Automatically Disaggregating the Total Electrical Load in Residential Buildings: a Profile of the Required Solution

H S Matthews[1], L Soibelman[1], M Berges[1], E Goldman[1]

[1] Carnegie Mellon University, Pittsburgh, PA, USA
hsm@andrew.cmu.edu (Lead author email address)

Abstract. Renewed emphasis on the cost and environmental impacts of electricity are leading to a greater awareness for the need to measure and manage loads in buildings of all types. In this paper we focus on the residential buildings sector by looking at best available technologies for the commercial sector and discuss how they can be best applied to residences, by looking for effective solutions that can identify and track electricity use by the most important loads.

1 Introduction

There has been a renewed emphasis on managing energy consumption due to both economic and environmental pressures in recent years. Home owners are faced with several challenges when considering ways to reduce their energy (electricity) consumption:

- Generally only aggregate consumption (e.g., electricity use for the whole house) is known.
- Energy conservation efforts are very difficult to plan and assess without disaggregated data by appliance or appliance type.
- Available data does not provide enough information to help answer questions like: “how can I limit my electric bill to $50 per month?”, “will I get any savings if I lower my thermostat by a few degrees?”

Our explorations have surprisingly not identified a broad range of existing solutions to address these challenges. Instead, we have identified many systems reporting aggregate data (i.e., measurements taken at the main electrical feed) with low cost, and some systems able to provide consumption reports for individual electrical outlets, designed to measure each appliance separately. A few other systems collect data at the electrical circuit level, measuring the current flowing through the outgoing wire of each circuit breaker.

Information has a price in this setting. The more information the system delivers, the more costly it is and vice-versa. Solutions that are tailored for tracking individual appliances, and hence deliver detailed information, have a high installation cost as they require more sensors to be placed and the implementation of a communication link between them. On the other hand, systems which only measure the total electrical load of a building require little effort for installation at the cost of less detailed information.

The ideal setup is one where the hardware/installation costs are low, and the amount of information obtained from it is high. This is precisely what we would like to achieve. This paper describes the challenges and requirements for automatically disaggregating the total electrical load into individual load components.

1.1 Previous Work

One primary type of load monitoring is referred to as Non-Intrusive Load Monitoring (NILM), because it allows tracking appliance-specific loads without intruding (placing sensors or other devices) into the consumer's property. Many researchers have worked on this problem. The solutions, in general, consist of the following steps:

- **Feature extraction**: voltage and current waveforms are captured, and different features are extracted from them (i.e., real/reactive power, harmonics, etc.). These calculations occur at different time scales, depending on the specific solution.
- **Event detection**: changes in the features extracted above, according to dynamic or static thresholds, are detected and flagged as events for later identification.
- **Pattern recognition**: flagged events are processed using a trained algorithm that matches the sample feature-set with known and labeled events previously recorded (where previously can mean before installation of the system or during the operation process). The set of known and labeled events are often called appliance signatures, because they are a collection of features that distinguish an appliance event from others.

Our literature survey suggests that there exist feasible solutions to the problem so far. These solutions have incrementally expanded the types of features used to identify individual appliances. Initially, Hart (1992) patented a method for disaggregating individual electrical loads from a residential building using normalized real and reactive power measurements, obtained at the main feed and sampled at 1Hz (Hertz). This work was later improved by Norford and Leeb (1996) by introducing transient event detection to account for appliances with similar characteristics in the real and reactive power signature space. The sampling rate was increased for this implementation, although there is no reference to the precise value they used. Finally, Laughman et al. (2003) describe using current harmonics (3rd and 5th) as another feature to better deal with continuously variable loads, and improve the disaggregation process.

Cole and Albicki (1998) developed a modified version of Hart's algorithm which incorporates a logical analysis to separate between simultaneous on/off events. Farinaccio and Zmeureanu (1999) developed a rule-based system to disaggregate the total electricity consumption. Murata and Onoda (2002) also built on top of Hart's algorithm and incorporated radial basis function networks to estimate the power consumption of inverter type appliances. A neural network algorithm implementation is discussed by Prudenzi (2002) to identify large domestic appliances using 15 minute sampled data. Baranski and Voss (2004) present a complex algorithm combining clustering, genetic algorithms, finite-state machines and heuristics to be used with data acquired from an optical sensor installed on the utility's meter, greatly reducing the cost of implementation, assuming 1 Hz sampling rate.

2 Technological Challenges

To understand the technological challenges posed by this problem, one first needs to understand the signals being studied. Particularly, it is necessary to understand what information is encoded at different frequencies.

2.1 Understanding the Phenomenon

Electricity in the United States and in many other parts of the world is delivered as alternating current (AC) at 60Hz. This means that the sinusoidal current and voltage waveforms have a period of 1/60 of a second. Different appliances in a residence draw different amounts of current, and can do so in different ways. If the load is completely resistive, such as a toaster oven or an incandescent light bulb, then the current and voltage waveforms are in phase (e.g., their peaks coincide). If, on the other hand, there are inductive or capacitive elements on the load as it is the case with AC motors and lamp ballasts, then there will be a phase shift between the waveforms.

Power on an AC circuit is a complex quantity which can be computed from the current and voltage waveforms. When these waveforms are in phase the resulting power is all positive, and it is said to be entirely consumed by the appliance. However, when there is a phase shift between the waveforms, the resulting power can be negative in some portions of the 60 Hz cycle (where the voltage and current have opposite signs), indicating that part of the power is consumed by the appliance, while another part of it is stored and later returned to the source. The effect of this phase shift on power consumption is depicted in Figure 1.

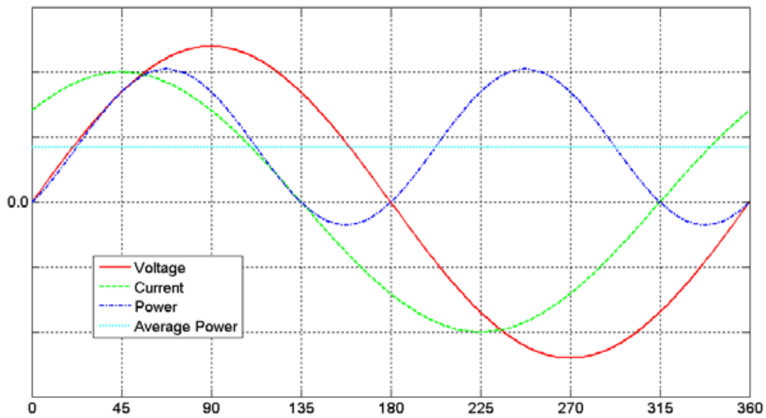

Fig. 1. Average (light blue) and instantaneous (dark blue) power computed from the voltage (red) and current (green) waveforms. It can be seen that the current waveform is out of phase with the voltage waveform, causing portions of the instantaneous power to become negative.

If the load is completely resistive such as a toaster oven or an incandescent light bulb, then the current and voltage waveforms are in phase. If, on the other hand, there are inductive or capacitive elements on the load as it is the case with AC motors and lamp ballasts, then there will be a phase shift between the waveforms.

The ability of a circuit to perform work at a particular time is known as real power. The power that is stored and returned is defined by the term reactive power. Apparent power is computed by multiplying the RMS values of the voltage and current waveforms. The ratio between real and apparent power is known as the power factor, and can be used to calculate the shift between the current and voltage waveforms.

To be able to characterize the shape of this signal in the time domain, we would need to sample it at more than five times the frequency of interest: a rule of thumb for accurately reproducing any waveform. Alternatively, if we were only interested in the frequency domain then a 120Hz sampling rate, two times the frequency of interest, would suffice according to the Nyquist-Shannon sampling theorem. This theorem states that to capture the highest frequency component of interest in a signal, one must sample at more than twice that frequency. With this last sampling rate, most of the traditional power metrics, such as root mean square (RMS) voltage, RMS current, real power, reactive power, and others can be computed.

Some appliances, such as the switching power supply in some personal computers, have a non-sinusoidal current draw and introduce new higher frequency components to the signal, known as harmonics, which are integer multiples of the fundamental frequency (60Hz). As an example, the 5^{th} harmonic would be at 300Hz.

To capture current harmonics, the sampling rate needs to be increased accordingly. Again, if only frequency domain information is needed, the Nyquist-Shannon sampling theorem tells us we would need to sample at twice the frequency of the highest harmonic we are interested in (600Hz for the above example). Since different appliances with a non-sinusoidal current draw have different harmonic characteristics, this information can turn out to be useful for our dis-aggregation purpose.

The start-up draw of different appliances is also characteristic for each appliance-type. Such transient events can be as short as a few milliseconds and as long as several minutes. To capture those with a short duration, a sampling rate in the order of kHz is needed. Norford and Leeb (1996) used 8 kHz for their application involving electric motors, for example.

There is information at an even higher level of the frequency spectrum. Patel et al. (2007) showed that the mechanical activation of a light switch generated electrical noise which could be captured and identified using a 100MHz sampling rate.

2.2 Sensors

As discussed previously, the amount of information we can obtain from the signal is directly related to the sampling rate used. Most of the commercially available power-metering technologies we have found in our investigation have an internal sampling frequency high enough to deal with short start-up transients. However, the external reporting rate is generally much slower (less or equal than 1Hz), as these solutions use custom algorithms to extract features out of the high-sampling frequency data, and produce a down-sampled summary in terms of RMS voltage, current, power, etc.

There is certainly some value in these down-sampled data sets. But whatever information is not computed from the raw values and cannot be obtained with the down-sampled data is lost. This is one of the main challenges: commercially available power monitoring solutions implement fixed algorithms that compress the raw data in a way that is appropriate for power quality or aggregate consumption monitoring (e.g., they measure at 50 kHz but only report an average at 1 Hz). These algorithms are not necessarily the appropriate ones needed for load disaggregation.

An alternative approach is to use a general purpose data acquisition board to capture the raw current and voltage waveforms at a reasonable sampling rate, and develop the algorithms for extracting the necessary information from these data. The hardware costs of this solution are much lower, but implementing the feature extraction algorithms would require extra effort.

2.3 Meter Inconsistencies

Assuming that the necessary data is available, there are still other challenges to face. Appliance signatures (collection of features associated with an appliance event) are closely related to the technology used to measure them, meaning that each different commercial power meter could yield different values for these features. To test this Figure 2 shows the results of using three different power meters (*EnerSure* by Trendpoint Systems, *WattsUp Pro* and the *Model 21-1850/CI Digital Power Meter* by Brand Electronics), and obtained average steady-state real power and power factor readings for a fluorescent lamp and a toaster, which are examples of a completely resistive load (toaster) and one with a reactive power component (fluorescent lamp with ballast).

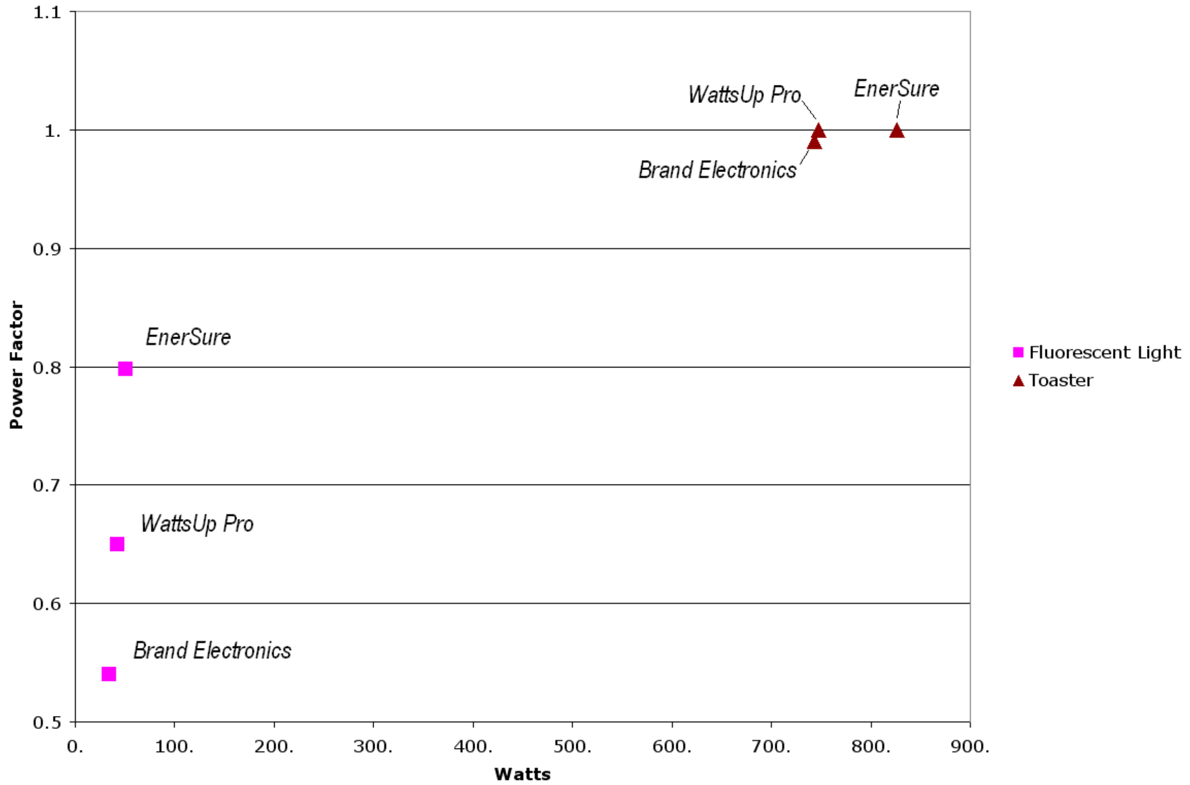

Fig. 2. Average steady-state consumption of a fluorescent lamp and a toaster measured using three different metering technologies (EnerSure by Trendpoint Systems, WattsUp Pro and 21-1850/CI Digital Power Meter by Brand Electronics). Horizontal axis shows true power, and vertical axis represents power factor.

Given that the measurements are not consistent, the generated signatures would be characteristic not only of the appliance but also of the metering technology, and so, generalizations become much more difficult.

Household appliances change their behavior with time (from malfunction, decay, etc.), and so does their signature. This can work to our advantage, if we are interested in not only disaggregating the individual loads from the main feed, but also in tracking their efficiency. Norford and Leeb (1996), for example, took advantage of this fact to implement a fault-detection algorithm for HVAC systems. However, this issue also makes the task of disaggregation more difficult, because appliance signatures may change over time but must still be recognized as belonging to a single appliance.

3 Pattern Recognition Challenges

So far we have discussed the technological challenges to extracting important features out of the aggregate electrical signal in a residential building. Just as a prism does to light, or a Fast Fourier Transform does to a time-domain signal, these extracted features provide us with a more complex characterization of the original signal that better represents some of the details that interest us.

The next step is to explain how to recognize the effects of individual appliances in this new representation of the signal. In other words, and maintaining the prism analogy, how can we infer which object is the source of the light coming into the prism?
Two other processes need to be defined before we answer this question. The first involves detecting a change in the signal, which could be related to a change in the consumption of an appliance or group of appliances at any given time. This is called event detection.

The second process, pattern recognition, involves matching the features associated with the detected change, to a set of known features in order to provide this event with a label.

3.1 Event Detection

For a human, especially for a trained one, detecting changes in extracted features is not a hard task. Visual inspection of the values is usually sufficient. However, we need to reproduce this ability algorithmically in order to automate the process.

There are many ways in which events can be defined. For example, events can be changes in the average real power exceeding a certain threshold, the appearance of a known startup transient shape, or any other suitable condition. After the selection of any of these as the preferred method, one still needs to define the parameters for it to work correctly. It might well be that the chosen event detection method is not a single one, but rather a combination of many.

Two obvious challenges arise for the event detection: defining events, and setting the appropriate threshold parameters for this definition. For example, if changes in the average real power that exceed a certain threshold are used for defining events then appliances that switch on and off very quickly will be masked by the averaging window, and will not be identified. Defining the size of this window, the threshold and other parameters is also critical and each combination of these leads has similar effects (excluding some appliances and including others).

3.2 Pattern Recognition

The objective of the pattern recognition algorithm is to assign a label to the collection features previously extracted from the current and voltage waveforms, associated with each event that is detected.

As we discussed in previous sections, many researchers have tried different implementations of this process (using neural networks, cluster matching, etc). Almost all of the examined solutions were based on off-line training of the algorithm. This training would result in what some call an “Appliance Signature Library”. Due to the particular sampling frequency, averaging techniques, etc., this signature library would be tied to the equipment used to obtain the measurements, and to some degree to the individual appliances that were used as exemplars.

Building a library that includes all possible appliances (type, make, model, year, etc.) is impractical if not impossible. It is conceivable then that the task of *identifying* new loads would be done on-line (during the operation of the system) either by applying an ‘unknown’ label to every new event detected and having then allowing the user to correct this, or by using a networked, distributed approach much as the Gracenote’s MusicID®[1] allows users to both download and upload music files metadata, which can then be associated with the unique audio “fingerprint” for each music file.

If, on the other hand, on-line *learning* is allowed, then the matching process would need to rely on historical data to function properly. Whenever a new event is detected, and no match for the features associated with it is found, the algorithm would assign a new label to the event (i.e., “Unknown001”) and store it. If this event is repeated, a match will be found and the same label would be applied to any future repetitions of it. It is reasonable to envision having the user modify these unknown labels and assign them appropriate values.

The importance of the event detection algorithm becomes clear. If, for instance, noise in the signal is detected as an event, many “new appliances” would be recognized by the pattern recognition algorithm.

4 Feedback

The ultimate goal of this event detection and classification is energy conservation through behavior modification. Therefore the information must be identified and presented in such a way that the user is both motivated and guided toward reducing and/or time-shifting their energy use. Numerous studies have shown 10-15% reductions in energy use with various kinds of feedback from household electricity metering systems (Socolow, 1978; McClelland & Cook, 1979; Hayes & Cone, 1977, 1981). However, most conservation programs either deliver specific advice to a general audience (e.g. the “energy saving tips” in utility bills) or standardized information regarding a specific home’s consumption (e.g. the instantaneous power level and daily cumulative energy usage.)

Having individual appliance energy use data will allow feedback that is both specific about the conservation actions to take and also reflects the effects of the user’s actions. Instantaneous gas-consumption gauges in cars have sometimes been offered as a model for a household energy meter, but electricity use behavior is much more complex than the acceleration of a car. Not only are there many different regulation points—dozens of appliances rather than a single pedal—but instantaneous consumption is not an adequate metric. Fluctuations due to time of day and the cycles of large loads such as the refrigerator could easily obscure the effects of behavior change. Averaging consumption over longer periods of time can smooth out some of the noise but at the expense of immediate feedback; of all the things the user did differently last week, which ones were responsible for the drop in consumption?

The risk of presenting detailed real-time feedback is of creating an interface so complicated that it serves only to intimidate the user without educating or assisting them. This suggests

[1] http://www.cddb.com/business_solutions/music_id/

that the detailed data should be filtered to provide the most salient morsels, such as a sudden change in the usage pattern of a major appliance. As another example, heuristics could be added to draw attention to low-level but high-duty-cycle appliances (such as phantom loads). However, more research is required to determine what information is most useful to present based on the disaggregated consumption data available at any given time.

5 Conclusions

We have presented the hardware and software challenges, as well as the relevance of feedback in the context of energy consumption management solutions. We argued that a system with a low installation cost, high information retrieval rate and a good user interface is the ideal solution, and presented guidelines for how to achieve such a setup. But we still have to answer the question of why these solutions are not readily available on the market already?

We believe that on the hardware side, the cost of computing power has kept these solutions from becoming widely produced and adopted. The signal processing and machine learning algorithms also need to achieve a better performance in order to make these solutions feasible. Lastly, but most importantly, there is a great need for improved user interfaces which effectively balance the tradeoffs between providing too much and too little information to the residents.

Disaggregating the individual loads from measuring the total current and voltage waveforms in a residential building has been shown to be theoretically possible, and some actual implementations of this idea show that it is indeed feasible. However, more work on the area needs to be done in order to make this type of solution a reality.

We provide broad suggestions that would help improve current solutions:

- Make use of other available sensor data (i.e., light intensity measurements, temperature, etc.), and correlate it with the electricity data in order to improve disaggregation. For example, we could envision correlating the change in light intensity caused by turning on a light bulb on a room with the change in power consumption to infer that this specific change in power was due to a light bulb.
- Take advantage of existing computing power in modern residential buildings to lower the cost of the hardware required to process the signals and run the pattern recognition algorithms.
- Improve the learning algorithms by utilizing user-feedback through an interface. This feedback mechanism should be non-intrusive and must be designed to easily blend with the resident’s normal daily activities.

6 Acknowledgements

The authors would like to gratefully acknowledge the support of the Robert Bosch LLC Research and Technology Center. Additional support was provided by the Pennsylvania Infrastructure Technology Alliance (PITA).

7 References

Baranski, M. & Voss, J., 2004. Genetic algorithm for pattern detection in NIALM systems. In *Systems, Man and Cybernetics, 2004 IEEE International Conference on.*, p. 3462-3468 vol.4. Available at: http://ieeexplore.ieee.org/iel5/9622/30424/01400878.pdf [Accessed October 1, 2007].

Cole, A. & Albicki, A., 1998. Data extraction for effective non-intrusive identification of residential power loads. In *Instrumentation and Measurement Technology Conference, 1998. IMTC/98. Conference Proceedings. IEEE.*, p. 812-815 vol.2.

Farinaccio, L. & Zmeureanu, R., 1999. Using a pattern recognition approach to disaggregate the total electricity consumption in a house into the major end-uses. *Energy and Buildings*, 30(3), p.245-259.

Hart, G., 1992. Nonintrusive appliance load monitoring. *Proceedings of the IEEE*, 80(12), p.1870-1891.

Hayes, S. C. & Cone, J. D., 1977. Reducing residential energy use: Payments, information, and feedback. *Journal of Applied Behavior Analysis*, 10, 425-435.

Hayes, S. C. & Cone, J. D., 1981. Reduction of residential consumption through simple monthly feedback. *Journal of Applied Behavior Analysis*, 14, 81-88.

Laughman, C. et al., 2003. Power signature analysis. *Power and Energy Magazine, IEEE*, 1(2), p.56-63.

McClelland, L & Cook, S., 1979. Energy conservation effects of continuous in-home feedback. *Journal of Environmental Systems*, 9, 169-173

Murata, H. & Onoda, T., 2002. Estimation of power consumption for household electric appliances. In *Neural Information Processing, 2002. ICONIP '02. Proceedings of the 9th International Conference on.*, p. 2299-2303 vol.5.

Norford, L.K. & Leeb, S.B., 1996. Non-intrusive electrical load monitoring in commercial buildings based on steady-state and transient load-detection algorithms. *Energy and Buildings*, 24(1), p.51-64.

Patel et al., 2007. At the Flick of a Switch: Detecting and Classifying Unique Electrical Events on the Residential Power Line (Nominated for the Best Paper Award). UbiComp 2007: Ubiquitous Computing. Available at: http://dx.doi.org/10.1007/978-3-540-74853-3_16 [Accessed October 4, 2007].

Prudenzi, A., 2002. A neuron nets based procedure for identifying domestic appliances pattern-of-use from energy recordings at meter panel. In *Power Engineering Society Winter Meeting, 2002. IEEE.*, p. 941-946 vol.2.

Socolow, R. H., 1978. The Twin Rivers program on energy conservation in housing: Highlights and conclusions. In Socolow, R. H. (Ed.). *Saving Energy in the Home: Princeton's Experiments at Twin Rivers* (pp. 2-62). Cambridge, MA: Ballinger Publishing Company.