Implementation of Learning-Based Dynamic Demand Response on a Campus Micro-Grid*

Sanmukh R. Kuppannagari, Rajgopal Kannan, Charalampos Chelmis and Viktor K. Prasanna
Ming Hsieh Department of Electrical Engineering
University of Southern California, Los Angeles, California
{kuppanna,rajgopalk,chelmis,prasanna}@usc.edu

Abstract

Demand Response (DR) allows utilities to curtail electricity consumption during peak demand periods. Real time automated DR can offer utilities a scalable solution for fine grained control of curtailment over small intervals for the duration of the entire DR event. In this work, we demonstrate a system for a real time automated Dynamic DR (D^2R). Our system has already been integrated with the electrical infrastructure of the University of Southern California, which offers a unique environment to study the impact of automated DR in a complex social and cultural environment including 170 buildings in a "city-within-a-city" scenario. Our large scale information processing system coupled with accurate forecasting models for sparse data and fast polynomial time optimization algorithms for curtailment maximization provide the ability to adapt and respond to changing curtailment requirements in near real-time. Our D^2R algorithms automatically and dynamically select customers for load curtailment to guarantee the achievement of a curtailment target over a given DR interval.

1 Introduction

The reliable operation of a power grid requires the constant matching between fluctuating load and supply. Demand Response (DR) is a technique whereby customers are asked to curtail their demands, typically during peak demand periods - when the grid exhibits very high demands potentially surpassing the generation capacity. Given the availability of bi-directional smart meters, we have developed a system for performing DR optimization in real time which we call Dynamic Demand Response (D^2R).

We have implemented D^2R on our campus micro-grid to demonstrate its large scale feasibility and identify and resolve the challenges associated with practical deployment. Our D^2R technique uses learning of occupant energy strategy preferences (at fine grained scales ranging from buildings to floor levels within buildings) to make accurate electricity consumption predictions and individual curtailment recommendations using only a small subset of consumption data.

*This work has been funded by the U. S. Department of Energy under Award Number DE-OE0000192, the Los Angeles Department of Water and Power (LADWP) and the U.S. National Science Foundation under grant number ACI 1339756.

2 D^2R Implementation

Buildings on our campus are instrumented with smart metering and control that can be used to implement a large number of advanced energy curtailment strategies such as resetting temperature set points, reducing air flow, duty cycling. A simplified control and data flow diagram of our D^2R implementation is shown in Figure 1. The micro-grid utility initiates a DR event using OpenADR messages and provides a curtailment target γ to be achieved over a given DR interval T, typically 4 hours. The Policy Engine (PE) module provides campus wide curtailment strategy policy recommendations based on the analysis of historical consumption data and curtailment maximization customer (building) selection algorithms. Smart meter data is aggregated over 15 minute intervals into a consumption database. State-of-the-art data-driven models are then used by the PE module to predict energy consumption values over each 15 minute period comprising the entire DR interval T for each building across campus. This information is then provided to the optimization module in the form of a discrete time varying curtailment matrix. The outputs of the PE are sets of buildings-strategy pairs at 15 minute intervals that will achieve the required curtailment target γ over T.

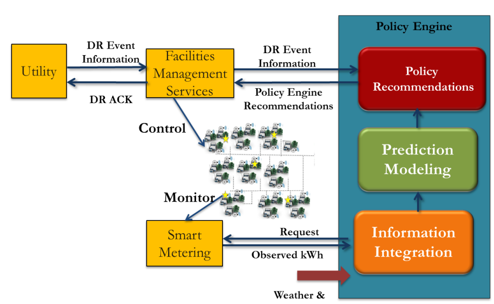

Figure 1: Control and Data Flow For D^2R Implementation

Our Policy Engine (PE) is composed of an influence model based demand prediction engine that feeds into a building-strategy selection optimization module as shown in Figure 2.

Historical data from the energy consumption database before the DR day is used in conjuction with time series forecasting techniques such as ARIMA and Lasso-Granger (for learning temporal dependencies among multiple timeseries) to learn occupant energy strategy preferences and load profiles. A significant challenge is to ensure accurate prediction in the absence of high quality consumption data. To ensure quality data we have developed several techniques: interpolation methods for estimating intermittent missing data and sophisticated influence based learning models for estimating systematically unavailable data over larger periods. We have shown that only a small subset ($\approx 7\%$) of the meters are required in real-time to make predictions for buildings across the campus micro-grid [Aman *et al.*, 2015].

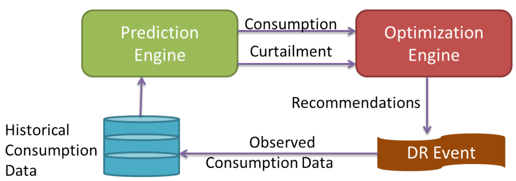

Figure 2: Policy Engine

The results of our adaptive customer prediction models are fed into our optimization module which consists of fast approximation algorithms for ILP based fair energy curtailment maximization [Kuppannagari *et al.*, 2016]. Given a targeted curtailment value γ over DR interval T, curtailment matrix M, the optimization module selects sets of building-strategy pairs for each 15 minute interval to achieve the curtailment value γ over T. Note that optimizing for γ over interval T without considering per-period curtailment values might lead to aggressive curtailment in some periods and low curtailment in others. This could be unsustainable to the utility as demands in periods with low curtailment may exceed the generation capacity. Therefore, we have defined and implemented the notion of Sustainable DR (SDR) in which the targeted curtailment value γ is distributed proportionally (as $\{\gamma_t\}$) across interval T [Kuppannagari *et al.*, 2016]. Additionally, we have included the notion of fairness (via building curtailment budgets) and strategy switching overhead as constraints in our optimization algorithm. We have shown that this problem is NP-hard and have developed fast polynomial time approximation schemes (PTAS) as well as bounded randomized rounding heuristics with provable error bounds i.e. deviation from the curtailment target γ. We have shown that our SDR algorithms achieve results with a very low absolute error of 0.001-0.05 kWh range [Kuppannagari *et al.*, 2016].

3 Implementation Challenges

Real time quality data availability is a key challenge that we have addressed earlier (Section 2). Real world challenges such as changing environmental conditions may affect our prediction accuracy, buildings may have to be dropped out of DR due to unresponsiveness or thermal comfort violation. We address such cases by performing dynamic customer reselection to offset deviations in the curtailment target. Another challenge is scalability while maintaining curtailment accuracy which we have addressed by developing fast polynomial time approximation algorithms with bounded errors. Finally, extending our implementation to a large city environment will face human behavioral challenges. Determining the right customer incentives to obtain a reasonable compliance rate is an open challenge that needs to be addressed.

4 Innovations, Benefits and Importance

Traditionally, DR programs communicate the schedule to customers much ahead of time and are very inefficient due to inaccuracies over long prediction horizons. We have addressed these inefficiencies using our innovative real-time D^2R techniques which can be dynamically adjusted to manage supply/demand imbalances over short timescales (15 minutes). Using D^2R, we have achieved curtailment values as high as 1.2 MW in a single DR event in our campus micro-grid.

A key innovation of our D^2R technique is the development of influence based machine learning techniques to predict consumption even with partial or missing data. Another key innovation is the development of fast polynomial time optimization algorithms for customer selection which provide near optimal results.

By addressing the various challenges, we have implemented a DR program with very low curtailment errors (due to prediction inaccuracies and errors due to the approximation algorithms for SDR) which are $\leq 10\%$ as opposed to 30% in other techniques [Kwac and Rajagopal, 2014]. Hence, our state-of-the-art implementation can be used as a model to implement city scale efficient automated DR programs as is being currently undertaken in our city.

5 Details of the Demo

We will demonstrate a delayed implementation of a live DR event that occurred earlier. The consumption values until 12 noon of the DR day will be used by our prediction engine. A simulated curtailment request with a timestamp of 12 noon will be sent to the PE which is implemented as a webserver. We will show the consumption predictions and strategies recommended by the PE over 15 minute intervals. Finally, we will evaluate the success of our DR algorithm by measuring the curtailment error.

References

[Aman *et al.*, 2015] Saima Aman, Charalampos Chelmis, and Viktor K Prasanna. Influence-driven model for time series prediction from partial observations. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.

[Kuppannagari *et al.*, 2016] Sanmukh R. Kuppannagari, Rajgopal Kannan, Charalampos Chelmis, Arash S Tehrani, and Viktor K Prasanna. Optimal customer targeting for sustainable demand response in smart grids. In *The International Conference on Computational Science*, 2016.

[Kwac and Rajagopal, 2014] Jungsuk Kwac and Ram Rajagopal. Targeting customers for demand response based on big data. *arXiv preprint arXiv:1409.4119*, 2014.