Unsupervised Approximate Power Trace Decomposition Algorithm

Kyle Anderson and José M.F. Moura
Electrical and Computer Engineering Department
Carnegie Mellon University
{kda, moura}@ece.cmu.edu

Mario Berges
Civil and Environmental Engineering Department
Carnegie Mellon University
marioberges@cmu.edu

***Abstract*—We propose the Approximate Power Trace Decomposition Algorithm for decomposing a power signal into components based on power consumption levels. Such a decomposition provides users with more detailed information about their energy consumption without attempting the difficult problem of disaggregating all of the devices in a building. Much of the NILM literature has focused on the problems of event detection and not treated the full problem of energy disaggregation; we have relaxed the problem by moving away from individual appliance disaggregation and provided a solution that makes estimating the energy used by different classes of devices feasible. We present preliminary results using phase A of the BLUED dataset.**

I. Introduction

Over the past decade, Non-Intrusive Load Monitoring (NILM) has attracted renewed attention from both industry and academia. Despite the revival of interest in this technology, there is still a lack of cohesive performance metrics that are accepted and there is relatively little work that deals with actual energy disaggregation. Much of the literature deals primarily with detection and classification of transients but not with the important and most useful step of actually reporting energy consumption.

One of the main challenges to providing a complete device-level disaggregation on the energy level is the sheer number of devices present in homes. This number may be tempered by work that has shown that in the US 12 devices account for 80% of average household electrical loads [1], but the task of disaggregation on this many devices, even were they the only ones active, would still be challenging. The situation in real deployments is made more difficult by the fact that they must be separated from the larger number of devices in the home that have less overall energy consumption.

These challenges have caused the authors to consider relaxations to the NILM problem that would create a problem that is more easily solvable while still providing valuable information to the user. To this end, we propose the Approximate Power Trace Decomposition Algorithm (APTDA) that, instead of providing an estimate of appliance-level energy consumption, provides estimates of the energy consumed by devices operating in different ranges of power consumption. Reasonable ranges can be learned from the power data and require no user feedback or input.

The organization of the paper is as follows. Section II defines the NILM energy disaggregataion problem and introduces the device power-consumption based relaxation we propose. Section III gives a detailed overview of the Approximate Power Trace Decomposition Algorithm (APTDA). In section IV we present preliminary results for the APTDA algorithm and provide a brief discussion of how they might be used and of challenges in establishing metrics for verifying our method. Finally, in section V, we conclude and discuss challenges that remain to be addressed by our method as future work.

II. NILM Problem and Relaxation

In general, the non-intrusive load monitoring problem is to decompose the aggregate energy consumed over a time period of interest into energy consumed by relevant components during that time. The overall energy consumed by electrical devices in a home is a superposition of their individual energy consumptions. So, we consider a disaggregation model that is additive, namely:

$$E^{[1,T]} = \sum_{k=1}^{K} E_k^{[1,T]} \tag{1}$$

This equation decomposes the total energy consumed during the time period $[1,T]$, $E^{[1,T]}$, into K components, $E_k^{[1,T]}$. What these components are and how many to choose depends on the application and the end goal. In the remainder of the discussion, we simplify notation by dropping the time interval $[1,T]$ with the understanding that when energy is mentioned it must refer to energy consumed over a fixed period of time.

There are many potentially interesting decompositions, and we discuss a few here. In the full appliance disaggregation problem, K is equal to the number of electrical devices in the installation and each E_k represents the energy consumed by each single device. A second possibility is an activity-based decomposition where each E_k represents the energy consumed by a different activity such as cooking, watching television, heating/cooling the home, etc. Next, we could have a room-based decomposition where K is the number of rooms in the home and each E_k represents the power consumption of an individual room. One last example we mention is a person-based grouping, where each E_k represents the energy consumed by a particular individual in the home. We note that not all of the scenarios presented here fit perfectly into

k	Power consumption range (Watts)
0	Background
1	$0 - 105$
2	$105 - 720$
3	$720+$

TABLE I
POTENTIAL POWER CONSUMPTION RANGES FOR A POWER CONSUMPTION-BASED DECOMPOSITION.

the model in (1), given that they are not all strictly additive, but the main point of potential alternative decompositions is still valid.

Apart from these examples, there are certainly many other conceivable decompositions that may be of interest to consumers, some more feasible to achieve than others. We are particularly interested in decompositions that exploit the nature of the power signal and make few assumptions about what devices, rooms, or people are or are not present in the home.

With this in mind, we propose an energy decomposition by power consumption level. We consider K ranges of power consumption and define each E_k to represent the energy consumed by appliances that consume power in the k^{th} range. For example, the power consumption ranges could be as in table I.

Using the ranges in table I, then, E_0 represents the total energy used by background devices, E_1 represents the total energy used by devices that consume power in the range of 0-105 Watts, E_2 represents the total energy used by appliances that consume power in the range of 105-720 Watts, and E_3 represents the total energy used by devices that consume power at a rate of greater than 720 Watts.

This approach is advantageous for many reasons. First, the power ranges can be learned from the data without any training or input from the users. Second, while the ranges do not specifically tell the user which devices are in each range, the information can be used to suggest to the user which devices may be in each range based on standard power consumption levels for various appliances. Third, the energy usage of each component of the decomposition can be compared over different time intervals, for example from one week or month to the next. Fourth, there is no need to have knowledge of the number of devices in the building: the number of ranges to use, K, can also be learned from the data.

We can see in figure 1 what the APTDA algorithm produces as its output. A piecewise-constant approximation, $\overline{P}$, is made of the power signal and this $\overline{P}$ is decomposed by the APDTA algorithm into the four signals shown in figure 1. The top signal, $\overline{P}_0$ is the estimate of background power consumed in the building during the week. The following signals, $\overline{P}_1$, $\overline{P}_2$, and $\overline{P}_3$, are estimates of the power consumed by devices operating in each of the 3 ranges from table I. The sum of the K decomposed signals (in our example $K = 4$), as shown in (2), is equal to $\overline{P}$. We refer to $\overline{P}_0$ as the background

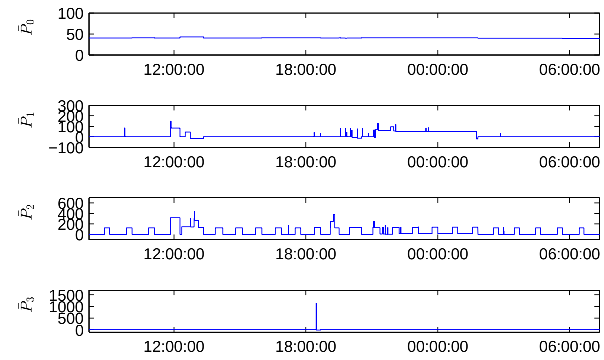

Fig. 1. These plots show a portion the final decomposition of the APTDA algorithm on phase A of the BLUED dataset.

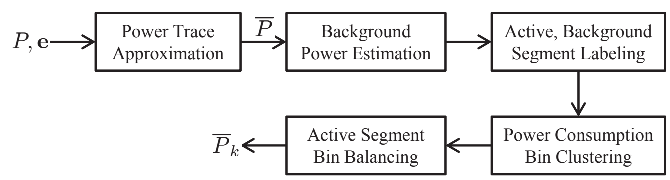

Fig. 2. Block diagram for the APTDA algorithm.

component and $\overline{P}_1, \ldots, \overline{P}_{K-1}$ as active components.

$$\overline{P} = \sum_{k=0}^{K-1} \overline{P}_k \qquad (2)$$

III. APTDA OVERVIEW

The APTDA algorithm takes as input the power trace of a building and a set of event locations and outputs a decomposition of the power signal into components based on how much power devices in different ranges use, as discussed in section II above. In this section we provide a brief sketch of the algorithm using the block diagram of figure 2 as our guide.

In the first stage of the APTDA algorithm, "Power Trace Approximation", the original power signal, P, and a set of event locations, $\mathbf{e}$, are used to compute a piecewise-constant, energy-preserving approximation of the power signal, $\overline{P}$. Figure 3 shows an example input and output for this stage taken from phase A of the BLUED dataset [2]. The event locations serve as starting points to create the endpoints of the piecewise-constant segments in $\overline{P}$.

The next two stages in the algorithm, "Background Power Estimation" and "Active, Background Segment Labeling," deal with the separation of background and active power consumption by devices. The background power level in a building is the amount of power consumed by devices that are always left on or in a standby mode. By "active power" we mean the power consumed by appliances whose turning on and off is observable in the power signal. The background power level is easily identified visually as the lowest level to which the power signal always returns. To compute the level, however, we use

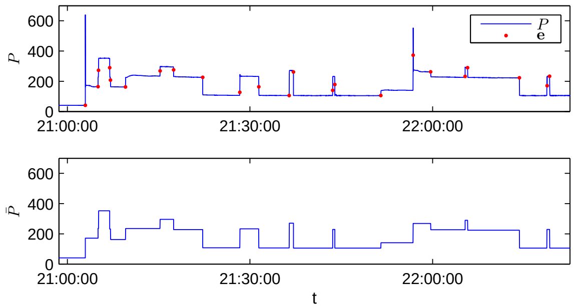

Fig. 3. Top: A portion of the original power signal P with event locations $\mathbf{e}$. Bottom: The piecewise-constant approximation $\overline{P}$.

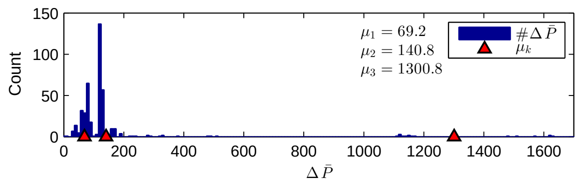

Fig. 4. Histogram of the positive changes in $\overline{P}$ from phase A of the BLUED dataset. The red triangles denote the cluster centroids, μ_k, found by applying the k-means algorithm to all of the $\Delta\overline{P} > 0$ for $k = 3$.

an approach similar to [3], where percentiles are used. For phase A of the BLUED dataset the background power level remains consistent throughout the week and is approximately 40 Watts.

The background power consumption is a useful quantity for the consumer to know but it is also important in the APTDA algorithm for distinguishing which portions of the power signal represent consumption only due to background devices and which portions represent power consumption due to both background and active devices. Recalling (2) above, we can see that during background segments the active components, $\overline{P}_1, \ldots, \overline{P}_{K-1}$, must be identically zero since there is no active power consumption. This means that each active segment can be treated separately since the active components must begin from zero and return to zero during each active segment.

The next stage, "Power Consumption Bin Clustering," deals with learning the power consumption ranges and beginning to form the active components $\overline{P}_1, \ldots, \overline{P}_{K-1}$. Each change in $\overline{P}$ denotes activity corresponding to an electrical device in the building. If the change is positive then we assume that a device has turned on (or if it was already on, entered a state in which it consumes more power). If the change is negative, then a device has turned off (or decreased its power consumption). Figure 4 shows the histogram of all of the positive changes in $\overline{P}$ on phase A of the BLUED dataset.

A k-means clustering is performed on these $\Delta\overline{P} > 0$. The centroids of these clusters, μ_k, determine the power consumption ranges for the active components. Specifically, the limits of the power ranges for the active components are determined by the midpoints between the μ_k. Thus, with the

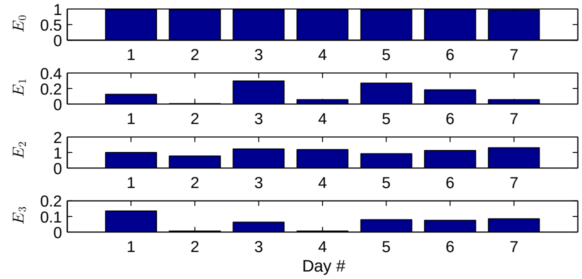

Fig. 5. Energy in kWh consumed by each component per day in phase A of the BLUED dataset.

μ_k as in figure 4, the ranges for the $\overline{P}_k$ are 0-105 Watts, 105-720 Watts, and greater than 720 Watts. Once the power ranges have been determined, each of the $\Delta\overline{P}$ is assigned to the range that its absolute value belongs to. We also refer to the power consumption ranges as "power consumption bins."

The final stage of the algorithm, "Active Segment Bin Balancing," corrects the assignments performed in the previous stage so that each active component satisfies certain criteria that must be imposed in order to ensure that a reasonable decomposition is obtained. Examples of these criteria include, among others, non-negativity in each active component and net positive energy consumption of each component during each active segment.

To recap, the APTDA algorithm begins by forming a piecewise constant estimate, $\overline{P}$, of the original power trace, P, and a given set of events, $\mathbf{e}$. After estimating the background power consumption of the building, $\overline{P}$ is split into segments of either active or background power consumption. Next, the distribution of all the changes in $\overline{P}$ is analyzed in order to establish relevant power consumption ranges. Finally, these ranges are used to process each of the active segments and decompose $\overline{P}$ into relevant components useful for tracking the total energy consumed by devices operating in each power range.

IV. Results and Discussion

The output of the APTDA algorithm can be used to view the breakdown of energy consumed by different components during selected time intervals. Figure 5 shows how much energy each component uses per day of phase A of the BLUED dataset. As expected, the background component, E_0, is very consistent during the week. Component E_2, those devices consuming between 105 and 720 Watts, consumes much more energy than either components 1 or 3. In these results, the ground truth events provided with the BLUED dataset were used and the power was computed at a rate of 1Hz.

We have also done analysis using events from the modified generalized likelihood ratio (GLR) event detector described in [4] and [5]. Using power from phase A of the BLUED data sampled at 1Hz, we varied the parameters of the GLR detector for a total of 284 combinations. In the BLUED dataset on phase A there are 872 events throughout the week and the detectors return between 724 and 3,014 events. Different

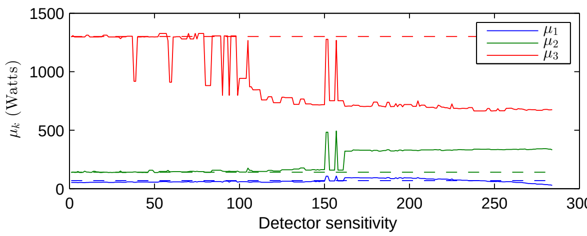

Fig. 6. The centroids of the clusters obtained by k-means ($k = 3$) clustering on the $\Delta\overline{P}$ for event detectors ranging from low to high sensitivity. The dashed lines show the cluster centroids when the ground truth events were used.

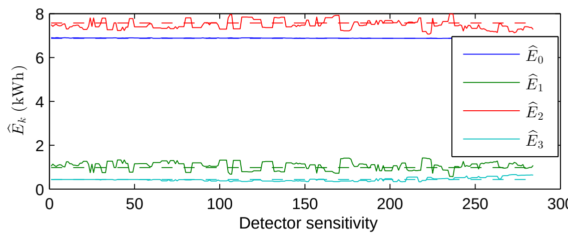

Fig. 7. The energy reported for the entire week by the APTDA algorithm run with event detectors ranging from low to high sensitivity and with the power ranges forced to be the same as those found using the ground truth events.

sets of events create different power traces $\overline{P}$ and this affects the centroids found by the k-means clustering, which in turn means that the power ranges used may be different.

Figure 6 shows how the cluster centroids change depending on the sensitivity of the event detection algorithm. We see that the least sensitive detectors maintain cluster centroids much closer to the centroids found when using the ground truth algorithm (the dashed lines). We note that the number of clusters was fixed at 3 to make comparisons easier.

The variability of the cluster centroids makes it impossible to compare the energy attributed to each component directly because the power ranges are not the same. However, when the ranges learned from the ground truth edges are forced on all of the event detectors we see in figure 7 how the energy reported for each range compares. Here, the dashed lines represent the energy reported by using the ground truth events, and the hat on the $\widehat{E}_k$ denotes that the energies pertain to forcing the algorithm to use the power ranges of the ground truth edges. Note that the reported energy for each component is very robust to the various event detectors.

One of the primary challenges of evaluating the APTDA algorithm is that it reports quantities that are not directly observable. The $\overline{P}_k$ are not estimates of actual quantities but are constructs for grouping devices of similar power consumption together. Furthermore, it is possible that a single device may have modes that operate in different power consumption ranges. In this case we can also see that it is not enough to simply sum the individual consumption of devices in each range to obtain meaningful ground truth for evaluating the algorithm.

With this said, one of the main strengths of the APTDA algorithm lies in its usefuleness for comparing results over different time periods rather than its absolute decomposition. By comparing consumption day to day, week to week, or month to month, users can see how their consumption pattern changes for the different consumption ranges. This can help them to infer and identify which devices may be responsible for reduced or increased consumption.

V. Conclusions and Future Work

We presented the Approximate Power Trace Decomposition Algorithm that decomposes a power signal into components based on power consumption levels. Such a decomposition is useful for providing users with more detailed information about their energy consumption without attempting the difficult problem of disaggregating all of the devices in a building. Much of the NILM literature has focused on the problems of event detection and not treated the full problem of energy disaggregation; we have relaxed the problem by moving away from individual appliance disaggregation and provided a solution that makes estimating the energy used by different classes of devices feasible.

There are many ways that we would like to improve and extend this work. The results presented here are preliminary but show potential. Moving forward, we hope to analyze the method's sampling frequency sensitivity to determine the lowest acceptable sampling rate. We would also like to explore creating a joint event detection and power trace approximation algorithm that does not assume fixed events but jointly solves both problems. Presently, the algorithm does not account for changes in background power consumption, but this is an important problem for making the algorithm applicable in real settings. Finally, we also plan to incorporate reactive power, Q, to provide additional information for the decomposition.

Acknowledgment

The authors would like to gratefully acknowledge the support from the National Science Foundation (NSF) grant #09-30868. The opinions expressed herein are those of the authors and not of the NSF.

References

[1] D. Carlson, H. S. Matthews, and M. Berges, "One size does not fit all: Averaged data on household electricity is inadequate for residential energy policy and decisions," *Energy and Buildings*, vol. 64, pp. 132–144, September 2013.

[2] K. Anderson, A. Ocneanu, D. Benitez, D. Carlson, A. Rowe, and M. Berges, "BLUED: a fully labeled public dataset for Event-Based Non-Intrusive load monitoring research," in *Proceedings of the 2nd KDD Workshop on Data Mining Applications in Sustainability (SustKDD)*, Beijing, China, Aug. 2012.

[3] S. Pattem, "Unsupervised disaggregation for non-intrusive load monitoring," in *Machine Learning and Applications (ICMLA), 2012 11th International Conference on*, vol. 2, Dec 2012, pp. 515–520.

[4] K. Anderson, M. Berges, A. Ocneanu, D. Benitez, and J. M. F. Moura, "Event detection for non intrusive load monitoring," in *Proceedings of the 38th Annual Conference on IEEE Industrial Electronics Society (IECON)*. Montreal, Canada: IEEE, Nov. 2012.

[5] M. Berges, E. Goldman, H. S. Matthews, L. Soibelman, and K. Anderson, "User-centered non-intrusive electricity load monitoring for residential buildings," *Journal of Computing in Civil Engineering*, vol. 25, no. 6, pp. 471–480, 2011.