Nonintrusive Disaggregation of Residential Air-Conditioning Loads from Sub-hourly Smart Meter Data

Krystian X. Perez[1], Wesley J. Cole[1], Michael Baldea[1,2], and Thomas F. Edgar[1,3*]

[1]McKetta Department of Chemical Engineering, The University of Texas at Austin, Austin, TX 78712
[2]Institute for Computational Engineering and Sciences, The University of Texas, Austin, TX 78712
[3]Energy Institute, The University of Texas, Austin, TX 78712
Email: krystian.perez@utexas.edu

Abstract — **The installation of smart meters has provided an opportunity to better analyze residential energy consumption and energy-related behaviors. Air-conditioning (A/C) use can be determined through non-intrusive load monitoring, which separates A/C cooling energy consumption from whole-house energy data. In this paper, a disaggregation technique is described and executed on one-minute smart meter data from 19 houses in Austin, TX, USA, from July 2012 through June 2013. 19 houses were sub-metered to validate the accuracy of the disaggregation technique. The R^2 value between the predicted and actual A/C energy use for the 19 houses was 0.90. The accuracy of A/C disaggregation on five-minute data was found to be comparable to one-minute data. However, fifteen-minute data did not yield accurate results due to insufficient granularity.**

Index Terms *Nonintrusive Load Monitoring, Disaggregation, Residential Energy, Air Conditioning, Smart Meter*

I. Introduction

Approximately 15% of primary energy use within the United States is due to heating, ventilation and air conditioning (HVAC) [1]. The HVAC load placed on the grid by residential consumers is highly variable and strongly influenced by weather and human activity patterns. Meeting such fluctuations in demand is challenging for grid operators and requires excess capacity from generation facilities to be available. Understanding how residential air-conditioning (A/C) usage for individual houses and entire neighborhoods are affected by external factors can lead to increased uniformity of energy demand on the grid.

Smart meters allow the derivation of valuable information about residential A/C energy use through non-intrusive load monitoring (NILM). Sub-metering specific circuits such as the A/C is expensive. NILM mitigates this problem by employing algorithms to extract the A/C usage from the whole-house measurements provided by the smart meter data.

The NILM method was first described by Hart in 1970 [2] and many advances have been made since. The technique used in this paper is different from previous methods in that it is only concerned with the A/C use. One of the challenges in Hart's method is that it used clustering to match changes in power of on and off events, which required previous knowledge of devices to accurately pair behaviors to devices. By focusing exclusively on A/C energy, this new technique can be accurately tuned for any house to separate just A/C energy use. Rather than pair on and off power events, this technique uses edge detection and k-means clustering to find key parameters on A/C behavior. The parameters are then used to identify A/C on and off events. Since A/C is a dominant feature in the energy profile, the on/off events of an A/C unit can be detected. A statistical analysis evaluates the accuracy of utilizing a general disaggregation algorithm by comparing actual vs. estimated A/C use. Additionally, the accuracy of disaggregating 1-,5- and 15-minute data are compared using the NILM technique described in the paper.

II. Method

A. Data

Whole house energy usage was taken from 19 single-family houses in the Mueller neighborhood in Austin, TX from July, 2012 to June, 2013. Each house had been metered with an eGauge power monitor that reported whole-house power consumption in watts on one-minute time intervals. The 19 houses had additionally been sub-metered with an additional meter to directly measure power used for the A/C system during the full year. The Mueller neighborhood consists mostly of newer (since 2007), green-built houses [3]. The houses were equipped with electric A/C cooling units and natural gas heating systems.

B. Disaggregation Outline

This paper uses a NILM to disaggregate the A/C cooling energy consumption from the 1-minute whole-house energy consumption data. In this technique, the magnitude of change in load that signals the A/C turning on or off is found, which is then used to identify on and off events of A/C use. A decision flowchart for the disaggregation process is seen in Figure 1. The algorithm functions as an edge-detection algorithm. Similar techniques which has been used in other disaggregation applications [2].

First, energy use data for each house was separated by day from midnight to midnight. The difference in energy between each time step (ΔE_i) was calculated and stored as seen in equation 1 where ΔE_i is the change of energy between time *i* (the current time step) and *i + 1*.

$$\Delta E_i = E_{i+1} - E_i \tag{1}$$

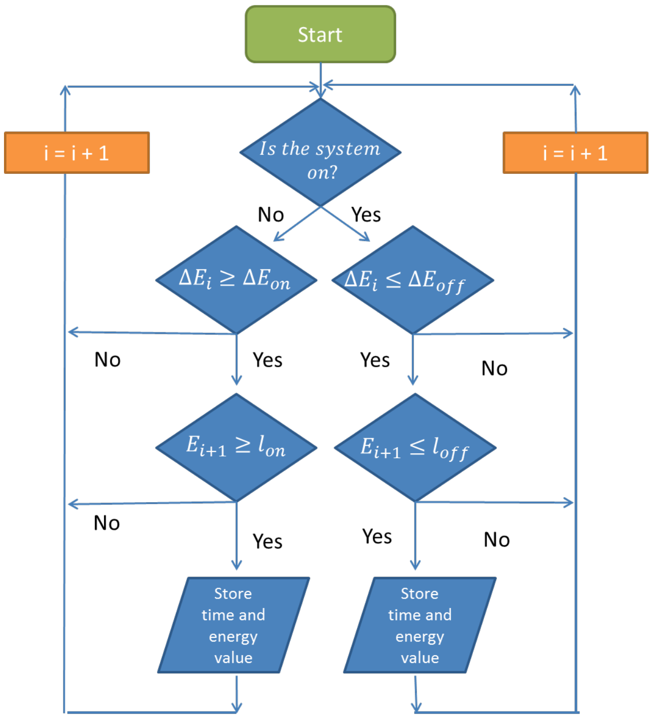

Figure 1. Decision flowchart to demonstrate how A/C cycles are determined and stored. ΔE_i is the change in energy at time step *i*. ΔE_{on} and ΔE_{off} are the changes in magnitude large enough to indicate the system has turned on or off. l_{on} and l_{off} are limits on how much power is used by the A/C system.

The value of *i* in the study was 1 minute unless otherwise stated. If the HVAC system begins in the "off" state, ΔE_i is compared to ΔE_{on}, which is the minimum energy change to signal an "on" event.

If the difference in energy was large enough to signal that the A/C unit turned on ($\Delta E_i \geq \Delta E_{on}$) and the energy at time *i + 1* was greater than a lower limit ($E_{i+1} \geq l_{on}$), then the time stamp was stored as the system turning on. The sum of the change in energy value at time step *i* with the average of the ΔE values before and after the signal was also stored ($\Delta E_i + \frac{\Delta E_{i-1}+\Delta E_{i+1}}{2}$). The ΔE_i value represents the amount of energy increased due to the A/C system turning on. The average of the ΔE value before and after the time step was added for events where the system ramped on between time steps. Otherwise the algorithm would underestimate energy because of the offset.

The stored value incorporated power consumed by the air-handling unit (the blower fan and controls). The algorithm stepped forward in time until ΔE_i was less than the condition for the "off" signal ($\Delta E_i < \Delta E_{off}$) and E_{i+1} was smaller than the upper limit ($E_{i+1} \leq l_{off}$), meaning the decrease in the energy consumption was strong enough to indicate that the system turned off and the total energy was below the lower limit threshold. Then the algorithm would step forward in time until another on signal was found. The decision process was repeated until the end of each day. The total estimated A/C energy consumption was calculated by multiplying the total duration of time during which the A/C unit was on with the average of the A/C energy values over a day. Therefore the power required by the A/C was assumed to be constant whenever the A/C was on. In reality, the amount of energy consumed by the A/C slightly fluctuates throughout the day as a result of inefficiencies due to changes in outside temperature [4]. Power draw will increase at higher temperatures, such as the afternoon, when the COP decreases. The averaged value was used because it is resistant to noise that would otherwise decrease the estimated value. However, as a result the resulting estimated energy value may be lower than if there was a correction factor dependent on outdoor temperature that adjusted for a changing COP.

In Austin, during early morning hours and warm weather the dominant feature in the energy profile is the sporadic cycling of the A/C unit. The premise of the algorithm is that during months where A/C is used (and most likely the dominant consumer), there is a time during the day when the A/C load is a dominant feature in the electricity load. That time period can be used in the algorithm as a training period to determine key parameters, such as the magnitude of A/C spikes that indicate that the A/C unit has turned on (ΔE_{on}). The parameters derived from this training period are then applied to the rest of the data in order to identify and extract information about the A/C energy usage.

III. Results and Discussion

A. Validation from Sub-metered houses

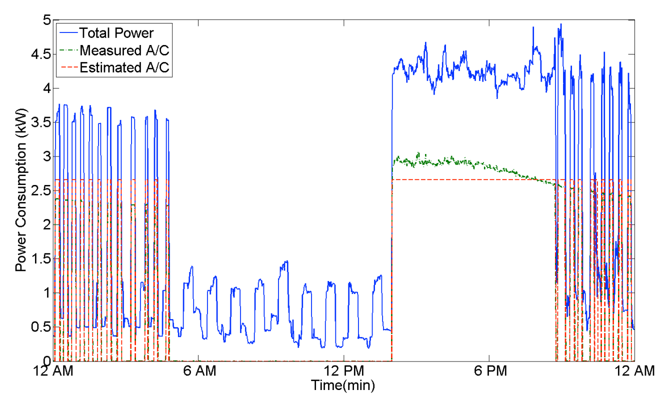

Figure 2. Sample day of disaggregated energy for a house. Measured A/C energy denotes the values given from the sub-circuited eGauge on the A/C system. Estimated A/C is the disaggregated energy.

The 19 houses that had the A/C sub-metered were used as a benchmark to evaluate the accuracy of the algorithm. The estimated daily A/C energy consumption was compared with the measured daily energy consumption. Figure 2 shows the estimated values of the disaggregation technique compared with the measured values reported by the sub-meter for one house for one day. The estimated energy aligns very closely with the measured A/C energy for most days. It was found that throughout the time period of May through October, the estimated values are very close to measured values. During the rest of the year, the algorithm was less accurate because other loads in the profile appear similar to the A/C. Thus the algorithm detected false events and would overestimate A/C energy.

The algorithm does not differentiate changes in energy due to the A/C or other appliances if they are the same magnitude. Therefore, when A/C is not a dominant load the load is still perceived as A/C. The degree to which the estimated A/C energy is different from the measured energy depends on how many loads mimic the A/C. Generally speaking, most A/C units in this study consumed about 2.5 kW of power when on. It appears that only high loads such as an electric dryer are of the same size. In the case of a highly efficient unit that drew substantially less power, such as 1 kW, the algorithm likely would not successfully discriminate between that and other similarly sized loads. For some houses, accuracy in estimated A/C use during winter months was the same as summer and for others the difference was substantial. Therefore only cooling months were taken into consideration. The limitation to this algorithm is that it is specific to A/C loads during cooling months.

Table 1. House ID for sub-metered houses. R-squared value is that given from the fit of the data to the parity plot. The CV(RMSE) values (in %) are also given.

House ID	R^2	CV(RMSE)	House ID	R^2	CV(RMSE)
A	0.92	22.4%	K	0.74	64.6%
B	0.99	7.2%	L	0.96	22.9%
C	0.95	16.3%	M	0.92	30.6%
D	0.95	13.1%	N	0.82	30.3%
E	0.97	10.5%	O	0.98	12.4%
F	0.95	17.4%	P	0.69	44.1%
G	0.99	7.5%	Q	0.97	14.2%
H	0.93	25.3%	R	0.95	15.2%
I	0.87	25.8%	S	0.97	8.9%
J	0.99	9.5%			

Observing all the houses that had the A/C sub-metered, a parity plot of May through October of estimated A/C energy consumption vs. measured A/C energy consumption exhibited high R^2 (coefficient of determination) values and relatively low CV(RMSE) (coefficient of variation of the root mean square error) values as shown in Table 2. R^2 is used to describe the variation in the linear relationship between measured and modeled A/C use while the CV-RMSE is used to quantify variation of residuals (the differences between predicted values and observed values.) ASHRAE Guideline 14 adopted CV-RMSE to evaluate prediction uncertainty of energy inverse models, which would also include the energy disaggregation model presented in this work [5]. According to the standard, “typically models are declared to be calibrated if they produce …CV(RMSE)s within ±30% when using hourly data.” House N, while having a high R^2 also had a regular electric vehicle charging load that resembled an A/C unit turning on for a sustained time. Thus the algorithm effectively tracked the variation in A/C consumption, but consistently over-estimated the amount of A/C energy. We note that house P also had lower values of these goodness of fit metrix. A closer inspection of the data showed low A/C use during morning hours and thus a poor training period to identify A/C loads. However, houses with low A/C use in the morning could be screened prior to disaggregation. Overall the results on the sub-metered houses indicate that the disaggregation technique accurately differentiates the A/C energy from the overall energy profile. On average for the 19 houses, the total estimated A/C energy value differed from the measured value by less than 14% (2.1 kWh) of the measured energy consumption.

B. *Various Sampling Times*

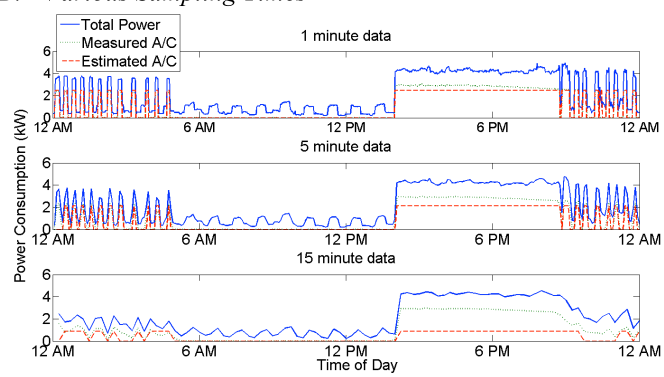

Figure 3. Comparison of sample times on the ability of the algorithm to estimate A/C power consumption for a single day. The top plot uses a one minute sample time, the middle plot uses five minute sample time, and the bottom plot uses fifteen minute sample time. As sample time increases, it becomes more difficult to disaggregate the A/C load.

In this analysis, the data set considered energy use measured in 1-minute time intervals. However, most utilities have smart meters that report energy usage in 15-minute time intervals. As energy use data becomes coarser, it becomes more difficult to identify and separate individual loads (see Figure 3). A brief analysis evaluated the sampling time to which A/C loads can still be reliably disaggregated. Figure 3 is a comparison of the energy profile for 1, 5 and 15-minute data. 5-minute data is visibly coarser than the 1-minute data, but still retains a degree of granularity that makes it possible to recognize which peaks are the A/C. The 15-minute data is even coarser, making it difficult to even visibly distinguish A/C peaks from the base load. For 5-minute data, it may still be possible to count the number of cycles if the cycle times are shorter than 5 minutes. In the 15-minute profile it was extremely difficult to accurately measure cycle numbers. The 5-minute data does have advantages. For example, taking the average across 5 minutes, anomalies, such as random spikes, would not be seen in the profile. Lower frequency data have lower hardware storage requirements but this savings comes at the cost of granularity as highlighted below.

Figure 4 describes the error differences between 1, 5 and 15-minute data. The disaggregation technique was scaled to match 5 and 15-minute data in terms of the training period. The techniques then followed the same algorithm to cluster changes in energy to define the parameters that indicate on and off events. They then stepped through the algorithm just like in the 1-minute technique. The resulting absolute error values were found and then compared with the 1-minute technique as seen in Figure 4. Overall, both 1 and 5-minute data had lower absolute error values while the 15-minute data had significantly higher values. Although for some houses, the 5-minute data had a smaller absolute error. In the 1-minute data scheme, sometimes a very busy energy profile made it hard to distinguish A/C events from other appliances, especially if a system ramped on between sampling times. The 5-minute profile smoothed out the profile to remove most hard to distinguish events. On days with low A/C usage, the 5-minute profile is especially accurate because of this smoothing effect. Depending on the required amount of detail, 5-minute time intervals can be accurate enough to make control decisions while reducing the amount of data stored. It also gives a fairly accurate overall picture of energy use. 15-minute data does not have high enough resolution to discriminate between A/C and non-A/C events. When considering the benefit vs. cost of storing smart grid data, utilities should consider flexibility in new installations so that the benefits of more discrete data can be fully explored.

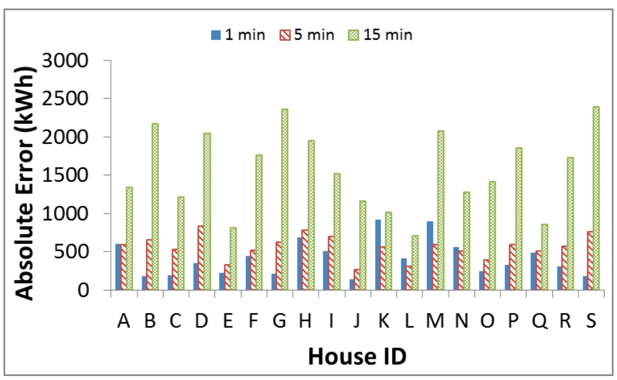

Figure 4. Comparison of R^2 and absolute error values for 1, 5 and 15-minute data.

IV. CONCLUSIONS

This paper has shown that a generic NILM algorithm was able to accurately disaggregate A/C usage from total energy profiles. The average R^2 value of measured vs. estimated A/C energy values for 19 sub-metered houses was 0.90 with an average CV(RMSE) of 25.4. This technique requires a training period in which the A/C is the dominant load, but once the parameters have been estimated, they will likely remain valid until the A/C unit undergoes significant change.

A comparison between 1, 5 and 15-minute sampling times revealed the accuracy of 5-minute data in disaggregating A/C loads to be comparable to 1-minute data. However, 15-minute data did not yield accurate results due to insufficient granularity of data.

ACKNOWLEDGMENT

This work was supported by the Pecan Street Research Institute (a 501(c)3 nonprofit public-private partnership in Austin, Texas), the United States Department of Energy, and the University of Texas at Austin. This material is also based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE-1110007. Any opinion, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

[1] U.S. Energy Information Administration, “Annual Energy Review 2011.” 27-Sep-2012.

[2] G. W. Hart, “Nonintrusive appliance load monitoring,” *Proc. IEEE*, vol. 80, no. 12, pp. 1870–1891, 1992.

[3] J. D. Rhodes, C. R. Upshaw, C. B. Harris, C. M. Meehan, D. A. Walling, P. A. Navrátil, A. L. Beck, K. Nagasawa, R. L. Fares, W. J. Cole, H. Kumar, R. D. Duncan, C. L. Holcomb, T. F. Edgar, A. Kwasinski, and M. E. Webber, “Experimental and data collection methods for a large-scale smart grid deployment: Methods and first results,” *Energy*, vol. 65, pp. 462–471, Feb. 2014.

[4] M. Kim, W. V. Payne, P. A. Domanski, S. H. Yoon, and C. J. L. Hermes, “Performance of a residential heat pump operating in the cooling mode with single faults imposed,” *Appl. Therm. Eng.*, vol. 29, no. 4, pp. 770–778, Mar. 2009.

[5] ASHRAE, “ASHRAE Guideline 14-2002 Measurement of Energy and Demand Savings.” ASHRAE, 27-Jun-2002.