Advanced Demand Response Solutions for Capacity Markets

Rim Kaddah, Daniel Kofman, Fabien Mathieu, Michal Pioro

► To cite this version:

Rim Kaddah, Daniel Kofman, Fabien Mathieu, Michal Pioro. Advanced Demand Response Solutions for Capacity Markets. 11th International Conference on Innovations in Information Technology (IIT'15), Nov 2015, Dubai, United Arab Emirates. 10.1109/INNOVATIONS.2015.7381509 . hal-01225293

HAL Id: hal-01225293

<https://hal.science/hal-01225293>

Submitted on 6 Nov 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Advanced Demand Response Solutions for Capacity Markets

Rim Kaddah*, Daniel Kofman*, Fabien Mathieu[§] and Michal Pioro^{†‡}

* Telecom Paristech, 23 Avenue d'Italie, Paris, France

Emails: {rim.kaddah,daniel.kofman}@telecom-paristech.fr

[§] Alcatel-Lucent Bell Labs, route de Villejust, Nozay, France

Email: fabien.mathieu@alcatel-lucent.com

[†]Institute of Telecommunications, Warsaw University of Technology, Poland

[‡] Department of Electrical and Information Technology, Lund University, Sweden

Email: michal.pioro@eit.lth.se

Abstract—The Internet of Things (IoT) paradigm brings an opportunity for advanced Demand Response (DR) solutions. Indeed, it enables visibility and control on the various appliances that may consume, store or generate energy within a home. It has been shown that a centralized control on the appliances of a set of households leads to efficient DR mechanisms; unfortunately, such solutions raise privacy and scalability issues. In this paper we propose an IoT-based DR approach that deals with these issues. Specifically, we propose and analyze a scalable two levels control system where a centralized controller allocates power to each house on one side and, each household implements an IoT-based DR local solution on the other side. A limited feedback to the centralized controller allows to enhance the performance with little impact on privacy. The solution is proposed for the general framework of capacity markets.

I. INTRODUCTION

The growing deployment of intermittent renewable energy sources at different scales (from bulk to micro generation) advocates for the design of advanced Demand Response (DR) solutions to maintain the stability of the power grid and to optimize the usage of resources.

DR takes advantage of demand flexibility. The level of gain depends on the granularity of visibility and control on the demand. The Internet of Things (IoT) paradigm enables implementing DR at the finest granularity (individual appliances), and deploying IoT-based solutions becomes feasible, both from technological and economical points of view.

The introduction of capacity markets in several countries provides incentives for: the flexibility end users could provide through DR mechanisms; the deployment of flexible generators (for which the energy cost is higher than the average).

In this paper, we focus on DR solutions for keeping power consumption below a certain known capacity limit for a well-defined period of time. A possible application is for utility companies, which are interested in limiting the cost of the capacity certificates they have to acquire in the capacity market (for securing supply). Such cost reduction is facilitated by keeping power consumption below known thresholds.

In [3] the authors propose and analyze several IoT-based DR mechanisms. They show that fine-grained visibility and control on a set of households at an aggregation point enables to maximize users perceived utility. However, this approach may cause scalability as well as privacy problems. On the other hand, they consider two levels control systems where a central controller allocates available capacity to households based on some static information (e.g. type of contract). Then, local controllers leverage IoT benefits for local optimization, without any feedback to the central controller. The drawback of such approach is that it significantly reduces the total utility perceived by the users.

In this paper, our main contribution is the proposition and evaluation of an intermediate approach, based on two level systems with partial feedback from the local controllers to the central entity. The feedback sent has little impact on privacy. Proposed solution enforces fairness by considering two levels of utility for each appliance (i.e., vital and comfort). We compare the performance of the proposed scheme with the two cases studied in [3] (fully centralized solution and two level system with no feedback). Results are analyzed for homogeneous (all households have the same characteristics) and heterogeneous scenarios. We show that for both cases, the proposed algorithm outperforms the scheme with no feedback. It runs in a limited number of iterations, which ensures good scalability and limited requirements in terms of communication resources.

The paper is organized as follows: Section II presents the related work. The system model and allocation schemes are introduced in Section III and IV respectively. In Section V, we study the performance of the proposed control scheme and compare it to two benchmark control approaches through a numerical analysis of the model. Conclusion and future work are presented in Section VI.

II. RELATED WORK

Most proposed DR approaches can be classified in 2 groups, namely, incentive based and direct control based. Incentive based approaches aim to induce a targeted behavior of users through dynamic prices. Authors in [5] propose a dynamic pricing scheme based on a distributed algorithm to compute

*[§] This author has carried out the work presented in this paper at LINCS (www.lincs.fr)

optimal prices and demand schedules. In this work, we are interested in direct load control of heterogeneous appliances in the context of capacity markets. Direct load control has the advantage of providing tight consumption guaranties. Such mechanisms have been mainly proposed for providing system services (for example, real time following of a flexibility demand curb) and not in the context of capacity markets. While detailed appliance models are proposed in pricing papers (see for example [5]), most previous work on direct control focus on specific types of appliances. For instance, the authors in [7] propose an online control of deferrable ON-OFF loads. A wide range of proposals focus on Thermostatically Controlled Loads (TLCs) ([6], [8], [10], [4]).

Recent work has proposed schemes that are capable of taking into account flexibility of any generic appliance ([11], [12]). Authors in [12] propose a customer reward scheme that incentivizes users to accept direct control of loads. They propose a greedy algorithm (maximizes utility slot by slot) based on the utility that each appliance declare for each slot. It's shown in [3] that such an allocation has low performance since decisions are taken with no view on the global time period. It also has obvious privacy issues. As stated in the introduction, the present work builds on conclusion of the analysis in [3]. Most closely related work to our proposal is the scheme presented in [11] which is very similar to [5] if prices are interpreted as control signals. The authors in [11] propose to solve a similar problem but their approach requires convergence of the algorithm to produce an allocation that does not violate total capacity constraint. In our proposal, stopping the algorithm before convergence (in case of communication delays or loss for certain homes) will degrade performance but still propose a feasible solution. The authors in [11] do not discuss scalability and communication requirements in terms of number of iterations required. They also assume convex utility functions. Moreover, due to the type of information provided to the centralized controller, our solutions better guaranties privacy. Finally, they don't consider fairness. We introduce the concept of vital and comfort utility for each appliance. On the one hand, this enables to better model the real utility for end users and, on the other hand, we intrinsically introduce a certain level of fairness.

III. SYSTEM MODEL

We consider an aggregator in charge of allocating power to a set of H households under a total capacity constraint $C(t)$. t represents time slots. We suppose that during a defined time period (measured in slots), in absence of control, predicted demand would exceed available capacity. We call this period a DR period. We denote by DE_a and DE_h the functional groups in charge of decision taking at the aggregator side and at the user h side (one per home), respectively. DE_a is in charge of allocating power to each household (C_{ht}), under the total power constraint. For each house h , DE_h has two main roles: the collection of information on variables monitored at user premises (state of appliances, local temperature, etc.); the enforcement of control decisions received from DE_a (e.g.

by controlling the appliances). More details will be given in Section IV when introducing the considered allocation schemes.

System Parameters and Exogenous Variables	
H	Number of homes
A	Number of classes of appliances
$P_m^a(h)$	Minimum power consumed by appliance a in home h
$P_M^a(h)$	Maximum power consumed by appliance a in home h
$\pi_{v/c}^a(h, t)$	Preference coefficients
$C(t)$	Available power capacity at time slot t
$L(h)$	Power limits for home h
t_M	DR period duration in time slots
$T_m(h)$	Minimum acceptable indoor temperature for home h
$T_P(h)$	Preferred indoor temperature for home h
$T_M(h)$	Maximum acceptable indoor temperature for home h
$T_0(h)$	Initial indoor temperature for home h
$F(h), G(h)$	Coefficients for temperature dynamics in home h
$T_e(t)$	Exterior temperature at time t
Control Variables and Controlled Variables	
U_{ht}^a	Utility of appliance a in home h at time t
U_{vht}^a	Vital component of U_{ht}^a
U_{cht}^a	Comfort component of U_{ht}^a
X_{ht}^a	Power consumed by appliance a in home h at time t
x_{ht}^a	= 1 when appliance a in home h at time t is active and = 0 otherwise
T_{ht}	Temperature of home h at time t
C_{ht}	Capacity limit allocated for home h at time t
g_{ht}	Subgradient of the utility function at point C_{ht} for home h at time t

TABLE I: Table of notation

A utility function is defined for each controlled appliance to express the impact of its operation on user's satisfaction. We assume electrical appliances are classified among A classes. Appliances of the same class have similar usage purposes (e.g., heating) but may have different operation constraints. Appliance of class a at home h operates within a given power range $[P_m^a(h), P_M^a(h)]$.

Following [3], a specific utility function is modeled for each class of appliances based on usage patterns and criticality, users' preferences and exogenous variables (e.g. external temperature). So, the utility of an appliance is expressed either directly as a function of its consumption or as a function of some monitored variables (see Section V for an example).

In the present work, we introduce two levels of utility per appliance, vital and comfort. The first one expresses high priority targets of high impact on users' wellbeing and the second one expresses less essential preferences.

For notation, we write utilities as vital/comfort pairs: $U_{ht}^a = (U_{vht}^a, U_{cht}^a)$ denotes utility of appliance a at time t for home h . When controlling a set of appliances, policies target to satisfy comfort only if vital needs cannot be further covered for all appliances. Control decisions are based on the lexicographical order comparison of utility values: For two values U_{ht}^a and U'_{ht}^a , we say $U_{ht}^a > U'_{ht}^a$ iff $U_{vht}^a > U'_{vht}^a$ or ($U_{vht}^a = U'_{vht}^a$ and $U_{cht}^a > U'_{cht}^a$). Utilities can be summed using element-wise addition.

We denote by $\pi_v^a(h, t)$ (resp. $\pi_c^a(h, t)$) the maximal vital (resp. comfort) utility associated to appliance a at home h and time t . These values, which we call *preference coefficients*, represent how the importance of appliances is modulated depending on the preferences of users.

We assume that each house has a power limit $L(h)$ sufficient to achieve a maximal utility.

The optimization problem considered in this paper consists in maximizing the total utility (using the lexicographic total order) of users under system constraints. We notice that fairness is introduced through the lexicographic ordering of vital and comfort utility values (no comfort power is allocated to any house if not all vital needs are covered). We do not directly focus on revenues but expect that reaching maximal users' utility leads to maximal gains for all involved players. Utility companies can provide better services for a given total allocated power, which should translate into higher revenues, or reduce the expenses in the capacity market for a given level of service, which should reduce total costs. End users can save money due to attractive prices they get for participating to the service and adjusting energy consumption to their predefined policies. Notation is summarized in Table I.

IV. ALLOCATION SCHEMES

We present here two reference schemes that will be used for benchmarking purposes, along with our proposed solution.

A. Global Maximum Utility

The centralized global optimization is formulated by equations (1).

$$\max_{X_{ht}^a, x_{ht}^a} \sum_{t=1}^{t_M} \sum_{h=1}^H \sum_{a=1}^A U_{ht}^a \quad (1a)$$

s.t.

$$\sum_{h=1}^H \sum_{a=1}^A X_{ht}^a \leq C(t), \quad \forall t \quad (1b)$$

$$P_m^a(h)x_{ht}^a \leq X_{ht}^a \leq P_M^a(h)x_{ht}^a, \quad \forall t, \forall h, \forall a \quad (1c)$$

$$x_{ht}^a \in \{0, 1\}, \quad \forall t, \forall h, \forall a. \quad (1d)$$

This problem was studied in [3], which shows the importance of fine grained information from homes on the performance of the control.

One can solve (1) if all the informations about appliances and their utility functions are transmitted by the repartitors DE_h to the aggregator DE_a , which can then compute an optimal global solution and notify the repartitors accordingly.

While being optimal with respect to the utilities (by design), this allocation, called *GM*, has two major drawbacks. First, it requires to compute the solution of a complex problem, which may raise scalability issues. Second, fine grained information harvesting may cause privacy related issues which can affect the acceptance of the control scheme by users. Thus, it may be preferable to store information locally at homes with a local intelligence. This leads to the following scheme.

B. Local Maximum Utility

This control scheme, denoted *LM*, considers only one way communication from DE_a to DE_h (no feedback is transmitted from DE_h to DE_a), and decision is made at both levels.

First, DE_a allocates power to homes proportionally to their subscribed power, so the power allocated to home h is

$$C_{ht} = \frac{L(h)}{\sum_i L(i)} C(t).$$

Then, at each home h , DE_h decides the corresponding allocation per appliance by solving the restriction of (1) to h , using C_{ht} instead of $C(t)$.

By design, *LM* is scalable (only local problems are solved) and private information disclosure is kept to a minimum. The drawback is that the corresponding allocation may be far from optimal [3].

C. SubGradient decomposition

We aim at achieving a reasonable trade-off between the centralized solution *GM*, which provides maximum performance in terms of total utility value, and the local solution *LM*, which enforces scalability and privacy.

To do so, we propose a simple primal decomposition, denoted *SG*, of the global *GM* problem into a master problem, described in (2), and subproblems, described in (3).

Master problem

$$\max \sum_{h=1}^H U_h \quad (2a)$$

$$\sum_{h=1}^H C_{ht} = C(t), \quad \forall t \quad (2b)$$

$$C_{ht} \geq 0, \quad \forall h \quad \forall t. \quad (2c)$$

Subproblems

For each home h , the following mixed integer linear problem (MILP) is solved:

$$U_h = \max \sum_{t=1}^{t_M} \sum_{a=1}^A U_{ht}^a \quad (3a)$$

$$\sum_{a=1}^A X_{ht}^a \leq C_{ht}, \quad \forall t. \quad (3b)$$

We briefly describe the main steps of *SG*: *SG* needs to be bootstrapped with an initial power allocation. Then, for up to K_{MAX} iterations, DE_a transmits to each repartitor DE_h the current allocation proposal $\{C_{ht}\}_t$. DE_h then solves the corresponding subproblem (3) and sends back the total utility U_h feasible, along with the subgradients associated to the current solution. Using the values reported by the repartitors, DE_a then tries to propose a better solution. In the end, the best found solution is used.

We now give the additional details necessary to have a full view of the solution.

1) *Initial allocation*: The first allocation is based only on homes' static maximum power limit. Following [3], we use a round-robin strategy: we allocate to some houses up to their power limit until the available capacity $C(t)$ is reached; we cycle with time the houses that are powered. The interest for *SG* of such an initial allocation (e.g. compared *LM*) is that it gives an initial diversity that will help finding good subgradients.

2) *Subproblem and subgradient computation*: The subproblem (3) is solved by DE_h as in *LM*, using the values C_{ht} proposed by DE_a . Reporting U_h is straightforward once the local solution is computed. For the subgradients, let g_{ht} denotes a subgradient of the utility function at point C_{ht} . The value g_{ht} can be found in two ways: either analytically or by taking optimal value of the dual variables of (3b). In this paper, we use an analytical computation of g_{ht} based on the utility functions.

3) *Finding better solutions:* To update the current solution at the k -th iteration, DE_a does the following:

- It first computes a value $\alpha_k g_{ht}$, where α_k is a parameter that depends on the iteration number. This value represents potential increase of C_{ht} .
- It then adjusts the new values of C_{ht} so they stay positive and fit the capacity constraints.

For the first phase, the step size α_k for each iteration k is a crucial parameter. Thus, choosing appropriate step sizes is key to speeding up resolution. Intuitively, α_k should be chosen to make the allocation update (dictated by $\alpha_k g_{ht}$) useful for high consumption appliances during the first iterations. Then, α_k should decrement with k so that the update is able to modify allocations corresponding to low consumption appliances.

For the adjustment phase, it is important to deal with cases where allocation update $\alpha_k g_{ht}$ is larger than available capacity $C(t)$ or maximum subscribed power $L(h)$ of home h , so we first cap $\alpha_k g_{ht}$ at the minimum between power limit of the smallest home ($L_m := \min_h L(h)$) and system capacity $C(t)$. We therefore define $\beta_{kht} = \min(\alpha_k g_{ht}, L_m, C(t))$.

Then for each t , we remove some positive common value λ_t to the C_{ht} to keep the sum of the allocations equal to the total capacity $C(t)$. To avoid houses with low C_{ht} to be badly impacted (in particular to avoid negative allocations), a subset I_t of the houses will be “protected” so that their values cannot decrease. In details, we do the following, starting with $I_t = \emptyset$:

- We compute λ_t such that the values

$$C'_{ht} = \begin{cases} C_{ht} + \max\{\beta_{kht} - \lambda_t, 0\} & \text{if } h \in I_t, \\ C_{ht} + \beta_{kht} - \lambda_t & \text{otherwise,} \end{cases} \quad (4)$$

sum to $C(t)$. See [9], [1] for more details.

- We protect (e.g. add to I_t) all houses that get a negative value C'_{ht} .
- We iterate the steps above until all C'_{ht} are positive. DE_a then proposes C'_{ht} as a new solution to investigate.

Remarks: While the solution described here applies to a 2-level hierarchy (DE_a , DE_h), it can be generalized to M levels to take into account static maximum capacity of different aggregation points on a hierarchical distribution network: considering an aggregation point m at a certain level, the subgradient for m can be obtained by adding up subgradients from its children.

Also note that the proposed scheme can run asynchronously in the sense that it does not require all houses to communicate simultaneously. In fact, as soon as two homes respond, reallocation can be made based on the sum of the power for responding homes without having to wait for others to respond.

V. NUMERICAL ANALYSIS

We now propose to evaluate the performance of our proposed solution for a specific use case.

A. Parameters and settings

To study the performance of the control schemes for several values of capacity, we choose the following system parameters:

Class	$[P_m^1(h), P_M^1(h)]$	$[P_m^2(h), P_M^2(h)]$	$F(h)$	$G(h)$
1	[50, 1000]	[1000, 4000]	0.0017	0.075
2	[50, 500]	[1000, 2000]	0.0008	0.0365

TABLE II: Classes of houses

Fig. 1: Utility of light power

Fig. 2: Utility of T_{ht}

- The size of the system is $H = 100$ houses.
- We select a slot duration of 5 minutes.
- The DR period is set to $t_M = 100$ slots (≈ 8 hours).
- We consider two types of appliances ($A = 2$): lighting ($a = 1$) and heating (index $a = 2$). Utility functions for both appliances have a vital and a comfort component.
- Vital light utility is fully obtained as soon as the minimal light power $P_m^1(h)$ is reached, while comfort utility linearly grows from $P_m^1(h)$ to $P_M^1(h)$ (See Fig.1).
- For heating, vital utility linearly grows until the minimum tolerable temperature $T_m(h) := 15^\circ C$ is reached, while comfort utility linearly grows from $T_m^1(h)$ to the preferred temperature $T_P(h) := 22^\circ C$ (See Fig.2).
- We suppose a constant external temperature $T_e(t) = 10^\circ C \forall t$ and an initial temperature $T_0(h) = 22^\circ C \forall h$.
- We suppose constant preference coefficients during the whole period: $\pi_v(h, t) = \pi_c(h, t) = 1 \forall h \forall t$.
- Temperature in homes evolves according to a simplified conductance/capacity model that leads to the following dynamics:

$$T_{ht} = T_{h(t-1)} + F(h)X_{ht}^2 + G(h)(T_e(t) - T_{h(t-1)}).$$
- Two types of houses are considered (See Table II), with class 2 having better energetic performance than class 1 (less light power required to achieve full utility and better insulation).

We suppose that the total available power is constant over the DR period, $C(t) = C$. We analyze the model for different values of C , ranging from low (only one type of appliances can be used) to full capacity (all appliances can be used).

While this model is rather simple (two types of appliances, constant values for C , π and T_e), we believe that the knowledge required to compute good solutions is sufficient to capture the trade-off between the efficiency of an allocation and the privacy of the users.

Fig. 3: Average utility per home over the DR period as a function of the available capacity (homogeneous case)

For the decomposition problem, we fix the maximum number of iterations to $K_{MAX} = 100$ iterations (Suppose it is a desired communication constraint). In the present, α_k is chosen non-summable diminishing and specifically equal to $500000/(\sqrt{k}+1)$. As stated in section IV-C, the order of magnitude of α_k is crucial to make capacity updates useful for homes and to speed up performance improvement. We choose to divide by the square root of k rather than k to slow the decay of α_k . The computation of g_{ht} is done by taking the highest slope of the utility functions with regard to X_{ht}^a at the solution of the local optimization problem (For heating, g_{ht} is computed by multiplying the slope of the utility function by the power coefficient $F(h)$ in the considered temperature dynamics).

The numerical analysis of the various presented mixed integer linear problems has been carried out using IBM ILOG CPLEX ([2]).

In the following, we discuss two cases: homogeneous and heterogeneous. For the homogeneous case, all houses belong to class 1 and for the heterogeneous one, we suppose 50 houses of class 1 and 50 houses of class 2.

B. Results on the homogeneous case

Figure 3 presents the main results on the homogeneous case. It displays the average utility per home over the DR period as the function of the available capacity C , for the three considered schemes (GM , LM and SG). For better readability, vital and comfort utilities are displayed separately.

Note that with the chosen parameters, the maximal feasible utility (vital and comfort) is 2. Another value of interest for vital utility is 1.75, which corresponds to situations where all houses are able to achieve vital light ($P_m^1 = 50$ W) but none

has the power necessary for heating ($P_m^2 = 1000$ W) so there is no control of temperature. Because the requirement for vital light is very low, it can be seen as a worst case situation.

Using a static allocation, LM struggles more than the other schemes for rising the vital utility above that threshold. It can only start to use heat for $C = 10^5$ (1000 W per house). Maximal vital utility is reached for $C = 105 \times 10^3$ (1050 W per house) and maximal utility (vital and comfort) necessarily requires $C = 2 \times 10^5$ (2000 W per house).

Obviously, GM , the optimal solution, is able to achieve better utilities. In particular, it achieves maximal vital utility even for very low capacities (down to 3×10^4), thanks to its ability of finding a working rolling allocation that allows all houses to use heat for a sufficient part of the period.

As expected, our proposal, SG , stands in-between these two opposite schemes. It is able to improve the vital utility of houses for values below $C = 10^5$, even if it fails to perform as good as GM . With respect to the comfort utility, it performs on par with LM even in situation where it devotes resources on heating (for vital utility) while LM does not.

It should be noted that the homogeneous case is a kind of worst case for SG . Actually, by design, if all homes have the same $\alpha_k g_{ht}$ for a certain t , SG struggles to break ties between the sets I_0 and I_1 . This is the reason why SG does not outperform LM for very low capacities. Also, the algorithm consumes many iterations to reach its best solution (up to 20 in our experiments). As we are about to see, SG performs better in a heterogeneous case.

C. Results on the heterogeneous case

The results for the heterogeneous case are shown in Figures 4 (class 1) and 5 (class 2).

We first discuss GM . For vital utility, the results are pretty much similar for both classes to the homogeneous case, with maximal value obtained even for low capacities (down to 3×10^4). For the comfort utility, however, one notices that GM leads to better values for class 2 compared to class 1. This is due to the fact that class 2 houses have better energetic performance, so once vital utility is ensured for all, it is more gainful to allocate energy to class 2.

The same reason explains the poor performance of LM . Let us remind that the static allocation is proportional to the maximum power $L(h)$ of homes. So for a given capacity, class 1 homes get more power than class 2 ones. As a result, while performance of class 1 is satisfactory, performance of class 2 is terrible. In particular, the capacity required for class 2 houses to achieve maximal vital utility is very high: $C = 1.7 \times 10^5$, which corresponds to 1700 W per house (regardless the class).

Lastly, we observe that compared to the homogeneous case, the performance of our solution SG is now closer to GM than to LM . In particular, SG manages to take advantage of the heterogeneity to reach high vital utility values more quickly than in Figure 3. Regarding comfort utility, it stays below GM values but manages to give descent values for both classes, which gives a clear advantage over LM (especially regarding the handling of class 2 houses).

(a) Vital utility

(b) Comfort utility

Fig. 4: Average utility per home as a function of the available capacity (heterogeneous case, class 1)

Moreover, *SG* converges faster on the heterogeneous case: the scheme takes between 3 to 8 iterations to find the best allocation over the k_{MAX} iterations.

VI. CONCLUSION

We propose an IoT-based demand response approach, which we name *SG*, that relies on a 2 level control scheme. Intelligence (decision taking) is split between a centralized component and a set of local controllers (one per home). The proposed control approach enables reaching good performance in terms of utility perceived by users while keeping privacy and providing scalability. Moreover, priority is provided to critical needs and fairness among households is introduced.

We show that the approach outperforms schemes where the central controller takes decisions based on the available total capacity and on static (contract based) information about the households. Results for the proposed use cases show that the proposed scheme requires a limited number of iterations to render effective solutions. Moreover, the proposed solution is robust in the sense that the system will keep working even in periods where the proposed algorithms have not converged and in cases where information is delayed or lost.

Future work will encompass a study on the power allocation algorithms for the *SG* scheme considering i) a broader range of classes of appliances and ii) more general cases for the available capacity curb. We will also study the effect of communication impairments on the global performance and on fairness. Finally, we will analyze the cost savings under realistic cost models, looking for solutions that will target minimizing the total expenses a provider will incur in the

(a) Vital utility

(b) Comfort utility

Fig. 5: Average utility per home as a function of the available capacity (heterogeneous case, class 2)

Capacity market while keeping a predefined level of service.

REFERENCES

- [1] M. Held, P. Wolfe, and H. P. Crowder. Validation of subgradient optimization. *Mathematical programming*, 6(1):62–88, 1974.
- [2] IBM Ilog CPLEX optimizer. <http://www-01.ibm.com/software/commerce/optimization/cplex-optimizer/>.
- [3] R. Kaddah, D. Kofman, and M. Pioro. Advanced demand response solutions based on fine-grained load control. *Intelligent Energy Systems (IWIES), 2014 IEEE International Workshop on*, pages 38–45, 2014.
- [4] G. Karmakar, A. Kabra, and K. Ramamritham. Coordinated scheduling of thermostatically controlled realtime systems under peak power constraint. *IEEE 19th RealTime and Embedded Technology and Applications Symposium (RTAS)*, 2013.
- [5] N. Li, L. Chen, and S. H. Low. Optimal demand response based on utility maximization in power networks. *IEEE Power and Energy Society General Meeting*, 2011.
- [6] J. L. Mathieu, M. Kamgarpour, J. Lygeros, and D. S. Callaway. Energy arbitrage with thermostatically controlled loads. *European Control Conference (ECC)*, 2013.
- [7] S. Meyn, P. Barooah, A. Bušić, Y. Chen, and J. Ehren. Ancillary service to the grid using intelligent deferrable loads. *arXiv preprint arXiv:1402.4600*, 2014.
- [8] A. Nayyar, J. Taylor, A. Subramanian, K. Poolla, and P. Varaiya. Aggregate flexibility of a collection of loads π . *IEEE Conference on Decision and Control*, 2013.
- [9] M. Pioro and D. Medhi. *Routing, Flow and Capacity Design in Communication and Computer Networks*. Morgan Kaufmann Publishers, 2004.
- [10] B. M. Sanandaji, H. Hao, K. Poolla, and T. L. Vincent. Improved battery models of an aggregation of thermostatically controlled loads for frequency regulation. *arXiv preprint arXiv:1310.1693*, 2013.
- [11] W. Shi, N. Li, X. Xie, C.-C. Chu, and R. Gadh. Optimal residential demand response in distribution networks. *IEEE journal on selected areas in communications*, 2014.
- [12] C. Vivekananthan, Y. Mishra, G. Ledwich, and F. Li. Demand response for residential appliances via customer reward scheme. *IEEE Trans. Smart Grid*, 5(2):809–820, 2014.