

# Integrated Platform for Automated Sustainable Demand Response in Smart Grids

---

Vasileios Zois  
Department of Computer  
Science  
University of Southern  
California  
Email: [vzois@usc.edu](mailto:vzois@usc.edu)

Marc Frincu  
Department of Electrical  
Engineering  
University of Southern  
California  
Email: [frincu@usc.edu](mailto:frincu@usc.edu)

Viktor Prasanna  
Department of Electrical  
Engineering  
University of Southern  
California  
Email: [prasanna@usc.edu](mailto: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.

# Integrated Platform for Automated Sustainable Demand Response in Smart Grids

Vasileios Zois

Department of Computer Science  
University of Southern California  
Email: vzois@usc.edu

Marc Frincu

Department of Electrical Engineering  
University of Southern California  
Email: frincu@usc.edu

Viktor Prasanna

Department of Electrical Engineering  
University of Southern California  
Email: prasanna@usc.edu

**Abstract**—Demand Response(DR) is a common practice used by utility providers to regulate energy demand. It is used at periods of high demand to minimize the peak to average consumption ratio. Several methods have been proposed over the previous years on how to formulate and deal with the problem of excess demand. Following these methods automated systems for initiating and regulating demand response events have emerged. In this paper we present an automated system for providing an estimated demand response schedule of participating customers. We quantify the achieved energy reduction using information about the baseline consumption and the consumption during DR. Our goal is to provide a sustainable reduction to ensure the elimination of peaks in demand. The proposed system includes an adaptation mechanism for when the provided solution does not meet the DR requirements. We conducted a series of experiments using consumption data from a real life micro grid to evaluate the efficiency as well as the robustness of our solution.

**Keywords:** *dynamic demand response, sustainable reduction, automated demand response, real time adaptation, scheduling.*

## I. INTRODUCTION

Reliable energy distribution has been the cornerstone of the energy industry. Utility providers are concerned with meeting the energy demand while ensuring the viability of the distribution network. Over the last years the traditional power grids have evolved to complex cyber-physical systems [7], [17] consisting of bi-directional smart meters that report energy consumption in real time. The collected data can be used to predict future consumption [4] or deal with periods of high demand [20] through Demand Response (DR) [20], [23] techniques.

DR is well known paradigm used by utility providers to shape customer load. A variety of techniques have been employed to minimize consumption based on direct control [21] or customer voluntary participation [4]. While both paradigms have been used extensively they fail to eliminate demand peaks as they may shift the original load to other less busy periods of the day.

To deal with this scenario we have introduced the notion of Sustainable DR (SDR) [24]. A DR event is said to be sustainable if it achieves a consistent reduction through the whole DR period. We define as consistent the reduction of low deviation between the observed values of consecutive intervals. This can be formally described by:




Fig. 1: Sustainable reduction compared to peak minimization.

$$S = \frac{1}{n} \cdot \sum_{j=2}^n |L_j| + \max(L_j) - \min(L_j) \geq 0 \text{ given:}$$

$$R = \{[x_1, f(x_1)], [x_2, f(x_2)], \dots, [x_n, f(x_n)]\} \quad (1)$$

$$\forall j \in [2, n], L_j = \frac{f(x_j) - f(x_{j-1})}{x_j - x_{j-1}}$$

The equation shows the reduction vector  $R$  described by the reduction values  $f(x_i)$  for each sample  $x_i$ . If  $S \rightarrow 0$  then the achieved reduction is said to be *highly sustainable*. A visual representation of SDR compared to peak minimization is depicted in Fig. 1.

Achieving a sustainable reduction is a hard combinatorics problem because there are numerous combinations of customers to be utilized. An approximation algorithm needs to be employed in order to solve the problem quickly. However the need for real time adaptation is imperative as customer behaviour is hard to predict. The proposed system provides an initial schedule which can be efficiently updated(e.g., **dynamic SDR** –  $SD^2R$ ) to meet the given reduction target.

In this paper we address these issues and present an automated system used to suggest and enact  $SD^2R$  schedules. The system has the ability, based on real-time consumption data, to quickly adapt in order to compensate for any prediction inaccuracies. It is currently being used as a support tool in the controlled micro-grid of the University of Southern California (USC) campus by the Facility Management Services (FMS).

## II. RELATED WORK

There have been numerous attempts to deal with consumption demand. Utility providers can either compensate by buying extra power at high prices [5] or employ DR strategies. The latter is a well known concept divided into two categories which include direct control and voluntary participation [2]. Arguments in favour of both techniques have yield several different solutions driven by specific use cases. These address residential buildings [6], offices [10] as well as large industrial facilities [19] and data centers [14]. In this paper we focus on the USC campus microgrid which includes a mixture of various building types including residential, offices, libraries, and mixed spaces. Our work is based on directly controlling the building equipment to achieve a specified reduction target.

Previous work in the domain of load manipulation includes attempts to minimize peak demand by shifting it to less busy hours of the day [13], [22] or optimizing load consumption while minimizing costs from the customer perspective [16]. The above methods rely on cooperative customer action and has the main drawback of not ensuring the sustainability of the DR event. In contrast, we aim on finding a group of customers the participation of who ensure the sustainability of the DR event. 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 [9].

A significant contribution in the domain of ADR is the DRAS (Demand Response Automated System) [11]. This system is designed with the goal of eliminating human intervention when scheduling DR events. It is used to broker the communication between the utility providers and directly controlled equipment. After a bidding procedure the clients to participate in the DR event are selected. Our system is designed to work in cooperation with DRAS by automating the selection of equipment/buildings to be controlled during DR for each participating client. In doing so it also provides an estimated reduction based on that selection.

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

## III. INTEGRATED PLATFORM OVERVIEW

The FMS at USC operates an integrated platform for ADR which manages the DR activities on the USC campus. In addition the micro-grid is part of the LA DWP Smart Grid Regional Demonstration project. The platform comprises of several independent modules: the *DRAS server*, the *Integrated Building Control* (IBC), and the *Policy Engine*. Requests for DR events come to the IBC from the DRAS server in the form of messages following the OpenADR specifications [3]. Based on them IBC sends to the Policy Engine an XML message containing the list of buildings and strategies to be used, the targeted curtailment value in kWh and the DR event period. The engine will reply with a subset of building-strategy pairs




Fig. 2: DR event timeline.




Fig. 3: Overview of the existing integrated platform for ADR.

which IBC can use in the DR. A number of HVAC strategies and their combinations can be currently used (cf. Fig. 5): Global Zone Temperature Reset (GTR), Variable Frequency Drive Speed Reset (VFD) [18] and equipment Duty cycling. Based on the received list the IBC will send OpenADR command and control messages to the building automation system currently installed in 36 on campus buildings. Power consumption for each of the buildings is monitored at 1 minute intervals and aggregated for ease of use in 15 minute kWh values which are stored on an FTP database for the Policy Engine to use. The updates are in real-time in order to offer the engine the most accurate view of the system and to allow it to efficiently adapt the DR strategies to recent changes in the buildings' consumption patterns.

The IBC constantly communicates with the engine in order to sustain and achieve the curtailment target. For this it relies on three types of messages: *Far*, *NEAR*, and *DURING*. The first two are sent immediately following the DRAS request, respectively 15 minutes ahead of the DR event in order to decide the initial set of building-strategy pairs to be used. The *NEAR* message is designed to capture any possible changes in the set determined by *Far*, e.g., due to baseline adjustments. The *DURING* message is sent on an hourly basis during the DR event to update the building-strategy pairs based on the so far achieved curtailment and an estimated achievable target for the remaining period. Figure 2 shows a typical DR event.

## IV. POLICY ENGINE MODULE

The policy engine's functionality is split in two main parts dealing with the communication with the FMS's integrated building control and with the actual selection process. Figure




Fig. 4: Policy Engine components & data flow.

4 shows the main components and their interaction. The engine is deployed as a REST service [8] which is periodically polled by FMS. Two independently running components deal with the communication: the *Data clean-up* and *Message Processing*. The former is responsible for retrieving real time and historical consumption data that will be used later in the selection process. Since this data comes directly from the buildings, it is not cleansed and filtered by an MDM (Meter Data Management) component. As a result our component takes the role of a simplified MDM. The latter is responsible for receiving and processing XML based request messages as well as constructing and transmitting a response message. They include information about the duration of the DR event, the participating buildings along with the corresponding strategies that should be considered and the curtailment target.

The selection process is handled by a single *Selection Policy* component. The component consults a curtailment matrix containing the predicted reduction for the whole day for each building-strategy pair. The values in this matrix are computed based on the difference between the buildings' consumption in the absence of DR as predicted by a baseline and the actual consumption during the DR events. Given that numerous ADR strategies can be used each line corresponds to a building-strategy pair and contains 96 values corresponding to 15 minute interval readings. Depending on the baseline the matrix values can vary significantly impacting as a result the selection process and its accuracy.

Because we target  $SD^2R$  we require the selection module to be able to periodically adapt the selection of buildings to the reality provided by the real time consumption data. For this the component will periodically (i.e., whenever FMS sends a *DURING* message) estimate, based on real time consumption values, the achieved curtailment and the remaining target. The remaining target will then be used to make a reselection of building-strategy pairs to be used next. Two choices exist: add/remove building-strategies and modify the strategies of already used buildings. Given building specific mechanical and practical constraints there are limitations on the strategies which can be selected to replace existing ones. Figure 5 depicts the state diagram of the allowed changes for the case of the USC micro-grid.




Fig. 5: Allowed transitions between DR strategies. The strategies are techniques used to control directly HVAC units and include: Global Zone Temperature Reset (GTR), Variable Frequency Drive Speed Reset (VFD), Duty Cycling (Duty).

#### A. Selection Algorithm

The selection algorithm provides an approximate solution to the problem of finding a sustainable reduction schedule [24]. It is formulated based on the change making problem [15]. Buildings(customers) are grouped according to a calculated estimated reduction derived from the curtailment matrix. Using a default value set corresponding to the US coin set (i.e.,  $C = \{1, 2, 5, 10, 25, 50, 100\}$ ) a number of corresponding bins are constructed. The bin ranges are scaled to incorporate the estimated reduction values. The scaling factor(unit value =  $u$ ) is computed using a variety of developed heuristics which are out of the scope of this paper. The buildings are then distributed into their corresponding bin. At this point each building is paired with the specific strategy that minimizes the error from the corresponding bin value (i.e., the upper bound of the bin range). This ensures that a strategy with the maximum reduction of the lowest deviation between consecutive intervals is selected. This strategy will provide us with a sustainable reduction. The building-strategy pairs are sorted per bin based on the Euclidean distance from the specified reduction denoted by the bin value. Finally these pairs are combined to achieve a given reduction through greedy indexing of the constructed bins.

In Algorithm 1 an overview of the selection procedure is presented. The reduction estimates used to group the buildings are called representatives. The representatives are utilized to select a suitable scaling factor called unit value. The given reduction to be achieved is denoted as  $M$ . The default value and the scaled value set are denoted as  $\tilde{c}$  and  $c$  respectively. The algorithm then continues to make a selection in the way described previously. Findings that verify the accuracy as well as the efficiency of the algorithm are part of previous work presented in [24].

---

**Algorithm 1** Change Making Scheduler

---

**Input:** Curtailment vectors for each customer-strategy pair.  
**Output:** List of customer-strategy pair.

```

1: representatives  $\leftarrow$  customers.representatives()
2: v  $\leftarrow$  calc_unit_value(representatives)
3: for i  $\leftarrow$  0 to c.size do
4:   c[i]  $\leftarrow$   $\tilde{c}$ [i]  $\cdot$  v
5: end for
6:  $\tilde{M} \leftarrow M/v$ 
7: bins  $\leftarrow$  distribute( $\tilde{c}$ , customers)
8: for i  $\leftarrow$  1 to buckets.size do
9:   sort(bins[i])
10: end for
11: for i  $\leftarrow$  c.size to 1 do
12:   j  $\leftarrow$  0
13:   while  $M - \tilde{c}$ [i]  $\geq 0$  do
14:     while c[j]  $-$  bins[i].customer[j].reduction  $\geq 0$  and
         j  $\leq$  bins[i].length do
15:       result.add(bins[i].customer[j])
16:       c[j]  $\leftarrow$  c[j]  $-$  bins[i].customer[j].reduction
17:       j  $\leftarrow$  j + 1
18:     end while
19:      $\tilde{M} \leftarrow \tilde{M} - c[i]$ 
20:   end while
21: end for
22: return result

```

---

## V. EXPERIMENTS & RESULTS

We conducted a series of experiments using consumption data acquired from the USC campus micro-grid. We focused on comparing the predicted reduction to the actual achieved reduction. We also conducted experiments to test the effectiveness of the adaptation mechanism, i.e., the sustainability of the curtailment (cf. Eq. 1). Finally we present the results which measure the execution time of our *Policy Engine* module.

### A. Experimental Setup and Evaluation Methods

The dataset we used consists of consumption data for 33 buildings for a period of one year. It also includes the observed consumption during DR events where building equipment are directly controlled using different strategies. The time series data are sampled at a fixed rate of 15-min granularity. The consumption data during regular operation was used to establish the baseline consumption for each building. This baseline was predicted using Southern California Edison (CASCE) [1]. The actual consumption during a DR event was predicted using the ARIMA model [12]. The reduction achieved for each building was measured as the point difference of these two time series. The designated period of the DR event was 1-5PM which is the period of peak demand at USC [4]. However the curtailment matrix consists of 96 intervals which represent a whole day. Hence a schedule can be provided for an arbitrary period of the day given the existence of reduction data. For the experiments we used a variety of reduction targets ranging from 100 to 3,000 kWh. In the experiments we used




Fig. 6: Comparison of predicted to achieved reduction using an increasing window of previous events.




Fig. 7: Comparison of predicted to achieved reduction with stable window=2 size.

observed consumption information from previous consecutive events to predict the upcoming events. The events to which we compared our prediction are presented in consecutive order according to the date they were submitted.

### B. Achieved Reduction

The effectiveness of the policy engine was evaluated by comparing the predicted reduction with the actual achieved reduction. The predicted reduction is calculated using a window of continuously increasing size that includes the previous DR events. The consumption of these events is predicted and compared with the baseline of the next immediate event. The actual reduction is the observed reduction of the next DR event not included in the prediction.

The results of the experiment can be seen in Fig. 6. The diagonal line represents the ideal scenario where the prediction is 100% accurate and matches the observed actual curtailment.

Points above the ideal line represent a higher actual reduction than the predicted one and in reverse points that are below the line represent a lower actual reduction than the predicted one. It can be seen that the actual reduction in most instances is lower than the predicted reduction. However the reduction target is always very closely approximated. This means that the difference between the actual and the predicted reduction is




Fig. 8: Comparison of predicted to actual reduction using real time DR adaptation.

due to the prediction accuracy. In fact the calculated reduction will present a higher error since we need to predict both the baseline and the consumption during DR. This increases significantly the overall error of several building combinations.

Intuitively we can suspect that external events (e.g., weather related factor) that characterize older DR events affect the behavior of each building and in turn the accuracy of the prediction method. This can be avoided when smaller historical windows are being used for the prediction. In this case the predicted reduction would be influenced only by adjacent events in time. We performed experiments using a stable window of size 2. This means that for the same events we utilized only the previous two DR events to determine the reduction. The results are presented in Fig. 7. It can be observed that in most cases the predicted reduction matched more closely the actual reduction. However the results are far from the ideal case for the last events. In our problem the accuracy of the prediction methods is very important as the selection procedure assumes accurate data to make a calculated suggestion.

As it was mentioned in Sect. IV the system has the ability to adapt on demand. This means that it can detect through real time monitoring that the achieved reduction is lower than expected and suggest a new schedule to make up for the missing reduction. This feature can help increase the achieved reduction overcoming the induced prediction error. We performed experiments using the difference between the predicted and actual reduction to test the adaptation mechanism. In Fig. 8 we present the results of these experiments. It can be observed that the quality of the results increased. The adaptation mechanism successfully compensates for the prediction error and produces better results in terms of the actual achieved reduction.

### C. Sustainability

As we stated in Sect. I a sustainable reduction is very important if we want to keep the curtailment across a longer time frame. Next we evaluate the sustainability of the provided solution using Eq. 1 which measures the slope variations between consecutive points. It ranks the sustainability of the provided solution independently of the targeted reduction. The




Fig. 9: CDF depicting the level of sustainability of the predicted reduction for a variety of reduction targets.




Fig. 10: CDF depicting the level of sustainability of the achieved reduction for a variety of reduction targets.

sustainability was measured for both cases of the window prediction methods described previously.

In Fig. 9 we present the level of sustainability of the predicted reduction. It is observed that the reduction achieved is highly sustainable since almost 90% of the solutions provided have a slope variation less than 1.

Although the results of the predictions are promising it does not follow up for the sustainability of the actual achieved reduction. In Figure 10 we present the CDF plot of calculated level of sustainability for the observed reduction. The graph shows a clear increase in the slope variations indicating a sustainability lower than the predicted one. This is a direct result of the low prediction accuracy which cannot be solved by on demand DR adaptation unless a highly accurate baseline is used. This is because adaptation is based on the prediction which provides a reduction more stable than the actual observed reduction. Achieving a sustainable reduction adaptation will have to utilize information about the real time evolution of the observed consumption data and make a decision based on them.

### D. Execution Time

Synthetic data were used to measure the execution time of each component of the selection procedure. Here we use 10 strategies per building which is a realistic scenario drawn from the available strategies of the real dataset used in our previous




Fig. 11: Percentage utilization of the overall execution time for each individual step of the selection procedure by the policy engine.

experiments. The results are presented in 11. It is clear that the execution is independent of the reduction target. It depends only on the size of the dataset. Roughly 90% of the execution time is spent on retrieving the reduction information. This includes updating the consumption values with the latest data from the FMS FTP server. As the size of the input increases with the number of buildings this becomes a serious bottleneck. The actual execution time however is relative low unless taking into account the retrieval time. The execution time can be further improved by pre-computing information which is regularly used in the selection procedure. This includes the representatives or even the constructed building groups, unless the curtailment matrix is frequently updated.

## VI. CONCLUSION

In this paper we presented an integrated platform for sustainable DR. The DR events are handled by providing a highly sustainable schedule of participating customers given a specific reduction target. It has been shown that the presented system has the ability to overcome the prediction error when utilizing the real time adaptation mechanism. However the adaptation cannot currently ensure a sustainable reduction because it is based only on the predicted reduction values. This indicates that the prediction methods are the ones limiting our system's effectiveness.

In our future work we need to deal with two aspects. The first is related to the efficiency of the selection procedure as our implementation will not scale well for a larger number of households. Moreover it is imperative to deal with the inaccuracies of the prediction models to achieve a highly sustainable DR event.

## ACKNOWLEDGMENT

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).

## REFERENCES

- [1] 10-day average baseline and day-off adjustment. Technical report, Southern California Edison, 2011.
- [2] 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.
- [3] O. A. Alliances. Open automated demand response, 2014. URL: <http://www.openadr.org/> (accessed May 23, 2014).
- [4] S. Aman, M. Frincu, C. Charalampou, 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).
- [5] 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.
- [6] 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.
- [7] A. Z. Faza, S. Sedigh, and B. M. McMillin. *Reliability modeling for the advanced electric power grid*. Springer, 2007.
- [8] 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.
- [9] M. Frincu and Z. Gima. A bottom-up approach to sustained curtailment and comfort for controlled demand response. In *IEEE Conference on Technologies for Sustainability*, page accepted, July 2014.
- [10] 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.
- [11] E. Koch and M. A. Piette. Architecture concepts and technical issues for an open, interoperable automated demand response infrastructure. *Lawrence Berkeley National Laboratory*, 2008.
- [12] S. Makridakis and M. Hibon. Arma models and the box-jenkins methodology. *Journal of Forecasting*, 16(3):147–163, 1997.
- [13] F. Mangiatordi, E. Pallotti, P. Del Vecchio, and F. Lecce. Power consumption scheduling for residential buildings. In *Environment and Electrical Engineering (EEEIC), 2012 11th International Conference on*, pages 926–930, May 2012.
- [14] K. Mares. Demand response and open automated demand response opportunities for data centers. *Lawrence Berkeley National Laboratory*, 2010.
- [15] S. Martello and P. Toth. *Knapsack problems*. Wiley New York, 1990.
- [16] 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.
- [17] K. Moslehi and R. Kumar. A reliability perspective of the smart grid. *Smart Grid, IEEE Transactions on*, 1(1):57–64, 2010.
- [18] N. Motegi, M. A. Piette, D. 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).
- [19] 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.
- [20] 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.
- [21] 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.
- [22] 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.
- [23] P. Siano. Demand response and smart grids—a survey. *Renewable and Sustainable Energy Reviews*, 30:461–478, 2014.
- [24] V. Zois, M. Frincu, C. Chelmis, M. R. Saeed, and V. Prasanna. Efficient customer selection for sustainable demand response in smart grids. In *Green Computing Conference (IGCC), 2014 International*, page accepted, November 2014.