require precise placement on the breaker panel, a key requirement in previous approaches. This is enabled through a self-calibration technique using a neural network that dynamically learns the transfer function despite the placement of the sensor and the construction of the breaker panel itself. We also demonstrate the ability to actually infer *true power* using this technique, unlike past solutions that have only been able to capture *apparent power*. We have evaluated our technique in six homes and one industrial building, including one seven-day deployment. Our results show we can estimate true power consumption with an average accuracy of 95.0% during naturalistic energy use in the home.

## Author Keywords

Energy monitoring; sustainability sensing; smart home; ubiquitous computing

## ACM Classification Keywords

H.5.m. Information interfaces and presentation (*e.g.*, HCI): Miscellaneous.

## INTRODUCTION

Energy conservation and eco-feedback research continues to be an important focus in the Ubicomp and HCI communities. Given that 28% of U.S. energy consumption is directly contributed by household activities [3], the home is a natural place to study. However, just obtaining whole home power consumption information in real-time by homeowners or even researchers is not a simple task. For instance, certain smart meters provide data at 15 minutes

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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [Permissions@acm.org](mailto:Permissions@acm.org).

UbiComp'14, September 13 – 17, 2014, Seattle, WA, USA  
Copyright 2014 ACM 978-1-4503-2968-2/14/09...\$15.00.

<http://dx.doi.org/10.1145/2632048.2636087>



**Figure 1: Left: Contactless power consumption sensor on the breaker panel; Right: An interface showing whole-home total power consumption in real-time.**

intervals, however gaining access to that information is challenging due to closed-source and often private protocols and application interfaces. The most common approach to date is to install commercially available CTs *inside* the breaker panel. However, safely installing this device requires hiring a trained electrician as it involves placing a sensor around the main electrical feed in the breaker panel. Most researchers and homeowners do not have the training or confidence to do such an installation. In fact, the National Electric Code (NEC) has strict rules on the requirement of professional installation of CTs. In addition, certain U.S. states disallow CTs to be installed inside the breaker panel at all, in which case an expensive pass through meter solution is required. The latter requires involvement of the utility company as an end-user cannot tamper with or alter the installation of an electricity meter.

To address this challenge, in previous work Patel *et al.* [12] introduced a contactless power consumption sensor that reduces such deployment burdens by offering a “stick on” sensor that goes on the outside of the breaker panel. This technique utilized magnetic sensors to sense the magnetic field induced by the 60Hz current flowing through the main lines inside the breaker panel. While this initial work was a step towards simple and easy to deploy non-intrusive power monitoring, there are some important limitations to consider. First, the existing approach requires the user to precisely position the sensor on the panel, which is a difficult task for an end-user to perform. Second, the approach assumed a linear transfer function between the magnetic sensors and the current, which limits its accuracy to a small current range. Third, the approach did not take

into account the small fields generated by the various branch circuits that may reside in the area directly behind the magnetic sensors. Fourth, the previous approach was only able to infer apparent power and not true power since it does not take into account the phase information between the voltage and current waveforms. This is an important limitation as it is not able to accurately infer power use of highly inductive loads, which now tends to constitute much of the power consumption in a modern home (CFLs, LEDs, HVAC, Computers, TVs, etc). In addition, researchers in the energy disaggregation community have limited utility with just the apparent power data.

In this work, we have significantly improved this previous work by introducing a self-calibrating technique that dynamically generates a multi-order transfer function between the magnetic sensors and the actual current across the entire range of power use in the home. Instead of mathematically modeling the transfer function *a priori*, we use a learning approach to generate this transfer function for each home, which is less susceptible to differences in breaker panel design and construction. In addition, it removes the need for precise placement of the sensor because it takes into account “interference” from any branch circuits. Prior work [12] has also assumed that a plug-in calibrator would draw known power loads to fit a transfer function. However, the drawback of this approach is that it assumes the calibrator is able to draw a large range of loads (~0-20 kW) depending on the size of the home and types of appliances present. This is impractical because of safe heat dissipation limitations in addition to being difficult to build it in a small form factor. In this work, we introduce a technique that uses a calibrator with a much smaller range (0-300 W) by leveraging the insight that we can use a home’s natural electrical activity throughout the day as a part of the self-calibration sequence.

The main contributions of this work are as follows:

- This approach has the ability to predict the phase angle between the current and voltage to infer true power. This is equivalent to predicting the waveform itself and not just the magnitude.
- A self-calibrating approach that does not require precise placement of the sensor on the breaker panel and uses the actual power use throughout a day for calibration.
- A neural network-based learning approach that dynamically generates a multi order magnetic sensor transfer function.

Through deployments in six homes and one industrial building we show that this new approach can predict RMS current and phase angles with an accuracy of 96.0% and 94.3%, respectively. Overall it can predict real power consumption with an accuracy of 95.0% in real-world naturalistic energy use. We also evaluated our technique by

placing the sensors in different non-ideal positions on the panel and achieved an accuracy of 97.4% across all the placements. This high accuracy is ideal for many applications such as energy disaggregation, activity inference, and eco-feedback while reducing the barrier to entry by greatly simplifying the installation process. Also note that a ~5% maximum error in our proposed system is much better than the commonly used IEEE C57.13 C-class CTs (rated for  $\leq 10\%$  error). Very expensive CTs with 1-2% error are used in specialized applications such as precision current measurements, and in those application scenarios we assume that the end user would have access to the current carrying inductors.

#### RELATED WORK

There are many commercially available sensors for measuring and showing appliance level energy use at each outlet, such as the Conserve Insight™, GreenSwitch, and Kill-A-Watt™. In case of whole house power consumption, some of the popular commercially available solutions are The Energy Detective and (TED®) and the PowerCost Monitor. Installing TED involves placing a CT around the main electrical feeds (mains) *inside* the breaker panel, which requires a professional installation due to high-voltage shock hazard. On the other hand, PowerCost can easily be installed by a homeowner without hiring an electrician, but requires either electromechanical meters or electronic meters with an exposed and compatible optical port. Hence it is constrained to specific type of meters with its update rate as well as performance dependent on the meter and exposed data ports of the same.

Because of such limitations, contactless solutions are emerging that try to infer power without having direct access to the mains. Cooley *et al.* [5] describes one such solution that measures the current at individual circuit breakers using a magnetic sensor placed on the face of the breaker switch itself. While this approach is promising, most electric codes do not allow anything to be placed on the circuit breakers for extended use because of the potential interference with its life-saving cutoff operation. In addition, a sensor would have to be placed on each circuit breaker to gather whole home power use or on the main circuit breaker, if present. Lorek *et al.* [11] also describes a similar magnetic field based approach where a magnetic sensor have to be placed on every breaker switch on the panel. In addition to requiring a number of sensors, this system also needs to be calibrated manually by the homeowner, which might be impractical and extremely difficult for a homeowner to perform.

Patel *et al.* proposed a solution that uses a pair of magnetic sensors placed on the face of the breaker panel (instead of the breakers) to sense the current flowing through the main bus bars [12]. A set of LEDs were used to help guide the user on placement. Similar to our approach, this system also used a load calibrator to create a transfer function. However, they assumed a linear transfer function and that

the calibrator could emulate the entire power range of the house. Despite the use of LEDs to help with placement, other branch circuits and stray wires impact the magnetic field under the sensors. In our investigations, we found that the state of the magnetic flux changes throughout the day as various appliances are used. This means that the LEDs only help if the breaker panel state remains the same after the initial installation. Also, the previous approach only inferred apparent power and didn't take into account the phase angle between current and reference voltage.

To overcome these limitations, we use a calibrator with smaller loads and leverage the natural household electrical activities throughout the day to generate a transfer function for the entire range of power use in the home. Because of the in situ dynamic model, our system is not limited to perfect placement of sensors. In addition, because of the capability of predicting absolute current waveform, we can also calculate phase angle in real time. To the best of our knowledge, this is the first attempt to solve the problem of calculating phase angle using a single set of magnetic sensors.

This work would allow researchers in the energy disaggregation community easy access to power data in a home without the need for professional installation. Many approaches have been developed that use power usage analysis to infer appliances use [1,7,8,14]. Approaches that use alternative ways to infer appliance could use our approach to associate the actual power use to each inferred appliance event [2,4,9,13].

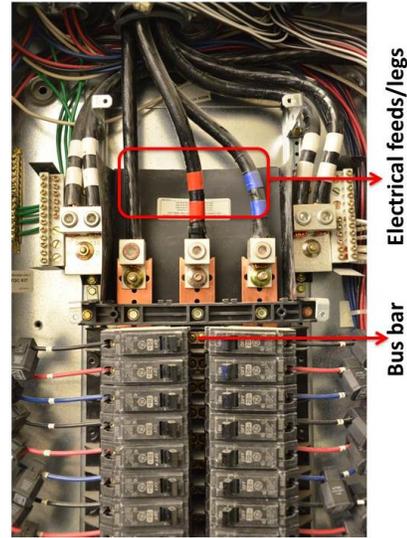
#### SYSTEM DESCRIPTION AND ALGORITHMIC DETAILS

Breaker panels in the U.S. comply with the General Electric "style" based on the guidelines from National Electrical Manufacturers Association (NEMA). Briefly, there is a front surface with an access door that covers the interior where main electrical feeds or lines are connected to the bus bar. In this work, we have focused on typical U.S. breaker panels such as the shown in Figure 2, however the approach presented in this paper can scale to varying designs of breaker panels and to those in other countries as our learning approach is dynamic and calibrates in-situ.

#### Theory of Operation

Our sensing approach involves computing the current consumption in the home by inferring the current being drawn through the main feeds/legs coming in the home at the breaker panel. In general, most homes in the U.S. have split single-phase electrical service where each leg is 180 degrees apart from each other. Industrial buildings usually have three-phase service where three phases are 120 degrees offset from each other. In either case, we need to predict the current flowing through all the legs. The field generated from the main legs allows us to estimate the current flow through each leg separately, which radiates a few centimeters from the wire and even through the layer of sheet metal. In the ideal situation, magnetic field scales linearly with the current. However this relationship is not as

simple in practice because of fields from all neighboring wires, reflected magnetic fields, and magnetic nonlinearities of the sheet metal.



**Figure 2: Interior of a typical U.S. three-phase breaker panel. Three main feeds are connected to the bus bar. Split-phase systems have two electrical feeds instead of three.**

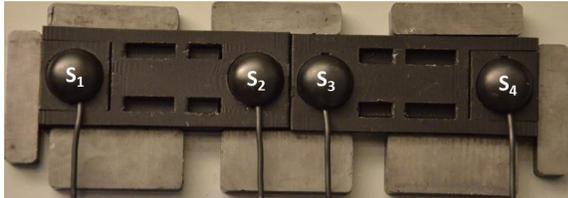
To sense the radiated magnetic field, we use a set of magnetic pickup sensors placed on the outside of a breaker panel. To convert this sensed magnetic field to current, we need to determine a transfer function, *i.e.*, given the sensed magnetic flux, determine what the underlying current flow in the main leg is that induces the flux. We infer this function using a calibrator plugged in an electrical outlet that draws a known amount of current (by powering a resistive load) at a given time while the sensor senses the change occurring in the magnetic field due to that current draw. In future, we imagine the calibrator embedded in an in-home energy display.

At first we create an initial transfer function using this collected calibration data which only works for a small range (limited to the range of loads the calibrator can provide) of magnetic field values. We store this range as calibrated region while keeping track of the present magnetic sensor values. These values change over time as appliances are used in a home. Every time these values reach an un-calibrated region, the calibrator is commanded to pull a small load. The difference in the observed magnetic field signal at that level is used to update the transfer function.

In the rest of this section, we describe the hardware and software of our prototype implementation. We primarily focus on the intuition behind using a machine-learning model and describe its use in creating a dynamic transfer function. It should be noted that during the rest of the description, we assume a home environment having two-phase power supply, however our system scales to a three-phase industrial setting as well.

## Hardware

Our prototype system consists of two components: a *Sensor unit* and a *Calibration unit*. Figure 1 shows a sample placement of the sensor unit on a breaker panel. The data from this unit is collected using a DAQ (National Instruments USB-6259) connected to a laptop. Figure 4 (a) shows a calibration unit installed in an electrical outlet. The calibration unit is also connected to the same laptop. The laptop sends the control logic to the calibration unit based on the captured data from DAQ.



**Figure 3:** The sensor unit consists of four magnetic pickup sensors surrounded by some permanent magnets.

### Sensor Unit

The sensor unit (see Figure 3) consists of four magnetic pickup sensors (we use RadioShack removable telephone pickups. Model no: 07C12) surrounded by permanent magnets. The sensors detect magnetic field generated from the 60 Hz current flow from the legs behind the panel (as well as some of the branch circuits). The surrounding magnets are placed to reduce the effect of magnetic nonlinearity of the sheet metal located in between the main lines and magnetic sensors. Briefly, the time difference between the actual current waveform and sensed magnetic waveform is depended upon the magnetic saturation and permeability of the material. Saturating the magnetic field reduces the nonlinearity induced by the sheet metal. In other words, the nonlinearity of the phase difference between the actual and sensed waveform reduces, which plays a key role during the phase angle calculation. We discuss the effect of this phenomenon later in the deployment and results section.

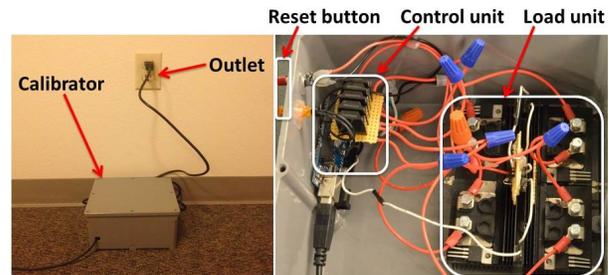
### Calibration Unit

Our system consists of two calibration units installed in two different outlets, each on a different leg. Each unit briefly draws a series of known loads (0-300 W) for automatic calibration described before. It also senses the line voltage for calculating true power.

The calibration unit consists of four high wattage resistors (we use Ohmite chassis mount resistors. Part no. TGHVLR100JE) connected in several series and/or parallel combinations through a set of relays. A microcontroller drives the relays to provide 25 W, 100 W, 200 W, and 300 W loads.

Figure 4 (Right) shows a zoomed in view at the calibration unit. It should be noted that in previous work, Patel *et al.* [12] also used a calibrator to model step power changes. However, they made an assumption of the calibrator pulling a load up to ~20 kW depending on the appliances present in

a home, which is impractical to build in practice because of the large and unsafe heat dissipation requirements. In contrast, our calibrator only pulls up to 300 W and leverages the actual electrical activities in the home to calibrate the entire range.



**Figure 4:** Left: Calibrator unit installed in one of the outlets; Right: Closer look at the calibration unit

## Software Implementation Details

The four magnetic sensors ( $S_1$ ,  $S_2$ ,  $S_3$ , and  $S_4$ ) shown in Figure 3 sense the magnetic field generated from each leg of the breaker panel. Our goal is to generate a transfer function that converts these magnetic fields into current waveforms flowing through only each of the main legs. Since we are interested in predicting real power, we need to predict both the RMS value of the current waveform and phase angle between the current and voltage waveform. Therefore we decided to predict the current waveform instead of just the RMS value.

Creating a transfer function to compute the current waveform given the magnetic flux is non-trivial because of challenges posed by fundamental characteristics of breaker panel and the sensed magnetic field:

- *Stray Magnetic Flux:* In addition to the main feeds, breaker panel also consists of other electrical wires going through all the breaker switches. There are also wires passing around the main lines and each of those radiates good amount of magnetic field depending on the current flowing through them. Figure 2 shows a sample illustration of this situation. The magnetic pickup sensors catch the magnetic field radiated from all them. However, we are only interested in the magnetic field radiated from the main lines. All the magnetic field radiated from the surrounding wires need to be eliminated during the prediction. In other words, the transfer function should be able to identify flux changes caused only by the two main legs.
- *Position of Sensor:* The amount of magnetic field sensors receive depends on the distance between the legs and sensors. Our goal is to create a placement invariant system that works on any position of the breaker panel. Hence the transfer function should accommodate any distance between the legs and sensors.
- *Isolating flux from each main leg:* Current flowing through each line contributes to the sensed magnetic

field of all four sensors. However, we don't know a priori how much each line contributes to each sensor. For example in Figure 3, the leftmost and rightmost magnetic sensors are probably influenced mostly by the leftmost leg and rightmost leg, respectively. But the ratio of influence is unknown. For the two middle sensors, the scenario is even more unpredictable. The transfer function should be able to figure out the ratio by which each leg influences each sensor.

- *Breaker panel wiring uncertainty:* Finally, despite having guidelines from NEMA and NEC, internal wiring of breaker panels varies a lot among each other depending on the skill of the electrician who installed it. The transfer function should be able to work with any breaker panel with any type of wiring.

Because of all these characteristics that vary across different breaker panels, same amount of electrical load induces different amounts of magnetic field in different panels. In fact even in the same panel with the same positioning of sensor unit, the relationship between load and magnetic field depends on the existing magnetic field inside the whole panel. For example, let us say that the baseline current through one leg is  $I_1$  and a positive change of  $I_{ch}$  amount results in a positive change of  $S_{ch1}$  in  $S_1$ . If the baseline current changes to  $I_2$ , the same positive  $I_{ch}$  change will cause a different amount of change  $S_{ch2}$ . Depending on how the magnetic fields radiated from different wires and their constructive or destructive interference, the value of  $S_{ch2}$  could even be negative despite a positive  $I_{ch}$  value.

In summary, the relationship between the electrical current and sensed magnetic field is nonlinear and depends on the existing baseline magnetic field and presence of other magnetic fields. To accommodate this variability and nonlinearity, we observed that we can create multiple polynomial equations for each "state" of the breaker panel. Thus, if we can define the state in terms of magnetic flux, we can build a function for each state. Such a problem is well suited for machine-learning techniques; when given states as inputs, we need to learn a function. We use a neural network as it essentially learns a polynomial function to predict output from input.

#### *Constructing Neural Network Model Using the Calibrator*

As we mentioned previously, the calibrator can cycle through a series of 25 W, 100 W, 200 W, and 300 W loads. Right before the calibrator turns on a load, the system starts tracking the sensor values. Once the load is on, it causes a change in the total current and in-effect the magnetic flux. This change in flux is recorded by the system. Therefore for each calibrator action, our system constructs a training instance for the neural network. Structure of such an instance is shown in Figure 5. First 8 columns of every row are features to the learning algorithm. 9<sup>th</sup> column is the output value that the algorithm will try to learn. The neural network consists of one input layer, one output layer, and two hidden layers having five neurons in each of the layers.

$S_{1p}$	$S_{1c}$	$S_{2p}$	$S_{2c}$	$S_{3p}$	$S_{3c}$	$S_{4p}$	$S_{4c}$	$I_{ch}$
----------	----------	----------	----------	----------	----------	----------	----------	----------

**Figure 5: Structure of a training row. First 8 columns are inputs and 9<sup>th</sup> column is output.**

Here  $S_{1p}$ ,  $S_{2p}$ ,  $S_{3p}$ , and  $S_{4p}$  are the RMS values of four sensors before the calibrator turns on the load.  $S_{1c}$ ,  $S_{2c}$ ,  $S_{3c}$ , and  $S_{4c}$  are RMS values after the calibrator turns on the load.  $I_{ch}$  is the amount of current the calibrator added to the leg where it is plugged-in. Note that as the relationship between magnetic flux change and electrical current change depends on the existing magnetic flux present in the panel; we take both previous and current magnetic flux as our features instead of just taking the change of flux.

In our implementation, the calibrator turns on each load for five seconds. Therefore after five seconds when the calibrator turns off the load, the system captures a similar event and calculates a similar training instance for the turn-off event. It must be noted that we never have access to the absolute value of current going through each leg. The only a priori information is the amount of current change the calibrator is intended to cause to the leg. Consequently, our neural network will only be trained to predict the change in current value, not the absolute current value.

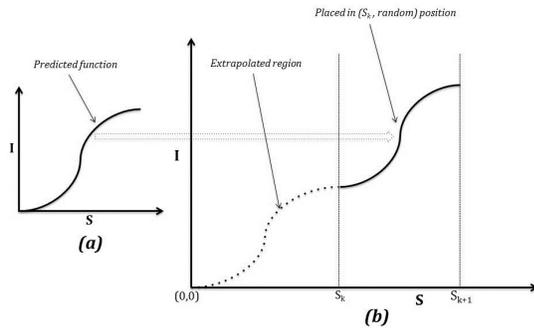
However, our goal is to predict absolute current waveform going through each leg. To achieve this, we use a geometric translation technique that leverages the prediction model and natural electrical activities in a home to create a transfer function that converts sensor values to current waveform. Below is the description of our technique. For the sake of simplicity, we assume one magnetic sensor instead of four and one leg instead of two.

#### *Creating the transfer function*

Once the system starts, the only information it knows is the current RMS magnetic field of the sensor ( $S_k$ ). Initially the calibrator pulls a series of 100 W, 200 W, and 300 W loads (3 times each) on top of this field. Hence the field value changes and the system keeps track of the maximum value of the sensor ( $S_{k+1}$ ). Now from  $S_k$  to  $S_{k+1}$ , it has 9 calibration events. For each event, there are 2 training instances (one for on event and one for off event) as mentioned before. Therefore it gathers 18 training instances from sensor value of  $S_k$  to  $S_{k+1}$  and uses these instances to train the neural network model described earlier.

This results in a function ( $F_k$ ) that can convert magnetic field change value from  $S_k$  to  $S_{k+1}$  to current change value  $I_{ch}$  (see Figure 6(a)). The goal is to find a function  $F$  that can convert any magnetic field value  $S$  to absolute current value  $I$ . Therefore the function  $F_k$  is placed into appropriate position of  $F$ . As already mentioned, the system never knows the absolute value of  $I$ , thus, a random y-axis value  $R$  is assumed and the function placed on ( $S_k, R$ ) position. Figure 6(b) shows an illustration of this approach.

Now the system does not know how the function looks like from  $0$  to  $S_k$  position. Therefore the function is extrapolated



**Figure 6: (a) Predicted function from a small range of sensor values ( $S_k - S_{k+1}$ ). (b) Placing the function in appropriate position.**

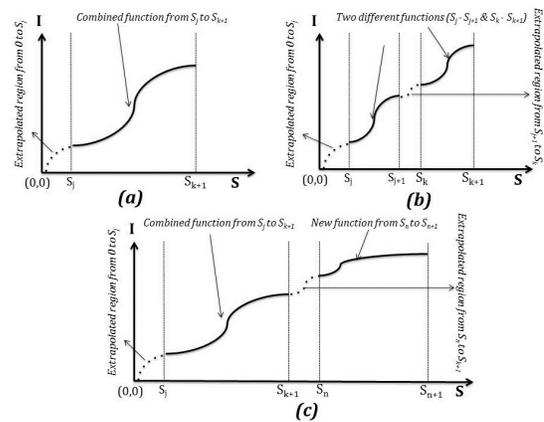
from  $(S_k, R)$  to  $(0, 0)$  as shown in Figure 6(b). Obviously this region does a poor job in translating  $S$  to  $I$ . Therefore the system waits until the value of  $S$  falls below  $S_k$ , following which the calibration process is reinitiated. Just as before, it results in a new function ( $F_j$ ) that can convert magnetic field value from  $S_j$  to  $S_{j+1}$  where  $S_j < S_k$ . Now if  $S_k < S_{j+1}$ , the system combines  $F_j$  with  $F_k$  and creates a new function that covers from  $S_j$  to  $S_{k+1}$  (Figure 7(a)). Otherwise it keeps two separate functions  $F_j$  and  $F_k$  that covers ranges from  $S_j$  to  $S_{j+1}$  and  $S_k$  to  $S_{k+1}$ , respectively (Figure 7(b)).

As evident in Figure 7(a), the new extrapolated region is from  $0$  to  $S_j$ . If the sensor value ever falls below  $S_j$  (e.g. during night when most of the appliances are off), the system will initiate a new calibration cycle for the new region. Consequently, the system will recreate the whole function from the new position to  $S_{k+1}$ .

On the other hand in Figure 7(b), the new extrapolated regions are from  $0$  to  $S_j$  and  $S_{j+1}$  to  $S_k$ . In case the value of sensor is  $S_m$  where  $S_{j+1} < S_m < S_k$ , the calibrator will trigger again and the system will create a new function from  $S_m$  to  $S_{k+1}$ . Note that this time the calibrator will only extrapolate from  $S_m$  to  $S_{j+1}$  as the system already has a function from  $S_j$  to  $S_{j+1}$ .

Similarly, if more appliances are turned on and the sensor value ( $S_n$ ) exceeds  $S_{k+1}$ , the calibrator will trigger and the system will create a new function from  $S_n$  to  $S_{n+1}$  (Figure 7(c)). Again the system will only extrapolate from  $S_n$  to  $S_{k+1}$  as it already has a prediction function from  $S_j$  to  $S_{k+1}$ . In other words, as time goes by, the extrapolated regions will shrink more and more and the system will have a better translation function from  $S$  to  $I$ .

In summary, the key idea is that as the system runs in a house, it captures the usual electrical activities that increasingly provide it with a wide range of sensor values to learn from. As more appliances are turned on and off, the system gets a chance to calibrate for more and more ranges and the prediction function gets increasingly accurate. We will discuss the calibration requirement more in the results section.

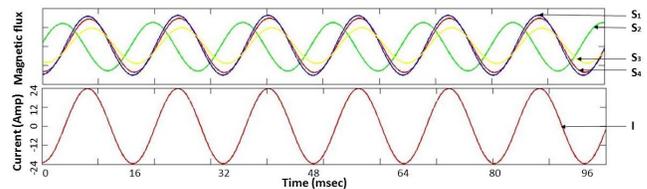


**Figure 7: (a) Combined function:  $S_j - S_{k+1}$ . Extrapolated region:  $0 - S_j$  (b) Two separate functions:  $S_j - S_{j+1}$  &  $S_k - S_{k+1}$ . Extrapolated regions:  $0 - S_j$  &  $S_{j+1} - S_k$  (c) New function:  $S_n - S_{n+1}$ . Extrapolated regions:  $0 - S_j$  &  $S_{k+1} - S_n$ .**

## ANALYZING TRANSFER FUNCTION

### Prediction Using the Transfer Function

Once the system starts, it creates a function  $F$  that takes four magnetic field values from four sensors ( $S_1, S_2, S_3, S_4$ ) and translates them into current waveform  $I$ . Figure 8 shows a sample output of the prediction function  $F$ .



**Figure 8: (Top) Four waveforms from four sensors  $S_1$  (blue),  $S_2$  (green),  $S_3$  (yellow),  $S_4$  (red); (Bottom) Predicted current waveform  $I$  through one leg.**

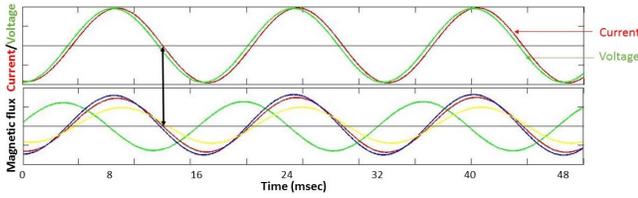
As shown in the Figure, the system is able to predict raw current waveform flowing through each leg. In other words, it can predict both the RMS current ( $I$ ) and phase angle ( $\theta$ ) between the line voltage and current waveform. Predicting this  $\theta$  is very important from an energy monitoring perspective as it can tell us the real power consumed by the household as oppose to apparent power.

### Phase Angle ( $\theta$ ) Prediction Analysis

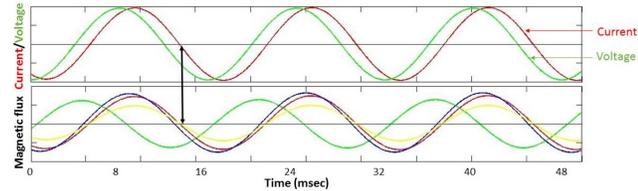
To predict phase angle  $\theta$ , we rely on the following hypothesis:

*“Any change into the phase of the current waveform will also be reflected into the sensor waveform.”*

Figure 9 shows an example of the validity of the hypothesis. From the top graph, we see that voltage and current waveform are closely in phase with each other ( $\theta$  is small). If we carefully inspect the bottom graph, we can see that two of the magnetic waveforms (blue and red) have the same phase characteristics (zero crossing rise and fall in almost same timestamps) as the current waveform. In other words, our prediction function will have more influence from these two sensors while predicting current waveform.



**Figure 9: (Top) Current and voltage waveform are almost in phase. (Bottom) Two waveforms (blue and red) are following the current waveform.**



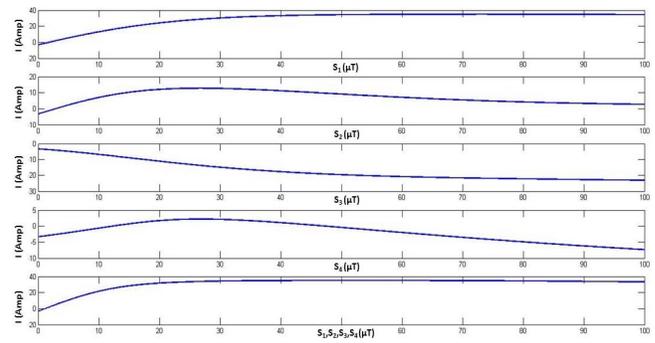
**Figure 10: (Top) Current waveform is lagging the voltage waveform. (Bottom) Two waveforms (blue and red) are still following the current waveform.**

Figure 10 (top) shows us a different scenario where the current waveform is lagging the voltage waveform by an angle  $\theta$ . From the bottom graph, we can see that blue and red waveforms are also following the current waveform. In other words, when current waveform is phased shifted by angle  $\theta$ , four sensor waveforms will also get phase shifted by some angles  $\theta_1, \theta_2, \theta_3,$  and  $\theta_4$ . Obviously these angles will be different from the original phase shift  $\theta$ . But the sensors that are mostly influenced by the current waveform will have a closer shift to the angle  $\theta$ . Therefore the difference ( $\theta_{diff}$ ) between original shift and sensed shift will be small.

However due to the presence of the sheet metal between the main lines and magnetic sensors, the phase difference ( $\theta_{diff}$ ) between the actual current waveform going through the lines and magnetic waveform sensed by the sensor becomes a nonlinear function of magnetic saturation and permeability of the material. The surrounding magnets built into our prototype sensor unit are intended to saturate the magnetic field and reduce the nonlinearity effect. As a result,  $\theta_{diff}$  becomes near constant and the transfer function predicts the phase angle with good accuracy. We discuss the accuracy of this phase prediction approach more in the next section.

### Visualizing the Transfer function

Mathematically, the transfer function can be expressed as follows:  $I = F(S_1, S_2, S_3, S_4)$ . As the function is 5 dimensional (4 inputs and one output), it's hard to visualize the effect of each sensor on the prediction output. Therefore we decided to decompose the visualization using results from each sensor. Figure 11 shows the result. For each of the first four plots, we change one sensor ( $S_1/S_2/S_3/S_4$ ) value from 0  $\mu$ T to 100  $\mu$ T linearly keeping all the other sensor values to 0  $\mu$ T. The bottom plot assumes all the four sensors increasing from 0  $\mu$ T to 0.05  $\mu$ T.



**Figure 11: Decomposed transfer function. Top four plots show predicted current ( $I$ ) based on just one sensor value ( $S_1, S_2, S_3,$  and  $S_4$ ). Bottom one shows the plotting of  $I$  based on all four sensors.**

Though the plotting is not ideal (in actual case, the current is predicted based on different combinations of all sensor values), it gives us certain interesting insight. As an example, after a certain field value, the predicted current values go down for all the sensors except  $S_1$ . This phenomenon is observed because of the presence of multiple magnetic waveforms inside the panel. As the phases of these waveforms are different and they are always changing based on load condition, there are constructive and destructive interferences in different locations inside the panel. Hence depending on the location of the sensor placement on the breaker panel, some sensor senses destructive interferences when there is a positive change in the current waveform and exhibits an inverse relationship between current and magnetic field.

If we inspect the bottom plot where all the sensor values are increasing, we find an interesting similarity between this plot and the topmost plot where only  $S_1$  is increasing. Although for all the other three sensors, the current ( $I$ ) is decreasing after a while; it is always increasing in case of the bottommost graph. Essentially, this behavior means that the transfer function is mostly influenced by  $S_1$ . In other words, this sensor reflects the current waveform more precisely than other sensors. The neural network learns that and increases the coefficient of this sensor more than other sensors. Therefore the amplitude and phase of the predicted current is mostly determined by  $S_1$ . This is one reason why a machine learning based approach is more appropriate for this kind of problem, since it would be nearly impossible to fit a single polynomial for these observations.

### IN-HOME DEPLOYMENT AND EVALUATION

To validate our technique, we conducted experiments in six different homes and one industrial building. Homes had two-phase wiring system and industrial building had three-phase system. We collected data from one house for a longer period, spanning seven days and from other places for a shorter period (spanning two days). This allowed us to show the general applicability of our system to a diverse set of breaker panels as well as the longer-term temporal stability of our solution. Table 1 shows the summary of the homes used in our evaluation.

ID	Panel type	Style/Built/ Remodeled	Size/Floors
H1 (*)	Two-phase	Apartment/ 1993/NA	550 sq. ft./ 1 flr.
H2	Two-phase	House/1972 /2002	1250 sq. ft./ 1 flr.
H3	Two-phase	Apartment/ 1931/1994	800 sq. ft./ 1 flr.
H4	Two-phase	House/1960 /NA	2220 sq.ft./ 1 flr.
H5	Two-phase	House/1987 /NA	1340 sq. ft./1 flr.
H6	Two-phase	House/NA/ NA	1452 sq. ft./1 flr.
I1	Three-phase	Industrial /2003/NA	NA

**Table 1: A summary of the homes showing the panel type, style, and size (\* 7 days deployment).**

It should be noted that all of our data collection sessions were performed under a *naturalistic setting* with the usual home appliances comprising of inductive, resistive, and other complex harmonics appliances. We did not give the residents any instructions on the use of their electrical appliances or requested any changes in their daily routines and household tasks. Once installed, our system ran in background for the entire data collection session with no human interaction at all.

#### Data Collection Procedure

Our system was packaged such that it could be rapidly setup in a home. The sensor unit was placed on a breaker panel using double-sided tape. To collect the ground truth, we installed commercially available high end transformer-based split-core CTs (99% accurate) inside the breaker panel prior to installing our sensor unit on the outside of the breaker panel. Both the sensor unit output and CT output were collected using the same DAQ device.

We used long extension cables to bring the two different outlets of different phases closer to the laptop. We then plugged in the two calibrators in the outlets. The calibrators and the data acquisition device were connected to a laptop. The laptop controlled the calibration unit, recorded all the data from the data acquisition device, and performed all the algorithmic processing in real-time. We just recorded the original and predicted waveforms for post-experiment analysis. The whole software portion was written in Matlab.

#### Deployment Results

##### Accuracy Measurements

For each of the deployments, we calculated the RMS current value in ampere, RMS line voltage in volts, and

phase angle of the current waveform with respect to voltage in degree every second. These quantities are recorded both for ground truth current waveform (measured from the CTs) and predicted current waveform (output from the algorithm). Finally we calculate the real power consumption for both of them in every second as follows:  $P = V_{RMS} \times I_{RMS} \times \cos \theta$ . It should be noted that during accuracy calculation, we only considered the accuracy of calibrated regions. However after a certain time, most regions became calibrated and all of the data were taken into consideration.

##### Results

Our system requires installing two calibrators in two different phases of a house. Then, based on the calibrator data, it creates two different functions  $F_1$  and  $F_2$  for two phases  $P_1$  and  $P_2$ . During the evaluation, we also considered the case of using just one calibrator in one of the phases. Therefore for each home, we calculated the accuracy for all three possible cases: using just one calibrator in  $P_1$ , using just one calibrator in  $P_2$ , and using both calibrators in both phases. It should be noted that during all of our deployments, both calibrators were installed in both of the phases all the time. But as we recorded both of the functions  $F_1$  and  $F_2$  for  $P_1$  and  $P_2$ , respectively; we just used  $F_1$  to predict current in both  $P_1$  and  $P_2$  and  $F_2$  to predict current in both  $P_1$  and  $P_2$ . Table 2 shows the summary of all the deployment results.

From Table 2, we see that the average accuracy across all the deployments while using two calibrators is 95.0%. This shows the robustness of our system in predicting real power across different breaker panels and placement. It should be noted that [12] achieved an accuracy of 97.36%, however it only calculated apparent power step changes in a controlled environment with a few fixed appliances while our experiments calculated accuracy with real power obtained in real environment with natural electrical activities for longer durations.

#### EVALUATION UNDER DIFFERENT CONDITIONS

##### Effect of Placement on Accuracy

As we claimed, our system does not rely on the precision of placement of sensors. In all of our deployments, depending on the structure of the breaker the accuracy remained unaffected. To further analyze the positioning effect on accuracy, we conducted an experiment in a controlled environment. Creating a controlled environment for this experiment was necessary, as we did not want the accuracy to get affected by different electrical conditions.

##### Controlled Experiment Environment

We placed the sensor unit at 6 different locations on the breaker panel (See Figure 12). For each of the locations, we maintained a controlled environment as follows:

First we made sure that the environment is electrically quiet and no appliances are being turned on or off, after which we measured the baseline power consumption (C) of the

ID	Deploy ment time (Hour)	Power Range (Watt - Watt)	Using both phase calibration (%)			Using Phase 1 Calibration (%)			Using Phase 2 Calibration (%)		
			I <sub>RMS</sub> (Amp)	cos $\theta$	Power (Watt)	I <sub>RMS</sub> (Amp)	cos $\theta$	Power (Watt)	I <sub>RMS</sub> (Amp)	cos $\theta$	Power (Watt)
H1 (*)	168	252 - 4952	98.1	96.8	96.2	90.3	88.4	89.1	86.7	86.4	86.1
H2	48	396 - 6840	95.6	97.7	96.7	89.2	84.7	86.6	91.3	85.9	87.7
H3	48	598 - 6673	96.9	94.3	95.8	92.9	89.3	90.3	91.7	88.8	89.3
H4	48	707 - 12373	97.2	95.3	96.0	90.4	85.5	87.4	85.3	81.0	84.9
H5	48	441 - 5567	94.2	93.9	94.0	86.6	84.0	85.7	87.2	82.5	84.7
H6	48	311 - 4110	93.3	90.8	91.2	87.4	82.1	83.1	88.1	86.4	86.7
I1	48	1920 - 5982	96.8	91.6	95.2	83.1	78.3	80.1	84.3	81.1	82.9
Aggre gate	456	252 - 12373	96.0	94.3	95.0	88.5	84.6	86.0	87.8	84.5	86.0

Table 2: In-home deployment results. The results are shown with two calibrators (one on each electrical phase) and a single calibrator (\* 7 days deployment).

environment. We then turned on a 300 W load from the calibrator 3 times on top of the baseline and created a prediction function that works from C W to C+300 W. After that we turned on a 100 W fan, which brought the baseline to C+100 W. As we already have a function that worked from C to C+300 W, the prediction function was expected to perform well for the current load condition.

After 10 seconds we turned the load off. We then turned on a 1300 W heater and went through the same procedure as described in the previous paragraph, that is, we calibrated the system from C+1300 W to C+1600 W. Finally, keeping the 1300 W load on, we turned on a 500 W rice cooker and repeated the same procedure to calibrate from C+1800 W to C+2100 W.

minimum accuracy was 96.3% with an average accuracy of 97.4%. This experiment confirmed that our approach works independent of sensor position on the breaker panel with high accuracy.

**Accuracy without Active Calibrator**

The longer our system runs in a house, the wider and more accurate our calibrated regions become. As it covers more of the house’s power consumption range, the calibration frequency further decreases as well. Thus, as long as the power consumption resides within the calibrated region, the calibrator can be turned off with little effect on the overall accuracy. To test this hypothesis, we performed an additional experiment.

First, we ran our system for 24 hours in a home with all the existing appliances and the system calibrated for the region between 247 W – 5344 W, yielding an overall accuracy of 95.7%. Next, we turned the calibrator off and introduced four new appliances (two bulbs of 125 W and 250 W, one fan of 100 W, and a heater of 700 W) with different load profiles than that was used during calibration. We turned on each of the appliances, both individually and in combination while keeping the total power consumption within the calibrated range. This resulted in a small drop in accuracy to 94.2%. This experiment confirmed that even with the calibrator turned off and new appliances being introduced, the overall accuracy does not significantly deteriorate as long as the consumption resides within the previously calibrated region. Moreover, this experiment also shows that the generated function does not overfit based on existing appliances. It is flexible enough to work with any new appliance as long as the total consumption does not exceed the calibrated region.

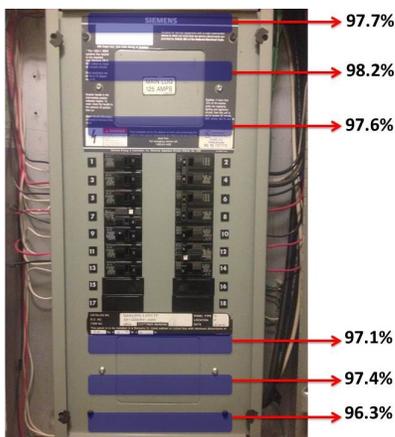


Figure 12: Accuracy on different positions. Min accuracy: 96.3%. Average accuracy: 97.4%.

**Accuracy Results**

Figure 12 shows the accuracy in each of the 6 positions. We observed that for all positions on the breaker panel, the

### Accuracy with Low Power Factor Loads

Low power factor loads such as Switch Mode Power Supply (SMPS) appliances consume power in higher order harmonics of 60 Hz power. As our sampling rate is 9.6 kHz, the magnetic sensors capture harmonic contents up to 4.8 kHz (79 harmonics). The key takeaway is that our magnetic pickup sensor is similar to CT in that both need to be sampled at a high enough sampling rate to capture 60Hz harmonics as desired by the end use case. The difference is that our sensor does not need to be around a wire and hence much easier to install, and for the same reason it needs to *learn* what the transfer function of the sensed magnetic field to the actual current flow is.

#### Analysis of energy in harmonics

We performed an experiment to understand how much power in a home can be attributed to the harmonics of a sensor where we could considerably reduce the 60 Hz. Our motivation for doing this analysis was to design a sensor where we could considerably reduce the engineering costs by reducing the sampling rate and data bandwidth requirements. We found that for a typical home over a period of a month, the 60 Hz harmonics contribute to only 0.15% of total power. This suggests that a simpler sensing system could be designed when only total power measurement is of concern to the end user, albeit with a ~0.15% loss in accuracy. Of course, the sensing system itself is capable of capturing all harmonics as mentioned above.

#### Deployment experiment with SMPS appliances

To further investigate the accuracy of our system with low power factor appliances, we performed a 7 day deployment in one of the homes with a bias towards SMPS appliances (two televisions, two laptops, an array of CFL bulbs, an active air conditioner, and frequent use of a microwave). This resulted in a small drop of  $I_{RMS}$  and  $\cos\theta$  accuracy, yielding 95.9% and 90.0%, respectively. The total power accuracy was 92.2%. This experiment further confirms our claim that the system also works with low power factor loads.

### DISCUSSION AND CONCLUSION

We have presented a practical approach to automatically calibrate a stick-on real power meter that can be installed by the homeowner without manual calibration. This system does not require professional installation service of opening the breaker panel and installing CTs around the main electrical feeds. To the best of our knowledge, the approach presented here is the first non-contact one to calculate phase angle between the current and voltage waveform. This allows for the calculation of real power as opposed to apparent power. In addition, our technique is also independent of sensor placement, which greatly reduces the installation effort required from end users. Most importantly, unlike the previous approach [12], our system is capable of predicting real-time, absolute power consumption in a home.

In-home deployments show an average accuracy of 95.0% across seven different places running naturalistic appliances. This result is encouraging and can be applied to a variety of scenarios such as energy disaggregation [2,8,13] and eco-feedback research [6]. This is especially true given the low installation burden, that is, the researchers are not even required to open their breaker panel. The current state of the art requires the users to buy and install CTs inside the breaker panel. Some of those systems<sup>1</sup> do not measure line voltage and therefore can be less accurate. Also the error of most CTs is 5-10% as defined by IEEE<sup>2</sup>. Moreover, their rated error is at a low current level. At higher current (*e.g.*, >2A, which is usually expected in a whole-home scenario), the errors are much higher.

To assess the energy viability of using the self-calibration approach, we also calculated the energy dissipated by the calibrator across all of our deployments. On average the energy use per home is just 0.181 kWh for the calibration to converge on the full transfer function. In our current prototype, we calibrate the system each time the consumption falls into an un-calibrated region. In future, we can easily envision a future deployment that uses a threshold and calibrate only when the consumption falls outside the threshold region. This will reduce both the number of calibration cycle as well as the energy dissipation. Furthermore, our calibrator was designed to pull up to a 300 W load; however, we found through post processing that it may be possible only go up to 50 W thus reducing the power needs during calibration.

For this prototype, we used four off-the-shelf magnetic pickup sensors. These pickup sensors are essentially inductors for all practical purposes. With the recent advent of ubiquitous circuit printing techniques [10], we envision printing an array of inductors on a flexible substrate as magnetic pickup sensors and sticking them on a breaker panel just like a sticker. Finally, to show the result of our technique to end-user, a display unit is needed to be install in some outlet of the home. We can foresee packaging the calibrator within this display unit, thus eliminating the need for an external calibrator unit.

### ACKNOWLEDGEMENTS

This work was partially supported by the National Science Foundation (Award: CNS-1253709).

<sup>1</sup> <http://www.smarthome.com/75505/Aeon-Labs-DSB09104-ZWUS-Z-Wave-Home-Energy-Monitor/p.aspx>

<sup>2</sup> [http://w3.usa.siemens.com/us/internet-dms/btlv/PowerDistributionComm/PowerDistribution/docs\\_MV/TechTopics/ANSI\\_MV\\_TechTopics91\\_EN.pdf](http://w3.usa.siemens.com/us/internet-dms/btlv/PowerDistributionComm/PowerDistribution/docs_MV/TechTopics/ANSI_MV_TechTopics91_EN.pdf)

## REFERENCES

1. Abbott, R. and Hadden, S. Product Specification for a Nonintrusive Appliance Load Monitoring System. *EPRI Report# NI-101*, 1990.
2. Gupta, S., Reynolds, M.S., and Patel, S.N. ElectriSense: Single-Point Sensing Using EMI for Electrical Event Detection and Classification in the Home. In *Proc. of UbiComp 2010*.
3. Bin, S. and Dowlatabadi, H. Consumer lifestyle approach to US energy use and the related CO2 emissions. *Energy Policy* 33, 2 (2005), 197–208.
4. Chen, J., Kam, A.H., Zhang, J., Liu, N., and Shue, L. Bathroom Activity Monitoring Based on Sound. (2005). In *Pervasive 2009*, 47–61.
5. Cooley, J.J., Member, S., and Vickery, D. A Retrofit 60 Hz Current Sensor for Power Monitoring at the Circuit Breaker Panel. .
6. Froehlich, J., Findlater, L., and Landay, J. The design of eco-feedback technology. In *CHI 2010*, 1999–2008.
7. Hart, G. Advances in nonintrusive appliance load monitoring. *Proceedings of EPRI Information and Automation Conference*, (1991).
8. Hart, G.W. and Member, S. Nonintrusive Appliance Load Monitoring. In *Proceedings of the IEEE*, 1992.80(12):p. 1870-1891.
9. Ho, B., Kao, H.C., Chen, N., et al. HeatProbe: A Thermal-based Power Meter for Accounting Disaggregated Electricity Usage. In *UbiComp 2011*.
10. Kawahara, Y., Hodges, S., Cook, B.S., and Abowd, G.D. Instant Inkjet Circuits: Lab-based Inkjet Printing to Support Rapid Prototyping of UbiComp Devices. In *UbiComp 2013*, 363–372.
11. Lorek, M.C., Chraim, F., Pister, K.S.J., and Lanzisera, S. COTS-based stick-on electricity meters for building submetering. *2013 Ieee Sensors*, (2013), 1–4.
12. Patel, S.N., Gupta, S., and Reynolds, M.S. The design and evaluation of an end-user-deployable, whole house, contactless power consumption sensor. In *CHI 2010*.
13. Patel, S.N., Robertson, T., Kientz, J.A., Reynolds, M.S., and Abowd, G.D. At the Flick of a Switch: Detecting and Classifying Unique Electrical Events on the Residential Power Line. In *UbiComp 2007*, 271–288.
14. Zeifman, M., Member, S., and Roth, K. Nonintrusive Appliance Load Monitoring: Review and Outlook. *57*, 1 (2011), 76–84.