Lawrence Berkeley National Laboratory
Lawrence Berkeley National Laboratory

Title
Quantifying Changes in Building Electricity Use, with Application to Demand Response

Permalink
https://escholarship.org/uc/item/6068k5nh

Author
Mathieu, Johanna L.

Publication Date
2011-05-23

DOI
10.1109/TSG.2011.2145010

Peer reviewed

Quantifying Changes in Building Electricity Use, with Application to Demand Response

Johanna L. Mathieu, Phillip N. Price, Sila Kiliccote, Mary Ann Piette

Lawrence Berkeley National Laboratory

April 2011

DISCLAIMER

Quantifying Changes in Building Electricity Use, with Application to Demand Response

Johanna L. Mathieu, Phillip N. Price, Sila Kiliccote, and Mary Ann Piette[*]

Abstract

We present methods for analyzing commercial and industrial facility 15-minute-interval electric load data. These methods allow building managers to better understand their facility's electricity consumption over time and to compare it to other buildings, helping them to 'ask the right questions' to discover opportunities for demand response, energy efficiency, electricity waste elimination, and peak load management. We primarily focus on demand response. Methods discussed include graphical representations of electric load data, a regression-based electricity load model that uses a time-of-week indicator variable and a piecewise linear and continuous outdoor air temperature dependence, and the definition of various parameters that characterize facility electricity loads and demand response behavior. In the future, these methods could be translated into easy-to-use tools for building managers.

Keywords: data visualization, demand forecasting, demand response, energy efficiency, energy management, load management, regression analysis

1 Introduction

Building managers often look for opportunities for energy cost savings through energy efficiency, electricity waste elimination (through anomoly detection, changing operational schedules, etc.); peak load management; and participation in Demand Response (DR) tariffs and programs, which incentivize reduced electricity consumption during peak hours or when grid reliability is jeopardized. Analysis of fifteen-minute-interval whole-building electric load data is a good starting point for discovering opportunities to reduce energy costs through building energy management. These data are usually available to commercial and industrial (C&I) customers; however, they are not commonly used by building managers because the raw data are difficult to process and interpret. Therefore, more effective methods are needed to translate electric load data into actionable information.

In this paper, we present methods for analyzing 15-minute-interval electric load data from C&I facilities. Specifically, we describe graphical representations of electric load data, a regression-based electricity load model, and the definition of various parameters that characterize electric load as a function of time (i.e. the "load shape") and DR behavior. Our technical contributions are as follows:

1. We describe new ways of visualizing electric load data;

2. we introduce a time-of-week indicator variable into regression models of electric load;

3. we avoid the use of change-point models, but still capture a nonlinear relationship between outdoor air temperature and load, by dividing temperatures into many intervals and then fitting a piecewise linear and continuous temperature dependence;

4. we define new parameters to characterize electric load shapes and DR behavior;

5. and we apply the modeling methods to evaluate DR effectiveness.

In Fig. 1, we present a framework for using electric load data to assess opportunities for DR, and evaluate the effectiveness of the DR strategies that were implemented. Traditionally, building managers develop DR strategies after only minimal analysis of their load data. They primarily work with their utility or aggregator to adapt DR strategies, such as those in [1], that have worked well in other buildings. Therefore, they may miss building-specific DR opportunities or may implement DR strategies that are ineffective in their buildings. Building managers do not typically approach DR strategy development systematically, in part because there are too few DR methods and tools [2].

Therefore, we propose that building managers use tools incorporating data analysis methods such as load shapes, standardized load parameters, parameter plots, and load prediction to analyze their facilities' current and historic load shape. These methods are described in Sections 3 and 4. With knowledge resulting from this analysis and knowledge of DR strategies that have worked well in other buildings, a building manager is able to develop a list of informed questions that help direct an evaluation of building operations, controls, systems, and end-uses. This process is explained in Section 5. The results of the evaluation allow the building manager to identify potential DR strategies that are specific to his or her building.

The effectiveness of DR strategies that have been executed is determined using load prediction models to estimate what load would have been on a DR event day if a DR event had not occurred (a literature review is presented in Section 4.1). There are few tools available to building managers that automate load prediction. Also, additional data analysis methods could enhance the interpretation of load prediction results. Therefore, we propose that building managers use tools incorporating load prediction, DR residuals, and DR parameters to evaluate DR effectiveness. These methods are detailed in Sections 4 and 6.

As Fig. 1 shows, developing DR strategies should be an iterative process: identify a DR opportunity; implement and evaluate it; and, if desired, modify it to improve the response or move on to the next opportunity. Improving a building's DR strategies benefits the building by reducing its energy costs. It also benefits the grid [3]. Of course, these benefits must be weighed against DR costs [4].

A similar framework to that in Fig. 1 could be applied to assess opportunities for any form of building energy management, but we focus on DR because (1) more DR estimation tools and methods are needed [2], and (2) the use of DR is expanding. In many utility service territories, C&I facilities are called to shed or shift load on hot summer afternoons when the electricity grid is stressed with high loads due to residential and commercial cooling. In the future, facilities may be dispatched at any time, especially those participating in wholesale energy, capacity, and ancillary service markets (such as spinning reserve [5], non-spinning reserve [6], and regulation/load following [7]), and as DR is used to support the integration of intermittent renewable energy resources [7, 8]. Therefore, new analytical methods are needed to assess opportunities for and evaluate the effectiveness of "any day, any time" DR.

2 Data Sources

2.1 Electric Load Data

Fifteen-minute-interval electric load data from dozens of C&I facilities (peak load >200 kW) were obtained from the Pacific Gas and Electric Company (PG&E), a California investor-owned utility. Table 1 shows the C&I facilities used in examples throughout this paper. All facilities studied in this research participated in the Critical Peak Pricing (CPP) Program, a DR program previously offered by PG&E. DR events, generally triggered by high outdoor air temperature predictions, were called on up to twelve summer days per year (referred to as 'DR days' in this paper). They were announced by 3pm the business day before the event. In exchange for lower partial-peak and on-peak energy prices on non-DR days, energy prices on DR days were raised to three times the normal rate from 12-3pm (moderate price period) and five times the normal rate from 3-6pm (high price period).
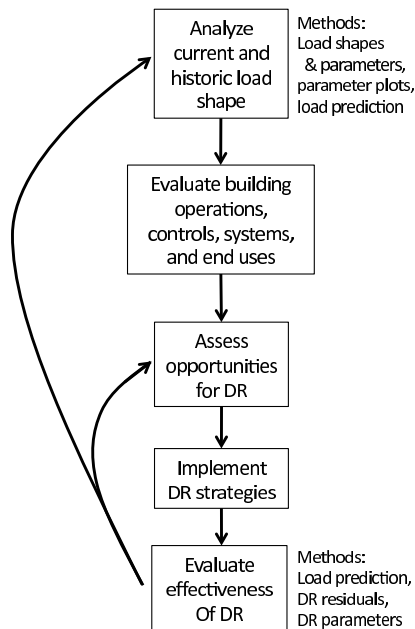
Figure 1: Framework to assess and evaluate DR opportunities.

Table 1: Facilities used in examples.

Facility	Type	Location	Weather Station(s)	Peak Load* (kW)	Load Factor*	Peak Load Intensity* (W/m²)
County Bldg	Office	Martinez	Concord	543	0.33	44.6
Furniture Store	Retail	East Palo Alto	Palo Alto & Hayward	1344	0.50	48.2
Bakery	Industrial	Oakland	Oakland	620	0.61	66.1

*Computed for May - Sept 2008, with 15-minute-interval data.

These facilities used the Open Automated Demand Response (OpenADR) communication specification to send DR event notification and/or electricity prices from the utility to facilities [9]. Using OpenADR allowed the facilities to automatically respond to DR events using pre-programmed control strategies such as HVAC set point changes, light switching and dimming, and industrial process shedding [1].

2.2 Temperature Data

For each facility, outdoor air temperature data from the nearest National Oceanic and Atmospheric Administration weather station were obtained from the National Climatic Data Center [10]. The data were collected by the weather station approximately hourly. Many of the temperature data are missing. We interpolate the temperature data linearly to assign an approximate temperature to every 15-minute-interval. When six or more hours of data are missing, we instead leave the temperature vector blank for those hours. In some cases, when the temperature data for a given site were particularly spotty we have filled the holes with data from a backup weather station (another nearby NOAA weather station). Table 1 shows the weather station(s) from which data were acquired for each of the facilities used in the examples.
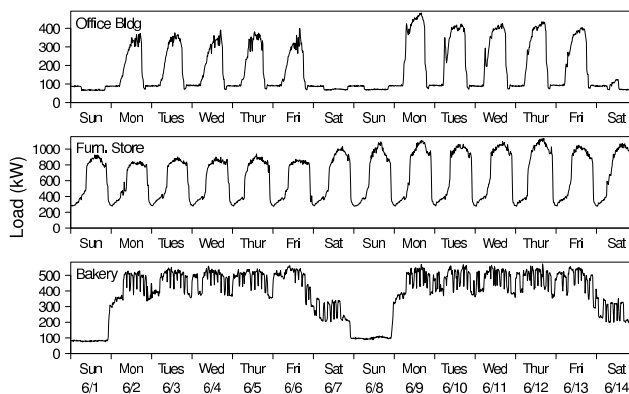
Figure 2: Electric load versus time.

3 Visualizing Electric Load Data

It is possible to learn a significant amount about the performance of a building over time, and as compared to other buildings, by plotting electric load data in various ways [11, 12, 13]. Analyzing electric load data graphically can lead to more insights into building characteristics, operations, and use than can be gained by only analyzing summary statistics. In this section, we describe several ways of visualizing electric load data. We first plot time series data and describe what can be learned from these plots. We then define and plot a small set of parameters that are useful for describing load variation from one day to the next. These "parameter plots" may help identify aspects of load shapes more easily.

3.1 Load Shapes

Fig. 2 shows electric load versus time for the three example facilities. Striking differences between facilities are immediately apparent including differences in operating hours, daily load shape regularity, the magnitude of daytime versus nighttime loads, and the variation in load from one 15-minute-interval to the next (i.e. the smoothness of the load shape). In addition, day-to-day changes within facilities can be observed. For instance, all of the facilities have higher loads in the second week shown, likely due to higher outdoor air temperatures (and therefore more need for cooling) during the second week.

   Many building load shapes share some features (Fig. 3). Most buildings have a clear *base load*, attained during the night, below which the power consumption rarely falls. In the early morning, the HVAC system switches from nighttime to daytime operation, and, if the building interior warmed/cooled overnight, the HVAC system may turn on at high power to cool/warm the building. This results in a short-lived load spike called the *morning start-up*. As the morning continues, load increases (*morning ramp-up* with increased occupancy and, in the cooling season, with increased outdoor air temperature. At some point the building reaches its *peak load* for the day. Peak loads can be computed over any time interval (e.g., daily, seasonally, yearly), and are generally more variable than base loads. In the afternoon or evening, the HVAC system switches back to nighttime operation and the power consumption quickly decreases, a phenomenon called the *evening setback*. In some buildings, some excess over the base load persists into the evening leading to an *evening shoulder*.

   Plotting time series load data and/or overlaying data from different time periods can be useful for noticing and characterizing changes in load shapes and their features. However, there are limitations to this approach:

1. For most facilities, energy consumption is a function of weather; however, we are generally interested in understanding changes in energy consumption that are not caused by weather variation. To deal with this issue, we present our method for weather-normalizing load data in Section 4.2.
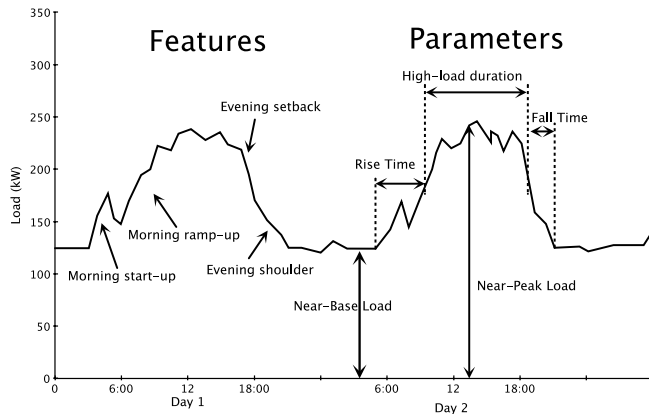
4

Figure 3: Load shape features and parameters.

Table 2: Load shape parameter definitions.

Load Shape Parameter	Definition
Near-Base Load (kW)	2.5^{th} percentile of daily load.
Near-Peak Load (kW)	97.5^{th} percentile of daily load.
High-Load Duration (hrs)	Duration for which load is closer to near-peak than near-base load.
Rise Time (hrs)	Duration for load to go from near-base load to start of high-load period.
Fall Time (hrs)	Duration for load to go from end of high-load period to near-base load.

2. Some phenomena may be difficult to recognize in plots of time series load data. For instance, noticing a tendency for load to increase gradually over a long period might be difficult, since this small trend will often be superimposed on seasonal variation and other features. Overlaying data separated by a year may reveal that load is higher now than a year ago, but will not reveal whether the change was gradual or abrupt.

3. Graphical approaches to understanding and comparing load shapes are useful only when people are able to devote time and effort to using them. Automated methods can potentially reduce this effort, though fully automating the *analysis* of time series data such as those in Fig. 2 is a daunting task.

3.2 Load Shape Parameters

As with any time series, load data invite the calculation of a wide variety of summary statistics. It is useful to distinguish between two kinds of summary statistics: those that summarize various aspects of the load and its variability, and those that summarize the amount of load or load variability that is not related to weather, i.e. weather-normalized summary statistics. In this section, we consider summary statistics that are *not* weather-normalized.

We recommend five parameters that are useful for describing load shapes (Fig. 3 and Table 2). The value of each of the parameters can be calculated for each day and these values can be summarized (e.g., mean and standard deviation of each parameter).

We define the "near-base" and "near-peak" load since the base and peak load summarize extreme events, which may differ qualitatively from events that are actually of interest. For instance, the power consumption during a power outage should not be considered a building's base load. As for the peak load, consider two buildings with the same power consumption profile that each have power usage spike for a full 15 minutes.
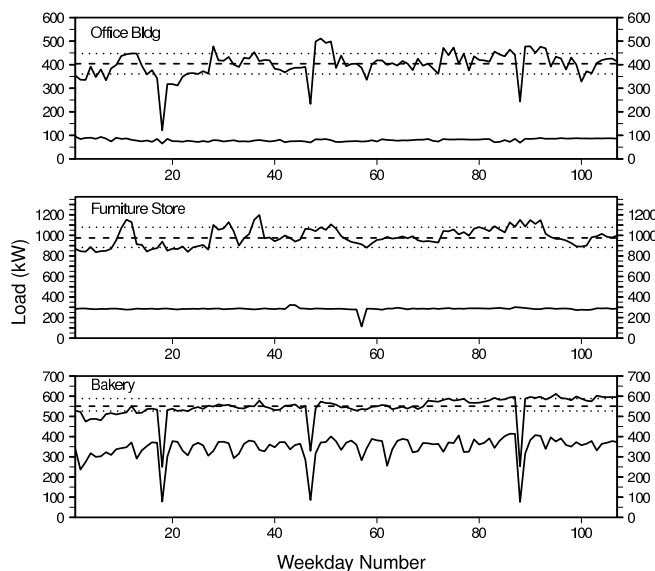
5

Figure 4: Near-base (lower curve) and near-peak load (upper curve) for each weekday from May to Sept 2008. Dotted and dashed lines show the 15^{th}, 50^{th}, and 85^{th} percentiles of the near-peak load. Significant dips in load for the office building and bakery are holidays (Memorial Day, Independence Day, and Labor Day).

The first building's spike occurs during a single 15-minute-interval, but the second building's spike is spread across two 15-minute-intervals. The measured peak load for the first building will be higher than that of the second building despite that the fact that they both consumed the same amount of power for the same amount of time. To help deal with these issues, we recommend using the 2.5^{th} and 97.5^{th} percentile of daily load instead of the minimum and maximum.

We define three time intervals, high-load duration, rise time, and fall time, to characterize how the load changes throughout the day. In practice, it is difficult to find definitions of these time intervals that yield consistent, easily interpretable results. The definitions in Table 2 work well if the load shape is something like Fig. 3, but it do not produce useful numbers in some cases, such as for load shapes that do not vary substantially over time or load shapes that have multiple extreme maxima and minima during each day.

3.3   Parameter Plots

Plotting the parameters defined in Table 2 can help us to recognize phenomena that we would be likely to miss if we only analyze plots of 15-minute-interval load data. For example, in Fig. 4, we plot each facility's near-base and near-peak load for weekdays from May to Sept 2008. Examining these plots reveals several things: (1) the office building's and furniture store's near-base loads were relatively constant, while the bakery's near-base load varied day-to-day; (2) the bakery's near-peak load increased over the course of the summer; and (3) the bakery's near-peak load did not vary significantly day-to-day, unlike that of the office building and furniture store. While all of this information *could* be obtained from a plot of 15-minute-interval load data, the many layers of information present in such a plot make it very difficult to identify these trends.

Fig. 5 shows an example of what one can learn from plotting a facility's high-load duration over time. Analyzing the plot we learn that, in early 2008, the furniture store's operating hours were extended by more than an hour each day, and became more uniform day-to-day. On an initial inspection of the store's 15-minute-interval load data we did not notice this change (though in retrospect it is visible). In comparison, when we inspected Fig. 5, we immediately noticed the change and when it occurred. More examples of useful
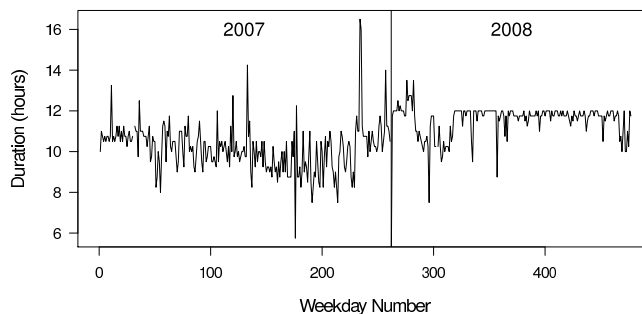
Figure 5: High-load duration for the furniture store for each weekday from Jan 2007 to Oct 2008.

parameter plots can be found in Price [14].

4 Predicting Electric Loads

In this paper, we focus exclusively on understanding past electric load. We use statistical models that quantify the electric load as a function of time-of-week and outdoor air temperature. Following standard statistical terminology, we call the output of the model a "prediction" of the electric load, even if the events occurred in the past. Predictions of future load are referred to as "forecasts."

Electric load prediction is useful for comparing how a facility is currently performing to how it has performed in the past. Specifically, load prediction is used to: (1) understand changes in a facility's electricity consumption patterns from one time period to the next; (2) quantify the effectiveness of DR strategies; (3) quantify the effectiveness of energy efficiency retrofits; and (4) perform anomoly detection (by finding times when the building is not behaving as it has behaved in the past). In each case, the predicted load is compared to the actual load. Importantly, the predicted load is computed under the same key conditions as those that lead to the actual load. For example, using actual weather data to compute the predicted load allows us to 'weather-normalize' our prediction; the remaining differences between the predicted and actual load are not weather-dependent.

Weather-normalization is especially important for facilities with significant temperature-dependent loads (e.g., cooling loads, electric heating loads). For example, in Fig. 6 we show the furniture store's actual and predicted load for three days in July in 2006 and 2009. Predictions are computed using the load prediction method that will be introduced in Section 4.2. To give a sense for model accuracy, the top plots compare actual and predicted load for each year. The bottom left plot compares actual load data across years. From this plot, we learn that the facility used significantly less energy in 2009 than 2006. However, we would like to know how much of the difference is due to changes in equipment, operations, and use, and how much is simply due to weather. The bottom right plot shows weather-normalized predictions. Specifically, predictions from 2006 and 2009 are shown, as well as predictions that use the 2009 model but 2006 temperatures. Comparing the gray and thin black lines we can see the portion of the savings not due to weather, while comparing the thin and thick black lines we can see the portion of savings due to weather. Some, though not all, of the difference in daytime load is due to weather, while almost none of the difference in nighttime load is due to weather.

We next present a brief review of existing load prediction methods. Then, we propose a load prediction method that uses a time-of-week indicator variable and a piecewise linear and continuous outdoor air temperature dependence. In addition, we briefly discuss the sources of error associated with load prediction.
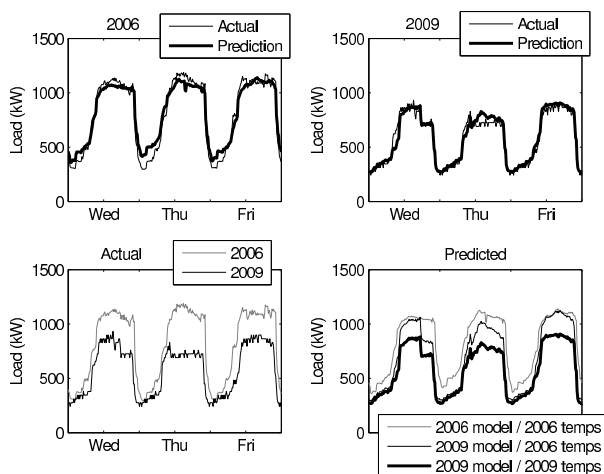
Figure 6: The importance of weather-normalization. The top plots compare actual and predicted load for each year, the bottom left plot compares actual load data across years, and the bottom right plot shows weather-normalized predictions. '2009 model / 2006 temps' refers to a load prediction made using a baseline model created with 2009 load data and 2006 outdoor air temperature data, etc.

4.1   Existing Methods for Load Prediction

Electric utilities use simple methods to predict the electric loads of facilities that participate in DR programs. The predictions are called baselines because they provide a baseline against which demand reductions are computed. California electric utilities use methods such as averaging the electric use profiles of the three days with the highest energy usage out of the last ten business days, or averaging the electric use profiles of the last ten business days [15, 16]. To adjust for weather and other conditions on DR days, predictions are sometimes multiplied by a morning adjustment factor: if a building is using, say, 10% more electricity on the morning of a DR day than on other mornings, then its predicted afternoon load is increased by 10% as well. Unfortunately, this approach has a serious problem: if a facility shifts some of its load to the morning on a DR day, the result of the morning adjustment is an overestimate of the amount of energy the building would have consumed during the DR period. Essentially the building is credited twice for the same energy shift: once by using less energy than it would have if it did not participate in DR, and once by overestimating the amount of energy that it would have used.

More advanced load prediction methods have been developed for a variety of applications including estimating the effectiveness of energy efficiency retrofits and forecasting utility-scale electric loads. Claridge [17] discusses many approaches for using historical electric load data to model the electricity consumption of C&I facilities including linear regression models, calibrated simulations, Fourier series models, and neural network models (e.g., [18]). Granderson *et al.* [19] also describe several other methods for residential load prediction including non-linear models such as locally-weighted regressions, "bin" models in which load predictions are based on the average load for time periods that share the same bin as current conditions (e.g., weather and time-of-day), and nearest-neighbor models in which the current load is predicted to be the same as it was when previous conditions were closest to current conditions. Taylor *et al.* [20] compare several methods for forecasting utility-scale electric loads including exponential smoothing models, ARIMA models, neural network models (e.g., [21]), and regression with principal component analysis.

In this paper, we use linear regression models because – when constructed appropriately – they provide a good fit to load data in most buildings, their results are easy to interpret, they are easy to modify, and they present negligible computational burden. In addition, regression methods have performed well when compared against other load prediction methods [22, 23, 24]. In 1986, Fels [25] introduced the PRInceton
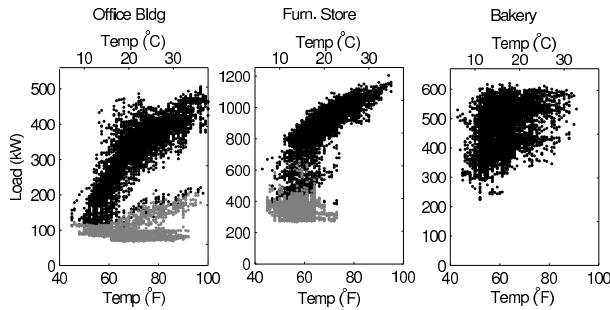
8

Figure 7: Scatter plot of load versus temperature, including both occupied mode load (black) and unoccupied mode load (gray). Data shown are from May-Sept 2008.

Scorekeeping Method, or PRISM, a regression-based load prediction method to standardize the measurement of energy conservation savings. Heating degree-days and monthly electricity consumption are related through a simple piecewise linear regression model. Other innovative regression-based load prediction methods followed including methods using finer resolution (e.g., daily and hourly) electric load data, change-point models [26], and multiple linear regressions [27, 28]. Kissock *et al.* [26] developed regression models specifically for commercial buildings, while Kissock and Eger [29] developed models for industrial buildings.

4.2   Load Prediction Method

We have developed a linear regression-based load prediction method that includes two novel features: (1) a time-of-week indicator variable, and (2) a piecewise linear and continuous outdoor air temperature dependence derived without the use of a change-point model or assumptions about when structural changes occur. Both of these of these features will be discussed in depth at the end of this section, after our load prediction method is introduced.

A facility's electric load is usually a function of both temperature and time-of-week, as shown in Figs. 7 and 8.[1] Therefore, we consider both temperature and time-of-week in our regression model. Our method is as follows: we divide a week (Monday-Friday) into 15-minute-intervals (indexed by i), e.g., the first interval is from midnight to 12:15 on Monday morning, the second interval is from 12:15 to 12:30, and so on. A different regression coefficient for each time-of-week, α_i, allows each time-of-week to have a different predicted load. Additionally, we expect a temperature dependence like that in Fig. 9. When the outdoor temperature is high, cooling load will increase with temperature, and when the outdoor temperature is low, heating load will increase as temperature decreases (even when electricity is not used as the heat source electricity will be required to run pumps and fans when the building is heating). For some range of moderate temperatures, the load may be insensitive to temperature because neither cooling nor heating is needed (the temperature is said to be in the "dead-band"). Sometimes the outdoor air temperature may be so high that the cooling capacity cannot achieve the desired indoor temperature set point, at which point load is at the maximum possible air conditioning load (maxed-out). Additional change-points are also possible (e.g., in facilities with two chillers, the second of which only turns on when the first is operating near capacity).

This nonlinear temperature effect can be modeled with a piecewise linear and continuous temperature-dependent load model. For each facility, we divide the outdoor air temperatures experienced by that facility into six equally-sized temperature intervals.[2] For example, if the minimum temperature experienced by the facility were 50°F (10°C) and the maximum temperature experienced by the facility were 110°F (43.3°C),

---

[1]Temperature and time-of-week are correlated: the highest temperatures generally occur in the afternoon and the lowest temperatures generally occur at night. Therefore, both time-of-week and temperature effects are superposed on each plot, and Fig. 7 shows *total load* versus temperature, not *temperature-dependent load* versus temperature.

[2]Any number of intervals could be used, but we recommend using at least twice the expected number of change-points, but not so many as to cause over-fitting problems.
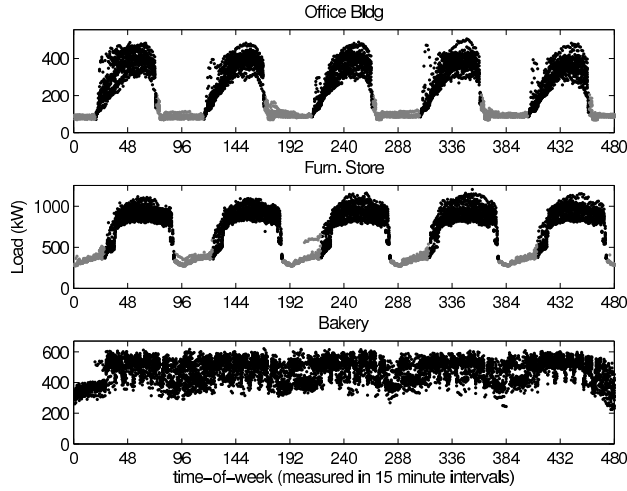
Figure 8: Scatter plot of load versus time-of-week, including both occupied mode load (black) and unoccupied mode load (gray). Data shown are from May-Sept 2008, so for each time-of-week there are 21-22 data points.
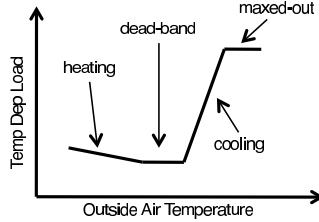

Figure 9: Temperature-dependent load.

the temperature intervals would be 50-60°F (10-15.6°C), 60-70°F (15.6-21.1°C), 70-80°F (21.1-26.7°C), 80-90°F (26.7-32.2°C), 90-100°F (32.2-37.8°C), and 100-110°F (37.8-43.3°C). A temperature parameter, β_j with $j = 1...6$, is assigned to each outdoor air temperature interval.

To achieve piecewise linearity and continuity, the outside air temperature at time t (which occurs in time-of-week interval i), $T(t_i)$, is broken into six component temperatures, $T_{c,j}(t_i)$ with $j = 1...6$. Each $T_{c,j}(t_i)$ is multiplied by β_j and then summed to determine the temperature-dependent load. Let B_k ($k = 1...5$) be the bounds of the temperature intervals. Component temperatures are computed using the following algorithm:

1. If $T(t_i) > B_1$, then $T_{c,1}(t_i) = B_1$. Otherwise, $T_{c,1}(t_i) = T(t_i)$ and $T_{c,m}(t_i) = 0$ for $m = 2...6$, and algorithm is ended.

2. For $n = 2...4$, if $T(t_i) > B_n$, then $T_{c,n}(t_i) = B_n - B_{n-1}$. Otherwise, $T_{c,n}(t_i) = T(t_i) - B_{n-1}$ and $T_{c,m}(t_i) = 0$ for $m = (n+1)...6$, and algorithm is ended.

3. If $T(t_i) > B_5$, then $T_{c,5}(t_i) = B_5 - B_4$ and $T_{c,6}(t_i) = T(t_i) - B_5$.

An example computation is shown in Table 3.

The temperature parameters β_j are only used when a facility is operating in occupied mode since one would expect a facility's response to temperature would change at night. The start and end of the occupied mode are manually determined by looking at average load profiles on non-DR days. In Figs. 7 and 8, data

10

Table 3: Example of component temperature computation, for $B_1 = 10, B_2 = 20, ...B_5 = 50$ (in arbitrary temperature units).

T	$T_{c,1}$	$T_{c,2}$	$T_{c,3}$	$T_{c,4}$	$T_{c,5}$	$T_{c,6}$
2	2	0	0	0	0	0
18	10	8	0	0	0	0
32	10	10	10	2	0	0
47	10	10	10	10	7	0
58	10	10	10	10	10	8

from occupied mode and unoccupied mode are differentiated (the bakery is never in unoccupied mode on weekdays). Analyzing Fig. 7, it is clear that the office building and the furniture store exhibit different behavior in different modes. For all facilities, occupied load, L_o, is estimated as follows:

$$\hat{L}_o(t_i, T(t_i)) = \alpha_i + \sum_{j=1}^{6} \beta_j T_{c,j}(t_i). \tag{1}$$

To predict load when the building is in unoccupied mode, we use a single temperature parameter, β_u since we expect most facilities in our data set to be operating at or near the dead-band at night.[3] Unoccupied load, L_u, is estimated as follows:

$$\hat{L}_u(t_i, T(t_i)) = \alpha_i + \beta_u T(t_i). \tag{2}$$

The parameters α_i for $i = 1...480$, β_j for $j = 1...6$ and β_u are estimated using non-DR day load and temperature data with ordinary least squares. Each of the 487 parameters is physically meaningful: power use varies in each 15-minute-interval in a week and varies as a function of outdoor air temperature. We use 15-minute-interval data from May through September, so approximately 20 data points are available to estimate each α_i; hundreds or thousands to estimate each β_j; and thousands to estimate β_u. Applying the model to data from a much shorter interval, such as four or six weeks, would likely run into problems from "over-fitting," with parameter values being overly influenced by stochastic variability in the data.

To test how well the load prediction method works we plot predictions on top of actual load data in Fig. 10. We also include scatter plots in Fig. 11. As can be seen in both figures, the prediction method works well for the office building and furniture store, but does not work as well for the bakery. The accuracy of the prediction is a function of how well the explanatory variables (time-of-week and outdoor air temperature) capture the power consumption of the facility. In the case of the bakery, we would expect that power consumption would be a function of the timing of industrial production processes, which is not captured in the model.

For the same facilities, we plot actual and predicted temperature-dependent load in both occupied and unoccupied model in Fig. 12. The bakery's electric load is less correlated with temperature than the office building's or furniture store's electric loads.

To our knowledge, our load prediction method differs from existing methods in two ways. First, we use a time-of-week indicator variable, α_i. Most methods we have seen compute separate regressions for each time-of-day (e.g., [30] and Model 5 in [15]), but do not capture day-to-day load variation. Adding a weekday/weekend indicator variable [31] or day-of-week indicator variable [24] helps with this problem but effectively only shifts the daily load shape up or down as a function of the day-of-week; it does not allow for other day-to-day changes in the load shape, such as shorter operating hours on certain days of the week, or days of the week with consistently higher peaks. A time-of-week indicator variable solves this problem, improving model accuracy.

---

[3]In climates where facilities transition between heating, dead-band, and cooling at night, it is advisable to use more temperature parameters to more accurately capture the load of the facility in unoccupied mode.
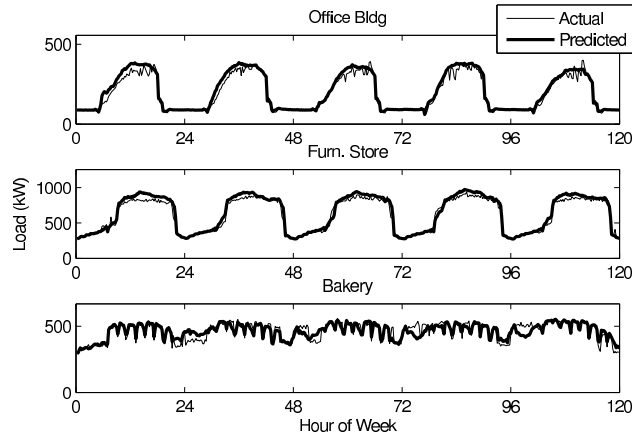
Figure 10: Actual versus predicted load time series for Monday-Friday, June 2-6, 2008.
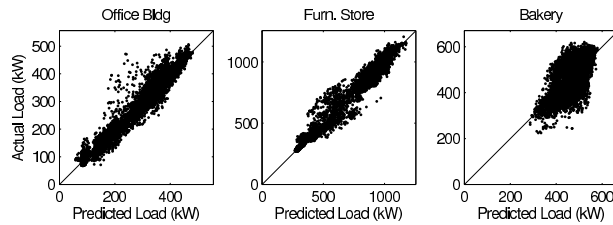
Figure 11: Actual versus predicted load scatter plots. Data shown are from May-Sept 2008. Significant outliers visible in the plot for the office building result from the model's inability to accurately predict the morning start-up.

Figure 12: Actual (dots) and predicted (line) temperature-dependent load. Data shown are from May-Sept 2008. The gray lines are piecewise linear and continuous functions of temperature, as defined in Section 4.2.

12

Second, our method is different because we avoid the use of change-point models. In much of the energy efficiency literature, change-point models are used to determine the outdoor air temperatures at which the building transitions from heating to the dead-band and fro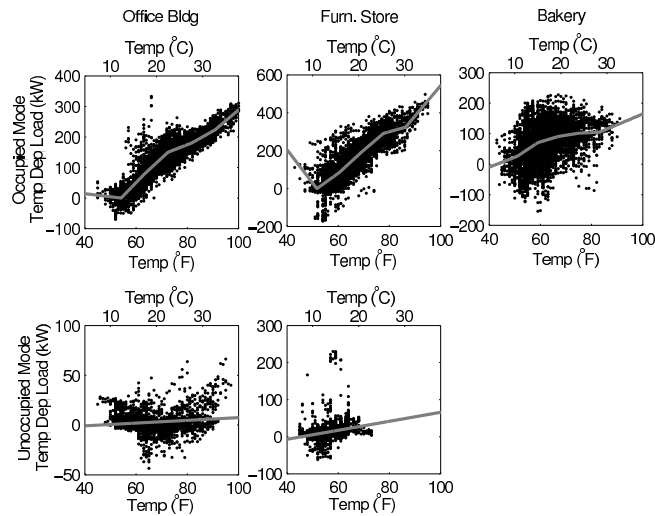m the dead-band to cooling (e.g., [26]). Our piecewise-linear modeling approach avoids the complexities of change-point models (such as the need for iterative regression) with no significant drawbacks. We divide the range of outdoor air temperatures experienced by the facility into six temperature intervals and allow the slope of the load versus temperature profile to be computed separately in each temperature interval, with the constraint that the predicted load must vary continuously as a function of temperature. We are able to do this because we are using high resolution load data (15-minute-interval) and therefore have sufficient data to determine the linear temperature dependence in each temperature interval.

We explored several variations on our load prediction method, including adding parameters associated with other weather data, such as humidity, to the model. However, since all of the facilities analyzed in this study are located in low-humidity climates, we did not include humidity in our final model. We also experimented with different numbers of temperature intervals, and allowing the temperature coefficients to vary with time interval. Most of these changes did not substantially improve the model fit. Surprisingly to us, even adding predictive variables such as the overnight outdoor temperature (which we expected to help predict the load during the morning start-up) did not improve the fit very much in most cases. In the end, we settled on the model discussed above. Others are possible and may be better for some buildings or some situations.

4.3 Load Prediction Error

Reddy *et al.* [32] enumerate the many sources of error associated with using regression analysis to model building electric load. Our regression model residuals are autocorrelated and heteroscedastic and the regression parameters, α and β, are correlated. Therefore, is very difficult to compute robust confidence intervals on the regression parameters (simply reporting the standard error associated with each regression parameter estimate, as in [25], underestimates the level of uncertainty). However, uncertainty on load predictions (as opposed to the regression parameters) can be approximated with the standard error, which can be computed at each interval, i. Computing robust confidence intervals on predictions made on days that were not used to build the regression models (e.g., DR days) is more complicated and is the subject of current research.

5   Identifying DR Opportunities

In Table 4, we show how the methods presented in Sections 3 and 4 can help building managers generate a list of informed questions that will help them identify opportunities for DR. Specifically, we list five data analysis methods and explain each method's relevance to DR. We also describe what the building manager should look for when examining a plot or evaluating a parameter. Lastly, we list some example questions that may be generated by applying each method. While the example questions in Table 4 are fairly general, a building manager could generate much more specific questions after data analysis. For example, in examining Fig. 4, one might ask: Why did the bakery's near-peak load increase during the summer, and could the additional load be curtailed during DR events or shifted outside of the DR period?

To answer these questions, the building manager will need to use the results of the data analysis to guide an evaluation of the facility's operations, controls, systems, and end uses (Fig. 1). For example, to figure out why the bakery's near-peak load increased during the summer the building manager will first need to determine when the daily near-peak load occurs. Then, he or she can determine which loads run during those hours, how much power those loads consume, and how that has changed during the summer. With this knowledge, the building manager can determine if the additional load is curtailable or shiftable, or if it is essential to critical building operation.

The benefit of applying these methods is that the building manager is able to focus his or her evaluation on issues relating to DR (though the same methods could be used to generate questions related to energy

Table 4: Partial list of methods for helping to identify DR opportunities.

Method	Relevance to DR	What to Look For?	Example Questions
Plot of daily load shape	Changes in electric load from hour-to-hour affect how much load a facility is able to shed.	How does load vary throughout the day?	Which equipment runs during each hour of the day? Which could be shed, shifted, or limited during a DR event?
Plot of near-peak and near-base load over time	Changes in electric load from day-to-day affect how much load a facility is able to shed.	How does load vary day-to-day?	Which equipment runs on the days with the highest loads? Which could be shed, shifted, or limited during a DR event?
Plot of high-load duration over time	Facilities can shed more load during hours when they are consuming more.	How long is the high-load duration, and how variable is it day-to-day and seasonally?	Which building systems/operations affect high-load duration and its variability?
Mean and standard deviation of the rise/fall time	Some facilities might have a hard time participating in DR while they are powering up/down.*	How long are the rise time and fall time, and how variable are they day-to-day?	Which building systems/operations affect rise/fall time and its variability?
Using load prediction to compare weather-normalized load shapes from two or more time periods	Variability in estimated DR sheds is a function of unmodeled load variability (discussed in Sec. 6).	How variable is the electric load over time (with weather held constant)?	Which loads are variable? Are the variable loads controllable (i.e. could they be shifted outside of DR periods)? Is load/shed predictability a requirement for the DR program?

*This is particularly relevant for facilities participating in DR to support the integration of intermittent renewable energy resources, such as wind and solar. These resources ramp in the morning and evening so facilities may be dispatched as they are powering up/down.

efficiency, electricity waste elimination, or peak load management). The answers to the questions not only help building managers develop DR strategies but also help them pair their facilities with the right DR programs, since all DR programs have different requirements (e.g., minimum shed required, length of time to hold shed, predictability of shed, allowable times-of-day or days-of-week of DR events, frequency of DR events).

6 Evaluating DR Effectiveness

In this section, we explain how the load prediction method described in Section 4.2 can be used to quantify the effectiveness of DR strategies. In addition, we define the 'DR Residual' and a small set of parameters useful for characterizing the DR Residual.

To estimate the effectiveness of a facility's DR strategies, we initially added parameters to our regression model to estimate the load shed during a DR event; however, this method did not allow us to understand shed-to-shed variability. Therefore, to analyze the effectiveness of a facility's DR strategies, we employ a 'predict and subtract' method:

1. Use the load prediction method described in Section 4.2 to develop a baseline model of electric load for the facility. Use only data from non-DR days to make the baseline model.

2. Acquire outdoor air temperature data from the DR day. Use the baseline model and the DR day temperatures to predict what load would have been on the DR day if the DR event had not been called.
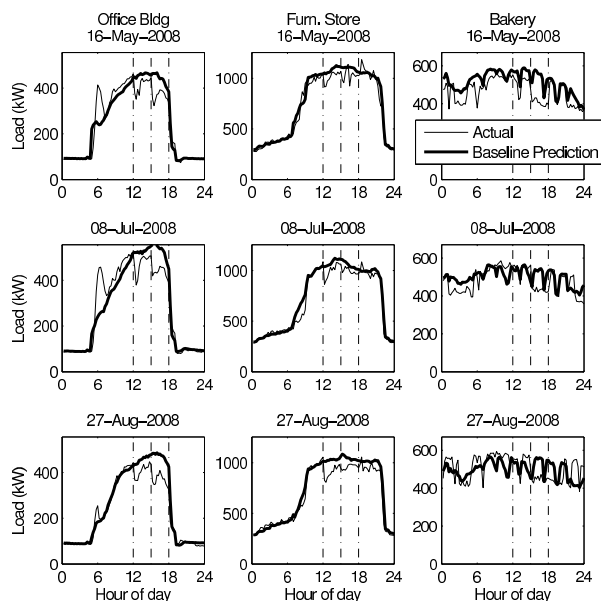
Figure 13: Actual and predicted load on three DR days (rows) for three facilities (columns). Vertical lines show the moderate and high price periods.

3. For each time interval in the day, subtract the baseline prediction from the actual electric load on the DR day. We call the resulting vector the *DR Residual*. For facilities participating in ancillary services markets, the DR residual is called *pseudo-generation* [6].

We applied this method to data from each of the facilities in Table 1. Since all DR events were called on non-holiday weekdays in May - Sept 2008, each baseline model was constructed only with load data from non-holiday weekdays during the same period. Though all of these facilities participated in a traditional DR program (DR events were only called hot summer afternoons), the same methodology could be used to evaluate "any day, any time" DR.

In Fig. 13, for each facility, we show the actual and baseline-predicted load for three DR days. In each case, the difference between the actual load and the baseline-predicted load is due to (1) the facility's DR strategies and other DR event-related behavior change, and (2) unmodeled load variability (i.e. explanatory variables in the model do not explain all of the components of the total load, resulting in model error as shown in Figs. 10 and 11). Despite automation of DR strategies (described in Section 2.1), for each facility, there is variability in estimated DR sheds. Shed variability results from unmodeled load variability, and, possibly, factors such as occupant behavior change on DR days; changes in building operations, controls, equipment, and end-uses from event-to-event; and changes in shed capacity as a function of outdoor air temperature, occupancy, total facility load, and other variables. To fully understand shed-to-shed variability, we need to compute prediction error on DR days, which is a subject of current research.

Because of load/shed variability from one DR event to the next it is instructive to look at plots of averages. In Fig. 14, for each facility, we plot the average baseline-predicted load, actual load, and DR residual for all eleven DR events days in 2008. The average DR residuals tell us many things including: (1) the office building and furniture shed load during both the moderate and high price periods, while the bakery only sheds load during the high price period; (2) the baseline model is unable to capture the office building's morning start-up; and (3) the furniture store's load rebounds after the DR event.

A small set of parameters can be used to characterize DR residuals (Table 5 and Fig. 15). The list is based on that developed by Mathieu *et al.* [33]. *Average demand shed* is important because it tells us, on
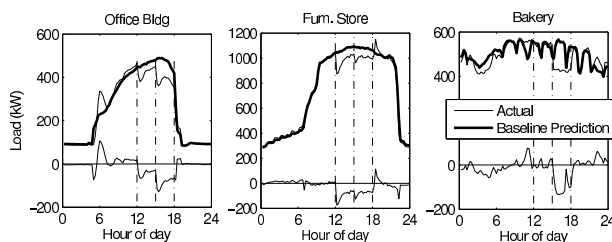
Figure 14: Average actual load, predicted load, and DR residual. Vertical lines show the moderate and high price periods.

Table 5: DR residual parameters and definitions.

DR Residual Parameter	Definition
Average demand shed (kW)	Baseline minus actual average load during DR event.
Intrashed variability (kW)	Standard deviation of demand shed during DR event.
Residual ramp time (min)	Duration for load to drop to average demand shed during DR event.
Rebound (kW)	Actual minus baseline average load in hour after DR event.
Daily peak demand (%)	Actual divided by baseline-predicted daily peak demand.
Daily energy (%)	Actual divided by baseline-predicted daily energy use.

average, how much load is shed during the DR event. *Intrashed variability* is a way for us to capture some of the dynamics of the shed. If intrashed variability is small relative to the average demand shed, the shed was held steady during the DR event, while, if it is large, the shed bounced around, increased, or decreased during the DR event. *Residual ramp time* tells us how quickly the facility sheds load, and *rebound* tells us how the facility behaves after the DR event. *Daily peak demand* and *daily energy* tell us how the DR event affects the facility's daily power and energy use.

These parameters can be computed for a certain DR event or an average DR event. In Table 6, for each facility, we give values for each parameter for each DR day shown in Fig. 13 and for the average DR day shown in Fig. 14. There is error associated with each parameter value, and the values for the means are more certain than the values for the individual events. For each facility, intrashed variability is relatively high as compared to average demand shed. Certain DR strategies, such as changing HVAC set points, often lead to more intrashed variability than strategies such as switching off lights or industrial processes. Both the county building and the furniture store can shed load within 15 minutes of the start of the DR period, while it takes the bakery 15-30 minutes. While all of the facilities use less *total* energy on most DR days, both the furniture store and the bakery may have higher *peak* demand on days when they participate in a DR event. Though a facility's DR day peak demand generally occurs outside of the DR event period (e.g., during the hour after the DR event, as is the case for the furniture store), these peaks affect demand charges.

7   Conclusion

We have presented methods for analyzing 15-minute-interval electric load data from C&I facilities. Applying these methods could help building managers ask the right questions to discover opportunities to reduce electricity bills through DR, energy efficiency, electricity waste elimination, and peak load management. We have focused on DR because of a lack of existing methods and tools for identifying DR opportunities and estimating DR effectiveness, and because the use of DR is expanding.

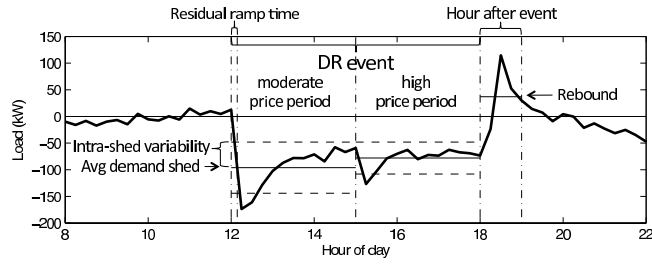A number of key findings from this study are as follows:

Residual ramp time   Hour after event

Load (kW)

150
100
50
0
−50
−100
−150
−200

DR event
moderate
price period   high
price period   ← Rebound

Intra-shed variability
Avg demand shed →

8   10   12   14   16   18   20   22
Hour of day

Figure 15: Several DR parameters shown on the furniture store's DR residual (from Fig. 14).

Table 6: DR parameter values.

Facility	Event Date	Average Demand Shed* (kW)	Intrashed Variability* (kW)	Residual Ramp Time* (min)	Rebound (kW)	Daily Peak Demand (%)	Daily Energy (%)
County Bldg	May 16	44 / 74	27 / 34	0-15 / 0-15	2	99	98
	July 8	56 / 93	39 / 28	0-15 / 0-15	5	93	99
	Aug 27	51 / 103	15 / 19	0-15 / 0-15	-15	92	93
	Mean**	42 / 88	21 / 25	0-15 / 0-15	-1	95	95
Furniture Store	May 16	73 / 67	33 / 62	0-15 / ***	50	105	97
	July 8	80 / 59	47 / 27	0-15 / ***	17	97	96
	Aug 27	124 / 83	50 / 42	0-15 / ***	-40	97	95
	Mean**	96 / 78	48 / 30	0-15 / ***	37	103	97
Bakery	May 16	— / 129	— / 34	—— / 15-30	-7	95	90
	July 8	— / 82	— / 40	—— / 15-30	-4	104	96
	Aug 27	— / 79	— / 37	—— / 15-30	-68	105	101
	Mean**	— / 102	— / 39	—— / 15-30	4	101	97

* Calculated for the moderate / high price period, except for the Bakery which did not shed load during the moderate price period.

*Calculated over all eleven events that occurred in the summer of 2008.

**The average demand shed during the moderate price period is higher than that during the high price period so it does not make sense to calculate this.

17

1. Plotting time series electric load data is useful for understanding electricity consumption patterns and changes to those patterns, but results may be misleading if data from different time intervals are not weather-normalized.

2. Parameter plots highlight key features of electric load data and may be easier to interpret than plots of time series data.

3. A time-of-week indicator variable (as compared to time-of-day and day-of-week indicator variables) improves the accuracy of regression models of C&I facility electric load.

4. A piecewise linear and continuous outdoor air temperature dependence can be derived without the use of a change-point model (which would add complexity to the modeling algorithm) or assumptions about when structural changes occur (which could introduce inaccuracy).

5. DR shed variability is a function of unmodeled load variability, and, possibly, other factors such as occupant behavior change on DR days; changes in building operations, controls, equipment, and end-uses from event-to-event; and changes in shed capacity as a function of outdoor air temperature, occupancy, total facility load, and other variables.

6. For measurement and evaluation accuracy, DR parameter values should be estimated at the end of the summer using all available data to build the baseline model, instead of directly after an individual event using only data before the event to build the baseline model.

7. DR parameters values computed for an individual event are far less certain than values computed for an average event (e.g., the average of all events in a summer); therefore, it is good practice to report values associated with means along with values associated with individual events.

To assess and evaluate "fast DR," such as facilities providing ancillary services [6], which occur on time scales of seconds to minutes, high resolution (e.g., 5 minute and 4 second interval) electric load data are needed. Though we present methods for analyzing 15-minute-interval electric load data, the same methods could be used to analyze higher resolution data.

In addition, some of the methods presented here could be used by building managers for real-time feedback control of DR resources. For example, at specific time intervals during a DR event, a baseline model built with data from non-DR days before the DR day could be used together with actual DR day temperature data to predict load. Subtracting the load prediction from the actual load results in the real-time DR residual, which can be used for feedback control [6]. A drawback to this approach is that the baseline model is built with fewer data, all from before the DR day. An existing barrier to this approach is that, while load and temperature data are often available in real-time, many existing building control systems are unable to handle these data in real-time.

The methods presented here could be translated into easy-to-use tools for building managers, helping them determine which DR strategies to implement and which DR programs to enroll in, and then evaluating the effectiveness of the DR strategies they have implemented. These tools could be integrated into real-time operating platforms to assist building managers in all forms of building energy management. Future work should include developing methods to quantify error in DR predictions, including an analysis of shed variability.

Acknowledgment

References

[1] N. Motegi, M. Piette, D. Watson, S. Kiliccote, and P. Xu, "Introduction to commercial building control strategies and techniques for demand response," Lawrence Berkeley National Laboratory, Tech. Rep. LBNL-59975, 2007.

[2] FERC, "National action plan on demand response," Prepared with the support of The Brattle Group, GMMB, Customer Performance Group, Definitive Insights, and Eastern Research Group, Tech. Rep., Jun. 2010. [Online]. Available: http://www.ferc.gov/legal/staff-reports/06-17-10-demand-response.pdf

[3] G. Heffner, "Demand response valuation frameworks paper," Lawrence Berkeley National Laboratory, Tech. Rep. LBNL-2489E, 2009.

[4] S. Kiliccote, M. Piette, G. Wilker, J. Prijanonda, and A. Chiu, "Installation and commissioning automated demand response systems," in *Proceedings of 16th National Conference on Building Commissioning*, Newport Beach, CA, 2008.

[5] J. Eto, C. Goldman, G. Heffner, B. Kirby, J. Kueck, M. Kintner-Meyer, J. Dagle, T. Mount, W. Schultz, R. Thomas, and R. Zimmerman, "Innovative developments in load as a reliability resource," in *Proceedings of the IEEE PES Winter Meeting*, vol. 2, 2002, pp. 1002–1004.

[6] S. Kiliccote, M. Piette, G. Ghatikar, E. Koch, D. Hennage, J. Hernandez, A. Chiu, and J. Goodin, "Using open automated demand response communications in demand response for wholesale ancillary services," in *Proceedings of Grid Interop Forum*, Denver, CO, 2009.

[7] D. Callaway, "Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy," *Energy Conversion and Management*, vol. 50, pp. 1389–1400, 2009.

[8] H. Wellinghoff, D. Morenoff, J. Pederson, and M. Tighe, "Creating regulatory structures for robust demand response participation in organized wholesale electric markets," in *Proceedings of ACEEE Summer Study on Energy Efficiency in Buildings*, Pacific Grove, CA, 2008.

[9] M. Piette, G. Ghatikar, S. Kiliccote, E. Koch, D. Hennage, P. Palensky, and C. McParland, "Open automated demand response communications specification (Version 1.0)," California Energy Commission, PIER Program, Tech. Rep. CEC-500-2009-063, Apr. 2009.

[10] NOAA, "NNDC climatic data online." [Online]. Available: http://www7.ncdc.noaa.gov/CDO/dataproduct

[11] J. Haberl, J. Kissock, R. Belur, and R. Sparks, "Improving the paradigm for displaying complex building energy consumption data." in *Proceedings of ASME International Solar Energy Conference*, Washington, DC, Apr. 1993, pp. 475–485.

[12] J. Haberl, R. Sparks, and C. Culp, "Exploring new techniques for displaying complex building energy consumption data," *Energy and Buildings*, vol. 24, no. 1, pp. 27–38, 1996.

[13] J. Haberl and M. Abbas, "Development of graphical indices for viewing building energy data: Part I-II," *Journal of Solar Energy Engineering*, vol. 120, pp. 156–167, 1998.

[14] P. Price, "Methods for quantifying electric load shape and its variability," Lawrence Berkeley National Laboratory, Tech. Rep. LBNL-3713E, 2010.

[15] K. Coughlin, M. Piette, C. Goldman, and S. Kiliccote, "Statistical analysis of baseline load models for non-residential buildings," *Energy and Buildings*, vol. 41, no. 4, pp. 374–381, Apr. 2009.

[16] M. Goldberg and G. Agnew, "Protocol development for demand response calculation- findings and recommendations," California Energy Commission (KEMA-XENERGY), Tech. Rep. CEC 400-02-017F, 2003.

[17] D. Claridge, "A perspective on methods for analysis of measured energy data from commercial buildings," *Journal of Solar Energy Engineering*, vol. 120, p. 150, 1998.

[18] S. Karatasou, M. Santamouris, and V. Geros, "Modeling and predicting building's energy use with artificial neural networks: Methods and results," *Energy and Buildings*, vol. 38, no. 8, pp. 949–958, 2006.

[19] J. Granderson, M. Piette, G. Ghatikar, and P. Price, "Building energy information systems: State of the technology and user case studies," Lawrence Berkeley National Laboratory, Tech. Rep. LBNL-2899E, 2009.

[20] J. Taylor, L. deMenezes, and P. McSharry, "A comparison of univariate methods for forecasting electricity demand up to a day ahead," *International Journal of Forecasting*, vol. 22, no. 1, p. 116, 2006.

[21] D. Park, M. El-Sharkawi, R. Marks, L. Atlas, and M. Damborg, "Electric load forecasting using an artificial neural network," *IEEE Transactions on Power Systems*, vol. 6, no. 2, pp. 442–449, 1991.

[22] J. Kreider and J. Haberl, "Predicting hourly building energy use: The great energy predictor Shootout-Overview and discussion of results," *ASHRAE Transactions*, vol. 100, no. 2, pp. 1104–1118, 1994.

[23] J. Haberl and S. Thamilseran, "The great energy predictor shootout. II: measuring retrofit savings," *ASHRAE Journal*, vol. 40, no. 1, pp. 49–56, 1996.

[24] R. Ramanathan, R. Engle, C. W. Granger, F. Vahid-Araghi, and C. Brace, "Short-run forecasts of electricity loads and peaks," *International Journal of Forecasting*, vol. 13, no. 2, pp. 161–174, 1997.

[25] M. Fels, "PRISM: an introduction," *Energy and Buildings*, vol. 9, no. 1-2, p. 518, 1986.

[26] J. Kissock, T. Reddy, and D. Claridge, "Ambient-temperature regression analysis for estimating retrofit savings in commercial buildings," *Journal of Solar Energy Engineering*, vol. 120, p. 168, 1998.

[27] S. Katipamula, T. Reddy, and D. Claridge, "Multivariate regression modeling," *Journal of Solar Energy Engineering*, vol. 120, p. 177, 1998.

[28] S. Katipamula, "Great energy predictor shootout II: modeling energy use in large commercial buildings," *ASHRAE Transactions*, vol. 102, pp. 397–404, 1996.

[29] J. Kissock and C. Eger, "Measuring industrial energy savings," *Applied Energy*, vol. 85, no. 5, pp. 347–361, 2008.

[30] H. Akbari, "Validation of an algorithm to disaggregate whole-building hourly electrical load into end uses," *Energy*, vol. 20, no. 12, pp. 1291–1301, 1995.

[31] L. Pedersen, J. Stang, and R. Ulseth, "Load prediction method for heat and electricity demand in buildings for the purpose of planning for mixed energy distribution systems," *Energy and Buildings*, vol. 40, no. 7, pp. 1124–1134, 2008.

[32] T. Reddy, J. Kissock, and D. Ruch, "Uncertainty in baseline regression modeling and in determination of retrofit savings," *Journal of solar energy engineering*, vol. 120, p. 185, 1998.

[33] J. Mathieu, A. Gadgil, D. Callaway, P. Price, and S. Kiliccote, "Characterizing the response of commercial and industrial facilities to dynamic pricing signals from the utility," in *Proceedings of ASME 2010 4th International Conference on Energy Sustainability*, Phoenix, AZ, May 2010.

Biographies

Johanna L. Mathieu received a B.S. in Ocean Engineering from MIT in 2004, and an M.S. in Mechanical Engineering from UC Berkeley in 2008. She is currently a PhD Candidate in Mechanical Engineering at UC Berkeley. Her research involves modeling buildings in order to quantify the effectiveness of demand response and modeling/controlling aggregated systems of appliances to support the integration of intermittent renewable energy resources.

Phillip N. Price has B.A.'s in physics and mathematics from Oberlin College, and a PhD in physics from the University of Kentucky. He has worked at Lawrence Berkeley National Laboratory since 1992, and has a wide variety of professional interests, including model-measurement comparison for complicated systems. He was elected a Fellow of the American Physical Society in 2003. He is also active in the environmental movement, with a special interest in protecting endangered species and sensitive habitat.

Sila Kiliccote is the Deputy Group Leader of the PIER Demand Response Research Center and a Program Manager in the Building Technologies Department at Lawrence Berkeley National Laboratory. Her areas of interest include characterization of building loads and demand reduction, demand responsive lighting systems, building systems integration and feedback for demand-side management. She has an M.S. in Building Science from Carnegie Mellon University and a B.S. in Electrical Engineering from the University of New Hampshire.

Mary Ann Piette is the Deputy of the Building Technologies Department and the Director of the PIER Demand Response Research Center at the Lawrence Berkeley National Lab. She is a member of the NIST Smart Grid Architecture Committee. She has received several awards related to programs to automate demand response, and she received the Benner Award for contributions to making building commissioning "business as usual." Ms. Piette has a B.S. in Physical Science and an M.S. in Mechanical Engineering from UC Berkeley, and a Licentiate in Building Services Engineering from the Chalmers University of Technology in Sweden.