New Approach to Analyzing Hourly Energy Usage Data to Obtain Fast, Accurate Savings Estimates

Robert Van Buskirk, Ph.D. and William C. Miller, Ph.D.[1]
Lawrence Berkeley National Laboratory, Berkeley CA

ABSTRACT

The paper demonstrates the use of hourly energy usage data to determine energy use changes in residential households in weeks rather than months or years. This can be applied to the determination of gross energy efficiency program savings with reduced cost. Particular applications could be in the area of program diagnostics ("Are savings occurring?"), program finance ("Are savings sufficient to achieve targeted energy bill reductions?"), or even overall program gross savings. This work should stimulate further innovation in the use of hourly or higher resolution usage data.

As automatic meter infrastructure (AMI) spreads, increasing the amounts of hourly usage data (AMI data) available. The variations in usage in hourly data embody the same factors that separate metering or data acquisition (e.g. surveys) illuminated when analysis was conducted with only monthly data. AMI data may embody this information less perfectly than direct data acquisition, but the information it does contain comes is almost immediately available at no additional cost.

This paper considers AMI data as basic time series and separates the noise in that series from evidence in a change in the pattern of usage using standard statistical techniques applied in a novel way. By applying a moving time window, step-wise changes in energy use patterns can be sensitively detected in very noisy data. This opens the possibility of being able to determine in a short time frame whether an energy efficiency action produced expected results rather than the one to two years now required.

Once a change has been identified, straightforward ranking techniques may be able quantify the change to a high degree of accuracy. In the case of non-weather dependent uses, this can be done with several weeks of data pre-change and post change. More complex techniques can account for the impact of temperature changes.

Introduction

The increasing availability of facility-level energy usage data at hourly or greater frequency (high-resolution data) opens new possibilities for analyzing energy usage and energy efficiency savings. Among the possibilities are developing new analytic approaches to address traditional questions, and developing new tools to analyze traditional or new issues. The work in this paper adds to these to these two areas and illustrates the opportunities the availability of high-resolution usage data provides.

---

[1] The authors may be reached at *RDVanBuskirk@lbl.gov* or *WCMiller@lbl.gov* respectively.

The opportunity posed by high-resolution data is timely as there is increasing interest in approaches to assessing energy efficiency savings that provide quicker, less costly but still usefully accurate answers. This interest arises not only among the historic users of energy efficiency savings calculations (regulators, EE program providers and energy service industry participants), but others such as those pursuing energy management or upgrades on their own.

The energy efficiency program industry can use faster (weeks rather than months or years) and accurate determination of energy use changes in residential households. A key potential application for using available hourly data is the determination of gross energy efficiency program savings with reduced cost. Particular applications could be in the area of program diagnostics (are savings occurring), program finance (are savings sufficient to achieve targeted energy bill reductions), what events are confounding clear determination of savings or even overall program gross savings.

Historically determining "gross" energy efficiency savings used the monthly energy usage available. Invariably, additional information was required because the information content of monthly energy usage is extremely aggregated. For example, in a typical whole-house regression analysis, monthly energy usage, weather, appliance holdings, number of occupants and their characteristics, and behavioral patterns were used. In some cases, additional information was obtained by separate metering activities. For example, lighting "loggers" have been used to determine actual lighting schedules in commercial lighting retrofits. Similarly, meters on air conditioning units to determine actual run times and KW used for commercial HVAC upgrades.

Today increasing amounts of hourly usage data including from advanced meter infrastructure (AMI) are becoming available. The variations in usage observed in hourly usage data are the result of the same factors that separate metering or data acquisition was designed to determine. While the AMI data embodies this information less perfectly than direct metering, what it does contain comes at no additional cost above obtaining the AMI data, and is available as rapidly as the AMI data is available.

Characterizing the Energy Use Signal

This paper explores the analysis of residential AMI data from a new analytical perspective in an attempt to increase its utility in estimating the impacts of efficiency measures or energy efficiency actions. In this new perspective, household electricity use is treated as comprising of two relatively distinct components that are characterized by differences in time scale. We define the short time scale signal as those energy use variations that consist of hourly, daily, and day of week fluctuations that respond to short-term behavior patterns of the household occupants. In contrast to these short-time variations, the long-term signal is defined by a series of step-wise changes in occupancy, behavior or equipment. We wish to measure empirically the longer term, step-wise changes in energy use.

A key difficulty of measuring the impact of any energy efficiency intervention comes from the challenge of measuring the impact of an efficiency measure amidst a background variation of idiosyncratic behavior. Behavior and occupancy have very complex patterns, and create relatively large, short-time-scale energy use changes that create a large natural variability in household energy use.

In our new perspective, most or all of the long term energy use changes in a household are occurring as step-wise changes: both the energy efficiency intervention and the long term component of the background variation are assumed to go through a distinct change at a particular point in time.[2] These step-wise changes may occur due to either equipment, or changes in occupancy or behavior (e.g. someone changes job location and commuting schedule, or adds a morning run to their daily routine). The approach is to measure the energy use changes for ALL major step-wise changes in energy use with good accuracy, and then focus on the efficiency impact measurement problem of interest. This is because if we have accurate measurement of all of the major energy use changes in a household, then the problem of identifying the energy efficiency measure impact reduces to a problem of identifying which of the several observed step-wise changes in energy use corresponds to the intervention of interest.

In this paper we demonstrate this measurement technique for example data. We begin the demonstration by initially describing some illustrative examples of how with a relatively short time series of hourly data it is possible to accurately measure step-wise changes in energy use. We follow the illustrative examples with a demonstration of how moving window averages can be used to detect step-wise energy use changes in noisy data. In the first part of the demonstration we show the application technique with a simple data model, and then we apply the technique to a more realistic, longer-term data series. We then illustrate the calculation of on-peak vs. off-peak energy use changes by comparing pre-change and post-change load frequency distributions. We conclude with a summary and discussion of next steps.

An Illustrative Example

In this section we illustrate how step-wise changes in energy use can be identified in hourly energy use data.

Figure 1 shows an example hourly residential electricity use data taken from the publicly available data from the End-use Load and Consumer Assessment Program (ELCAP). [1] We initially illustrate AMI data analysis methods with historical ELCAP data because it is a standard, publicly available data set of hourly data raising minimal issues with respect to privacy and confidentiality. ELCAP data has some differences with more current smart meter data in at least two aspects: First, it represents an area of the country where customers have relatively large electrical loads that usually includes the use electric resistance heating. Secondly, the measurement resolution of the load data is an order of magnitude lower that typical modern AMI data (~100 watts compared to 10 watts).

We can see if figure 1 a clear example of a step-wise change in energy use behavior in the context of very noisy hourly variations.

[2] Note that a continuous trend in energy use can often be approximated as a series of step-wise changes, so that the assumption of step-wise energy use changes at a distinct point in time may not be much of a restriction.

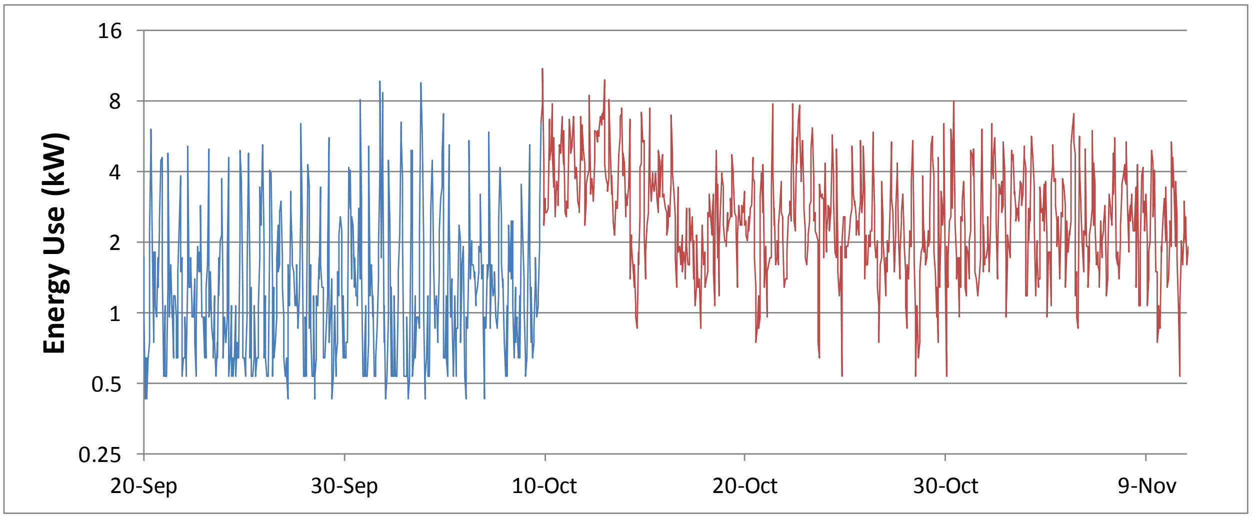

Figure 1: Example hourly data illustrating a pre- vs. post- changes in electricity use.

How precisely can load changes be measured?

We now discuss what happens when a change occurs that is substantially smaller than the change illustrated in figure 1. In the figure 1 example, previous to the transition the standard deviation of the natural log of the energy use is 0.69, with 473 data points indicating a standard error of measurement of the mean natural log of energy use is $0.69/(473)^{1/2} = 3\%$. This represents a pre-transition measurement error of 36 watts out of a median load of approximately 1200 watts. After the transition, the 800 data points have a standard deviation of the natural log of 0.51, indicating a standard error of $0.51/(800)^{1/2} = 2\%$ of 2500 watts or an error of about 50 watts. We can safely say that the event that occurred on the morning of October 8 resulted in an increase of electricity use of approximately 1300 watts.

In a state like California, average electricity use rates are approximately 15 kWh/day, substantially less than those in the households in the ELCAP study. [2] Assuming that the variance of the log of electricity use scale proportional to the use rate and that the averaging period is three weeks or more, it should be possible to measure step-wise transitions in energy use in Californian households to a standard error of less than 20 watts.

Example: A Step-wise change in load

Next, we illustrate how relatively small step-wise changes in energy use--while perhaps not evident by looking directly at time series data--can be very clearly seen by plotting load duration curves for pre- and post-change hourly data. To construct the illustration we simulate[3] three-week pre- and post-change load comparisons. Specifically we take a log normally distributed baseline load of mean 15 kWh/(24 hours) and a standard deviation of the log of 0.3

[3] We illustrate the concepts with simulated data to avoid privacy confidentiality costs/complications inherent in using AMI data from actual customers.

and add an incremental load of mean 90 watts and standard deviation 20 watts. Figure 2 illustrates the three-week load duration curves for two simulated loads with and without the incremental load.

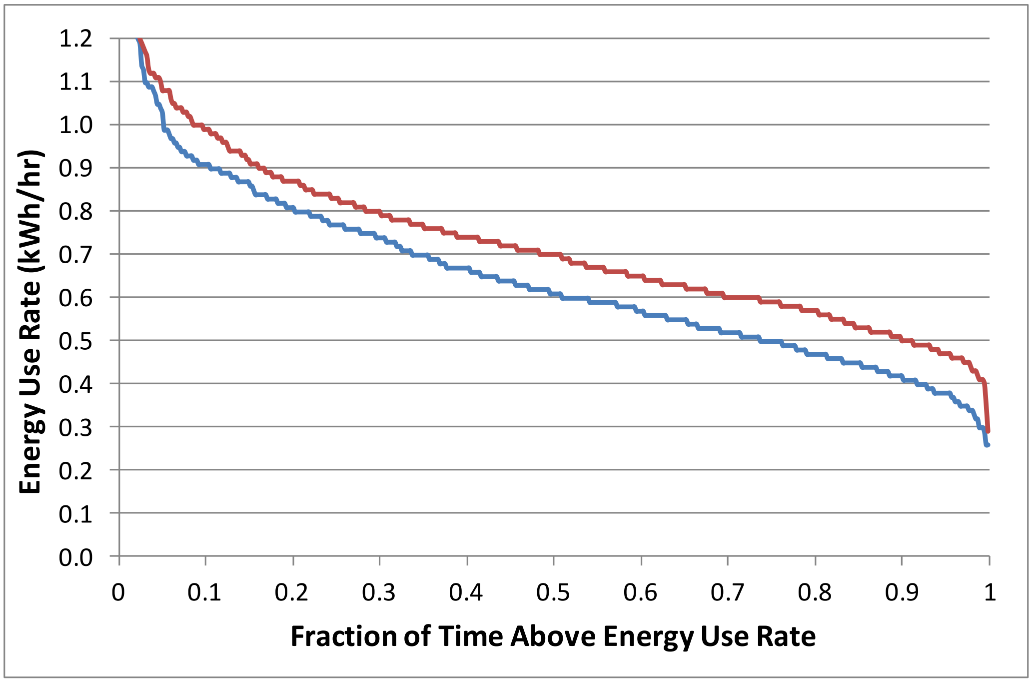

Figure 2: Pre-base-load-removal (red) and post-base-load-removal (blue) load duration curves.

What this particular example illustrates is that if one has pre-transition and post-transition periods clearly defined with fairly consistent average behavior in each of the pre-transition and post-transition periods, then it may be possible to create a clear measurement of the energy use differences even when the incremental load is relatively small. This example also illustrates that comparison of pre-transition and post-transition load duration curves may be an effective technique for measuring the energy use change at the transition.

Systematic Detection of Step-wise Changes in AMI Data

In this section, we illustrate a method that can be used for the systematic detection of step-wise changes in AMI data through the use of running-window calculations of average energy use differences.

Simplified Statistical Model of Data

We start our demonstration of the change detection method with a simplified statistical model of the AMI data. In this simplified model we assume that the AMI data consists of serially correlated, log-normally distributed random noise and generate synthetic data using a random number generator.

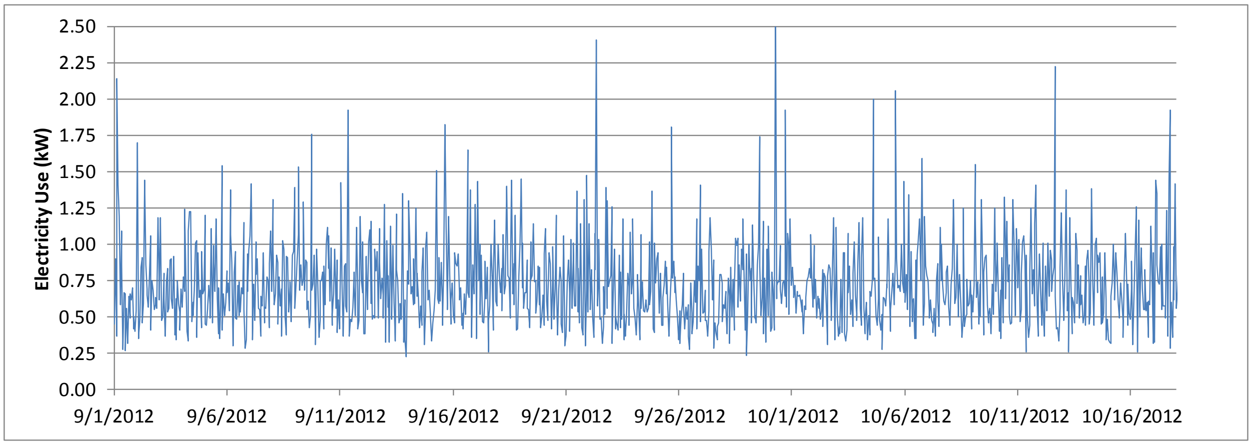

Figure 3: Synthetic energy use data generated from a statistical simulation model.

Figure 3 illustrates the synthetic data generated by the statistical model described above, where the statistical properties of the data correspond to the annual average statistical properties of AMI data from a residential household.

Calculation of Step-wise Change in Energy Use

We now illustrate with the synthetic data how to perform a detection calculation for a step-wise change in energy use. To synthesize a step-wise change in energy use, we add to our synthetic data after a particular date (specifically 6/30/12) a constant energy use change, and try to detect the change in the synthetic data as illustrated in Figure 4.

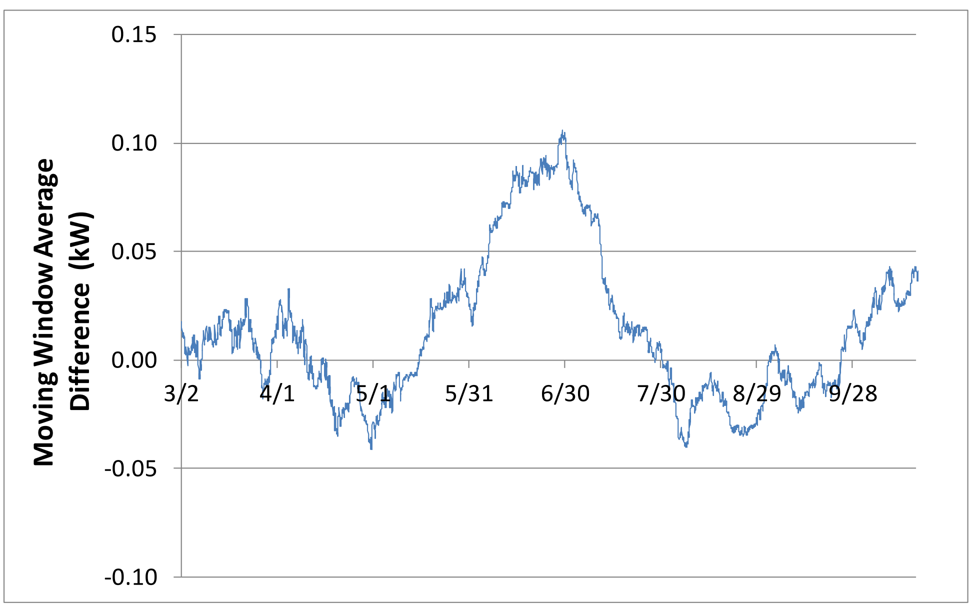

Figure 4: Moving window average detection of 0.10 kW change in energy use at 6/30 in synthetic data for a window size of 4 weeks.

The calculation that we use to measure energy use change in the synthetic data is a comparison between a forward and backward moving window average. To make this comparison, we calculate for each date, the average energy use 4 weeks ahead of the date, and the average energy use 4 weeks behind the date. When we reach a step-wise change in energy use, there should be a peak in this moving window difference of forward and backward averages that equals the size of the energy use step. This is exactly what we see in the calculation on synthetic data shown in Figure 4.

Figure 5 shows a recalculation of figure 4 with a range of window sizes. Note that the random numbers used in figure 5 are different than those in figure 4, so that pattern of noise is different. What we see in figure 5 is that as we change the size of the computational window, that the peak in the calculation of the moving window average difference is very stable at the transition date 6/30. Also while there is some error in the estimate of the energy use difference for small window sizes, as the window size gets larger, the error in the estimate of the energy use difference (i.e. the deviation of value at 6/30 from the actual step size of 0.1 kW) gets very small with increasing window size (and appears to be just a few watts).

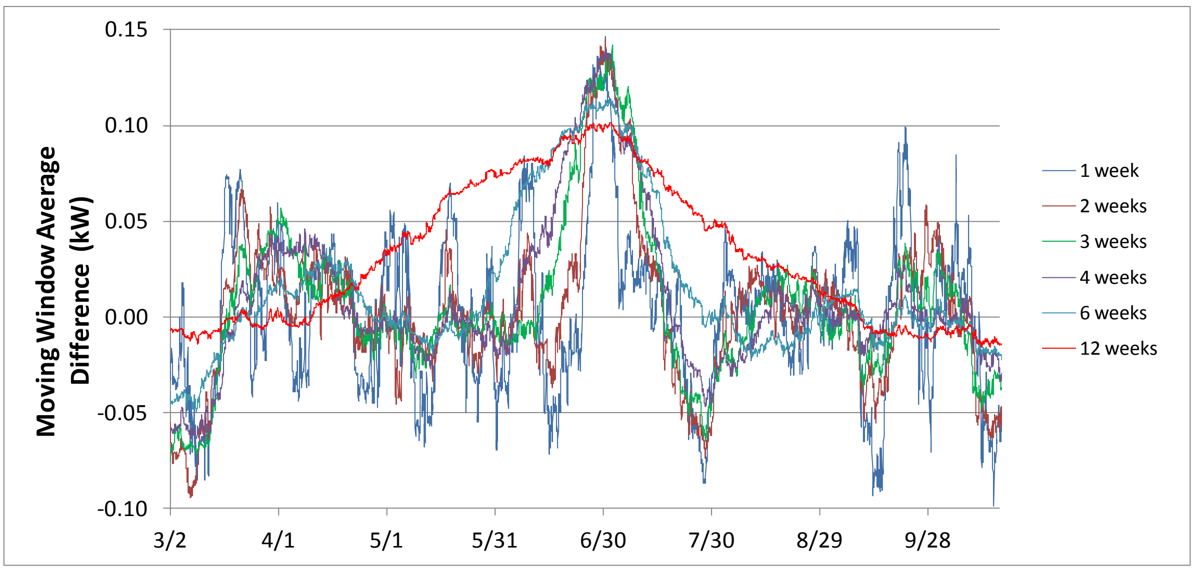

Figure 5: Moving window average detection of 0.10 kW change in energy use at 6/30 in synthetic data for a window sizes ranging from 1 week to 12 weeks.

Application to Real-world Data

In this section we illustrate the application of the change-detection calculation from the previous section to an example of real world data. Figure 6 shows the average difference between a forward window average and a backward window average for different moving window sizes for approximately 100 weeks of hourly residential load data.

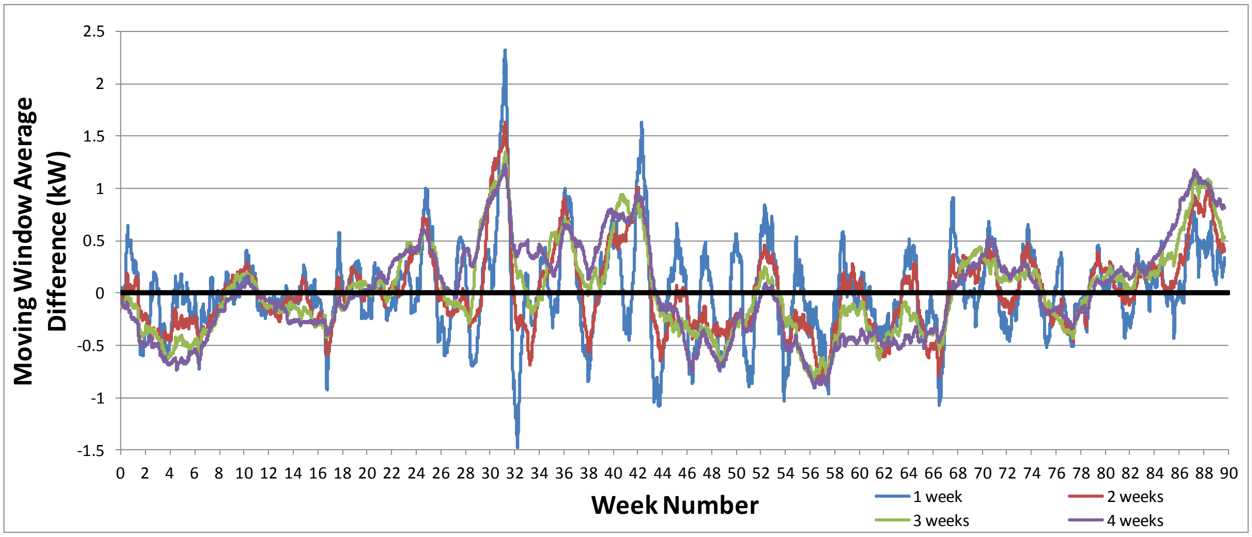

Figure 6: Moving window average energy use change for an example household from the ELCAP dataset. Window size for averaging varies from one week (blue curve) to four weeks (purple curve).

As we can see from figure 6, we can detect several different energy use changes occurring on different time scales. We can see very large changes in energy use week to week that probably correspond to the use of electric resistance heating more or less intensely during the winter. During the interim spring and summer months, there are more infrequent and usually smaller changes in energy use that probably correspond to changes in occupancy and changing various miscellaneous pieces of equipment that are in use.

We use figure 6 to develop a simple step-wise energy use change model for the AMI data. Specifically we examine each peak in the plot and determine the time and magnitude of 33 distinct energy use changes that range in magnitude from 0.3 to 2.3 kW. We illustrate the results of this model determination in figure 7.

The upper plot in figure 7 illustrates the cumulative energy use change represented by the 23 distinct changes detected in figure 6. Increases are likely caused by the moving in of more energy-intensive renters, the adoption of electric resistance heaters during the winter, or temporary changes in occupancy. Energy use reductions are caused by the removal of electric heaters, renters moving out or temporary changed in occupancy. The lower plot in figure 7 replicates the calculation in figure 6 for the noiseless step-wise energy use model shown in the upper portion of figure 7. The general pattern of figure 6 is replicated though the qualitative nature of the plot including the small-scale statistical fluctuations and smoothing of the energy change peaks is not seen.

But even though the simple step-wide energy use change model provides some general similarity to the empirical curves shown in figure 6, there are some distinct qualitative differences in the shapes of the different energy use change peaks that we would like to explore further.

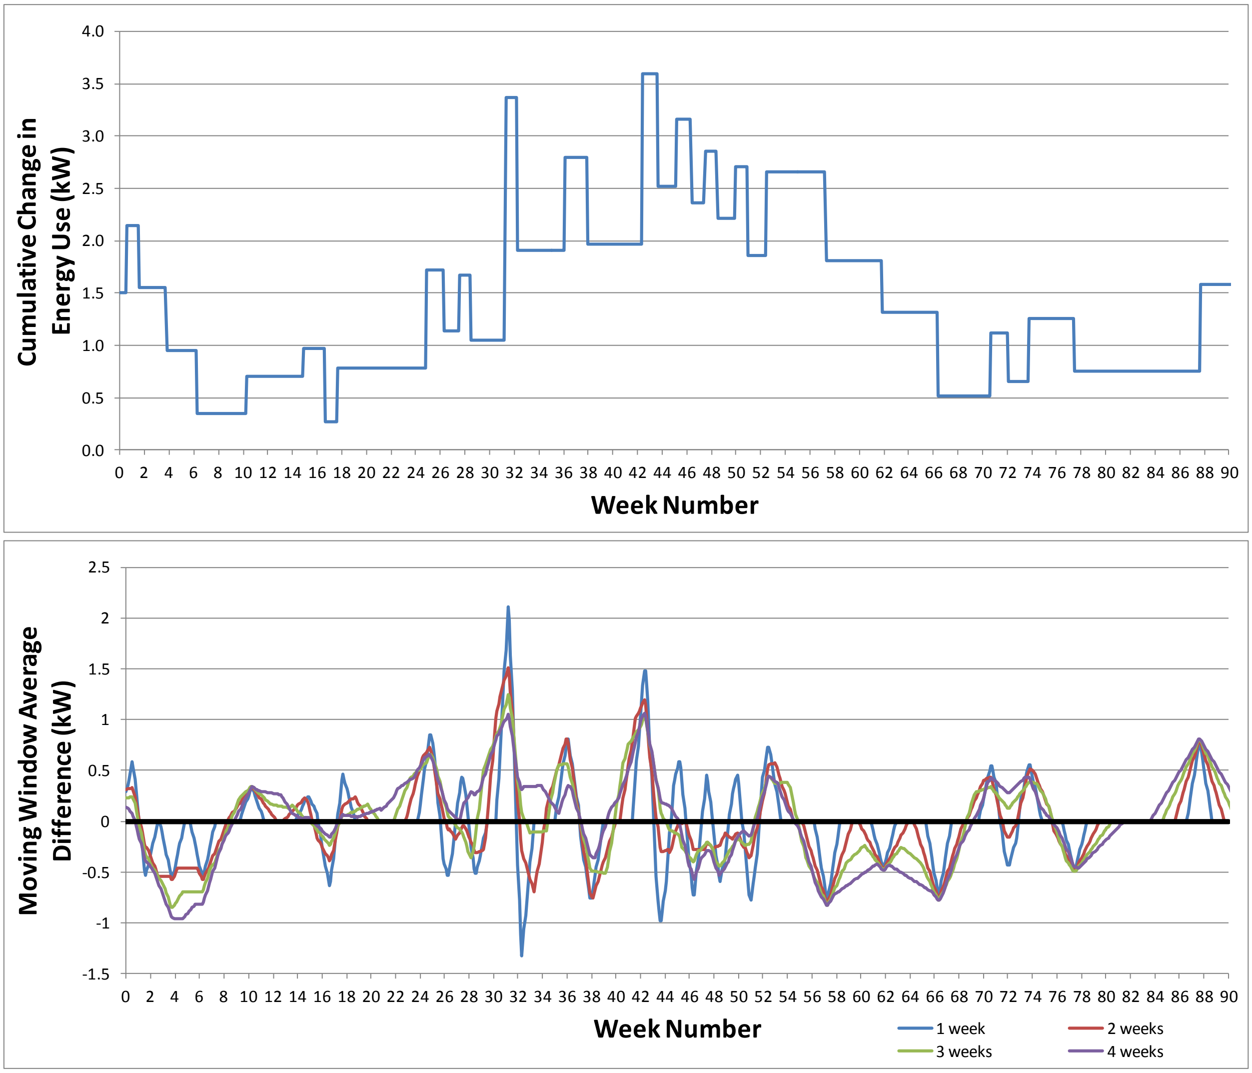

Figure 7: Step-wise change model of long-term energy use variations. The upper plot shows the cumulative step-wise energy use changes detected in figure 6, and the lower plot shows the moving window calculation of the average difference on the cumulative energy use difference.

While the procedures above demonstrate how changes in long-term energy use can be identified, we confirm their efficacy by the following experiment. The simple step-wise energy use change model developed above is modified by adding a noise term to the signal. If this doesn't conceal the long-term signal changes, then the method that isolated these changes has some evidence of robustness. Accordingly, we add a noise term to the cumulative energy use change signal by adding a random number of between 0 and 2 kW. The cumulative energy use plus noise term is shown in the upper graph in figure 8.

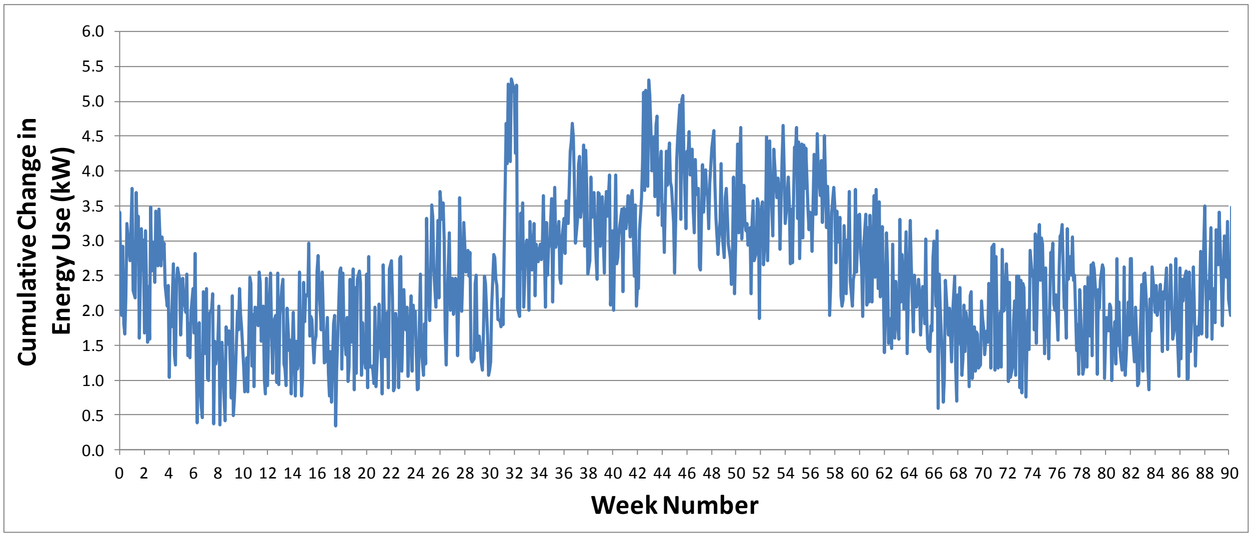

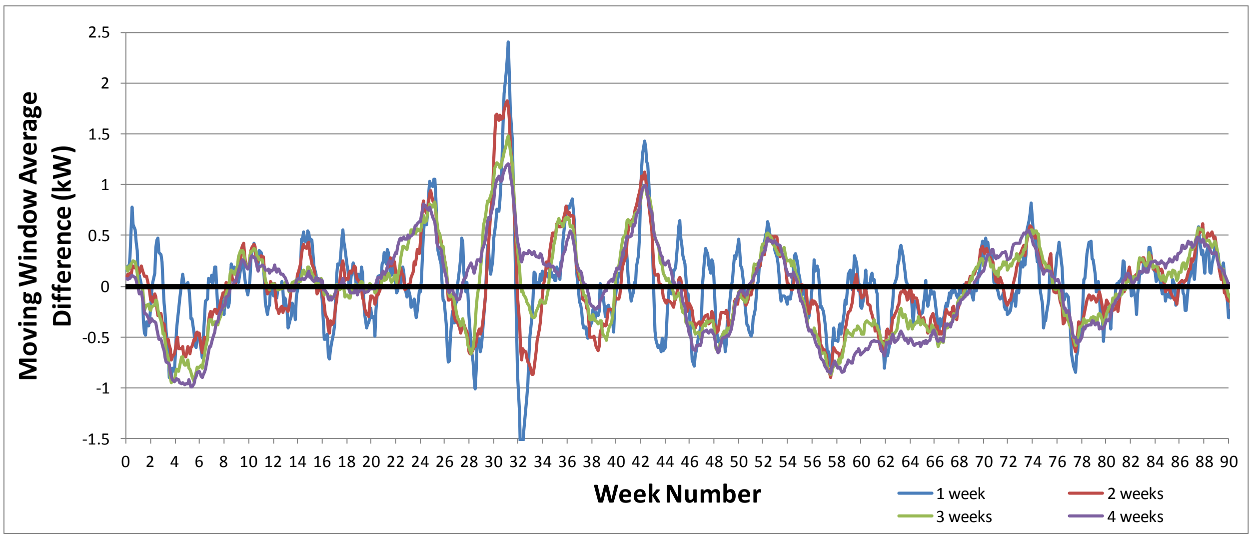

Figure 8: Step-wise change model of long-term energy use variations with noise added. The upper plot shows the cumulative step-wise energy use changes shown in figure 7 with a noise term (a random number between 0 and 2) added. The lower plot shows the moving window calculation of the average difference on the cumulative energy use difference.

The lower graph in figure 8 shows the result of a moving window difference in averages calculation for different moving window sizes. The basic features of the figure 6 are replicated. Each of the major deviations from zero shown in figure 6 is also present in the lower graph of figure 8. There is a noisy portion of the signal in both graphs with the smaller and higher frequency noise being smoothed to a greater extent with larger moving window sizes. While some minor details of the two graphs differ (specifically some of the specified energy use change peaks are distinctly wider in the 1-week window calculation in figure 6), the general qualitative and quantitative similarity gives us confidence that the type of statistical model for detecting individual household energy use changes is useful.

Determination of Energy Use Difference Distribution Details

Once the time and magnitude of energy use differences have been determined, it is fairly straightforward to characterize the distribution of energy use differences at the transition point. Note that in figure 2, when we plot two different load duration curves for an energy use transition (e.g. a dying server), this implies that we can calculate the energy change as a function of whether the load point is at peak load (i.e. percentile 0) or at base load (percentile 100%). In figure 9 we illustrate this calculation for the week 17 and week 70 transitions seen in figure 6.

To calculate the frequency distribution, we collect all of the data between the current and previous energy use change date, and calculated the corresponding energy use cumulative frequency distribution (as illustrated in figure 2). We then calculate the difference in the cumulative distributions for the post vs. pre-transition periods to calculate the energy difference as a function of frequency.

We see that for the week 17 transition, that the difference in energy use is concentrated at peak load, with only a small energy use change at base load. This can occur if there is a change in occupancy that involves very little change in equipment.

In contrast for the week 70 transition, there is an approximately constant difference in energy use. Because of the low resolution of the ELCAP data the curve is not smooth, but with higher resolution more current AMI data, such curves much smoother. The combined impact is a nearly constant 100 watts of electricity use increase.

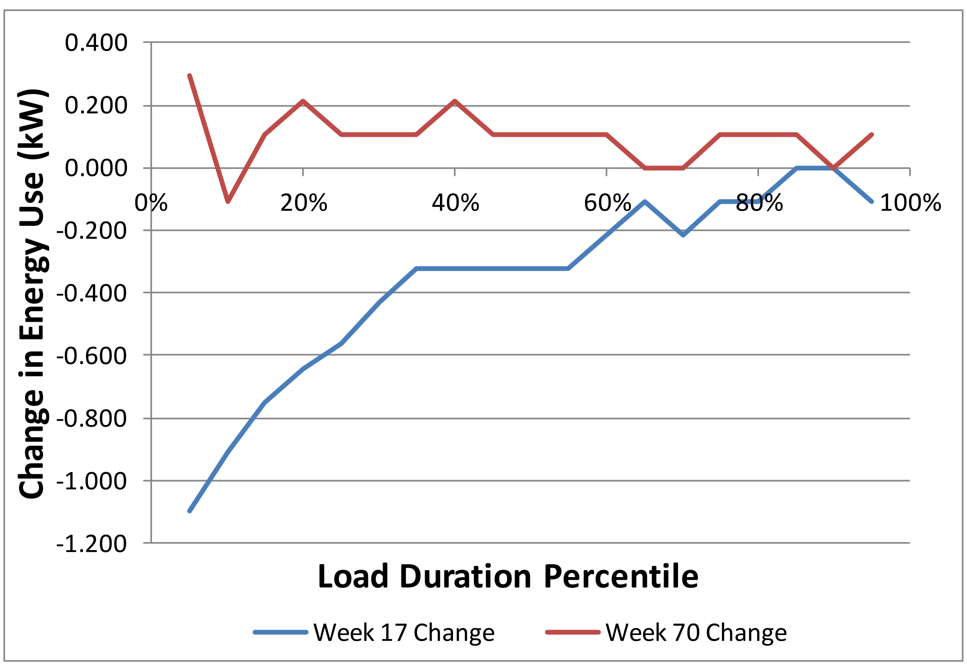

Figure 9: Distribution of energy use change for two transitions found in figure 6. The upper, red curve is the energy use change distribution for the week 70 change, and the lower blue curve is the distribution for the week 17 change.

Summary and Conclusion

In this paper we have illustrated the possibility of using AMI data to quickly and accurately detect electricity use changes in an individual household. We have considered using AMI data as basic time series, where the data consists of a series of step-wise energy use changes plus noise.

We applied a simple moving window calculation of average difference (between a forward window and a backward window) to detect points at which significant energy use changes occur. This allows us to decompose the AMI signal into a series of step-wise changes plus high frequency noise.

We then show how we can compare the load duration or load frequency functions before and after the detected energy use transitions to calculate the on peak vs. baseline distribution of energy use change.

This new method holds the potential of allowing for a new, cheaper, more accurate and much more rapid method of energy efficiency program impact detection. In this new method, AMI data would be decomposed into a series of step-wise energy use transitions, and then by mapping the time of an energy efficiency intervention and an observed energy use change, one can obtain a more precise and confident correlation between efficiency measure and household energy use at the individual household level. The method also holds promise for examining changes in energy use not induced by an energy efficiency program which are among the factors typically confounding traditional energy savings calculation approaches.

REFERENCES

[1] Pratt, R. G., et al. *Commercial equipment loads: End-Use Load and Consumer Assessment Program (ELCAP)*. No. DOE/BP/13795-26. Pacific Northwest Lab., Richland, WA (USA), 1990.

[2] Brown, Richard E., and Jonathan G. Koomey. "Electricity use in California: past trends and present usage patterns." *Energy Policy* 31.9 (2003): 849-864.