

Estimating Reduced Consumption for Dynamic Demand Response

Charalampos Chelmis, Saima Aman, Muhammad Rizwan Saeed, Marc Frincu, Viktor K. Prasanna

University of Southern California, Los Angeles, CA

Email: {chelmis, saman, saeedm, frincu, prasanna}@usc.edu

Abstract

Growing demand is straining our existing electricity generation facilities and requires active participation of the utility and the consumers to achieve energy sustainability. One of the most effective and widely used ways to achieve this goal in the smart grid is demand response (DR), whereby consumers reduce their electricity consumption in response to a request sent from the utility whenever it anticipates a peak in demand. To successfully plan and implement demand response, the utility requires reliable estimate of reduced consumption during DR. This also helps in optimal selection of consumers and curtailment strategies during DR. While much work has been done on predicting normal consumption, reduced consumption prediction is an open problem that is under-studied. In this paper, we introduce and formalize the problem of *reduced consumption prediction*, and discuss the challenges associated with it. We also describe computational methods that use historical DR data as well as pre-DR conditions to make such predictions. Our experiments are conducted in the real-world setting of a university campus microgrid, and our preliminary results set the foundation for more detailed modeling.

Introduction

With the rapid integration of advanced metering infrastructure, Smart Grids enable real-time implementation of dynamic demand-side control. Demand Response (DR) is a key load management technique which provides a cost-effective alternative to traditional supply-side solutions, meant to address demand increase during times of peak electrical load. Demand Response offers several advantages: it prevents blackouts, reduces the need for new generation units, and increases overall reliability of the electricity grid. Accurate estimation and evaluation of consumption reduction achieved by participants during curtailment is therefore critical to DR programs.

Figure 1 provides a conceptual diagram for consumption reduction and curtailment calculation during DR. A naive approach to determine the extent of curtailment during DR is to estimate what the consumption would have been in

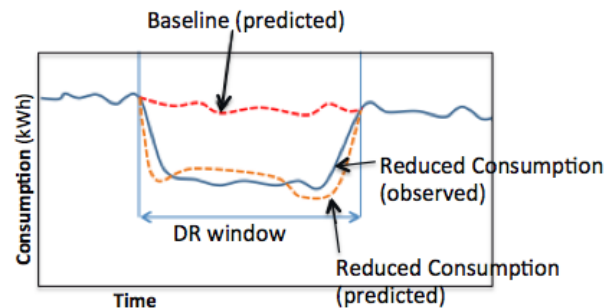


Figure 1: Conceptual diagram of *reduced consumption prediction* during Demand Response

the absence of DR (denoted as Baseline (predicted) in Figure 1), and compare such counterfactual prediction with the observed electric consumption during DR (denoted as Reduced Consumption (observed) in Figure 1). The amount of computed curtailment (i.e., area between observed Reduced Consumption and Baseline) depends on the accuracy of the baseline model used. Utilities have so far focused on developing more accurate baseline models. Different baseline models result in different curtailment estimates. In fact, selecting a reasonable baseline is non-trivial and may lead to misinterpretation of curtailment estimates. The task is further complicated both by the dynamic nature of DR, and its dependency on multiple factors, such as occupancy, ambient temperature, and weather conditions. More importantly, being able to estimate the reduced consumption during the DR window in advance can be beneficial for planning purposes, enabling utilities to (i) adapt to low generation (particularly with the introduction of unreliable renewable generation sources), (ii) intelligently target consumers for DR participation, (iii) select appropriate curtailment strategies¹, and

¹The USC campus data used in our study includes a variety of Fully-Automated Demand Response strategies (Piette, Kiliccote, and Dudley 2012): Global Zone Temperature Reset (GTR) (Motegi et al. 2006), Variable Frequency Drive Speed Reset (VFD) (Motegi et al. 2006), Equipment Duty Cycling (Duty), and their combinations. Such strategies directly reduce the heating, ventilation, and cooling (HVAC) loads, which make up a significant portion of the overall energy consumption of buildings.

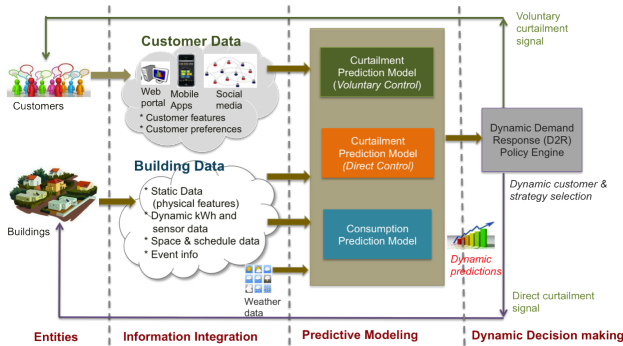


Figure 2: D2R Framework

(iv) maximize profit by avoiding expensive energy sources during peak hours.

In this paper, we formalize the open problem of *reduced consumption prediction* during DR and describe the computational methods we use to address this problem. Statistical, machine learning, and data mining techniques that work well for consumption prediction, e.g., auto-regression models, are of little use for reduced consumption prediction. For example, generic time series models such as ARIMA (Box and Jenkins 1970) postulate that short-term future consumption values can be estimated by consumption patterns in the immediate past. However, due to the abrupt change in consumption profile during DR, time series models, which are unable to capture sudden changes at the temporal boundaries of DR events, are deemed inappropriate.

To address this problem, we propose a method to estimate directly curtailed consumption during DR based on historical observations from past DR events. Specifically, we introduce a novel averaging model for direct curtailment prediction. We show that, despite its simplicity, our model a) provides a good fit to curtailed consumption during DR for customers who exhibit low consumption variability, b) its results are easy to interpret, and c) it presents negligible computational burden. Our experiments are conducted as part of the Dynamic Demand Response (D^2R) framework implemented at the University of Southern California campus in Los Angeles (Figure 2). Estimates of reduced consumption are used by the D^2R policy engine for optimal customer selection during the planning phase of future DR events.

Background and Previous Work

Research in automated demand response has received increased attention motivated by the need to reduce black outs and to enable dynamic electricity pricing. It has spanned areas such as communication and building control strategies for curtailment (Piette, Kiliccote, and Dudley 2012), motivating customers for participation in DR programs (Akasiadis and Chalkiadakis 2013), as well as for developing baseline models that make counterfactual predictions for baseline load on DR days (Mathieu, Callaway, and Kiliccote 2011). The latter are used to measure electricity consumption reduction during DR. For baseline models, utilities gen-

erally use simple averaging of time of day values from recent or similar days. Often, such predictions are multiplied by a morning adjustment factor to adjust for weather and other conditions on DR days. Regression (Mathieu, Callaway, and Kiliccote 2011) and time series methods have also been used as baselines. A comparison of various baseline models is presented in (Coughlin et al. 2009), where it is shown that baseline prediction accuracy is dependent on load variability and weather sensitivity.

Our focus is on prediction of reduced consumption during DR, which has received little attention so far. Due to limited data availability, small number of DR events per customer, and diversity of customer types (e.g., residential versus office buildings), reduced consumption estimation is a challenging task. Existing approaches assume that consumers who enroll in curtailment programs would always comply when asked to curtail (Lou et al. 2013; Simao et al. 2013). Our experiments suggest that this assumption does not always hold. For example when curtailment strategies, such as global temperature reset (GTR) or HVAC duty cycling are applied, they cannot violate occupants' thermal comfort limits, and hence may be aborted or modified midway during the DR window. This results in variable reduced consumption over DR events, even for the same customer. Our approach does account for variability through factors such as time of day, strategy used, and outside air temperature. More importantly, our focus is short-term reduced consumption forecasting during DR, which is essential for dynamic (near real-time) adaptation of customer selection programs to maintain or achieve a desired reduction in electricity demand in a dense urban area.

Problem Formulation

Our goal is to predict reduced consumption during DR using: 1) historic reduced consumption data from past DR experiments, and, ii) weather forecast data. We focus on 15-minute granularity consumption prediction for the DR window just before the start of DR. Such predictions are useful in selecting strategies and consumers to target for DR.

Let us denote the observed *reduced* consumption for j^{th} 15-min interval of i^{th} day by $r_{i,j}$. Let I be the number of DR days for a building and J be the total number of intervals per day for which DR is carried out. The observed reduced consumption for all the time points can be expressed as an $I \times J$ matrix, denoted as $R = [r_{i,j}]_{I \times J}$. Similarly, let K be the total number of intervals per day before DR begins. We denote consumption for k^{th} 15-min interval before DR on i^{th} DR day by $n_{i,k}$. The observed consumption for all time points before DR can be expressed as an $I \times K$ matrix, $N = [n_{i,k}]_{I \times K}$. We associate a vector of weather conditions with each consumption entry before DR. Particularly, for L unique weather attributes (such as temperature, humidity, etc.), each time point is associated with vector $\omega_{i,k} = \langle \omega_{i,k}[1], \omega_{i,k}[2], \dots, \omega_{i,k}[L] \rangle$, where $\omega_{i,k}[l]$ denotes the value for the l -th weather condition for time-slot (i, k) . Observed weather conditions for all time points before DR can be expressed as matrix $\omega = [\omega_{i,k}]_{I \times K}$. Our task is to estimate reduced consumption during DR given historical reduced consumption matrix R , consumption matrix N ,

and weather attributes matrix Ω .

Reduced Consumption Prediction

For normal consumption, observations made in the immediate past are usually a good indication of short-term future. However, for reduced consumption due to DR, this is not true as there is a sudden drop in the time series. In such a scenario, the historical observations (for the same curtailment strategy) are better predictors of future. We describe two ways in which historical patterns can be incorporated in prediction models.

Historical Average Model. We introduce historical average model that uses previous values for the same time on *similar* DR days for the same \langle building, strategy \rangle combination to forecast future reduced consumption. We consider past DR events per building to be similar if the same strategy was deployed. This results in 4-10 events per \langle building, strategy \rangle combination. When only a small number of past DR events is available per building, the average can be taken over all DR days. In this model, the prediction outputs are adjusted to account for indirect factors, such as weather conditions, on DR days (Coughlin et al. 2008). Particularly, we use a multiplicative factor defined as the ratio of average kWh usage of the first three of the four hours before the event to the average kWh usage for the same three hours from the average of past similar days. To adjust our forecast, we multiply our predicted curtailed consumption by the adjustment factor for each 15-min interval during DR.

Weighted Average Model. In the historical average model, we defined similarity in terms of weather conditions. Here, we consider two notions of similarity: (i) with respect to time, and (ii) Euclidean distance. We apply an exponential degradation function of a DR’s age in the computation of the average.

- *WtdAvTi*: Weights are selected to exponentially decrease with time according to recency, i.e., the time elapsed between the future DR event and historical DR events. Thus, observations closer to the DR day for which predictions are to be made, are given higher weights than those farther away from it. Intuitively, a building’s response to DR for a given strategy drifts with time, i.e., due to for example changing weather conditions or building characteristics such as better insulation, energy efficient lighting etc. Thus, events closer to each other are more likely to be similar.
- *WtdAvSi*: Weights are selected to be exponentially decreasing with decreasing similarity of pre-DR condition. Similarity between pre-DR condition on two DR days i and j is calculated as the Euclidean distance between the vectors N_i and ω_j . Intuitively, a building’s response to DR for a given strategy depends on *similar* conditions, which we capture in the form of indirect indicators N and Ω .

As in the case of historical average model, here too, we adjust our estimates by a *morning adjustment* factor.

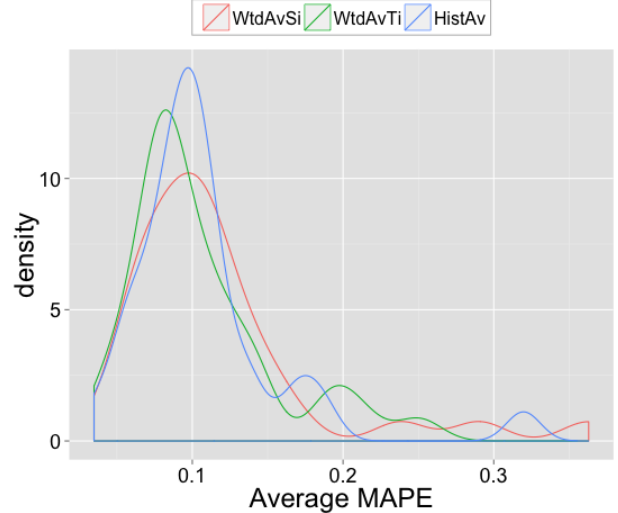


Figure 3: Density function of average MAPE for all buildings

Experiments

Dataset. We used data from DR experiments carried out at the University of Southern California campus as part of the DR demonstration project (Simmhan et al. 2011). Reduced and normal consumption data (in kWh) in 15-min intervals was collected from 35 buildings between Nov 2012 to Dec 2013. We focus on electricity consumption, measured in kWh over an interval of time. The dataset contains a diverse set of building types: academic buildings with teaching and office space, residential dormitories, and administrative buildings. The number of DR events across buildings is not homogeneous. Some buildings participated in more than 40 events, while others were rarely selected (less than 10 events). This results in 826 individual \langle building, strategy \rangle events overall. The choice of strategy is also heterogeneous. Building names have been obfuscated for privacy issues. Experiments were conducted while school was in session, allowing building responses to each strategy to be characterized during standard operation. Due to climate particularities DR events in the microgrid were conducted during the 1:00-5:00 PM time frame when demand peaks and temperature is high. In addition, hourly weather data was also collected from NOAA’s weather station at USC campus. Weather measurements include temperature and humidity, which were linearly interpolated to 15-min intervals.

Evaluation. We use leave one out cross validation for evaluating prediction performance. We compare model performance using the MAPE measure, which is defined as $\sum_i \frac{|O_{i,j} - P_{i,j}|}{O_{i,j}}$, where O is used to denote observed output and P is the predicted output.

Results. Figure 3 shows the density function of average MAPE errors for all buildings for all three averaging models. We observe that for majority of the buildings, the prediction

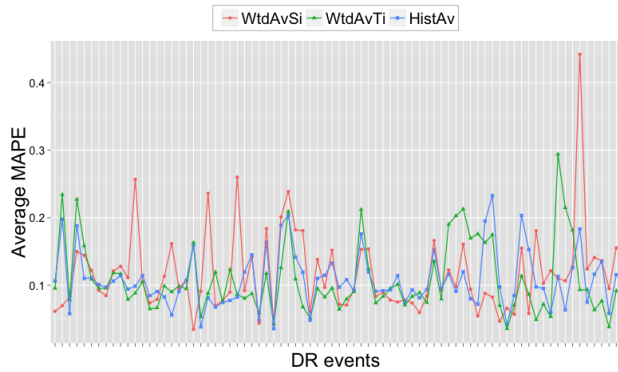


Figure 4: Average MAPE for all 78 DR event days between Nov 2012 to Dec 2013

errors is low (less than 15%), and in about half the cases, it is within 10%, indicating that averaging models provide a good estimate of reduced consumption during DR. The historical average model seems to perform slightly better than the other two models. Although simple, it derives its predictive power from errors being averaged out over the entire dataset. According to (Aman, Simmhan, and Prasanna 2014), acceptable threshold of prediction error for DR in individual buildings is 10%. Hence, those buildings whose performance does not meet this threshold require further investigation and more powerful models.

Figure 4 shows the performance of our models for individual DR events in our dataset. We found historical average model to perform best in 248 events. *WtdAvTi* was the best in 295 events, and *WtdAvSi* was the best in 283 events. These results suggest that an ensemble, i.e., a meta-model that would learn to select the best model for prediction for a given DR event, would perform favorably. We further plan to explore more advanced models in future work.

Discussion

To the best of our knowledge, we are the first to address the task of identifying and formally defining the problem of *reduced consumption prediction* during demand response window. While much work has been done on consumption prediction outside DR and on baseline prediction during DR, reduced consumption prediction has received little attention. However, accurate estimation of reduced consumption ahead of time can be beneficial in successfully planning and sustaining DR programs in Smart Grid. Reliable estimates of reduced consumption help the utility in: (i) deciding the duration of DR, (ii) intelligent selection of curtailment strategies, and (iii) targeting consumers for DR participation. We argue that traditional prediction approaches for consumption and baseline forecasting are inappropriate in this context. Our proposed strategies achieve good prediction accuracy by incorporating historical DR data as well as pre-DR condition data. Additionally, our proposed strategies are computationally inexpensive, making them favorable for scenarios with near real-time constraints, such as Dynamic

Demand Response (D^2R).

Acknowledgments

This material is based upon work supported by the United States Department of Energy under Award Number DEOE0000192, and the Los Angeles Department of Water and Power (LA DWP). The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof, the LA DWP, nor any of their employees.

References

- Akasiadis, C., and Chalkiadakis, G. 2013. Agent cooperatives for effective power consumption shifting. In *AAAI Conference on Artificial Intelligence*.
- Aman, S.; Simmhan, Y.; and Prasanna, V. 2014. Holistic measures for evaluating prediction models in smart grids. *IEEE Transactions in Knowledge and Data Engineering* PP(99).
- Box, G. E. P., and Jenkins, G. M. 1970. *Time series analysis, forecasting and control*. Holden-Day.
- Coughlin, K.; Piette, M. A.; Goldman, C.; and Kiliccote, S. 2008. Estimating demand response load impacts: Evaluation of baseline load models for non-residential buildings in california. *Lawrence Berkeley National Lab Technical Report, LBNL-63728*.
- Coughlin, K.; Piette, M. A.; Goldman, C.; and Kiliccote, S. 2009. Statistical analysis of baseline load models for non-residential buildings. *Energy and Buildings* 41(4).
- Lou, X.; Yau, D.; Nguyen, H. H.; and Chen, B. 2013. Profit-optimal and stability-aware load curtailment in smart grids. *IEEE Transactions on Smart Grid* 4(3).
- Mathieu, J. L.; Callaway, D. S.; and Kiliccote, S. 2011. Variability in Automated Responses of Commercial Buildings and Industrial Facilities to Dynamic Electricity Prices. *Energy and Buildings*.
- Motegi, N.; Piette, M. A.; Watson, D. S.; Kiliccote, S.; and Xu, P. 2006. Introduction to commercial building control strategies and techniques for demand response. Technical report, Lawrence Berkely National Laboratory.
- Piette, M. A.; Kiliccote, S.; and Dudley, J. H. 2012. Field Demonstration of Automated Demand Response for Both Winter and Summer Events in Large Buildings in the Pacific Northwest. *Lawrence Berkeley National Lab Technical Report, LBNL-6216E*.
- Simao, H.; Jeong, H.; Defourny, B.; Powell, W.; Boulanger, A.; Gagneja, A.; Wu, L.; and Anderson, R. 2013. A robust solution to the load curtailment problem. *IEEE Transactions on Smart Grid* 4(4).
- Simmhan, Y.; Prasanna, V.; Aman, S.; Natarajan, S.; Yin, W.; and Zhou, Q. 2011. Toward data-driven demand-response optimization in a campus microgrid. In *Proceedings of the Third ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, 41–42. ACM.