AARPA: Combining Mobile and Power-line Sensing for Fine-grained Appliance Usage and Energy Monitoring

Nirmalya Roy*, Nilavra Pathak*, Archan Misra‡
*Information Systems, University of Maryland Baltimore County
† School of Information Systems, Singapore Management University
nroy@umbc.edu, nilavra1@umbc.edu, archanm@smu.edu.sg

***Abstract*—To promote energy-efficient operations in residential and office buildings, non-intrusive load monitoring (NILM) techniques have been proposed to infer the fine-grained power consumption and usage patterns of appliances from power-line measurement data. Fine-grained monitoring of everyday appliances (such as toasters and coffee makers) can not only promote energy-efficient building operations, but also provide unique insights into the context and activities of individuals. Current building-level NILM techniques are unable to identify the consumption characteristics of relatively low-load appliances, whereas smart-plug based solutions incur significant deployment and maintenance costs. In this paper, we investigate an intermediate architecture, where smart circuit breakers provide measurements of aggregate power consumption at room (or section) level granularity. We then investigate techniques to identify the usage and energy consumption of individual appliances from such measurements. We first develop a novel correlation-based approach called CBPA to identify individual appliances based on both their unique transient and steady-state *power signatures*. While promising, CBPA fails when the set of candidate appliances is too large. To further improve the accuracy of *appliance level* usage estimation, we then propose a hybrid system called AARPA, which uses mobile sensing to first infer high-level activities of daily living (ADLs), and then uses knowledge of such ADLs to effectively reduce the set of candidate appliances that potentially contribute to the aggregate readings at any point. We evaluate two variants of this algorithm, and show, using real-life data traces gathered from 10 domestic users, that our fusion of mobile and power-line sensing is very promising: it identified all devices that were used in each data trace, and it identified the usage duration and energy consumption of low-load consumer appliances with $\sim 87\%$ accuracy.**

Keywords: energy, plug loads, green building, mobile application

I. Introduction

There is widespread interest in developing solutions that provide knowledge of the *fine-grained* usage and power consumption of everyday appliances (such as coffee makers and televisions) in residential buildings. Such interest is primarily driven by recent interest in energy-efficient building operations, especially as empirical evidence suggests that empowering consumers with greater awareness of their energy consumption patterns can result in 5-20% reduction in electricity usage [17][18]. However, we believe that, besides this energy-related benefit, the ability to precisely capture the usage profile of everyday consumer appliances also provides insight into an individual's *context*, at a fine granularity that existing approaches (typically based on mobile sensing [25]) simply cannot provide. For example, while past approaches such as [5], [15] can help classify activities such as "making dinner" or "watching TV", appliance monitoring can additionally indicate that the 'toaster was used today' (revealing details about the food items consumed) or 'the specific TV channel watched' [23].

This paper thus explores the technical feasibility of a vision where the sensing capabilities of body-worn pervasive devices are combined with the power-line sensing of appliance usage to provide significantly greater insight into the daily activities (formally called Activities of Daily Living or *ADLs*) of individuals, especially in their residential environments. While the empirical investigations carried out in this paper utilize smartphones (that may or may not be always carried around inside a home), an eventual embodiment will likely rely on wearable devices (e.g., smartwatches, smart-bracelets [26]) that are now gaining wider market acceptance and that a user will likely wear almost-continuously [27].

In this paper, we first use real-life measurement studies to develop an enhanced Correlation-Based Power Analytics algorithm, called **CBPA**, that applies correlation over both macroscopic and microscopic power consumption features, to identify *the total usage duration*, and *the total energy consumption*, of individual devices, from such circuit-breaker level aggregated data. While CBPA helps to successfully disaggregate room-level power data into individual devices in *some* practical cases of interest, its accuracy diminishes if the candidate set of possible low-load devices becomes modestly large. Accordingly, we then explore a joint sensor fusion approach, that combines mobile plus power-line sensing data, to first obtain a smaller, *filtered* set of candidate appliances whose cumulative power consumption is reflected in the reading of the smart circuit breaker. We provide two different variants of this ADL-driven approach, called Activity-Aware Room-level Power Analytics (**AARPA**), one rule-based and the other probabilistically-weighted, and then use real-world usage traces to establish their efficacy. Our work thus establishes how a *joint fusion of mobile-sensing* *based ADL recognition and room-level power-line consumption data* can provide a practical solution that (a) helps capture the energy consumption characteristics of low-load, commonly-used domestic appliances and (b) provides useful additional context about the lifestyle habits and context of an individual.

II. Related Work

Our work touches on several areas starting from context-aware power signature analysis to building energy management based on plug load meters.

Appliance Power Signature Analysis: Non-intrusive load monitoring (NILM) algorithm was initially proposed by Hart [6] for discerning individual appliance power consumption from total power measurements. The initial technique proposed a cluster analysis approach, over a two-dimensional signature space of real and reactive power. However, the data acquisition system required for the obtaining and storing reactive power measurements is costly. The heuristic end-use load profile algorithm proposed in [13] records the occurrence, timing, and magnitude of large spikes in powerline and disaggregates only relatively large loads, e.g., air conditioners, using a 15 mins sampling dataset, which inevitably limits the range of other consumer appliances that could be detected.

Green Building Energy Management using Plug Load Meters: A building energy auditing network based on Wireless AC plug-load meter [14] smart plugs has been proposed in [4]. The MIT plug power meter platform provides apparent power measurements for profiling a load over a short or long time scales [10]. A growing interest for building energy monitoring system has been noticed recently in industry as well due to the several startups, such as EnergyHub [1] and Greenbox [2]. A non-intrusive approach that employs machine learning on data collected from infrastructure sensors, such as magnetic sensors, has been proposed to infer fine-grained power usage in home [8]. PowerPedia enables users to identify and compare the consumption of the plug-level domestic appliances through a smart phone app [20].

Smartphone and Sensor based Energy Prediction: An iPhone App called Beware [3] provides the user information on energy consumption of entire home. It can detect the electricity consumption of different devices and notify the user if the devices use more energy than expected. Energy Lens [11] provides deeper real time visibility of plug-load energy consumption in buildings. It uses the mobile phone to provide a consumer with real-time energy analytics. [12] proposed an ad hoc sensor system that can monitor appliance power usage by exploiting multi-sensor fusion and unsupervised machine learning algorithms. In summary, plug-load meter based approaches can easily achieve a detailed device-level energy footprint but for a steep deployment, operating and maintenance costs [21]. Our approach is synergistic with studies in [5], which employed activity recognition, principally using infrastructural sensors, to estimate the *aggregate* energy consumption in a smart home environment. In contrast, we focus on using mobile sensing for more practical recognition of Activities of Daily Living (ADL), and focus on estimating the energy consumption at a finer, *individual appliance-level*, granularity.

III. Appliance Signature Analysis

Given our focus on identifying relatively low (or medium) load devices, we first present a brief empirical study of the power consumption characteristics of some typical devices. The goal here is to establish that even everyday appliances with seemingly similar power profiles often posses unique *temporal characteristics*, that we can hope to leverage while attempting to disambiguate the consumption of multiple devices. The absolute values of the instantaneous power readings are not of prime relevance, as: (i) precise readings, as well as temporal patterns, will clearly vary across device manufacturers and models, and (ii) individual absolute values are hard to tease out from a power consumption measurement that is a sum of a large number of individual appliances/devices. More specifically, we study the similarities and differences in the power consumption pattern of one *device pair* with relatively quite similar behavior:

Refrigerator vs. steam iron: While seemingly quite different in their operation (and role in a consumer's daily life), both exhibit a cyclic pattern of power consumption: the fridge due to duty cycling of its compressor, with the iron due to intermittent deactivation by the thermostat. Accordingly, our comparison deliberately looks at *worst-case* scenarios where device pairs have highly similar behavior–other pairings of these 4 devices will be much easier to identify and separate. The measurements are conducted using Moteware Smart Plugs (ACMe) [14], running in an TinyOS environment with an Ubuntu 12.04.3 LTS system (additional details of our measurement experiments will be provided in experimental section).

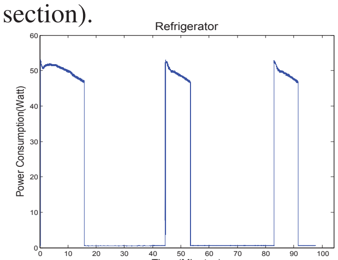

Figure 1. Refrigerator Power Consumption

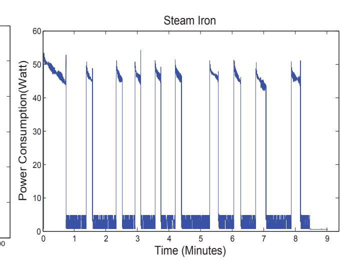

Figure 2. Steam Iron Power Consumption

Fig. 1 and 2 plots the power consumption pattern of the fridge against that of the steam iron. It should be clear that the consumption behavior exhibits the following similarities: (i) both devices exhibit a very *cyclical* consumption pattern, with each device being duty cycled for a significant fraction of the time (the fridge having its compressor turned off, while the iron having its heating coil turned off); (ii) both devices having significantly lower (at least 80%) lower power consumption during the “off” period of each duty cycle, and (iii) both devices have a longer transient duty cycle (with the fridge’s compressor being active to initially reach the pre-set temperature and the iron’s heating coil warming up to reach the desired temperature), as they ramp up from an initial state.

However, a finer-grained inspection of the patterns reveal some clear and insightful differences:

- The time period for the duty cycle is nearly constant for the fridge, while it exhibits irregular variation in the case of the steam iron. However, the “on” period for each duty cycle is fairly constant for both the fridge and the iron. This is really an artefact of the irregular pattern of human usage of the iron, in contrast with the fridge where human interaction is much less frequent and the behavior of the compressor is much more ‘regular’.
- Besides the variability, the average value of the time periods are markedly different. The time period for one cycle of the fridge is approx 38-40 minutes, while it is much shorter in the case of the iron (1-1.5 minutes).
- Finally, the power consumption during the “off” period of the duty cycle is quite different–it is essentially about 0.5W for the fridge, but around 3W for the iron.

The readings suggest that employing a technique that looks not just at initial transients, but over the regular operation cycle of each device, should be able to discriminate between these two devices, by taking advantage of their different temporal evolution patterns. Based on these insights, we now present our proposed CBPA algorithm, which seeks to utilize such temporal signatures in the power consumption characteristics of each individual device.

IV. Correlation-Based Power Analytics

The CBPA algorithm proposes to identify both the set of devices/appliances being used, and their individual power consumption, by exploiting both the *microscopic* and *macroscopic* features present in the time series of the total power consumption. *Microscopic* refers to the specific temporal nature of the signal waveforms and harmonics (e.g., as used in [9]), whereas *macroscopic* refers to power changes, etc as studied in [6]. However, capturing such microscopic features can itself present challenges–for example, to accurately capture harmonics, the Circuit Breaker needs to utilize a higher sampling rate, which in turn poses data transmission bandwidth and storage challenges. Similarly, monitoring the reactive power from the appliances is also a computationally intensive task [9]. To alleviate these problems, [9] proposed to use only transient signals for harmonic analysis. However, we are interested in not just detecting appliances via analyses of their transients, but also estimating their total energy consumption, implying that we need to analyze the steady-state operations as well.

To utilize the signatures present in both transient and steady-state phases of power consumption, our proposed CBPA utilizes a signal waveform analysis technique based on *Cross-correlation* [7], which captures the similarity of two waveforms as a function of a time-lag applied to one of them. Cross-correlation analysis is often used to detect the presence of a short-duration time series within a longer-duration signal. Mathematically, the cross-correlation (*xcorr*) sequence, defined between two jointly stationary random processes, x_n and y_n, with $-\infty < n < \infty$, is represented as:

$$R_{xy}(m) = E\{x_{n+m}y_n^*\} = E\{x_n y_{n-m}^*\} \quad (1)$$

where $E\{.\}$ is the expected value operator. *xcorr(x,y)* returns the cross-correlation sequence as a vector of size $(2 \times N - 1)$ vector, where x and y are length N vectors ($N > 1$). For continuous-valued signals, *xcorr* is computed via the *convolution* of two signals, by integrating as follows:

$$y(t) = x(t) \times h(t) = \int_{-\infty}^{\infty} x(\tau)h(t-\tau)d\tau \quad (2)$$

We then apply the correlation technique mentioned above to try and identify each individual appliance from the total load measurements. A key aspect of our methodology lies in our modeling and use of both the transient and steady-state phases: in particular, peaks in the correlation function often help to identify distinct ‘events’ associated with each appliance (e.g., a fridge’s compressor turning off or on); such peaks correspond to key transient characteristics. However, having isolated these peaks, we then model the steady state load as well to compute the energy consumed over the entire steady-state.

A. Generating Individual Power Signatures

In our approach, we first individually measure the load behavior (i.e., the characteristic power consumption curve) for each appliance, and thus create a device-specific *model* or signature of the power consumption. For obtaining the characteristic power consumption of a device or appliance, the device was plugged into an ACMe plug load meter, and monitored for a duration long enough to capture both transient and steady-state phases. We wrote a generic code to collect the relevant characteristic information from these measurements. For example, for the case of a mini-refrigerator, the fridge was left on for at least three compressor ON cycles. In this case, the characteristic information retrieved includes (i) the characteristic energy value when the compressor is ON, (ii) average energy consumption value when the compressor is OFF, (iii) time period of one cycle comprising both successive compressor ON and OFF states, and (iv) the energy consumption during the transient period of the compressor turning OFF. Similarly, generic codes are implemented for collecting characteristic power curves from other respective devices. To generate accurate characteristic curves, it is important to not only measure for a sufficiently long duration, but to also avoid usage artefact during the measurement phase.

B. *Feasibility of the CBPA Approach*

To test the possible use of the correlation-based CBPA approach, we utilize a real dataset collected using the Smart Plug over a 10 hours time period (614 mins) as shown in Fig. 3 from the appliances described below. To generate the ground truth about the device usage, we also collected the precise usage times of each device. a) Mini-Refrigerator is ON from 0 to 614 minutes, b) Fluorescent lamp is ON from 104 to 448 minutes, c) Mobile charger is ON from 145 to 344 minutes, d) Steam Iron is ON from 396 to 440 minutes, e) Table Fan is ON at level 1 speed from 416 to 494 minutes, at level 2 speed from 495 to 531 minutes, at level 3 speed from 532 to 573 minutes.

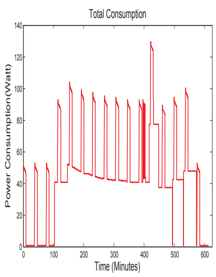

Figure 3. Aggregated Power Consumption

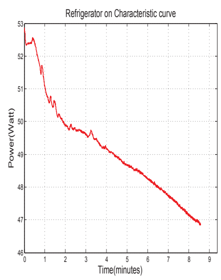

Figure 4. Refrigerator Characteristic Power Curve

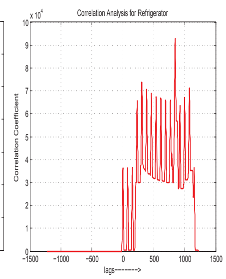

Figure 5. Refrigerator Correlation Analysis Curve

Fig. 3 plots the aggregate load, as measured by the plug load meter. To evaluate the effectiveness of our proposed CBPA technique, we start by trying to separate out the energy consumptions of refrigerator. For this we use correlation between the aggregate data set and characteristic curve generated by one of the load modeling function corresponding to refrigerator. Wherever there is a close enough match of aggregate data set curve (Fig. 3) and refrigerator characteristic power consumption curve (Fig. 4), peaks are generated at those points (Fig. 5). These peak points represent the location wherever refrigerator's Compressor ON state's consumption is present. Accordingly, we simply retrieve the consumption profile from the refrigerator's signature and align the consumption curve to those specific time instances; we then subtract the refrigerator's resulting *estimated* consumption pattern from the total power measurement to get the aggregated consumption of the *residual* devices, namely the steam iron, the mobile charger, the lamp and the fan. By similarly applying the correlation-based technique iteratively, we recover the energy consumption for both the steam iron and for the mobile telephone charger. Finally, we repeat the process to separate out the consumption patterns of the lamp and the fan. This approach proved to be fairly successful for this *somewhat arbitrary mix* of appliances, enabling us to recover the energy consumption and usage times for each appliance fairly accurately. As an illustration (to provide a unified view, detailed numerical results are deferred to experimental section), Fig. 6 shows the *ground truth* total energy consumption excluding the refrigerator, whereas Fig. 7 shows the total energy consumption excluding the *CBPA-computed* energy consumption of the fridge.

While this approach is successful in some cases, it turns out to be incapable of accurately estimating the energy consumption for many other combinations of appliances–in general, larger the set of possible appliances, the poorer the result. Motivated by these empirical observations, we next look into the AARPA approach of first using mobile sensing to infer an individual's ADL context, and using such context to reduce the set of appliances likely being used.

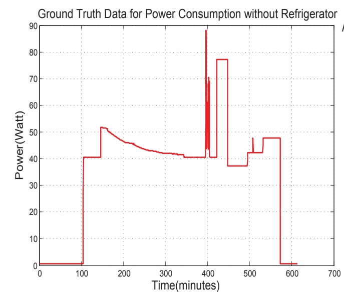

Figure 6. Total Power Consumption excl. Refrigerator (Ground Truth)

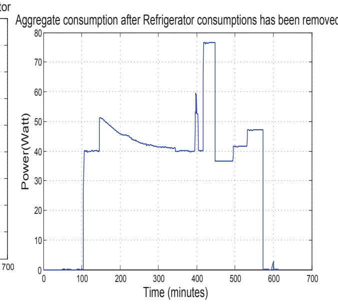

Figure 7. Total Power Consumption excl. Refrigerator (CBPA)

V. Activity-Aware Room-level Power Analysis

Our key idea is to leverage upon the large body of work on pervasive/mobile sensing, for inferring both low-level locomotive context [16], as well as high-level ADLs [15]. To be clear, this paper does not attempt to innovate in the domain of activity recognition, but instead borrows prior techniques that (i) apply a hierarchical classification model [15] to detect various ADLs and (ii) uses Wi-Fi measurement data to determine an individual's location at room-level granularity [19]. Mobile phones themselves have also been used as part of systems for energy attribution in green buildings (e.g., [11]), but more as an information presentation platform, rather than a sensing device.

A. *Appliance-Aware Activity Model*

As the second-stage of the AARPA framework, we then attempt to derive the set of *appliances* associated with each such ADL. More specifically, we use a-priori training data to associate each element $e, e \in \mathcal{A}$, with a set of appliances that the individual is likely to use (or, more generally, have the appliance be turned on for at least a portion of the overall activity duration) while engaging in that activity. Accordingly, the key element of AARPA involves a *mapping* from each element $e \rightarrow S(e)$, where $S(e)$ is a subset of the overall set of appliances $\mathcal{S}$= $\{microwave, toaster, steamiron, TV \ldots\}$. Once we have the smaller set $S(e)$, we then apply the CBPA technique described earlier, using this set $S(e)$ and the overall power consumption data obtained from the corresponding Smart Circuit Breaker. In this paper, we propose and explore two-variants of AARPA, that differ in the way the set $S(e)$, for each ADL e, is represented and used:

- *Rules based:* In this approach, called *RPA*, we collect the total set of appliances used (over the training phase) by an individual when engaging in a specific ADL. More specifically, the set $S(e)$ for any ADL e consists of all the appliances used during *any instance* of e, even if the appliance may have been used only in a small set of such instances, or for a very small duration.
- *Usage Weighted:* In this approach, called *WPA*, we compute the probability (or likelihood) $w(e,i)$ of a specific appliance i being used during a specific instance of ADL e, by observing and computing the average fraction of time that i is observed to be used during each instance of e. In this case, the CBPA algorithm is modified to use not just the absolute maximum correlation value described previously, but instead multiplies the correlation value of each appliance with its corresponding likelihood $w(e,i)$. The WPA technique can be viewed as a Bayesian analogue of RPA, with the identification of a specific appliance being weighted by the *a priori* likelihood of that appliance being used during a specific ADL.

Fig. 8 provides the high-level pseudocode for the overall AARPA technique, summarizing both the RPA and WPA variants. (For RPA, the $w(e,i)$ values are all set to 1.) Next we shall use empirical investigations to study the benefits and performance characteristics of these two AARPA variants, and compare them with the baseline version of CBPA (which includes in the candidate set *any* appliance attached to the circuit breaker from which the total power consumption details are obtained).

```
Procedure AARPA (input: aggregate power data (P)
output: appliance identification & power enumeration)
1. Empirically calculate appliance usage based weighting
   factor from trace data;
   1.1 For (e ∈ A) & (i ∈ S) {
          1.1.1 Compute w(e,i) = d_ei / Σ_{i=1}^{m} d_ei
   where w(e,i) is the usage weighting factor of e^th
   activity with i^th appliance and d_ei is the usage time
   duration of i^th appliance in association with e^th
   activity.
2. End-For
3. RP = P; // residual total power
4. For (i ∈ S) {
5.     Compute correlation δ_i = xcorr(RP,i);//convolution
6. End-for
7. Pick the appliance î with highest correlation based on
   δ_i × w(e,i)/sqrt(Σ w(e,i)^2);
   //correlation multiplied by a normalized weighting
   //factor
8. Subtract appliance î from S(e) //set-theoretic
   and its characteristic power from RP
9. Go to step 7 and repeat the process until S(e) = φ.
```

Figure 8. The AARPA Algorithm

VI. Experimental Setup and Results

In this section, we report on our experiments that investigate the benefit of the proposed (AARPA) techniques. Our experiments are conducted with real-life data traces of (i) appliance power consumption and (ii) smartphone sensor readings while participating in ADLs, collected from 10 users living in typical apartments and town home complexes.

A. Smart Plug Setup and Data Collection

The Moteware ACMe [14] was used as a plug load meter in this work. The ACMe is based around a Texas Instruments MSP430 microcontroller where power is monitored by an SPI-controlled Analog Devices ADE7753 power monitoring IC, attached with a Hall Effect sensor. The ACMe was plugged into an electrical wall socket; while an individual device was then directly plugged into this plug, an extension cord was used to connect multiple devices. The plug load meter measured and transmitted the power consumption data in watts (volt-amps), along with a timestamp, at a frequency of 16 Hz, both through a serial and wireless interface to the basestation, which generated a corresponding .csv file for subsequent analysis.

B. Android-based App Development and Data Collection

We designed an application to collect accelerometer and gyroscope data from an Android based Google Nexus smart phone device for monitoring the activity and appliance usage behavior of a typical user. It also asks the user to manually tag the semantic name of the location (such as bedroom, kitchen, living room etc) and the specific activity being performed to aid in labeling the data, as well as the (start, stop) times of individual appliances. The resulting data was stored in ARFF file format, for subsequent processing using the Weka toolkit [22]. We collected samples of ten users from six homes performing a variety of both kitchen-related ADLs (such as making breakfast, preparing dinner and washing dishes) and living-room related ADLs (watching TV, etc.). We collected samples for time periods between five to sixty minutes based on a specific activity, with sensor data collected at 80 Hz. While conducting each ADL, the users were free to place the smartphone in the on-body position of their choice.

C. Activity-Aware Power Signature Analysis

We investigate the issue of whether this activity-aware power signature analytics approach really helps to improve the detection and measurement of multiple appliances' power consumption. In particular, we experimented with 3 different strategies, which differ in whether or not, and how, they use of the additional room-level activity information to reduce the set of candidate appliances used in the CBPA algorithm. For each algorithm, we computed both the *duration error* (the difference between the ground-truth usage duration and that reported by AARPA), and the *energy* *consumption error* (the difference between the ground-truth and the AARPA output).

Exhaustive Power Analysis (EPA): In this approach, we consider all appliances in a particular location as members of the candidate set, without regard for their use during a specific ADL. For example, for the kitchen area, the set consists of all of the appliances: {microwave, toaster, coffee maker, boiler, hand mixer, and a grinder}, even if the user is "making breakfast" and has never used a hand mixer during this ADL. For this set of appliances, we note that the characteristics of the *hand mixer* were very similar to that of the *coffee-maker*, making disambiguation via CBPA very difficult. Overall, the EPA approach resulted in fairly high errors (see Table I for details), with average (across all appliances) duration and energy consumption errors of $\approx 35\%$ and 36% respectively.

Rule-based Power Analysis (RPA): In this approach, as explained earlier, the candidate set of appliances was defined a-priori for each separate ADL. For example, using ground-truth data about usage patterns, the "making breakfast" ADL is associated with the smaller appliance set: {microwave, toaster, coffee maker, and boiler}, and excludes the {hand mixer, grinder} devices. The resulting duration and energy consumption errors are lower than that achieved by EPA (about 31% and 22%) respectively.

Weighted Power Analysis (WPA): In this weighted-based approach, we additionally compute a *usage weight* for each candidate appliance (identified by the RPA process), as shown in Line 1.1.1 of Algorithm in Fig. 8, and use this weight to boost the correlation value. In this case, the average errors in duration and energy consumption are sharply reduced, to approx. 12% and 13% respectively, attesting to the promise of this approach.

Table I
PERFORMANCE METRIC

Power Analytics Methods	Start/End Time Error (%)	Total Energy Consumption (Joule) Error (%)
EPA	34.78	36.17
RPA	31.25	22.12
WPA	11.74	13.27

VII. CONCLUSION

In this work, we have advocated an intermediate approach for NILM-based monitoring of both the *usage episodes* and *energy consumption* of relatively low-load domestic appliances, that utilizes smart circuit breakers to measure the total power consumption at room (or sections) level granularity. We have developed a novel correlation-based analytics algorithm, CBPA, to identify the precise usage duration of, and the overall energy consumed by, appliances such as table fans and coffee makers based on both their steady-state and transient power characteristics. To overcome CBPA's limitations when the set of candidate appliances becomes moderately large, we then propose an analytics approach, called AARPA, that fuses pervasive/mobile sensing and high-level ADL recognition with such circuit breaker-level power readings. AARPA employs mobile sensing to first obtain a reduced set of candidate appliances likely to be used during an ongoing ADL episode, before applying the CBPA technique. Results from a set of 10 users show that a probabilistically-weighted variation of CBPA shows great promise in identifying the usage of such everyday appliances very accurately (every appliance usage episode was correctly inferred), and provides fairly accurate estimates (average error of around 13%) of both their usage duration and energy consumption.

ACKNOWLEDGEMENT

This work is supported partially by the NSF Award #1344990, and Constellation E^2: Energy to Educate Grant. We would like to thank Dr. Behrooz Shirazi for his valuable comments and feedback on this work.

REFERENCES

[1] EnergyHub, http://www.energyhub.com/

[2] Greenbox, http://www.getgreenbox.com/

[3] BeAware. 2010. http://www.energyawareness.eu/beaware/

[4] X. Jiang, M. V. Ly, J. Taneja, P. Dutta, and D. Culler, "Experiences with a High-Fidelity Wireless Building Energy Auditing Network", ACM SenSys 2009.

[5] C. Chen, B. Das, and D. Cook, "Energy prediction based on resident's activity", KDD 2010.

[6] G. W. Hart, "Nonintrusive Appliance Load Monitoring," Proceedings of the IEEE, Vol. 80, pp. 1870-1891, 1992.

[7] Steven Kay, "Fundamentals of Statistical Signal Processing", Prentice Hall

[8] Y. Kim, T. Schmid, Z. M. Charbiwala, and M. B. Srivastava. "Viridiscope: Design and implementation of a fine grained power monitoring system for homes". ACM UbiComp 2009.

[9] C. Laughman, et al., "Power Signature Analysis," IEEE Power & Energy Magazine, pp. 56-63, March/April 2003.

[10] J. Lifton, M. Feldmeier, Y. Ono, C. Lewis, and J. A. Paradiso, "A platform for ubiquitous sensor deployment in occupational and domestic environments", ACM IPSN 2007.

[11] J. J. Ortiz, Y. Noh, G. Saldanha, D. Su, and D. Culler, "Towards Real-time, Finedgrained Energy Analytics in Buildings Through Mobile Phones", ACM BuildSys 2012

[12] D. E. Phillips et. al, "Supero: A sensor system for unsupervised residential power usage monitoring", IEEE PerCom 2013

[13] J. Powers, B. Margossian, B. Smith,"Using a Rule-Based Algorithm to Disaggregate End-Use Load Profiles from Premise-Level Data," IEEE Computer Applications in Power, pp. 42-47, 1991.

[14] ACMe Calibration. Moteware. http://www.moteware.com/acmecal.php

[15] N. Roy, A. Misra, and D. Cook, "Infrastructure-assisted smartphone-based ADL recognition in multi-inhabitant smart environments", IEEE PerCom 2013.

[16] J. Kwapisz, G. Weiss, and S. Moore, "Activity recognition using cell phone accelerometers", *KDD 2010, pp. 10-18.*

[17] M. Chetty, D. Tran, and R. E. Grinter. "Getting to green: understanding resource consumption in the home," In UbiComp, pp. 242–251, 2008. ACM.

[18] J. Dobson and J. Griffin. Conservation effect of immediate electricity cost. feedback on residential consumption behavior. Proceedings of the 7th ACEEE Summer Study on Energy Efficiency in Buildings, 1992.

[19] G.V. Zaruba, M. Huber, F.A. Kamangar, and I. Chlamtac, "Indoor Location Tracking Using RSSI Reading from a Single Access Point", ACM/Springer Journal of Wireless Networks (WINET), vol. 13:2 April 2007.

[20] M. Weiss, T. Staake, F. Mattern and E. Fleisch, "PowerPedia - Changing Energy Usage with the Help of a Community-based Smartphone Application", Journal of Pervasive and Ubiquitous Computing, 2012.

[21] A.G. Ruzzelli, C. Nicolas, A. Schoofs and G.M.P. O'Hare, "Real-Time Recognition and Profiling of Appliances through a Single Electricity Sensor", In Proc. IEEE SECON 2010

[22] I. Witten, and E. Frank, "Data Mining: Practicial Machine Learning Tools and Techniques with Java Implementations," *Morgan Kaufmann, 1999.*

[23] U. Greveler, B. Justus, "Multimedia Content Identification Through Smart Meter Power Usage Profiles," *In Computers, Privacy and Data Protection CPDP, Brussels, Belgium, 2012*

[24] D. Cook, A. Crandall, B. Thomas, and N. Krishnan, "CASAS: A smart home in a box," *IEEE Computer, 46(6):26-33, 2013*

[25] N. Lane, E. Miluzzo, H. Lu, D. Peebles, T. Choudhury, and A. Campbell, "A Survey of Mobile Phone Sensing," *IEEE Communications, Sept. 2010*

[26] Huawei Smart Bracelet, http://www.huawei.com/us/index.htm

[27] Intel Make it Wearable, https://makeit.intel.com/