Enabling Automated Dynamic Demand Response: From Theory to Practice

Marc Frincu
University of Southern California
Dept. of Electrical Engineering
frincu@usc.edu

Charalampos Chelmis
University of Southern California
Dept. of Electrical Engineering
chelmis@usc.edu

Saima Aman
University of Southern California
Dept. of Computer Science
saman@usc.edu

Rizwan Saeed
University of Southern California
Dept. of Computer Science
saeed@usc.edu

Vasileios Zois
University of Southern California
Dept. of Computer Science
vzois@usc.edu

Viktor Prasanna
University of Southern California
Dept. of Electrical Engineering
prasanna@usc.edu

Acknowledgement

This material is based upon work supported by the Department of Energy under Award Number DE-OE0000192.

Disclaimer

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, the Los Angeles Department of Water and Power, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

Enabling Automated Dynamic Demand Response: From Theory to Practice

Marc Frincu
University of Southern California
Dept. of Electrical Engineering
frincu@usc.edu

Charalampos Chelmis
University of Southern California
Dept. of Electrical Engineering
chelmis@usc.edu

Saima Aman
University of Southern California
Dept. of Computer Science
saman@usc.edu

Rizwan Saeed
University of Southern California
Dept. of Computer Science
saeed@usc.edu

Vasileios Zois
University of Southern California
Dept. of Computer Science
vzois@usc.edu

Viktor Prasanna
University of Southern California
Dept. of Electrical Engineering
prasanna@usc.edu

ABSTRACT

Demand response (DR) is a technique used in smart grids to shape customer load during peak hours. Automated DR offers utilities a fine grained control and a high degree of confidence in the outcome. However the impact on the customer's comfort means this technique is more suited for industrial and commercial settings than for residential homes. In this paper we propose a system for achieving automated controlled DR in a heterogeneous environment. We present some of the main issues arising in building such a system, including privacy, customer satisfiability, reliability, and fast decision turnaround, with emphasis on the solutions we proposed. Based on the lessons we learned from empirical results we describe an integrated automated system for controlled DR on the USC microgrid. Results show that while on a per building per event basis the accuracy of our prediction and customer selection techniques varies, it performs well on average when considering several events and buildings.

Categories and Subject Descriptors

J.m [**Computer Applications**]: Miscellaneous

; H.4 [**Information Systems**]: Information Systems Applications

General Terms

smart grid, automated direct demand response

Keywords

Smart grid, demand response, direct control, microgrid, HVAC, curtailment prediction, customer selection

1. INTRODUCTION

With the increasing presence of smart appliances in our homes smart grids have a great opportunity to become ubiquitous in our society. Their unique advantages driven by the bidirectional smart meter communication channel between providers and customers make them suited for energy consumption optimization. Demand Response (DR) [40, 43] is a well known technique used by utilities to shape customers load especially during peak hours when the generation capacity is in danger of being exceeded. Utilities have at their disposal a variety of techniques including direct control [41], incentives [28] or voluntary programs [7] to reduce peak demand.

Direct control is a technique which lets utilities control the customers' appliances directly without relying on incentives or voluntary participation. For this reason it can be seen as intrusive and having a considerable impact on the customer lifestyles. Hence, it seems more suited for offices [7] and industrial complexes than for residential customers. Nonetheless its advantages including fine grained control of the amount of energy to curtail and the ability to efficiently target specific areas and time periods [14] makes it suited for further research into how it can be easily and with minimal impact integrated in complex microgrid environments consisting of heterogeneous buildings (e.g., offices, residential apartments, libraries, lecture rooms).

The USC microgrid, peaking at around 27 MW, is a unique environment to study the impact of direct control in complex social and cultural environments by offering a state-of-the-art control center capable of managing 170 buildings spread across two campuses totaling more than 50,000 sensors including smart meters, thermal, humidity, presence, and photosensitive devices. Together, they make the USC campus a truly "living laboratory" for advancing Smart Grid research and technology [36]. The USC Facility Management Services (FMS) owns and operates the campus electrical infrastructure which includes two substations and a 3 million gallon Thermal Energy Storage (TES) system complementing the existing chilled water system. The TES reduces the need for additional electrical generating facilities by chilling the water at night and using it in the air conditioning system during daytime when electricity is most expensive. FMS has more than 6 years worth of historical and real-time kWh consumption data, gathered from 33 DR enabled buildings, aggregated at 15 minute intervals. Combined with detailed information on the classes' schedule, buildings' occupancy, and weather data, it offers a unique opportunity to investigate the main challenges and possible solutions to adopting an efficient and reliable controlled DR program in complex dynamic environments. Currently FMS's focus is on HVAC based DR but upgrades to extend it to other DR techniques such as those based on the lighting system have been planned. The USC microgrid is also a test bed for the LA DWP Smart Grid Demonstration Project (SGRDP) [17] which involves more than 50,000 residential customers in voluntary and controlled (through smart home appliances) DR. LA DWP controls the DR activities on the campus microgrid by issuing curtailment notifications in advance and monitoring the progress and outcome of the DR events.

To achieve a customer-tolerated and utility-efficient controlled DR numerous factors need to be taken into consideration including: data privacy, customer satisfiability, reliability of the controlled DR action, and fast turnaround time of the decision process. To make things more complicate the DR model widely used nowadays needs to consider more realistic scenarios on when, for how long, by how much, and how (whom to pick) to reduce electricity consumption. We call this decision making process Dynamic DR (D^2R).

In this paper we present a functional software platform integrated with the FMS control center which leverages the information stored in heterogeneous data sources to increase the accuracy of the *curtailment prediction* and *customer selection* algorithms in a scalable, reliable, and secure platform for D^2R. Specifically we address the following aspects:

- An overview of the main challenges for D^2R and our approaches to solve them. Based on these results we have designed and implemented our D^2R system;
- The software architecture of the D^2R system and its main components;
- The system integration with the FMS and the LA DWP SGRDP;
- A comprehensive analysis of the results obtained from the real-life deployment of the system in terms of accuracy and efficiency. The novelty of the use case lies in the fact that we are trying to predict the curtailed consumption and to use it in our customer selection method.

The rest of the paper is structured as follows: Section 2 discusses in detail the main challenges D^2R adoption in real-life systems face. Section 5 revisits our main results to emphasize the experimental work we have done prior to implementing the system with respect to DR techniques, algorithms, and systems. Section 3 presents the system architecture and main components with emphasis on the integration between the proposed system, FMS and LA DWP. Section 4 analyzes the main results from the real-life experiments by looking into aspects such as accuracy, reliability, and timely delivery of the consumer selection decisions. Finally Section 6 wraps the overview and analysis, and lays out future research directions.

2. CHALLENGES AND OUR APPROACHES FOR D²R SYSTEMS

While the data deluge coming from various sensors (e.g., smart meters, weather stations, occupancy sensors) offers a fertile environment for researching D^2R techniques, enabling them on an actual smart grid or microgrid faces several challenges including information integration, privacy guarantees, predictive analytics, and dynamic decision making. Figure 1 depicts a typical D^2R system emphasizing these challenges and where they occur in the overall system. Information coming from the various heterogeneous data sources needs to be integrated under a common representation before being used by the consumption and curtailment prediction method. Finally the information obtained from the prediction techniques is used in the customer selection process. The entire process needs to ensure the privacy and security of the data obtained from sensors and information extracted from the prediction and selection mechanisms. In what follows we detail the challenges and our approaches for each of these steps:

2.1 Information Integration

Information available in a smart grid enables the design of novel D^2R techniques for finer control over the energy use. This information goes beyond details about the energy consumption readings available from the utility through smart meters, incorporating indirect influencers derived from customer activities, natural phenomena, and infrastructure behavior. For example, the current weather, the building occupancy fluctuations (e.g., weekdays vs. workdays vs. holidays), the building thermal properties (e.g., construction material), and the power grid infrastructure, can all affect energy consumption in a smart grid. Figure 2 shows how the different information sources are related and influence each other [45].

By integrating these various data sources D^2R algorithms can locate patterns among a large class of historical and real-time information to predict energy consumption and identify curtailment opportunities. However, in order to take full advantage of the diverse data influx, a holistic view of information across multiple domains is required. This integrated information model needs to be extensible to meet the organic and rapid growth of information sources in the smart grid domain, while being easy to interpret and manage the diversity of information and applications.

To address these challenges we proposed in our previous work [45] a semantic information model for the smart grid and demonstrated its use for DR in the campus microgrid, highlighting its extensibility, versatility, and usability.

2.2 Privacy Guarantees

Data from these diverse sources is used among others to predict energy consumption and compute billing invoices. Because some of the information comes from personal data, security and privacy issues have started to receive increas-

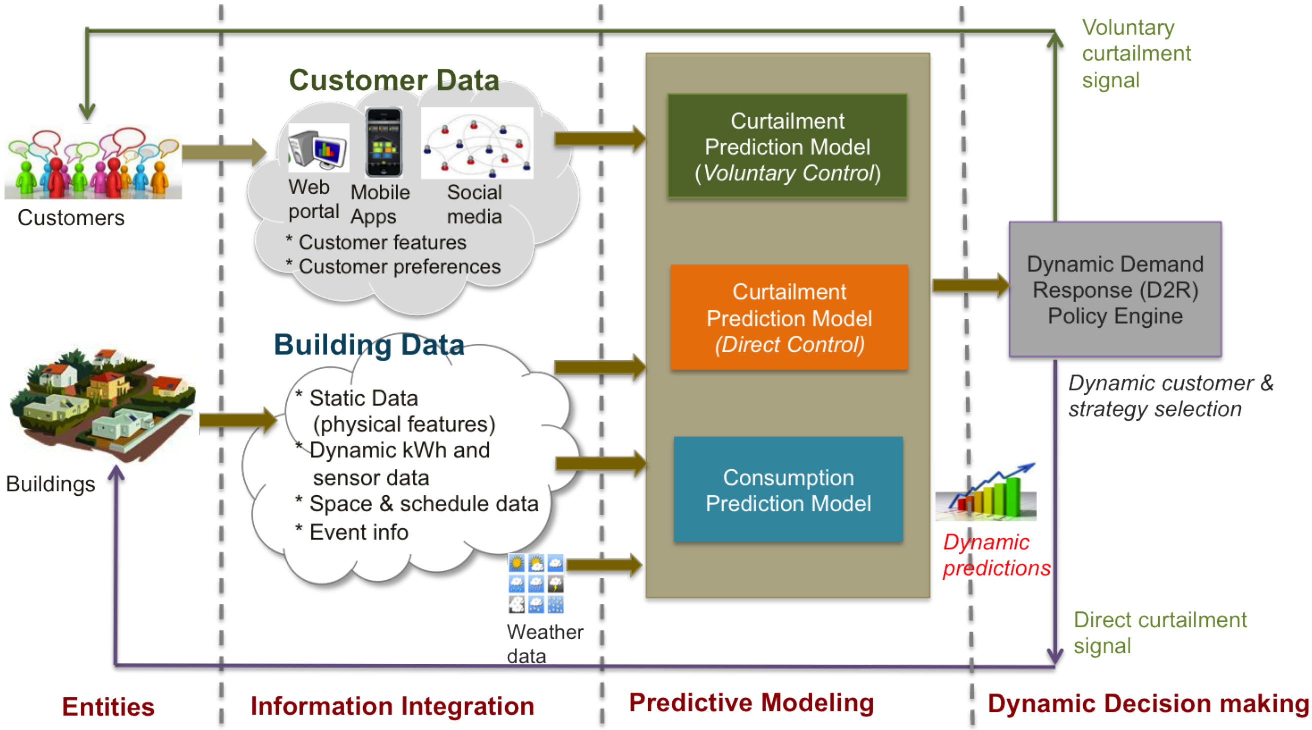

Figure 1: A Dynamic Demand Response (D^2R) System.

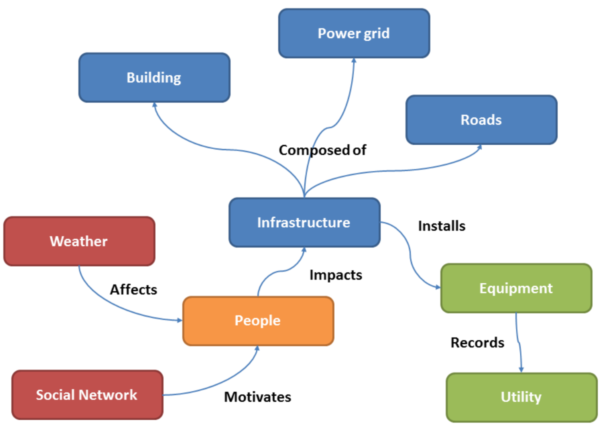

Figure 2: Relations between the various data sources in a smart grid.

ing attention [12, 31, 35]. Privacy aware data retrieval and processing is needed to meet statutory and consumer requirements. Adding too many privacy constraints, however leads to the system's inability to perform the intended task. A straightforward example comes from a simple DR use case where the future energy consumption of a consumer needs to be predicted based on historical data. Here, smart meter data needs to be sampled at an appropriate resolution (e.g., one meter value every 15 minutes). However, this granularity may cause privacy issues as it could be used to predict the consumer behavior at home [37] or the inadvertent release of fine grained meter data linking back to the consumer without his/hers approval. To avoid this utilities could decrease the sampling rate but this would impact the prediction accuracy by loosing information on short hourly consumption trends. Therefore one of the challenges in smart grid system engineering is to find a good trade-off between protecting a customer's privacy and being able to provide useful services. The first model for the trade-off assessment has proposed in our previous work [24] while in [25] we have proposed a model-driven solution for representing the smart grid system and use case communication patterns to assess their privacy impact using expert domain knowledge.

In addition, to ensure data privacy secure repositories are needed especially when outsourcing the storage using emerging technologies such as clouds which offer a reliable and cheap alternative to traditional storage servers. Existing solutions for clouds however offer basic key-based access control mechanisms and require extensive deployment of software infrastructures that limits their manageability. In addition relying on external services for ensuring security and privacy can lead to security breaches outside the control of the beneficiary. In [27] we have proposed Cryptonite, a solution that is controlled by the utilities, has a low key management overhead, and is compatible with existing storage service interfaces while not sacrificing performance and scalability.

2.3 Predictive Analytics

Predictive models for energy consumption provide insight into daily load requirements and peak demand periods being useful for planning DR events to curtail energy consumption. Utilities have traditionally used simple methods and statistical models that are easy to use and interpret [2, 3, 16]. However, the complexity and diversity of the data which impacts consumption motivates the design of advanced methods which can cope not only with the noise in the consumption data, but also with the missing data due to infrastructure, climate, or privacy issues.

Consumer clustering. While time series and regression tree based techniques can be trained on historical data to predict future consumption, they face two challenges when applied to large consumer pools: (1) the training time becomes costly [19], and (2) the variability in consumption data causes the prediction errors to increase [8]. Aggregating data from multiple consumers into a single "virtual consumer" helps address these aforementioned issues. Besides reducing the number of customers to train and predict for, it also reduces the temporal variability of the virtual customer, increasing the prediction accuracy [1]. Due to these advantages we have proposed a polynomial time incremental clustering technique for aggregating consumers [44]. The proposed method adds new customers to clusters only if the cumulative consumption prediction error of the aggregated cluster consumption is reduced by the new addition.

The efficiency of the aggregated consumption prediction depends however on the prediction method. Our previous studies have showed that the accuracy depends among others on the customer consumption pattern [8]. This means that an efficient system needs to decide and use the best method for a particular pattern. Doing an exhaustive search over all methods is time consuming and unsuited for D^2R where fast periodic decisions are needed. We have proposed a novel neural network classification method for predicting the best method based on the standard deviation of the historical time series [19]. The method has proven to be efficient both in the prediction outcome (84-94% accuracy) and in execution time (< 1s per consumer for predicting consumption once trained, making it suited for parallel execution).

Partial data. Complex infrastructures such as smart grids where large volumes of various data are trended at high velocities brings into picture the inevitable aspect of data quality. Due to infrastructure failures and privacy concerns data utilities may not have access to the complete set of data needed to take accurate decisions. For this reason it is imperative to develop advanced methods that rely only on the available data by extrapolating missing information from existing one. The rationale behind this is that some smart grid sensors may be correlated due to their spatial proximity. This is particularly true for the USC microgrid where there exists a strong correlation between dormitory occupancy and classroom schedules. Our empirical results have showed that recent (e.g., up to 4 hours old) real-time values of a few influential sensors are far more effective as predictors than the sensor's own older values. Extending this result we have also proven that we can use the same set of influential sensors for all all sensors and still get comparable results with the initial approach [6].

Curtailed consumption. The previous challenges concerning predictive analytics have focused on forecasting the consumption outside the DR event for planning purposes. Equally important in D^2R is the *prediction of the curtailed consumption* (i.e., consumption during DR) which helps utilities estimate the amount of energy shed and allows them to dynamically adjust the targeted consumers in case they under/over-shoot. Curtailed consumption estimation has received less attention with most of the utilities focusing on baseline prediction methods, e.g., the ISO models [2, 3, 16]. Reduced consumption prediction is particularly difficult because of the few number of events each consumer participates in thus resulting in small training data sets for prediction models. An additional challenge is that the DR events need to be time aligned and that selected past events need to take place in days similar with the current event. To increase prediction efficiency and account for dissimilarities a morning adjustment factor applied to the energy consumption is usually used [23]. Our studies [13] have shown for the USC microgrid case that averaging the reduced consumption of past DR events for a particular building-DR strategy pair achieves good estimates (<15%) for some of the pairs. However the complexity and impact of the different curtailment DR strategies used in a controlled environment [38] (e.g., Variable Frequency Drive – VFD [34], Global Temperature Reset – GTR [34], and Duty Cycling) requires more complex solutions relying on causal effects and accounting for more external factors.

2.4 Dynamic Decision Making

For D^2R to be successful accurate prediction techniques need to be complemented with an efficient customer selection. Typical DR events have a specific curtailment target C they need to meet over a given time period T. Given the dynamics and different consumption patterns of each customer the selected set needs to change over time to meet the specified target. Furthermore, utilities are sometimes interested in keeping the level of curtailment stable throughout the duration of the DR event, i.e., Sustainable D^2R (SD^2R). Reasons utilities may want this property include ensuring reliability in power distribution and cost reduction during peak hours. SD^2R as D^2R is a hard combinatorial problem due to the numerous feasible combinations possible.

The problem can be formulated as a binary linear programming optimization. Given n consumers and m DR strategies, let B be the matrix of building-strategy combinations, where each entry B_{ij} represents the predicted curtailment vector $< r_1, r_2, ..., r_k >$ obtained when applying DR strategy j on consumer i across the DR event time frame T. Each value $r_k \in \mathcal{R}$ is calculated at a 15 minute interval and represents the difference between the predicted baseline and the predicted consumption during the DR event. A positive curtailment estimate indicates that the strategy was successful in curtailing while a negative one indicates a predicted increase in consumption. It should be noted that the actual values depend on the efficiency of the methods used in the estimation [13].

Given the above notations the problem can be defined as $\min \quad x_{00}C + \sum_{i=1,n} \sum_{j=1,m} x_{ij}(-1)r_{ij}$ subject to the following constraints $\sum_{i=1,n} \sum_{j=1,m} x_{ij}r_{ij} \leq C$, $\sum_{j=1,m} x_{ij} = 1$ and $x_{00} = 1$. Here x_{ij} represents the participation matrix: $x_{ij} = 1$ if strategy j is used on building i, 0 otherwise. It should be noted that ideally $\sum_{i=1,n} \sum_{j=1,m} x_{ij}r_{ij} = C$ however due to the nature of r_{ij} this is usually not possible. To ensure the SD^2R at each time interval t a reselection occurs based on the new target $C_t = C - A_{t-1}$, where A_{t-1} represents the achieved curtailment in the previous interval computed as the difference between the predicted baseline and the actual observed curtailed consumption.

Integer linear programming problems are known to be $\mathcal{NP}$-hard. While for a small set of variables they may be able to offer a solution in a reasonable amount of time for a large set of customers a utility can have this becomes unfeasible due to the near real-time requirement of the D^2R. We have proposed in [47, 46] a fast and efficient heuristics based on a knapsack approach, i.e., we formulated the problem as a change-making problem. Our simulated experiments have shown a reduction in the number of selected customers as well as a high probability in achieving the targeted curtailment. This makes our solution reliable and suited for cases where human comfort is involved.

While this data driven approach has proven to work well for some cases, it should be noted that in some cases where a customer is not a single household but a larger entity such as a building, the curtailment may not be visible directly from the consumption data. For instance, a building may enter in DR by using GTR but other factors such as a lower outside temperature (leading to a smaller impact in curtailment) correlated with an increase usage of lights and IT equipment due to classes may hide the curtailment drop. To study this we need to analyze the correlation between the building and equipment level consumption during a DR event. Our analysis has shown that $\approx 50\%$ of the tested campus buildings exhibit a strong correlation between building and equipment indicating that in these cases a data driven approach can produce reliable consumer selection. For the other half a bottom-up approach the curtailment is derived from the mechanical properties should be used. In [20] we have proposed several heuristics that consider equipment and human comfort. Results have shown that the best method depends on the used DR strategy and that a certain level of human comfort can be achieved despite being a contradictory objective with respect to curtailment. The main drawback of the bottom-up approach is that it is only useful when the technical specifications of the equipment used in DR are known (e.g., how much cooling unit consumes when its fans are operating at 100%).

3. SYSTEM REALIZATION

While in Sect. 2 we have outlined the main challenges and our solutions for D^2R systems, in this section we focus on the actual system implemented in the USC campus microgrid. Section 4 will present the main results of over 400 DR real-life experiments on the system.

3.1 LA DWP - FMS - VSoE Integration

As mentioned in Sect. 1 the proposed system is part of the LA DWP SGRDP. The USC microgrid communicates with the LA DWP DRAS server [26] through OpenADR [5] messages. Internally, the software infrastructure consists of two main components: the Integrated Building Control (IBC) and the Demand Decision Support (DDS) module. The former is in charge of controlling buildings through the BACnet protocol [9], while the latter is responsible for the building-strategy selection during each DR event. Figure 3 gives an overview of the system's main components. A typical workflow begins with an LA DWP OpenADR message containing the request details for a DR event (e.g., datetime, duration, curtailment target). For the USC microgrid the duration is always set between 1:00PM and 5:00PM due to the specific LA climate which induces a peak load at that time. Once it has received the request the IBC forwards the details to DDS including additional details such as the list of buildings and associated DR strategies for the given event. DDS then processes the message and recommends based on historical and real-time data (i.e., consumption, weather, occupancy, etc.) which building-strategy pairs to use. This list is sent back to IBC which sends the selected buildings into DR by using the suggested strategies. Information on the achieved curtailment and estimated achievable target is sent back to the LA DWP DRAS every 15 minutes. To ensure SD^2R a building-strategy reselection is performed hourly. Finally, the outcome of the DR event is sent 15 minutes after its completion, by providing LA DWP with information on the achieved curtailment and the deviation from the initial estimate sent with the first DDS message. Next we detail the data flow and message types, the DDS, and the secure data repository.

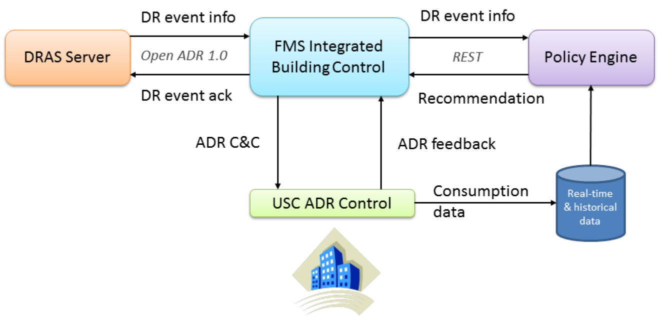

Figure 3: Overview of the complete automated D^2R system.

3.2 Data flow

The IBC constantly communicates with DDS in order to sustain the curtailment target. The communication protocol is based on REST [18] and contains two types of messages: *SELECT* and *INFO*. The former is used to initialize/adjust the buildings that participate in the DR event, while the latter is used to send back the current achieved curtailment together with the remaining estimate. Figure (cf. Fig. 4) depicts the message flow. Before the start of the DR event two SELECT messages are sent to initialize (FAR) the selection and to adjust ($NEAR$) it 15 minutes before the event start. The NEAR is intended to capture any changes due to extrinsic factors such as temperature fluctuations. During DR a series of INFO and SELECT messages follow at 15 minute intervals with the SELECT messages issued on the hour. Their role is to adjust the building-strategy selection to achieve SD^2R as explained in the next subsection. After the completion of the DR event a final INFO message is relayed back to the DRAS server through the IBC. It should be noted the IBC receives a third type of *ACTIVE* message which effectively triggers the start of the DR, however this message is not being relayed to DDS.

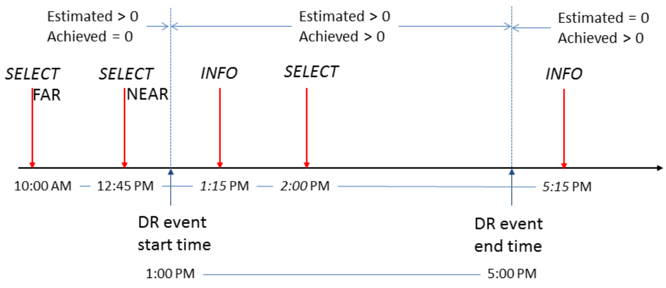

Figure 4: Message types .

3.3 DDS

At the core of the D^2R system lies DDS[1]. DDS is a Java library embedded with machine learning techniques for consumption and curtailment prediction as well consumer and DR strategy selection algorithms for SD^2R. The DDS is currently exposed to the IBC as a REST service which receives as input consumption data and DR requirements (e.g., duration, start time, available buildings and strategies) and returns the suggested list of building-strategies. Figure 5 depicts the overall functionality. At the core of the decision process lie three processes: the consumption baseline computation, the curtailment matrix setup, and the building-strategy selection.

The consumption baseline, used to determine the consumption as it would have been in the absence of the DR is utilized both by the selection policy to determine how much has been achieved and what is the remaining target, and the curtailment matrix to determine the estimated curtailment for the duration of the DR event. Our studies [13] have shown that the California Southern Edison [16] and fixed (i.e., the consumption value at the start of the DR) baselines are the least biased.

To accurately estimate the baseline consumption values just before the DR start should be used. However because the FAR message is usually sent a few hours in advance the estimates are likely to be off. The purpose of the FAR message is thus to give a rough estimate on the overall curtailment projection for the day. The refined selection is accomplished as a result of the NEAR message sent 1 hour ahead of the DR event.

The curtailment matrix is computed based on the curtailed consumption estimate given by our averaging method [13] (cf. Sect. 2.3) and the fixed baseline. Each row in the matrix contains sixteen 15 minute values (corresponding to the entire 1:00PM-5:00PM DR interval) for a single building-strategy combination. To ensure enough data is available for the averaging method we have performed more than 400 DR experiments between Nov 2012 - Dec 2014 1:00PM to 5:00PM on 33 USC buildings using VFD, GTR, Duty, or any combination supported by a particular building. Figure 6 depicts the distribution of various DR strategy tests per building. Based on these tests for a specific building-strategy we take all previous similar DR events associated with it and average the 15 minute curtailed consumption values recorded during DR. The 16 resulted values (4 readings per hour $\times$ 4 hours) are then subtracted from the fixed baseline for the upcoming scheduled event obtaining the estimated curtailment at 15 minute granularity. Once all building-strategy entries are computed the matrix is used by the policy selection engine to decide based on it and the so far achieved curtailment (0 if outside the DR for FAR and NEAR messages) which building-strategy pairs to select. During its decision making process the engine will first remove all negative curtailment matrix estimates as they indicate a possible increase in consumption rather than a decrease during the DR. Section 2.4 has addressed this possibility and while a bottom-up approach is possible [20] the current integrated system only considers the data driven approach based on the curtailment matrix.

During DR based on the achieved curtailment the engine may decide to add/remove a building-strategy pair or to change the strategy of an already selected building. In the latter case, due to mechanical limitations only certain changes are allowed. Figure 7 shows the allowed transitions between the DR strategies. It should be noted that the supported strategies also differ for each building. Finally, the selected building-strategy list is sent back to the IBC which adjusts the building in the DR accordingly.

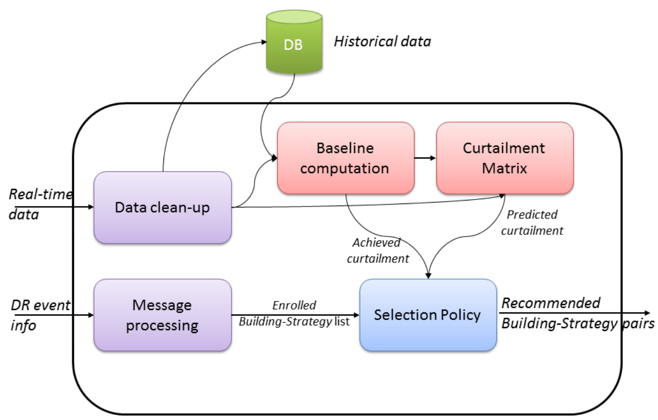

Figure 5: Overview of the D^2R decision system.

3.4 Data Storage

Due to the distributed nature of the automated D^2R system data is stored in several places based on the location of the service provider (i.e., LA DWP for DRAS, FMS for IBC, and DDS – which is stored in the Viterbi School of Engineering). FMS currently relies on plain text files to export the consumption data DDS requires. To ensure the security of these files, solutions such as our Cryptonite [27] discussed in Sect. 2.2 can be used. In the near future a migration

[1]`http://smartgrid.usc.edu/dds-javadoc/2.0/`

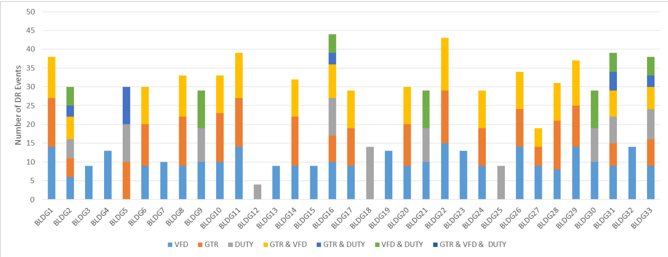

Figure 6: Strategy distribution for each tested building.

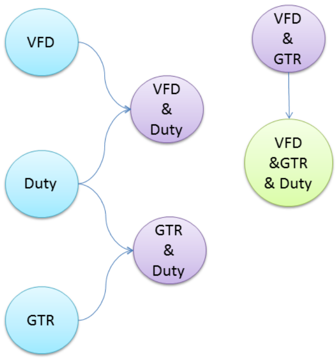

Figure 7: State transitions between DR strategies.

towards an ICONICS Hyper Historian [22] which offers high performance, robust, scalable, and secure consumption data storage compatible with the BACnet connectivity standard used by FMS is planned.

To manage the historical and real-time data sent by FMS with minimum network overhead DDS stores the real-time data in a scalable MongoDB repository [33] which offers application and file level encryption to secure the privacy of the data. As such it minimizes the impact on the network since all data is available near the server for the curtailment matrix computation and curtailed initial estimate and final observed values.

Besides consumption data some of the DDS consumption and curtailment forecast algorithms (e.g., regression tree) rely on additional information such as weather and classroom schedule. For the weather data we rely on the NOAA station located on the USC campus. Data coming from these additional sources is stored as plain text and updated every semester (for the schedule) or periodically through automated scripts (for the weather).

4. DEPLOYMENT AND EXPERIMENTS

The proposed system has been validated on 33 campus buildings each supporting one or more HVAC-based DR strategies (cf. Fig. 7). We have first analyzed the accuracy of our averaging model in predicting the curtailed consumption

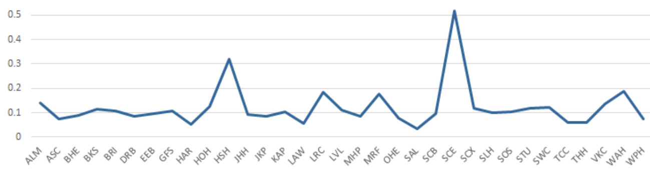

Figure 8: Average MAPE error per building for all events.

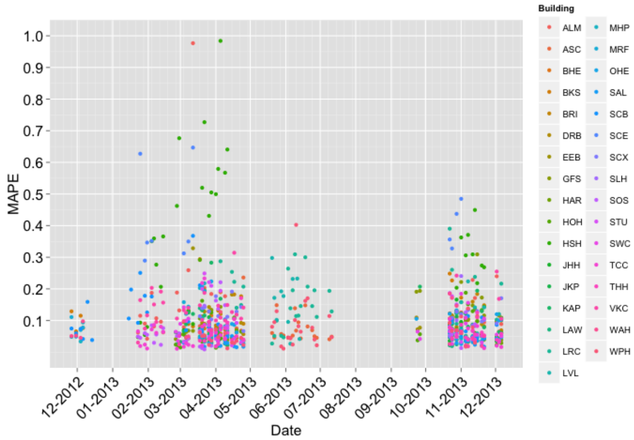

Figure 9: MAPE error per event event per building.

based on more than 400 DR events run between 1:00PM and 5:00PM from Nov 2012 to Dec 2014. These experiments are the same ones used for building the curtailment matrix discussed in Sect. 3.3. Cross validation was used to assess the MAPE error of the averaging error. Nine more DR experiments were run between Oct 24 and Dec 5 2015 to determine the accuracy of our DDS selection policy when relying on the curtailment matrix estimates. The accuracy was determined to be the MAPE of the total observed kWh consumption at the end of the DR end as compared to the initial estimate. Results are discussed in what follows.

4.1 Analysis

Figure 8 shows good results on average for the proposed curtailment method on a per building basis across all events. Around 49% of buildings had an error less than 10% while 6% were greater than 20%. However a detailed analysis as depicted in Fig. 9 shows a few outliers for certain events in Spring and Fall. For instance we have noticed few buildings which exhibit high variability with errors ranging between 10% to 50%-100% for certain events. Taken independently, these buildings are unsuited for DR. We are currently studying these buildings to determine the factors that differentiate them from other more predictable buildings. Initial results have not showed any conclusive results and we leave this analysis open for future work.

Results have shown us that given enough DR events our proposed method can provide on average good estimates in forecasting the curtailment of a future event. While this can be good for long term savings it means that for specific DR events the errors could still be high. Figure 10 depicting the errors for individual events shows a few of outliers

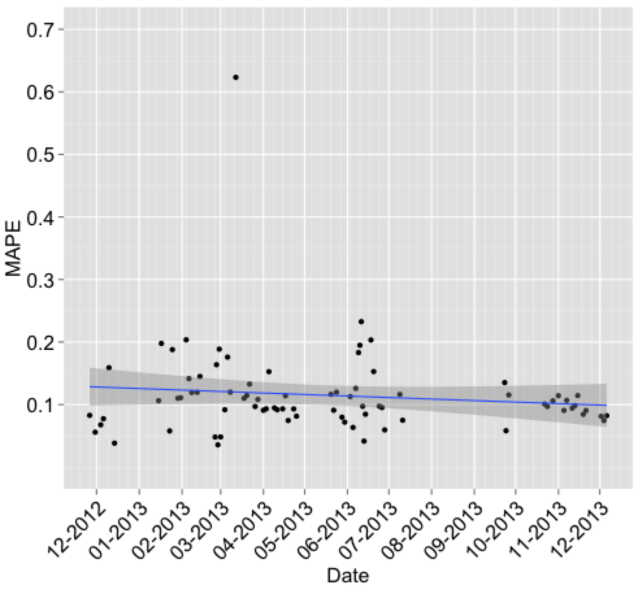

Figure 10: Average MAPE error per event for all buildings.

with errors above 20%. Complementing with the information from Fig. 9 it can be concluded that while individual buildings may experience fluctuations their aggregated results contribute to a low overall error per event.

The nine tests we performed to validate our hypothesis and cross validation experiments have shown that the method over predicts the curtailment by anywhere in the [12$, 46%] with three events being less than 16% off the initial estimate and five under 31%. Tests have shown however that as more tests are performed the accuracy of the initial estimate increases from 39.7% after the first five tests to 30% after all nine, with the last four having a deviation of 19.6%. This is due to the averaging model being able to rely on more relevant recent historical events.

The turnaround time of the building-strategy selection method from the moment it received the request and the time it sent back the recommendation is <1 minute. Simulated tests in our previous work [46] have showed that the method scales well with the number of customers making it suitable for larger customer pools.

5. RELATED WORK

Not many integrated DR systems for complex environments such as a campus microgrid exist. A recent paper discussed a research platform deployed on the University of California San Diego microgrid for developing large scale, predictive analytics for real-time energy management [10]. Contrary to our work which deals with DR by focusing on both consumption and curtailment prediction, the UCSD smart grid is focused on consumption prediction only to improve operational efficiency. As such they do not consider the complexity of D^2R and the challenges it raises.

There have been numerous attempts to deal with consumption demand. Utility providers can either compensate by buying extra power at high prices [11] or employ DR strategies. The latter is a well known concept divided into two categories which include direct control and voluntary participation [4]. Arguments in favor of both techniques have yielded several different solutions driven by specific use cases. These address residential buildings [15], offices [21] as well as large industrial facilities[39] and data centers [30]. In this paper we focus on a heterogeneous microgrid which includes a mixture of various building types including residential, offices, libraries, and mixed spaces. This environment offers a more realistic scenario for concepts such as smart cities. Our work is based on directly controlling the building equipment to achieve and sustain a specified curtailment.

Previous work aimed at load manipulation includes attempts to minimize peak demand by shifting it to less busy hours of the day [42, 29] or optimizing load consumption while minimizing costs from the customer perspective [32]. The above methods rely on cooperative customer action and has the main drawback of not ensuring the sustainability of the DR event. In contrast our system periodically reselects customers and DR strategies to sustains the DR. Utility providers can either rely on direct or voluntary participation as long the necessary consumption data from past DR events are available for the selection procedure. Maximizing human comfort plays an important role in a directly controlled environment an issue that we have addressed in our previous work [20].

To the best of our knowledge this is the first work to address the concept of automated SD^2R and to propose an automated equipment/building selection system for heterogeneous microgrids.

6. CONCLUSION

In this paper we have presented some of the main challenges in designing and implementing an automated system for controlled D^2R. Based on our results we have introduced a functional system deployed on the USC microgrid and validated by performing multiple on campus DR events. The USC microgrid is a heterogeneous environment equipped with an automated building control center making it suited for numerous controlled DR scenarios. Results have shown that on average the system achieves good accuracy in terms of predicting curtailed consumption during DR. Furthermore we have showed that by running multiple events the error between the estimated curtailment and final achieved one reduces significantly.

As future work we will continue our analysis on the possible factors that contribute to some buildings' high error variability. We are also looking at improving the curtailment prediction accuracy by designing ensemble models consisting of a combination of methods (one being the averaging method) each tested empirically to work best for specific cases.

Finally, we will focus on other controlled DR strategies besides HVAC, by looking at leveraging the information from the numerous sensors installed throughout buildings and the existing and planned energy storage solutions.

7. ACKNOWLEDGMENTS

This material is based upon work supported by the United States Department of Energy under Award Number number DE-OE0000192, 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. We also acknowledge the valuable contributions made over the years by Yogesh Simmhan (Indian Institute of Science), Zachary Gima (USC Mechanical Engineering), Fabian Knirsch (Salzburg University of Applied Sciences), Muhammad Usman Noor (USC Electrical Engineering), Prashant Nittoor (USC Computer Science), Ranjan Pal (USC Computer Science), and Qunzhi Zhou (USC Computer Science).

8. ADDITIONAL AUTHORS

Additional authors: Carol Fern (University of Southern California, Facility Management Systems, `carol.fern@usc.edu`) and Aras Ackbari (LA DWP contractor, `aras.ackbari@ladwp.com`).

9. REFERENCES

[1] Law of large numbers. URL: http://www.encyclopediaofmath.org/index.php /Law_of_large_numbers (accessed Dec 22, 2014).

[2] CAISO Demand Response User Guide, Guide to Participation in MRTU Release 1, Version 3.0. Technical report, California Independent System Operator (CAISO), 2007.

[3] Emergency Demand Response Program Manual, Sec 5.2: Calculation of Customer Baseline Load (CBL). Technical report, New York Independent System Operator (NYISO), 2010.

[4] M. Albadi and E. El-Saadany. Demand response in electricity markets: An overview. In *IEEE Power Engineering Society General Meeting*, volume 2007, pages 1–5, 2007.

[5] O. A. Alliances. Open automated demand response, 2014. URL: http://www.openadr.org/ (accessed May 23, 2014).

[6] S. Aman, C. Chelmis, and V. K. Prasanna. Towards prediction with partial data in sensor-based big data applications. In *Procs. 29th AAAI Conference*, page accepted, 2014.

[7] S. Aman, M. Frincu, C. Charalampos, U. Noor, Y. Simmhan, and V. Prasanna. Empirical comparison of prediction methods for electricity consumption forecasting. Technical Report 14-942, University of Southern California, 2014. URL: http://www.cs.usc.edu/assets/007/89887.pdf (accessed Mar 6, 2014).

[8] S. Aman, M. Frincu, C. Chelmis, M. U. Noor, Y. Simmhan, and V. K. Prasanna. Empirical comparison of prediction methods for electricity consumption forecasting. Technical report, University of Southern California, 2014.

[9] ASHRAE. Bacnet, 2015. URL: http://www.bacnet.org/ (accessed Jan 3, 2015).

[10] N. Balac, T. Sipes, N. Wolter, K. Nunes, B. Sinkovits, and H. Karimabadi. Large scale predictive analytics for real-time energy management. In *Big Data, 2013 IEEE International Conference on*, pages 657–664, 2013.

[11] A. Barbato, A. Capone, L. Chen, F. Martignon, and S. Paris. A power scheduling game for reducing the peak demand of residential users. In *Online Conference on Green Communications (GreenCom), 2013 IEEE*, pages 137–142, Oct 2013.

[12] A. Cavoukian, J. Polonetsky, and C. Wolf. Smartprivacy for the smart grid: embedding privacy into the design of electricity conservation. *Identity in the Information Society*, 3(2):275–294, 2010.

[13] C. Chelmis, S. Aman, M. Frincu, and V. K. Prasanna. Estimating reduced consumption for dynamic demand response. In *Procs. 29th AAAI Conference Workshops*, page accepted, 2014.

[14] C. Chelmis, S. Aman, M. R. Saeed, M. Frincu, and V. K. Prasanna. Estimating reduced consumption for dynamic demand response. In *AAAI Workshop on Computational Sustainability*, page accepted, 2015.

[15] J. Chen, B. Yang, and X. Guan. Optimal demand response scheduling with stackelberg game approach under load uncertainty for smart grid. In *Smart Grid Communications (SmartGridComm), 2012 IEEE Third International Conference on*, pages 546–551, Nov 2012.

[16] K. Coughlin, M. A. Piette, C. Goldman, and S. Kiliccote. Estimating demand response load impacts: Evaluation of baseline load models for non-residential buildings in California. Technical Report LBNL-63728, Lawrence Berkeley National Lab, 2008.

[17] L. DWP. Los angeles department of water and power smart grid regional demonstration, 2014.

[18] R. T. Fielding and R. N. Taylor. Principled design of the modern web architecture. *ACM Transactions on Internet Technology (TOIT)*, 2(2):115–150, 2002.

[19] M. Frincu, C. Chelmis, M. U. Noor, and V. K. Prasanna. Accurate and efficient selection of the best consumption prediction method in smart grids. In *Proc. IEEE International Conference on Big Data*, page in print. IEEE, 2014.

[20] M. Frincu and Z. Gima. A bottom-up approach to sustained curtailment and comfort for controlled demand response. In *Procs. 2nd IEEE Conference on Technologies for Sustainability*, page in print, 2014.

[21] I. Georgievski, V. Degeler, G. Pagani, T. A. Nguyen, A. Lazovik, and M. Aiello. Optimizing energy costs for offices connected to the smart grid. *Smart Grid, IEEE Transactions on*, 3(4):2273–2285, Dec 2012.

[22] I. Iconics. Hyper istorian, 2015. URL: http://www.iconics.com/Home/Products/Historians/Hyper-Historian.aspx (accessed Jan 5, 2015).

[23] C. Katie, P. Mary Ann, G. Charles, and K. Sila. Estimating demand response load impacts: Evaluation of baseline load models for non-residential buildings in california. Technical report, 2008. URL: http://drrc.lbl.gov/sites/drrc.lbl.gov/files/63728.pdf (accessed Nov 23, 2013).

[24] F. Knirsch, D. Engel, C. Neureiter, M. Frincu, and V. K. Prasanna. Model-based assessment for balancing privacy requirements and operational capabilities in

the smart grid. In *Proc. IEEE Conference on Innovative Smart Grid Technologies*, page accepted. IEEE, 2014.

[25] F. Knirsch, D. Engel, C. Neureiter, M. Frincu, and V. K. Prasanna. Model-driven privacy assessment in the smart grid. In *Proc. Int. Conference on Information Systems Security and Privacy*, page accepted. SCITEPRESS, 2014.

[26] E. Koch and M. A. Piette. Architecture concepts and technical issues for an open, interoperable automated demand response infrastructure. *Lawrence Berkeley National Laboratory*, 2008.

[27] A. Kumbhare, Y. Simmhan, and V. Prasanna. Cryptonite: A secure and performant data repository on public clouds. In *Cloud Computing (CLOUD), 2012 IEEE 5th International Conference on*, pages 510–517, 2012.

[28] X. Lou, D. Yau, H. H. Nguyen, and B. Chen. Profit-optimal and stability-aware load curtailment in smart grids. *Smart Grid, IEEE Transactions on*, 4(3):1411–1420, Sept 2013.

[29] F. Mangiatordi, E. Pallotti, P. Del Vecchio, and F. Leccese. Power consumption scheduling for residential buildings. In *Environment and Electrical Engineering (EEEIC), 2012 11th International Conference on*, pages 926–930, May 2012.

[30] K. Mares. Demand response and open automated demand response opportunities for data centers. *Lawrence Berkeley National Laboratory*, 2010.

[31] P. McDaniel and S. McLaughlin. Security and privacy challenges in the smart grid. *Security Privacy, IEEE*, 7(3):75–77, May 2009.

[32] A.-H. Mohsenian-Rad, V. Wong, J. Jatskevich, R. Schober, and A. Leon-Garcia. Autonomous demand-side management based on game-theoretic energy consumption scheduling for the future smart grid. *Smart Grid, IEEE Transactions on*, 1(3):320–331, Dec 2010.

[33] I. MongoDB. Mongodb, 2015. URL: http://docs.mongodb.org/manual/ (accessed Jan 5, 2015).

[34] N. Motegi, M. A. Piette, D. S. Watson, S. Kiliccote, and P. Xu. Introduction to commercial building control strategies and techniques for demand response. Technical report, Lawrence Berkely National Laboratory, 2006. URL: http://gaia.lbl.gov/btech/papers/59975.pdf (accessed Nov 4, 2013).

[35] C. Neureiter, G. Eibl, A. Veichtlbauer, and D. Engel. Towards a framework for engineering smart-grid-specific privacy requirements. In *Proc. IEEE IECON 2013, Special Session on Energy Informatics*. IEEE, 2013.

[36] U. of Southern California. Usc living lab and smart grid, 2014.

[37] O. Patri, A. Panangadan, C. Chelmis, and V. K. Prasanna. Extracting discriminative features for event-based electricity disaggregation. In *Proc. 2nd IEEE Conference on Technologies for Sustainability*, page in print. IEEE, 2014.

[38] M. A. Piette, O. Sezgen, D. S. Watson, N. Motegi, C. Shockman, and L. ten Hope. Development and evaluation of fully automated demand response in large facilities. Technical report, Lawrence Berkely National Laboratory. URL: http://drrc.lbl.gov/system/files/cec-500-2005-013.pdf (accessed Nov 4, 2013).

[39] M. A. Piette, O. Sezgen, D. S. Watson, N. Motegi, C. Shockman, and L. Ten Hope. Development and evaluation of fully automated demand response in large facilities. *Lawrence Berkeley National Laboratory*, 2004.

[40] F. Rahimi and A. Ipakchi. Demand response as a market resource under the smart grid paradigm. *Smart Grid, IEEE Transactions on*, 1(1):82–88, 2010.

[41] N. Ruiz, I. Cobelo, and J. Oyarzabal. A direct load control model for virtual power plant management. *Power Systems, IEEE Transactions on*, 24(2):959–966, 2009.

[42] A. Sepulveda, L. Paull, W. Morsi, H. Li, C. Diduch, and L. Chang. A novel demand side management program using water heaters and particle swarm optimization. In *Electric Power and Energy Conference (EPEC), 2010 IEEE*, pages 1–5, Aug 2010.

[43] P. Siano. Demand response and smart grids - a survey. *Renewable and Sustainable Energy Reviews*, 30:461–478, 2014.

[44] Y. Simmhan and M. Noor. Scalable prediction of energy consumption using incremental time series clustering. In *Big Data, 2013 IEEE International Conference on*, pages 29–36, Oct 2013.

[45] Q. Zhou, S. Natarajan, Y. Simmhan, and V. Prasanna. Semantic information modeling for emerging applications in smart grid. In *Information Technology: New Generations (ITNG), 2012 Ninth International Conference on*, pages 775–782, April 2012.

[46] V. Zois, M. Frincu, C. Chelmis, and V. Prasanna. Efficient customer selection for sustainable demand response in smart grids. In *2014 Green Computing Conference*, page in print, 2014.

[47] V. Zois, M. Frincu, and V. Prasanna. Integrated platform for automated sustainable demand response in smart grids. In *Intelligent Energy Systems (IWIES), 2014 IEEE International Workshop on*, pages 64–69, 2014.