Efficient Customer Selection for Sustainable Demand Response in Smart Grids

Vasileios Zois†, Marc Frincu‡, Charalampos Chelmis‡, Muhammad Rizwan Saeed‡, Viktor Prasanna‡
University of Southern California
Email: {vzois,frincu,chelmis,saeedm,prasanna}@usc.edu
Department of Computer Science†, Department of Electrical Engineering‡

***Abstract*—Regulating the power consumption to avoid peaks in demand is a common practice. Demand Response(DR) is being used by utility providers to minimize costs or ensure system reliability. Although it has been used extensively there is a shortage of solutions dealing with dynamic DR. Past attempts focus on minimizing the load demand without considering the sustainability of the reduced energy. In this paper an efficient algorithm is presented which solves the problem of dynamic DR scheduling. Data from the USC campus micro grid were used to evaluate the efficiency as well as the robustness of the proposed solution. The targeted energy reduction is achieved with a maximum average approximation error of $\approx 0.7\%$. Sustainability of the reduced energy is achieved with respect to the optimal available solution providing a maximum average error less than 0.6%. It is also shown that a solution is provided with a low computational cost fulfilling the requirements of dynamic DR.**

Keywords: scheduling, optimization, sustainability, demand response, real time, change making

I. Introduction

As demand for power increases so does the complexity involving safe and reliable energy distribution [14]. Recent innovations in power grids have provided smart tools, to help utility providers monitor and predict power demand [3]. Demand Response (DR) is a well known method employed by energy providers to control demand [2]. It is used to find the equilibrium between energy production and consumption which ensures reliable energy distribution.

Energy providers use different techniques to manipulate customer load which include: direct control [6], price incentives [13] as well as voluntary participation [3]. Their primary goal is to shift the energy consumption to different periods of the day thus eliminating peaks in demand. This can be achieved by employing a controlled energy reduction schedule of participating customers known as a DR event. Although this is a well accepted technique it might not produce the desirable results as it may be the case where peaks in demand are shifted to other periods of the day. An acceptable solution would be to evenly distribute consumption across a specified period of the day. This can be achieved by employing a sustainable DR event. DR event is said to be sustainable if it achieves consistent energy reduction for the scheduled time frame. It can be formally defined with:

$$
\begin{aligned}
&S = \frac{1}{n} \cdot \sum_{j=2}^{n} |L_j| + max(L_j) - min(L_j) \geq 0 \ \text{ given:} \\
&R = \{[x_1, f(x_1)],[x_2, f(x_2)], \dots [x_n, f(x_n)]\} \\
&\forall j \in [2, n],\ L_j = \frac{f(x_j) - f(x_{j-1})}{x_j - x_{j-1}}.
\end{aligned} \tag{1}
$$

R is the observed reduction described by a list of pairs denoting timestamp x_i and the reduction value of each observation during DR. The level of sustainability is measured with S. If $S \to 0$ then the achieved reduction is said to be *highly sustainable*.

Sustainable energy reduction is a hard combinatorial optimization problem. Customers deployed on top of the power grid are inherently unpredictable making the provision of a sustainable DR event even more difficult. This supports the need for an automated procedure that efficiently reacts and enacts dynamically(**dynamic DR**), to ensure the fulfillment of the DR requirements.

This paper addresses the problem of providing a DR schedule of participating customers based on past observations on their consumption behavior. The contributions of this paper can be summarized as follows:

- The proposed algorithm adheres to the time constraints of dynamic DR, providing a solution with low computational cost.
- The computed solution achieves a highly sustainable energy reduction in compared to the available optimal solution.
- The provided solution utilizes the minimum number of customers necessary to achieve the target energy reduction minimizing the level of intrusiveness on customers.

II. Related Work

The research presented so far deals with residential cases [9] mostly concerned with household appliances [11]. This can be considered as an unrealistic scenario as utility providers cannot sustain individual information for each unique appliance. The solution we propose needs to deal only with aggregated consumption observation for each individual customer. The observations are associated with different strategies which include direct control or voluntary participation through the use of incentives. The effectiveness of each strategy is embedded

978-1-4799-6177-1/14/$31.00 © 2014 IEEE

on the observed reduction values from past DR information which are considered to make a calculated decision when scheduling a DR event.

The problem of minimizing peak to average consumption ratio has been well studied. Several approaches have been proposed including dynamic programming [7], linear programming [15] and particle swarm optimization [17]. The primary goal of each approach was to directly control appliances which accounted for a high percent of the overall energy consumption. Although the initial approaches did not consider customer comfort the latter did. Our solution deals with customer comfort although not directly. It has been shown in related work that the reduction strategies can be optimized to maximize customer comfort [8]. The selection procedure considers customer behavior through consumption observations of past DR events assuming the customer comfort factor embedded into the data.

Alternate approaches on the demand regulation problem include game theoretic formulations constrained by real time pricing policies [5] or customer comfort levels [4]. Although these techniques rely on cooperative action and can be implemented in a distributed way they cannot ensure elimination of demand peaks. Customers' conflicting needs might still shift demand peaks in other periods of the day. Our goal is to provide a procedure that both solves the sustainable DR problem while also has the ability to quickly adapt to undesirable customer behavior.

A similar problem to ours has been studied in [18]. There the authors make a mixed-integer programming formulation and propose three approximate methods to deal with the case of a feeder failure. They assume customer compliance and deal with excess demand by balancing the load between working transformers and customers participating in DR. In contrast to this approach our solution considers only customers to handle the excess energy demand.

III. Problem Formulation

The notation used throughout this section are summarized in Table I. We formulate the problem as follows: we are given a set of n customers. Each one is associated with m available strategies. For each customer-strategy pair there exists a curtailment vector $(r_1, r_2, ..., r_T)$ representing a DR event of size T. The curtailment vector is calculated as the difference between the baseline consumption and the observed consumption during DR of past DR events. The accuracy of the predicted values depend on the prediction models [3] which are out of the scope of this work.

Let R (in kWh) be the targeted energy reduction of a DR event. Our goal is to find a subset of the participating customers to include in the DR event to curtail $\frac{R}{T}$ per interval. In this paper we deal with the simple case which requires for each customer to be paired with a single strategy from the available ones for the whole DR event time frame.The collective achieved reduction by the selected set of customers is described by:

B_{ij}	=	Curtailment Vector(CV) of customer i paired with strategy j.
$B_{ij}(t)$	=	Reduction value at interval t from the curtailment vector of the specified customer-strategy pair.
v	=	Unit value used for scaling original coin set to construct the bins.
$\tilde{C}$	=	Original coin set based on the US denomination. $\tilde{c}_i$ denotes the i-th coin.
C	=	Adjusted coin set for a specified unit value v. c_i denotes the i-th adjusted coin.
x_i	=	Number of customers used from i-th bin.
y_i	=	Number of customers residing in the i-th bin.
U	=	Set of representatives consisting of the reduction per interval of each individual customer.
u_i	=	The i-th representative corresponding to the i-th customer.
M	=	Targeted reduction per interval.
$\tilde{M}$	=	Dollar amount of the specified reduction scaled using v.
K	=	The set of bins being indexed using change making for target M.
$\tilde{K}$	=	The set of bins being indexed using change making for target $\tilde{M}$.
$\max_i u_k$	=	Max representative from the bin of range $(c_{i-1}, c_i]$.

TABLE I: Notation Table

$$\forall t \in [1, T] \ , \ B_{ij} \in A \ , \ \sum_{i=1}^{|A|} B_{ij}(t) \geq \frac{R}{T}. \tag{2}$$

Here A is the set of customer-strategy combinations selected to participate in the DR event. We aim to maximize the potential reduction constrained by minimizing the deviation between consecutive intervals. An energy reduction of lower value per interval can be sustainable, although it will not collectively achieve the given target. Our algorithm will provide a solution of the maximum possible sustainable energy reduction.

Our problem can be expressed as Integer Linear Programming Problem(ILP) [16]. It is known that ILP is $\mathcal{NP}$-hard. This means that for a small number of customers the solution is achievable with a low computational cost. The same is not true for a large number of customers. It can be also formulated as a 0-1 knapsack problem [12]. However dynamic programming does not fit well to the definition of this problem as we are dealing with real numbers as weights.

IV. Proposed Algorithm

The proposed solution is based on the change making problem [12] which addresses the question: *how a given amount of money can be made using the least amount of coins?*. The coins are the available customer-strategy pairs and their value is the predicted reduction per interval. Following this definition we group the customers into bins which are differentiated by their bin value. Each bin has a specific range defined by the coin set and the specified scaling factor called **unit value** v. The coin set(US coin set) provides an optimal solution in terms of the number of coins utilized. The bin range defined by coin $\tilde{c}_i$ will be $(\tilde{c}_{i-1} \cdot v, \tilde{c}_i \cdot v]$. The participating customers are grouped in the same bin if their reduction estimate, which we call **representative**, falls in the corresponding bin range. After the distribution step, customers are paired with the strategy that approximates most closely the bin value(upper bin range). The level of approximation is calculated using the euclidean distance of the corresponding curtailment vector from the bin value. The target reduction $M = \frac{R}{T}$, is normalized using the same unit value ($\tilde{M} = \frac{M}{v}$) giving the amount to be constructed. Finally the bins are indexed greedily from largest to smallest to create the given target. A pseudo-code describing the above procedure is presented in Algorithm 1.

Algorithm 1 Change Making Scheduler

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

$representatives \leftarrow customers.representatives()$
$v \leftarrow calc\_unit\_value(representatives)$
for $i \leftarrow 1$ **to** $c.size$ **do**
  $c[i] \leftarrow \tilde{c}[i] \cdot v$
end for
$\tilde{M} \leftarrow M/v$
$bins \leftarrow distribute(c, customers)$
for $i \leftarrow 1$ **to** $buckets.size$ **do**
  $sort(bins[i])$
end for
for $i \leftarrow c.size$ **to** 1 **do**
  $j \leftarrow 0$
  while $M - \tilde{c}[i] \geq 0$ **do**
    while $c[j] - bins[i].customer[j].reduction \geq 0$ **and** $j \leq bins[i].length$ **do**
      $result.add(bins[i].customer[j])$
      $c[j] \leftarrow c[j] - bins.customer[j].reduction$
      $j \leftarrow j + 1$
    end while
    $\tilde{M} \leftarrow \tilde{M} - c[i]$
  end while
end for
return $result$

A correct distribution of customers will always produce customer-strategy pairs that fit the constructed bin values. This will enable the change making algorithm to greedily produce a combination of pairs that achieve the overall target R. The representatives as well as the chosen unit value affect the approximation accuracy. A unit value is suitable if we can find at least one customer-strategy pair in each bin that achieves the bin value for all intervals in a DR event.

A. Representatives

The representatives are used to initially distribute each customer to a specified bin. If the distribution is done correctly then a bin value can be achieved by individual customer-strategy pairs. This would yield an accurate combination of bins which are going to be used to make the reduction target while taking advantage of the change making formulation.

Three methods were implemented for calculating the representatives. Their main properties are simplicity and low complexity. They are based on examining the curtailment vectors of each customer separately. The first method which is denoted as MAX, selects the maximum value from all intervals of all possible strategies for a customer. The second method called AVG, calculates an average of the reduction considering all intervals of all strategies. The third and final method called MAVG, selects the maximum value from the average reduction of each strategy available to a customer.

B. Unit Value

Selecting a suitable unit value must be a compromise between accuracy and execution time. We developed five different heuristics with different level of complexity and accuracy.

*a) **Greedy:*** The greedy approach uses as unit value the target reduction per interval($v = M$). This choice converts our problem to a simple 0-1 knapsack problem since all the customers under consideration will be in the first bin and selecting a subset of them will be constrained by the maximum weight/reduction target. Although simple this technique is known to be suboptimal when using a greedy algorithm. Motivated by this we developed simple heuristics that take into account reduction data patterns to provide a suitable unit value.

*b) **Minimum Goal Accumulated Bin Error (MGABE):*** This heuristic aims at selecting as unit value the representative that creates an adjusted coin set which minimizes the following two quantities. This is described in:

$$\min_{u_j \in U} \sum_{\forall c_i \in C} \Big(c_i - (\max_i u_k)\Big) + \Bigg(M - \sum_{\forall c_i \in K} (c_i \cdot x_i)\Bigg). \quad (3)$$

The first quantity refers to the accumulated error produced by the deviation of the estimated reduction of the customers inside a bin from the bin value. The second quantity is the deviation created by the difference between the target reduction and the selected combination of bin values to achieve M.

*c) **Minimum Accumulated Average Bin Error (MAABE):*** Another method is to consider minimizing only the quantity connected with the bin values. This heuristic considers the average reduction of the representatives that exist in each bin and tries to minimize their difference from the bin value. The representative that constructs the adjusted coin set which minimizes this quantity is selected. The heuristic is described formally with:

$$\min_{u_j \in U} \sum_{\forall c_i \in C} \Bigg(c_i - (\frac{1}{y_i} \cdot \sum_{\forall u_k \in (c_{i-1}, c_i]} (u_k))\Bigg). \quad (4)$$

*d) **Minimum Coin Error (MCE):*** Both of the previously presented heuristics consider all the bins when calculating the overall estimated error. It might be more advantageous to focus on specific bins which are related to the given target. This could potentially lower the computational cost. It can also provide a solution unaffected by reduction patterns irrelevant to our target. This heuristic estimates $\tilde{M}$ considering all the representatives as a potential unit value v. Using this information it pre-calculates the bins which are to be used in the end result. Then it selects v from the representative set, that minimizes the accumulated error induced only by the selected bins. It is formally described by:

$$\min_{u_j \in U} \sum_{\forall c_i \in \tilde{K}} \Big(c_i - (\max_i u_k) \Big). \tag{5}$$

This heuristic, is similar to the first part of Eg. 3. However it is calculated for only selected bins depending on the targeted reduction.

e) ***Unit Data Trend (UDT):*** A major drawback of selecting a unit value from the representative set is assuming that it contains a suitable one. In many cases this is not true as such a scenario depends strongly on the data distribution. UDT was developed to avoid this limitation. It uses the initial coin value set to form the bin values. The customers are distributed into their corresponding bins based on their estimates and then the following weighted average is calculated:

$$v = \frac{\sum_{i=1}^{7} (x_i \cdot \frac{\max_i u_k}{\tilde{c}_i})}{\sum_{i=1}^{7} x_i}. \tag{6}$$

The result from this calculation is used as the unit value. The weights represent the number of customers selected from bin i. The value corresponding to a specific weight is the maximum reduction observed in that bin divided by the original coin value of the bin.

C. Complexity

The size of the input is defined as $n \cdot m \cdot t$ where n is the number of customers, m is the number of available strategies and t is the number of intervals. A realistic scenario is having a constant number of intervals and strategies. From this we can safely presume that the size of the input is measured in the number of customers. The complexity is affected by specific initialization steps common to all of the developed heuristics.

The common operations include calculating the representatives (line 1), distributing customers into bins (line 7), sorting customer-strategy combinations in each bin according to their deviation from the bin value (line 9) and selecting the bins to index through change making (line 11). They present complexity of $O(n)$, $O(n)$, $O(n\ log\ n)$ and $O(n)$ respectively.

All the heuristics except for the greedy, add an extra computational cost. MGABE and MCE have the same complexity of $O(n\ log\ n)$ since both need to examine the whole representative set and calculate the error produced for each bin value. MAABE has the highest complexity since it needs to consider all the representatives and calculate the average value for each bin. Finally UDT calculates the weighted average an operation which is of linear complexity ($O(n)$).

V. Experiments & Results

The experiments evaluate the accuracy of the provided schedule in terms of the overall achieved reduction as well as the sustainability of the reduced consumption.

A. Data-set, Experimental Set up & Evaluation Methods

The data-set utilized contains consumption information from the smart grid deployed at the University of Southern California. The consumption measurements were sampled at 15-min granularity. Predicting consumption during DR was done using the ARIMA model [10]. The baseline consumption was determined by using the Southern California Edison(CASCE) [1] model. The data-set consists of 33 buildings which participated in 380 DR tests. The participating buildings include classroom, offices, dorm rooms as well as meeting halls. The DR events were conducted between 1:00 - 5:00 PM as this was the period of the day with the highest observed demand [3]. Although the strategies employed for each building are important for the overall achieved reduction they do not have any significance for the selection procedure. This is because the algorithm is just concerned about the reduction values which could be induced by any combination of strategies. For a more detail description the reader can reference [8].

The experiments consist of two parts. In the first part we evaluate the heuristics using all available buildings. In the second part we randomly select a subset of our original building set to test the robustness of each heuristic. Each building-strategy pair has equal probability to be chosen. The metrics used are the absolute percentage error in the first case and MAPE in the second case.

B. Overall Reduction

In Fig. 1 the absolute percentage error for various reduction targets is presented. The comparison is between the different techniques used to calculate the representatives. MAX surpasses for most heuristics the other two techniques. AVG and MAVG underestimate the potential reduction which results in placing buildings to bins of lower value limiting their maximum reduction. It has little effect on the greedy technique since it does not rely on combining multiple bins that face the previous issue. Reduction targets in the range of 500 - 1500 kWh present a lower accuracy error since they utilized a combination of bins. If the target is higher or lower respectively it relies on specific building pairs which are few and most of the time do not fit well the necessary reduction.

The results of Fig.1 denote overall MAX as the most suitable technique for calculating the representatives. In Fig 2 we compare the accuracy of each heuristic giving the CDF of the normalized absolute percentage error for MAX technique. The accuracy of the greedy heuristic is very promising although in several occasions is not consistent. The high accuracy is a result of the lower reduction values which fill the remainder to the target value as we greedily iterate from higher to lower reduction values. In several occasions such a combination will not be available yielding a lower accuracy.

MAABE provides the most consistent results in terms of accuracy. The unit value selected will balance the error distribution across all bins favoring bins of bigger value. This in combination with the change making algorithm, will provide an accurate solution since the overall approximation error

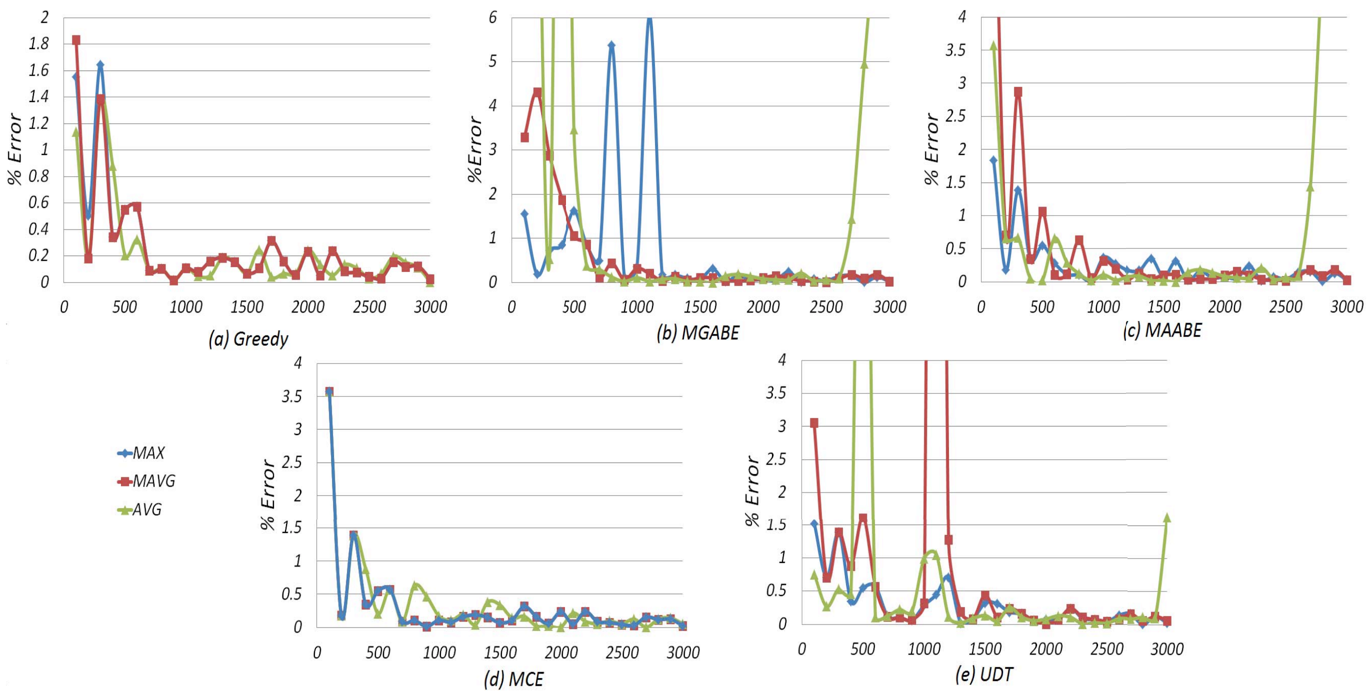

Fig. 1: The effect of representatives on the overall approximation error for a variety of targets for each one of the developed heuristics.

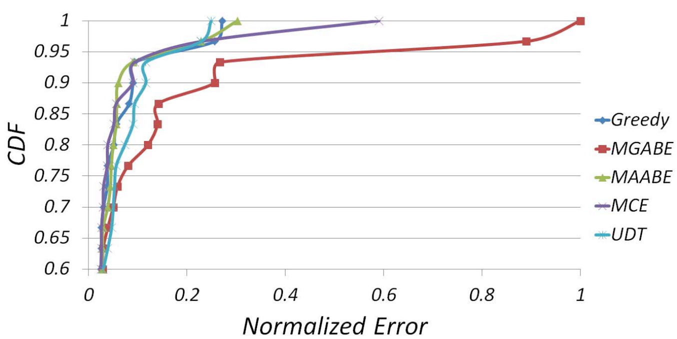

Fig. 2: CDF of absolute percentage error using MAX technique for the different reduction targets.

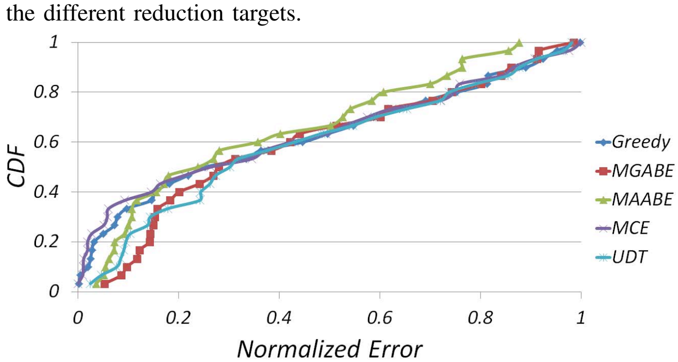

Fig. 3: Overall approximation error using randomly generated building-strategy sets.

will be reduced for larger reduction quantities denoted by the corresponding bin values.

MGABE is least effective maintaining larger inconsistencies in the observed approximation error. It is due to the minimization procedure which aims at fulfilling conflicting goals. A minimized error induced by the relative bin values does not

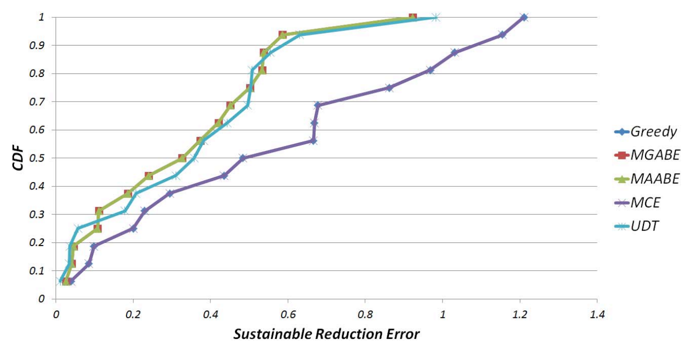

Fig. 4: Approximation error from the achieved sustainable load reduction in comparison to the optimal solution.

necessarily include a combination of them that minimize the deviation from the specified target.

MCE was developed as a local search heuristic. It provides a solution with low computational cost by focusing only on bins that are more likely to be used in the end result. The effectiveness of MCE can be evaluated better on a search space dense with potential solutions. It is shown in Fig 2 that the error evolves similarly to MAABE. Although there are occasions were it is higher. This is a direct result of its local search characteristics which are induced when examining only one representative from each bin, thus underestimating the potential combined reduction when more than one buildings are utilized.

UDT follows a similar procedure to that of MCE focusing on specific bins to produce an approximate unit value. It considers all the customer-strategy pairs needed for the given target reduction. The deviation between the customers under consideration will affect the unit value. In essence the constructed bin values will represent best the reduction with the largest weight or the one with the highest reduction value (Eq. 6). In such a scenario the combination of bins used will consist of some which might not yield the needed reduction as described by the bin value.

The presented results do not clearly distinguish a specific heuristic as the best one. In Fig. 3 we present the results of the experiments using the randomly generated building subsets. The results follow same pattern as before although the advantages and the disadvantages of the heuristics used are more clear. It can be sheen that for lower reduction targets all the heuristics produce a higher error than the greedy approach. This is a result of the change making algorithm which focuses on lower bins to *pay* for the reduction needed. In case there are less than the necessary buildings in them the approximation accuracy will be low. However for larger targets the accuracy of the heuristics is equal to or surpasses that of the greedy approaches. This is direct result of the change making combination which spreads to multiple bins and avoids the problem appearing for the lower reduction targets.

C. Reduction Sustainability

The sustainability achieved by each heuristic was evaluated against a given target of 3,0000 kWh. This choice was made to evaluate the robustness of the solution against a diverse combination of buildings utilized in the provided schedule. Achieving a sustainable reduction is not related to minimizing consumption. Our goal is to evenly distribute consumption for a specified period of the day. This will ensure the reliable operation of the power grid.

The heuristics that produce a highly sustainable reduction schedule are the ones that focus on finding building-strategy pairs that represent well each bin value. Those include MGABE, MAABE and UDT as the results of Fig 4 show. Although MCE is based on the same notion it does not produce similar results. This is caused by optimistically assuming that each bin contains the right amount of building-strategy pairs needed to be utilized for the given target. The greedy approach will be the worst in terms of sustainability as it does not consider matching the bin values when making a selection.

VI. Conclusion

In this paper we focused on solving the problem of Dynamic DR scheduling. We presented an approximate algorithm that achieves a sustainable energy reduction, while providing a low accuracy error with respect to the optimal solution. The provided method fulfills the requirements of dynamic DR, as it provides a solution with low computational complexity. The heuristics we developed differ in terms of accuracy for the achieved reduction, as well as the achieved level of sustainability. MAABE achieves the best accuracy but is computationally expensive. UDT and MCE can be used instead to lower the computational cost while maintaining an increased accuracy as the input size increases. MAABE is suitable for relatively small data-sets (e.g small industrial zones, university campuses) where solutions are few. UDC and MCE reduce the computational cost and should be used on large data-sets (e.g large residential neighborhoods) where it is expected for the search space to be dense with potential solutions.

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). We also want to thank for their initial contribution Yogesh Simmhan and Rongyang Liu.

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] 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).

[4] 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.

[5] 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.

[6] W.-C. Chu, B.-K. Chen, and C.-K. Fu. Scheduling of direct load control to minimize load reduction for a utility suffering from generation shortage. *Power Systems, IEEE Transactions on*, 8(4):1525–1530, Nov 1993.

[7] A. I. Cohen and C. Wang. An optimization method for load management scheduling. *Power Systems, IEEE Transactions on*, 3(2):612–618, May 1988.

[8] 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.

[9] N. Gatsis and G. Giannakis. Cooperative multi-residence demand response scheduling. In *Information Sciences and Systems (CISS), 2011 45th Annual Conference on*, pages 1–6, March 2011.

[10] S. Makridakis and M. Hibon. Arma models and the box–jenkins methodology. *Journal of Forecasting*, 16(3):147–163, 1997.

[11] 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.

[12] S. Martello and P. Toth. *Knapsack problems*. Wiley New York, 1990.

[13] 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.

[14] K. Moslehi and R. Kumar. A reliability perspective of the smart grid. *Smart Grid, IEEE Transactions on*, 1(1):57–64, 2010.

[15] K.-H. Ng and G. Sheble. Direct load control-a profit-based load management using linear programming. *Power Systems, IEEE Transactions on*, 13(2):688–694, May 1998.

[16] C. H. Papadimitriou and K. Steiglitz. *Combinatorial optimization: algorithms and complexity*. Courier Dover Publications, 1998.

[17] 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.

[18] H. Simão, H. Jeong, B. Defourny, W. Powell, A. Boulanger, A. Gagneja, L. Wu, and R. Anderson. A robust solution to the load curtailment problem.