Sustainable Energy, Grids and Networks 10 (2017) 75–83

Contents lists available at ScienceDirect

Sustainable Energy, Grids and Networks

journal homepage: www.elsevier.com/locate/segan

Demand response evaluation and forecasting — Methods and results from the EcoGrid EU experiment[☆]

Emil Mahler Larsen [a,*], Pierre Pinson [b], Fabian Leimgruber [c], Florian Judex [c]

[a] Danish Energy Association, Denmark
[b] Department of Electrical Engineering, Technical University of Denmark, Denmark
[c] Energy Department, Austrian Institute of Technology, Austria

ARTICLE INFO

Article history:
Received 28 September 2016
Available online 18 March 2017

Keywords:
Demand response (DR)
Real-time pricing
Demand forecasting
Smart grid

ABSTRACT

Understanding electricity consumers participating in new demand response schemes is important for investment decisions and the design and operation of electricity markets. Important metrics include peak response, time to peak response, energy delivered, ramping, and how the response changes with respect to external conditions. Such characteristics dictate the services DR is capable of offering, like primary frequency reserves, peak load shaving, and system balancing. In this paper, we develop methods to characterise price-responsive demand from the EcoGrid EU demonstration in a way that was bid into a real-time market. EcoGrid EU is a smart grid experiment with 1900 residential customers who are equipped with smart meters and automated devices reacting to five-minute electricity pricing. Customers are grouped and analysed according to the manufacturer that controlled devices. A number of advanced statistical models are used to show significant flexibility in the load, peaking at 27% for the best performing groups.

1. Introduction

Interest in Demand Response (DR) has grown in recent years as system operators look for new tools to meet the needs of a rapidly changing power system. Changes include increased production from renewables, tighter market coupling, and a surge in decentralised production and consumption from photovoltaics (PV) and electric vehicles. The changing needs of the power system can broadly benefit from DR in two ways: through emergency use, where a reliable reduction in demand is needed during infrequent critical periods, and through economic use, where demand exhibits continuous flexibility to bring down average costs in the power system.

There are many ways of changing consumption patterns, but dynamic tariffs in particular are gaining interest due to their potential to respond quickly to fluctuating production from renewable energy sources (RES) [1]. Indirect control is one such dynamic tariff that uses an incentive signal, e.g. a real-time price, to influence the load. Indirect control does not require an exact response from any one customer, but with a large number of loads that exhibit somewhat similar behaviour, a statistically likely response can be forecast [2]. The value of indirect control therefore depends heavily on being able to accurately foresee its response to the incentive signal, which has previously been proven complicated [3].

The challenge of determining how much DR there is in a load has previously been done using baseline profiling [4–7]. Baseline profiling requires a prediction of the load under non-DR conditions which is then subtracted from the observed consumption under a DR event. A key drawback of existing methods is the need for data from non-DR days, data for which will not always exist, or may be unreliable since new equipment and interaction with customers can make non-DR data unrepresentative. Existing baseline methods are also unsuitable for evaluating fast moving DR that is conditional on a wide range of historical prices and price forecasts. Existing methods typically look a load curtailment, while we consider both increases and decreases in consumption due to decreases and increases in real-time pricing. Finally, existing methods may be susceptible to overfitting, and may rely on just one or two dozen observations per parameter.

Demand forecasting literature is a well developed area that is useful in predicting the price-elasticity of a load, e.g. see [8,9], but

☆ This work was partly supported by the European Commission through the project EcoGrid EU (grant ENER/FP7/268199).

* Corresponding author.
E-mail address: eml@danskenergi.dk (E.M. Larsen).

http://dx.doi.org/10.1016/j.segan.2017.03.001

many modern approaches involve black box schemes like artificial neural networks (ANN), that obscure our understanding of the dynamics. Therefore, to get a full understanding of the controllable resources, we disaggregate the load into its constituent parts. Flexibility can then be interpreted in a useful way, so that it can be exploited for use in different services, or bid into a market. There are no existing methods that are appropriate for evaluating and integrating residential DR into a balancing market, in particular when data for non-DR days is unavailable, which was the task we had to achieve in the large scale smart grid experiment, EcoGrid EU.

We cultivate modelling approaches that determine how much DR a load is capable of delivering in terms of power and energy, and under what external conditions, e.g. ambient temperature. More specifically, DR is characterised in terms of peak response, time to peak response, energy delivered and ramping. Our primary motivation was to compare the performance of groups of houses with different hardware and software that receive real-time pricing. The tools developed were used to give feedback to hardware and software manufacturers so that their algorithms could be improved. Our second objective was to apply these attributes in the constraints of a balancing market so that the load could be controlled. With a balancing market scheduling DR, we sought to validate our method by comparing the observed response to the scheduled DR. Aside from the approach, an additional contribution are the state-of-the-art estimates of residential flexibility used in a five minute balancing market.

The paper is structured as follows: Section 2 investigates existing approaches for evaluating the success of DR, as well as previous work that was relevant when developing our own methodology. Section 3 introduces the experimental setup and the data gathered from the demonstration. Section 4 describes the models developed to analyse the DR activated during the demonstration. Section 5 presents DR analysis for different groups and results from real-time forecasting in the demonstration. Section 6 discusses uncertainty, future work and concludes.

2. Demand modelling

There are several lines or research that are relevant when assessing a DR program, including previous experimental studies, load forecasting research, and energy disaggregation research. Previous experimental analyses have looked at different types of DR, like critical-peak pricing (CPP) and time-of-use (TOU) pricing, and often consider human demographics and behaviour as an impact on DR. Forecasting literature has a widespread use in operation of power systems and offers a deeper insight into the statistical tools available, with a greater focus on weather conditions, calendar effects and economic variables. Energy disaggregation research is an up and coming area driven by new sources of data, like high resolution smart meter data from thousands or millions of customers.

2.1. Previous DR studies

The study of residential loads responding to prices goes all the way back to the 1970s and many fundamental aspects, like accounting for the time of day and ambient temperature in a statistical model, remain in use today. Studies have also included home type, size, income and smart thermostat ownership as model inputs [10,11], but residential DR studies have not been able to give concrete numbers in terms of power and energy the load is capable of delivering.

For medium and large commercial and industrial loads, baseline methods are widely used to determine financial settlement for participating customers. The baseline is simply the prediction of consumption under the assumption that no DR was present. The baseline is then subtracted from the observed consumption to determine the amount of load shed the customer was able to deliver during a critical period. Baseline models are created by regressing on historical data before DR events. This has been done with hourly interval data [4] and 15 min interval data [6]. In the latter case, it was observed that including parameters for load shed directly into the model did not give a reliable result, possibly because the model was too primitive or due to over parametrisation. Another approach is to average consumption for just 5–10 days before a DR event [5], yet such few observations may mean that this approach lacks robustness.

2.2. Forecasting approaches

Short-term load forecasting presents a number of useful tools that can predict how load changes with respect to price. Classical approaches to solving hour-ahead and day-ahead load forecasting problems include time-series methods like auto-regressive integrated moving average (ARIMA) models and exponential smoothing, also including geographical factors [12] and seasonal variations [13].

Recent advances in forecasting methods include spline-based methods [9], which avoid over-parametrisation by relying on a handful of splines to describe the baseload, although authors in [9] noted that some fidelity was lost during peak load periods. This work was applied to a price responsive load in New York, with parameters for price and, in theory, these parameters should allow a full evaluation of the DR volume, although this was not explored in practice.

Other modern advances in forecasting include multivariate state-space models [14], which feature time-varying regression coefficients that may be useful for analysing DR. Semi-parametric methods to predict the contribution of load from some non-linear variables [15] may also be useful for DR volume evaluation, although [15] did not apply the methodology to a price-responsive load.

Machine learning approaches like artificial neural networks (ANN) are also popular for forecasting, with positive results reported in [16]. The benefit from ANN includes being able to capture unspecified non-linear relationships between external variables like weather. It is likely that such an approach becomes increasingly valuable as demand becomes more non-linear and volatile with new external incentives like price and the growth of distributed energy resources (DERs). ANNs have, however, been criticised for leading to over-parametrised models [17] and do not necessarily outperform linear regression models [18]. From a DR evaluation perspective, ANN's black box form makes picking out price influences complex, especially when bidding a price response into an electricity market.

2.3. Energy disaggregation

Energy disaggregation has gained interest as automatic metering infrastructure becomes ubiquitous in many countries. Energy disaggregation tools can be used to see beyond the meter and uncover which devices are turned on despite only seeing a noisy, aggregated view of the load. The stated goal of disaggregation is to better understand the load and make well-targeted energy efficiency plans.

Of particular relevance to our study is the success in [19] of detecting air-conditioning use from 1-min interval smart meter data. Methods that rely on a dictionary of devices describing the real and reactive power each consumes have also previously been developed [20]. Grey-box, Markovian stochastic, Bayesian and logistic adoption models are also promising ways of identifying human behaviour and price-responsive devices in metering data [21]. However, disaggregation techniques are not conclusively proven with external influences or with variable speed devices.