

Discovering the Right Incentives for Demand Response Programs

Marilena Minou, Eugenia Kaskantiri, George D. Stamoulis,
Department of Informatics, Athens University of Economics and Business, Greece
{minou, kaskantiri, gstamoul} @aueb.gr

ABSTRACT

A Demand Response (DR) program can only be effective if it offers to users the proper incentives to participate and thus to modify their energy consumption patterns. In this paper, we focus on DR for residential environments. We propose a learning algorithm that helps the energy provider explore iteratively and discover for each user the minimum acceptable incentives that can motivate him to participate in DR on the basis of DR participation history and of profiling information possibly available. The provider can thus allocate incentives in the way that ensures the highest participation rate with the least possible total incentives, even when little information is available. We also deal with assessing, by means of a simple model, the effect of the provision of recommendations on users' participation in DR. We evaluate our algorithms for incentives' allocation and learning by means of simulations. Our results reveal interesting insights on the impact of profiling information on the allocation of the incentives for DR. The proposed algorithms and the environment implemented, when fed with appropriate values for certain parameters, can be employed to provide approximate evaluation of the performance of DR in practical cases.

Categories and Subject Descriptors

H.4 [Information Systems Applications]: Miscellaneous

Keywords

Demand Response, Incentives, Consumption Pattern, Budget Allocation, Learning Algorithm.

1. INTRODUCTION

DR programs constitute an efficient way to alleviate the peak demand problem in smart electrical-power grids. They encourage electricity end-users to adjust their consumption in response to DR events and signals issued by the energy provider. DR programs have been implemented both in industrial and commercial environments. Their successful penetration, particularly in the residential sector, can result in considerable savings, due to the fact that such environments account for a large portion of the total energy demand. However, the real success of such programs depends on offering adequate incentives for the participation and timely response of users to DR events, especially for critical peak rebate DR programs. Results from DR pilots indicate that the level of discomfort/inconvenience caused to users during a DR event due to modifications in their consumption pattern is a key factor that shapes DR participation; e.g. see the work of project WATTALYST available in [1] and references therein. In principle, users are

assumed to follow a particular consumption pattern according to their preferences. To be encouraged to participate in DR, energy providers offer various types of incentives to compensate users for the inconvenience caused to them. However, estimating the appropriate amount of the incentives needed to engage them to actively participate in DR is considered a major challenge. This is due to the type and amount of information that is necessary for the provider to carry out such an estimation, particularly information relating to demographic and consumption characteristics, such as profile of the household, its total consumption or consumption at the appliance level etc. The analysis of users' consumption patterns to obtain such information is a critical issue. Any request for reduction and/or shifting of power load in order to be successful should be consistent with the type of loads arising in each household, i.e. with the appliances used by each user and the constraints imposed by their operation. This of course implies that the provider, in order to acquire more detailed information on appliance usage and preferences, either employs the appropriate equipment, additionally to the smart meters, e.g. appliance level meters, or invests in different non-intrusive load monitoring (NILM) or load disaggregation systems and algorithms. These algorithms allow for the derivation of detailed information on appliance usage from data on the total consumption that is collected by a smart meter [2] and also for avoiding the additional cost of installing new infrastructure in both the provider and the user sides. While this profiling grants for a better and at once realisation as well as for a possibly accurate assessment of users' participation probability, users cannot be obliged to participate in DR and modify their consumption patterns, but can only be incentivized to reduce or defer consumption. The importance of incentives for successful DR programs is recognized and in fact has motivated several theoretical works in the literature; e.g. [3] deals with a similar problem with that addressed in the present paper. In this context, we propose a learning algorithm that helps the provider discover how to allocate DR incentives to ensure the highest participation rate (even when little information is available) while offering the least possible total incentives for achieving this participation rate.

Our algorithm utilises the available information in order to discover for each user the minimum acceptable incentives that motivate him to participate in DR. Indeed, to stimulate users' participation and/or incite them to follow the provider's requests for load curtailment or shifting, the provider is willing to dedicate a budget for providing incentives. In essence, the provider offers each user a reimbursement for his inconvenience in the form of monetary incentives. We assume that the provider aims to achieve the highest participation rate, which may amount to 100%. We develop a learning approach aiming to iteratively explore and exploit at the same time (in successive DR events) which incentive to offer in the next event based on the current estimates of users' participation probability, so that the highest participation rate is achieved and the least possible incentives for achieving this participation rate are offered by the provider. In fact, in the course of this process the provider may prefer to choose attaining a lower DR participation rate if this is considered more beneficial for him, e.g. according to

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.
e-Energy'15, July 14 - 17, 2015, Bangalore, India

© 2015 ACM. ISBN 978-1-4503-3609-3/15/07...\$15.00

DOI: <http://dx.doi.org/10.1145/2768510.2768519>

the required number of users participating in DR (or to the load that should be curtailed) or to the trade-off between the participation achieved and the total incentives offered. We also develop a simple model to study the impact in users' participation of recommendations. The recommendations are based on the information about users' consumption patterns and can address either the total consumption or the consumption of specific monitored appliances in each household. The evaluation of our work is done by means of simulations. To this end, we also develop a parameterized environment, which (when fed with appropriate values for certain parameters) can be employed for the approximate evaluation of DR performance in practical cases.

2. RELATED WORK

From the very broad literature on DR, our work is related mainly to learning of the DR incentives that should be offered per user and to a smaller extent to non-intrusive load monitoring (NILM). The literature on NILM is already extensive; e.g. see [2] (which was one of the first related articles), and [4] and [5] for some recent works. In this paper, we do not develop an approach for NILM, but we deal with how the provider can benefit from knowledge of consumption data (obtained through NILM) and/or profiling information in order to design more effective DR programs and allocate efficiently the DR incentives' budget. Regarding learning of the necessary DR incentives, [3] develops a related Multi-Armed Bandit (MAB) approach. The objective of that work is very similar to ours. However, in our work we only do learning without resorting to MAB, since at each DR event we offer some incentive to each of the users. Moreover, both in the model of [3] and in ours the user is characterized by a minimum incentive parameter, referred to as cost per unit reduction in [3]. The authors of [3] assume that the DR participation probability of a user is fixed provided that the offered incentive exceeds this threshold, while it equals 0 otherwise. In our model we assume that this probability varies with the incentive offered to the user according to a sigmoid function; see equation (3).

In this paper, we develop a methodology for an energy provider to exploit already available profiling data in order to dynamically discover for each user the minimum monetary incentives that should be offered to him to participate in DR, as well as to assess the impact of recommendations. These are proposed to users based on knowledge extracted from load disaggregation methods, or simply from profiling based on demographic data (see below). Thus, our methodology offers the provider a way of exploring a set of possible incentive allocations to choose from, each of which ensures specific levels of participation attained by offering the minimum total DR incentives using also the available profiling information. Our learning algorithm is iterative; each iteration corresponds to a DR event with simultaneous exploration and exploitation. Our methodology can be easily modified into a budget-limited problem, whereby the budget constraint is satisfied in all iterations.

3. THE MODEL

Consider a set of N users (and corresponding households) that are served by a single energy provider and are eligible for participation in DR. Each user $i \in N$ is characterised by i) a set of demographic characteristics, e.g. size of family, age, etc., ii) a consumption pattern, which is formed according to his needs and preferences and iii) the price of electricity. To simplify our analysis, we assume that each user $i \in N$ can only belong to one of two categories, to which users are categorized on the basis of the above profiling information; the categories are i) elastic (N_1), those who are willing to modify their consumption in order to benefit from a discount and/or reduced energy prices and ii) inelastic (N_2), those who are reluctant to participate in a DR program. How exactly this

categorization is done, falls beyond the scope of our work. Therefore, the provider should offer a higher incentive to inelastic users than to the elastic ones in order to engage them in DR if this is necessary due to the load that should be curtailed. To model user response to the incentives offered for DR, we assume that for each user i there is a minimum incentive value $t_{min,i,j}$ that triggers the user to participate in DR, yet not always, as explained later in detail. Moreover, for each user belonging to category $N_j, j = \{1,2\}$, this minimum incentive is drawn separately from the uniform distribution in the interval $[a_j, b_j], j = \{1,2\}$, which is the same for all users in $N_j, j = \{1,2\}$, while its expected value equals $\overline{inc}_{N_j} = mean(t_{min,N_j}) = \frac{a_j+b_j}{2}$. If the provider knows the classification of users by means of profiling information, then he can offer initially a different DR incentive inc_{N_j} to each category $N_j, j = \{1,2\}$. All users of N_j are then offered inc_{N_j} , which is expressed as

$$inc_{N_j} = \alpha * \overline{inc}_{N_j} \quad (1),$$

where the parameter α is an economic scaling factor that is used to relate the initial incentives to the average of the minimum incentive per category. This case is henceforth referred to as Approach 2. If the provider does not know the classification of users, then he is assumed to offer the same initial incentive to all users. This case is referred to as Approach 1. However, even under Approach 1, we take that the provider does have an estimate of the average \overline{inc} of the minimum incentive over both categories (e.g. by means of some profiling or historical information regarding user DR behaviour),

which in fact equals $\frac{\sum_{j=1}^2 \overline{inc}_{N_j} * N_j}{N}$, where $N = N_1 + N_2$ is the total number of users. Although this assumption is non-trivial, it is employed so that the provider makes a meaningful choice of initial incentives offered per user. Alternatively, if the provider does not have any estimate of the average \overline{inc} , then he can start with an arbitrary value of the initial incentives and employ the learning algorithm as is. Thus, in our model, after scaling \overline{inc} by the economic factor α , the provider offers the following initial incentive to all users of both categories:

$$inc = \alpha * \overline{inc} = \alpha * \frac{\sum_{j=1}^2 \overline{inc}_{N_j} * N_j}{N} \quad (2).$$

We have assumed that $t_{min,i,j}$ expresses the elasticity of users and in association with the incentives offered, this parameter can be used for the assessment of users' participation probability. In particular, when the incentive exceeds $t_{min,i,j}$ the participation probability p_i of this user should be close to 1; of course, the larger the incentive, the higher the participation probability. On the contrary, when the incentive is lower than $t_{min,i,j}$ this probability should be close to zero. Moreover, when provider gives recommendations to the user, the participation probability is considered to be higher than when no recommendations are used. Indeed, for a given amount of incentives a successful recommendation facilitates the user's planning of the consumption schedule based on the proposed reduction. Thus, the user can achieve the corresponding DR objective more easily and therefore more often than in the case when no recommendation is given. In the sequel, we assume that the participation probability of the users is given by the following formula:

$$p_i = \frac{e^{5(y_i-1)}}{e^{5(y_i-1)} + e^{5(1-y_i)}} * \gamma \quad (3)$$

The parameter γ accounts for the fact that the maximum participation probability is higher in the case when recommendations are offered. It is assumed to take the indicative,

values $\gamma = 1$ and $\gamma = 0.8$ for DR programs with and without recommendations respectively. These values are considered as representative of the positive impact of recommendations in the DR participation probability. The parameter y_i is the ratio of the incentive to the elasticity parameter of this user $i \in N$, i.e. $y_i = \frac{\text{incentive}}{t_{min,i,j}}$, where the variable *incentive* varies according to the approach implemented in each case. Therefore, the participation probability given by equation (3) has a sigmoid shape with a considerable increase from low to high values when the incentive offered to a user exceeds his elasticity parameter $t_{min,i,j}$. We choose to multiply $y_i - 1$ in the exponent by 5, as the resulting curve is steep but not very much, meaning that the participation probability does not increase or decrease very sharply. We have not taken $t_{min,i,j}$ as a strict threshold, to allow for some uncertainty in user participation. However, we still refer to this parameter as the minimum incentive.

4. EFFECTIVE INCENTIVE ALLOCATION

We present below three approaches $k = \{1,2,3\}$ each following a different strategy for providing incentives. The third approach is a learning algorithm aiming to assist the energy provider in deducing information about users' preferences in a dynamic manner and thus increasing the participation rate, by effectively allocating the incentives to be offered. The purpose of this algorithm is to grant the provider with additional knowledge concerning the trade-off between the amount of money used for incentives and the DR participation that can be achieved accordingly. We distinguish two cases with regard to the exploitation of the available information in the implementation of DR. In the first case, the provider applies a DR program without offering any recommendations on the actions to be taken by the users, while in the second case utilising the knowledge originating from the load disaggregation and profiling the provider offers recommendations regarding the load curtailment/shifting of certain of the appliances deduced, with the aim to examine to what extent the introduction of such recommendations leads to better results on user participation. Overall, the basic idea of the algorithm is to gradually exploit the available information in order to attain an efficient participation rate in conjunction with an efficient incentives allocation scheme. The approaches begin with the minimum information available to the provider. After each approach, the provider is assumed to enrich his knowledge of the monetary incentive preferences of each user, so that the maximum participation rate is achieved.

4.1 DR without recommendations

In this section, we consider the deployment of DR program without any recommendation and we run three distinct approaches.

4.1.1 Approach 1: DR with a single unified incentive

In this approach we consider that any prior information about the participation of users is either unknown to the provider, or ignored. Hence, the implemented DR program utilises the incentive defined in (2). The objective of the provider is to extract knowledge of the users' elasticity from their participation. After the DR program is executed, the provider identifies the set of users that participated in DR as $Z_1, Z_1 \subset N$ and estimates the total participation rate $PAR_1 = Z_1/N$.

4.1.2 Approach 2: DR using participation information and common incentive per category

Suppose that the provider by leveraging the demographic characteristics of users and the participation information from the previous approach generates better users' profiles with regard to their elasticity characteristics. Thus, utilising this information, the

provider deploys DR offering to each category the same incentive as in (1). In this case, we denote as $Z_{2,j}$ the set of users that participated in the DR and as $Par_{2,j} = Z_{2,j}/N_j$ the participation rate for each category. The total participation rate equals $PAR_2 = \sum_{j=1}^2 Z_{2,j}/N$.

Table 1. Mining the minimum threshold $t_{min,i,j}$ for $k = 3$

Step 1: Define the initial incentive to be offered either according to Approach 1 or according to Approach 2.
Step 2: Sort all users (Approach 1) and users in each category (Approach 2) in ascending order of $t_{min,i,j}$.
For each iteration r : Step 3: Examine users one by one (for more details refer to the text of Subsection 4.1.3). Step 3a: Reduce the incentive value for each of the participating users by δ : $inc_{new,3,r,1} = incentive - \delta$ Step 3b: Increase the incentive value for each of the non-participating users by δ' : $inc_{new,3,r,2} = incentive + \delta'$
Step 4: Set $inc_{3,i,j} = inc_{new,3,r,j}$ and mark user i as "Discovered". The algorithm terminates, when changes in the $inc_{3,i,j}, \forall j = \{1,2\}$ do not affect users' state, i.e. users do not change from participating to non-participating.
Step 5: Compute for each category $j = \{1,2\}$ the percentage of participation $Par_{3,r,j} = \frac{Z_{3,r,j}}{N_{3,j}}$ and total participation rate $PAR_{3,r,j} = \frac{\sum_{j=1}^2 Z_{3,r,j}}{N}$, where $Z_{3,r,j}$ is the set of users that participated in the DR.

4.1.3 Approach 3: Effective incentive allocation using learning of customized incentives

This approach can be defined as an extension of both the first and the second approach. We refer to them as Approach 3.1 and 3.2 respectively. The provider utilises the rate of participating users as input. The aim is to extract information concerning the minimum incentive $t_{min,i,j}$ of each user $i \in N_j$ in a dynamic way and employ it in a subsequent DR event. In particular, the approach consists of independent runs, each r corresponding to a DR event. (The first iteration is essentially an execution of either the first or the second approach.) We introduce two parameters δ and δ' that denote the amount of decrease and increase in the incentives offered. Their values are chosen to be quite small and fixed, so that there is limited dispersion of the resulting incentive values to be offered between the participating and non-participating users with similar values of $t_{min,i,j}$. Table 1 describes briefly the steps followed. In particular, given the outcomes of the previous approaches as starting points, at each subsequent r , the provider sorts the set of users (and in each category) in ascending order, so that users with the lowest thresholds to be investigated first, and the total amount of money spent to not increase rapidly. Then he reduces (resp. increases) the amount of incentives by δ to the participating users resp. (resp. δ' to the non-participating users). Reducing gradually the incentive of the users that participated in the first and second approach resp. allows for exploring (learning) their minimum incentive $t_{min,i,j}$ without affecting their participation. If in some iteration a user is not engaged with the new reduced incentive, then he is offered the same incentive for the next iteration as well, in order to confirm whether non-participation depends on the randomness of (3) or is due to the

low incentive offered. In such a case, users that do not participate for two subsequent times are marked as “Discovered” and in following iteration they are offered the incentive by which they participated the last time. Each initially non-participating user is given an incentive increased by $\delta' = 2\delta$ in order for the provider to gradually approach this user's $t_{min,i,j}$. The process continues until $t_{min,i,j}$ is indeed reached and the user is marked as “Discovered”. However, if a user does not participate when given the new incentive, we keep this incentive for next iteration. If this user still does not participate for this iteration, we increase his incentive by $\delta' = 2\delta$. Therefore, under this approach, the provider observes dynamically users' response to DR incentives offered until each user is given roughly the minimum amount of incentives that can lead him to a high participation probability. Thus, the algorithm both ensures consistently successful DR events in the intermediate iterations and improves gradually the participation rate, until the maximum participation is reached with the minimum total amount of DR incentives. Note that to maintain consistently successful DR events, we are more conservative in reducing an incentive value that proved to be effective than in increasing one that was not, for which a larger step is employed.

4.2 DR using recommendations

In this case, we consider DR programs, for which the provider utilises the consumption profiles stemming from load disaggregation (the details of such an algorithm are out of the scope of our paper) and offers recommendations regarding the load curtailment or shifting of specific appliances. The objective is to investigate whether the use of specific recommendations influences the participation rate of a DR program and to what extent. We employ again three such approaches that build on the same basis in terms of input data, methodology followed and the objectives to be served, as the approaches introduced in Section 4.1 albeit with minor differences. For this reason, we only describe them briefly, highlighting their differences.

4.2.1 Approach 1: DR with a single unified incentive

Again, as in Section 4.1.1, we assume that the provider has no information about users' participation. The DR program applied uses a single incentive that is unified for all users $inc_{1,i} = inc$. The difference lies in the fact that the provider can additionally offer recommendations regarding the actions of curtailment and/or shifting of the load for specific appliances. We evaluate the approach by means of independent runs, where each run is mapped to a DR event.

4.2.2 Approach 2: DR using profiling information and common incentive

Given the profiling information obtained by the demographic characteristics of users, the provider applies DR program and offers to all users of each category the average incentive as defined in (2), both when using recommendations and when not, aiming to examine whether the introduction of such recommendations has positive impact on the overall user participation and the participation in each type of users. Again the approach is evaluated by means of independent runs.

4.2.3 Approach 3: Effective incentive allocation using learning of customised incentives

The difference of this approach with the corresponding one of Section 4.1.3 lies in the fact that the provider wishes to discover the $t_{min,i,j}$ of each user but also the optimal choice of recommendation to accompany the DR message and achieve the highest participation rate. The approach is executed following the same methodology as

described in Table 1. At Step 3a, if a user does not participate in the DR event, at the next iteration, he is offered the same incentive $inc_{new,3,r,1}$ but supposedly with some other different recommendation. If he does not still participate for this iteration, he is categorised as non-participating and the new incentive offered is calculated according to Step 3b. Otherwise, we assume that the reason for his behaviour is an inappropriate recommendation. The same logic is followed at Step 3b.

5. EXPERIMENTAL EVALUATION

This section presents our evaluation results. We assume a set $N = 60$ users. N_1, N_2 are the subsets of elastic and inelastic users, for each of which the parameter $t_{min,i,j}$ is drawn from the uniform distribution within the range of $[2,6]$ and $[6,10]$ resp. For the decrement δ a small value of 0.1 is taken, while the increment δ' is taken $\delta' = 2\delta$. When evaluating the various approaches, we should keep in mind that users' behaviour in terms of participation is not fully predictable and depends on the probability defined in (3). Therefore, for all approaches, we measure the average participation rate for each of several iterations over a set of multiple independent runs of the experiment.

5.1 DR without recommendations

To experimentally evaluate our approaches, we study three cases of users' partition in the two categories, i.e. symmetric partition, asymmetric partition with more elastic users and asymmetric partition with more inelastic users. For conciseness, we present results for only one of the cases, namely the asymmetric partition with more inelastic users (75% of total users), as it is considered more challenging to identify $t_{min,i,j}$ for those users and achieve higher participation. However, the same evaluation approach is used for the other two cases and similar conclusions apply. Figure 1 depicts the results of the two static Approaches 1 and 2 resp. for a broad range of values of the economic scaling factor α , for each of which the average participation from 35 independent runs is depicted. We observe that in general the lower the value α and thus the initially offered incentives the less the (average) participation. Note that despite the stochastic modelling of users' participation in (3), the average participation rate is monotonic as intuitively expected.

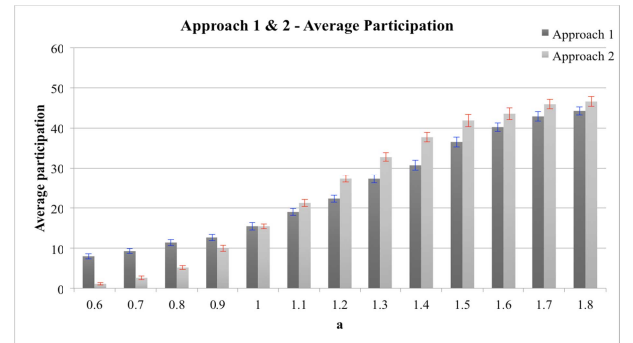


Figure 1. Approaches 1 & 2: Participation for each value of α .

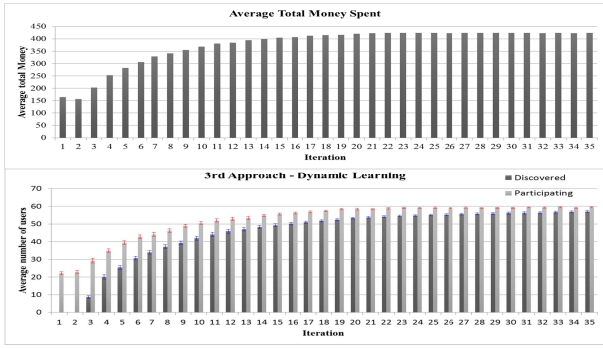


Figure 2. Approach 3.1, for $\alpha = 1.1$ and 35 iterations

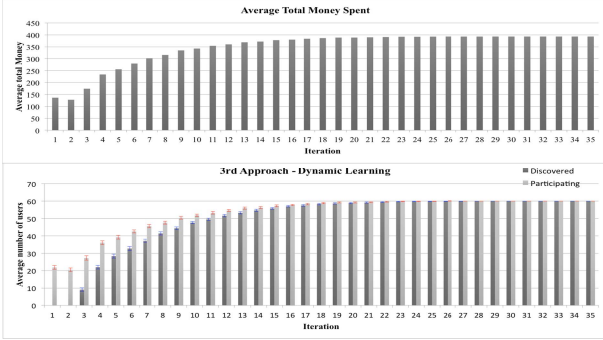


Figure 3. Approach 3.2 for $\alpha = 1.1$ and 35 iterations

The use of profiling information in Approach 2 in general leads to higher levels of participation than for Approach 1. Somewhat surprisingly though this does not apply for low values of the factor α , for which Approach 1 attains a higher average participation. This can be attributed to the fact that by offering the same incentives to all users, Approach 1 succeeds in inciting all elastic users. Approach 2 offers different incentives per category of users. Hence, for low values of α , these incentives may be lower than most users' thresholds, thus resulting in lower participation probability for Approach 2. Additional evidence for the above comparisons of participation under Approaches 1 and 2 is provided by also observing the associated confidence intervals, all constituting a small portion of the respective average participation; see Figure 1.

For the dynamic Approach 3, we examine and compare both possible ways of its implementation, i.e. as extension of the static Approaches 1 and 2. Thus, for both Approaches 3.1 and 3.2 we simulate 30 independent runs, with 35 iterations each for $\alpha = 1.1$. Figure 2 and Figure 3 indicate that both variations of Approach 3 converge to high average participation; Approach 3.2 achieves the highest possible participation (100%). The confidence intervals vary with the approach and the iteration but in general are relatively small. In addition, users are being discovered gradually as the algorithm converges. Actually, Approach 3.1 does not discover all users in contrast to Approach 3.2. However, the key difference lies in the total amount of money spent. For attaining a specific number of participating users, Approach 3.2, which employs classification in setting the initial incentives, utilises less money than Approach 3.1. This can be justified by the fact that the initial incentive per category of users in Approach 3.2 is closer to $t_{min,i,j}$; hence the participating users are offered from the start incentives that are more likely to be both effective and close to their real values of $t_{min,i,j}$. Therefore, it is really beneficial for the provider to utilize profiling information under Approach 3, enabling him to offer an appropriate initial incentive per category. In addition, by sorting users in

ascending order of $t_{min,i,j}$ under both Approaches 3.1 and 3.2, the algorithm avoids selecting at the beginning the inelastic users requiring a higher incentive to participate. In essence, in this manner the learning algorithm attempts to minimise the amount of money to be spent for incentives at each given level of participation rate. Thus, in its successive iterations, the algorithm produces a set of possible incentive allocations that attain different participation rates with nearly the minimum possible total amount of money for DR incentives. The provider can decide on when to stop trying to learn the required incentives of more users depending on his objectives for participation rate and/or total money to be spent.

5.2 DR using recommendations

In this case, users are again considered to be asymmetrically distributed with 75% of them being inelastic. For the two first approaches we perform again simulations of 35 independent runs for a broad range of values of the economic scaling factor α . Figure 4 reveals that with the use of recommendations too when the value of α increases the average participation raises similarly as in the results of Section 5.1. In the present case, the use of profiling information in Approach 2 leads to lower levels of participation compared to the Approach 1. We observe that the stochastic modelling in (3) results in smaller confidence intervals. Regarding Approach 3, we have run a set of 30 independent runs each with 35 iterations for each run for $\alpha = 1.1$ and for both variations of the approach. Figure 5 and Figure 6 indicate that the algorithm, either as Approach 3.1 or 3.2 converges really fast and thus reaches in a few iterations a very high average participation. It is noteworthy that Approach 3.2 results in discovering almost all users as opposed to Approach 3.1. Similarly to the outcomes described in Section 5.1, the use of the profiling information plays a key role also when combined with the use of recommendations, since it results in the extraction of more detailed and accurate knowledge of users' threshold and in a faster manner.

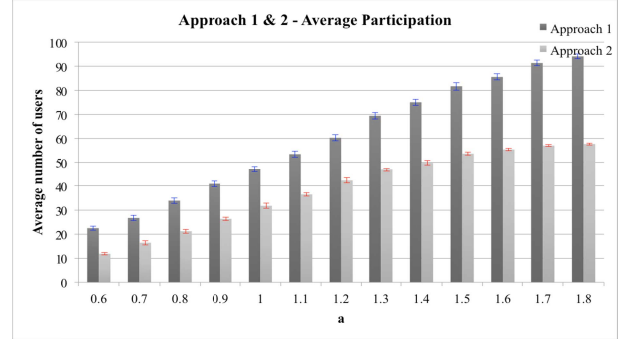


Figure 4. Approaches 1 & 2: Participation for each value of α .

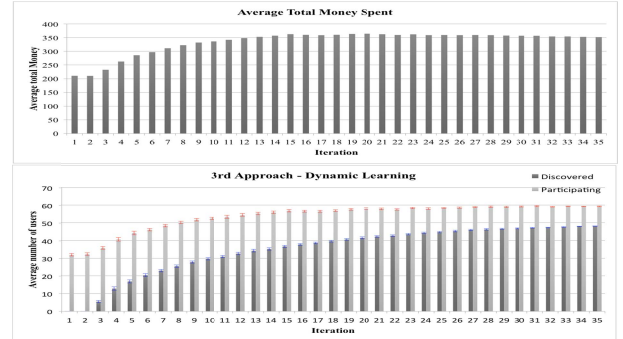


Figure 5. Approach 3.1 for $\alpha = 1.1$ and 35 iterations

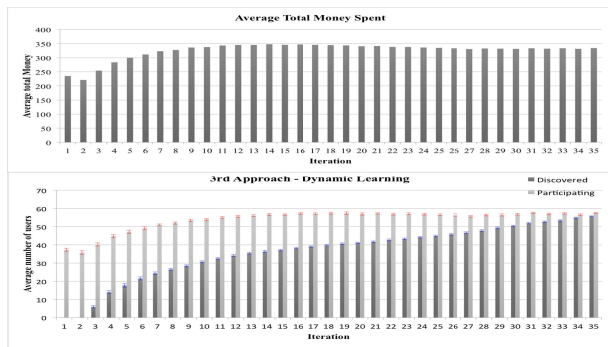


Figure 6. Approach 3.2 for $\alpha = 1.1$ and 35 iterations

We observe here as well the trade-off between the participation rate and the total money spent for this level of participation with Approach 3.2 ultimately requiring a smaller total amount for incentives than 3.1. Again, the appropriate selection rests with the provider given his objectives. For example, assume that the provider wishes to achieve participation only 58 out of 60 users, so he seeks to select the DR program that requires the least money. In general, we observe that both Approaches 3.1 and 3.2, either without or when using recommendations can fulfil the provider's objective. However, Approach 3.2 outperforms Approach 3.1 in both scenarios of applied DR with regard to the total amount of money spent. Thus, the provider should choose to implement Approach 3.2. In a practical case, the provider should estimate the value of the parameter γ in equation (3) for DR programs with and without recommendations respectively, and assess whether the DR budget reduction when he uses recommendations is worth the additional cost for deriving them.

6. CONCLUDING REMARKS

In this paper, we have introduced and evaluated three approaches (two static ones and a learning approach) that help the energy provider perform DR effectively exploiting information that are available from profiling and/or result from load disaggregation. In particular, the learning approach is applied in successive DR events and aims to explore and at the same time exploit the minimum acceptable incentives that motivate each user to participate in DR. Our study focused on two basic types of incentive-based DR programs offered to residential environments; namely, critical peak rebate DR programs with and without accompanying recommendations regarding the curtailment or shifting of the load of specific appliances. We have assumed two categories of users, namely elastic and inelastic. We have investigated different users' distributions between the two categories, but due to space limitations we have presented in detail the results for only one distribution; the conclusions for other distributions are similar. Our simulations reveal interesting insights on the impact of the use of profiling information on users' participation in DR programs and on the effective allocation of provider's budget for DR incentives.

In particular, in the case of the static Approaches 1 and 2, it turns out that using profiling information to offer customized incentives per category of users is beneficial except if the initial incentives are rather low. This implies that, in a practical case, if the budget for DR incentives offered to a specific set of users is relatively low, then a unified DR incentive should be offered to all users. Moreover, Approach 3 employs a learning algorithm in order to discover the minimum acceptable incentive for each user. After the first two iterations, and by starting with those users that are characterised by the lowest incentives accepted so far, the learning algorithm explores and discovers the lowest acceptable incentives

for an increasing subset of the users until the maximum participation is attained. Thus, after a few iterations, in each DR event the algorithm attains a DR participation rate by dedicating in DR incentives nearly the minimum of the corresponding amount of money required for achieving this rate. The specific DR program to be applied in subsequent events depends on the DR participation rate that the provider aims to achieve and on the amount of money he wishes to spend for this purpose. Thus, this selection depends on the significance for the provider of the participation of additional users in conjunction with the total incentives to be offered, given the provider's objectives on the load to be curtailed and operational constraints, if any. One should keep in mind, though, that applying DR by using Approach 3.2 and recommendations leads to better results concerning the total amount of money spent. Also, that for the static Approaches 1 and 2, the provider's budget plays a significant role on users' participation, since in general the higher the value of incentives offered the higher the participation. Under Approach 3, in all cases considered for the mix of users (elastic vs. inelastic) and the value of α , the participation raises in each DR event (iteration). Also, for all cases of user mix, the introduction of specific recommendations in the DR programs (made possible by means of using profiling information) leads to noticeable improvement of the convergence rate of the algorithm to the maximum participation. This should have been reasonably expected because we have assumed that the maximum participation probability γ in equation (3) is higher for DR programs with recommendations. Consequently, the total money that is spent in each run of a DR program is less in the case of using recommendations, and especially when employing Approach 3.2. It should also be noted that the introduction of parameter γ in our model offers an interesting new possible use of our algorithm. Indeed, by employing different values of γ close to real life conditions, the algorithm can serve as a tool for assessing the impact on the improvement of participation and of the trade-offs arising when recommendations are provided. This can be considered as an interesting direction of future work based on the results stemming from real-life trials, such as those conducted in the context of the WATTALYST project [1], through which the provider can appraise whether and to what extent the provision of more specific recommendations leads to higher participation of users in DR.

7. ACKNOWLEDGEMENTS

This research has been partly funded from the European Union's Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 288322 WATTALYST [1].

REFERENCES

- [1] WATTALYST project: www.wattalyst.org
- [2] Hart, G. W. 1992. Nonintrusive appliance load monitoring. *Proceedings of the IEEE*, 80(12), 1870-1891.
- [3] Jain, S., Narayanaswamy, B., & Narahari, Y. 2014. A Multiarmed Bandit Incentive Mechanism for Crowdsourcing Demand Response in Smart Grids. *In Twenty-Eighth AAAI Conference on Artificial Intelligence*.
- [4] Marceau, M. L., & Zmeureanu, R. 2000. Nonintrusive load disaggregation computer program to estimate the energy consumption of major end uses in residential buildings. *Energy Conversion and Management*, 41(13), 1389-1403.
- [5] Ardakanian, O., Keshav, S., & Rosenberg, C. 2011, August. Markovian models for home electricity consumption. *In Proceedings of the 2nd ACM SIGCOMM workshop on Green networking* (pp. 31-36). ACM.