

Effective Consumption Scheduling for Demand-Side Management in the Smart Grid using Non-Uniform Participation Rate

Tri Kurniawan Wijaya, Thanasis G. Papaioannou, Xin Liu, and Karl Aberer
School of Computer and Communication Sciences
École Polytechnique Fédérale de Lausanne (EPFL)
CH-1015 Lausanne, Switzerland
Email: {tri-kurniawan.wijaya, thanasis.papaioannou, x.liu, karl.aberer}@epfl.ch

Abstract—Periods of peak consumer demand in today’s electricity sector are expensive to satisfy and can be the source of power failures. One possible solution is the use of demand-side management (DSM) applying dynamic pricing mechanisms. However, instead of reducing peak loads, these mechanisms can lead to peak-shifting due to the herding effect of consumers’ load-shifting behavior. To overcome this problem, we explore strategies of assigning (non-uniform) participation rates to consumers. We use a generic method to find a near-optimal distribution setting for participation rates. Our method allows DSM designers to tune the system toward consumer convenience. This means less frequent consumption schedule changes, in the price of system performance. In addition, consumers do not need to reveal their detailed consumption schedules (hence, their privacy is preserved). Using experiments, we show the impact of the herding effect and evaluate the effectiveness of the proposed solution. We thereby demonstrate price fairness for consumers. Finally, we apply our solution to a more realistic environment – one where consumers change their consumption behavior every day.

I. INTRODUCTION

Today’s electricity sector faces significant challenges. These include rapidly growing consumption, aging infrastructure, rising prices, and sustainability. Smart grids hold the promise of connecting electricity producers, consumers, and prosumers (who act as both producer and consumer) to a more efficient, reliable, and sustainable power supply system [1]. In order to make this vision a reality, smart grids will have to incorporate several key components, e.g., smart meters for monitoring real-time usage, enhanced communications for real-time control and data flow, appropriate sensors for ensuring system stability, and intelligent energy management systems as a control mechanism for the benefit of all the actors connected to the grid.

Demand-side management (DSM) in a smart grid is a mechanism allowing producers to reduce peak loads by encouraging consumers to make more informed decisions about reducing their bills. For example, Figure 1 shows that peak electricity consumption in a group of Irish residential homes occurs in the evening. Maximum electricity production capacity must be calculated to satisfy peak loads. Lower peak loads mean lower requirements for electricity plant and overall electricity generation, leading to reduced costs and carbon emissions. One significant DSM mechanism is dynamic pricing: in order to




Fig. 1. Hourly average consumption of residential homes in Ireland in November 2009. Peak hours occur in the evening. Data obtained from CER smart-metering trial [12], using 782 residential homes belonging to the control group, with no missing value.

flatten out demand, electricity is priced more expensively at known periods of high consumer demand [2], [3], [4]. On site micro-storage capacity for electricity consumers has also been shown to reduce their costs [5], [6]. Other approaches that have been considered include: service curves, where consumers select their future capacity contracts according to their optimal consumption schedules [7]; demand side bidding, where consumers participate in an auction to fulfill their demand [8], and where residential consumers participate in the electricity market to reduce price volatility [9]; and decision support tools for maximizing consumer benefits from distributed energy resources [10]. A simulation framework to study demand side participation has also been built [11].

This paper focuses on a DSM mechanism where dynamic pricing drives consumers to adjust their consumption behavior. We particularly address the problem caused by the *herding effect*, when the majority of consumers reschedule their electricity consumption towards the same time slots. Dealing with the herding effect is essential, because instead of reducing the peak load, this phenomenon leads to the creation of a new peak, or peak-shifting [3], [13]. Figure 2 is a simple illustration of the herding effect. Our approach benefits from the existence of two-way communication channels between energy companies and consumers in the smart grid. In addition, we preserve consumers’ privacy, because consumers do not need to reveal their detailed consumption schedules.

Our paper’s contribution to solving the herding effect problem can be summarized as follows:




Fig. 2. An example of a herding effect. Initially, the peak hour is at t_1 with cost per unit 8 (assuming total cost is quadratic on load). Afterwards, A and B shift their consumption to t_2 (which previously had cost per unit 0). However, this makes t_2 the new peak (peak-shifting). Cost can be reduced if only A or B (but not both) makes the shift.

- **Non-uniform participation rates.** Applying a strategy of non-uniform participation rates to consumers maximizes the effectiveness of the dynamic pricing DSM mechanism. Furthermore, it is more realistic than uniform participation rates because the approach accommodates a population with different levels of willingness to participate in a DSM program.
- **Realistic consumption model.** We propose a consumption model, which is able to express different usage groups to avoid meaningless scheduling result.
- **Laziness property.** Non-uniform participation rates also enable us to vary strategy settings. We provide a framework for adjusting system laziness (for less frequent schedule changes) in the price of system performance.
- **Experimental evaluation.** We evaluate the price fairness of our proposed solution experimentally, as well as its effectiveness in dynamic environments that are closer to real-world conditions (daily variations in consumer consumption).

The remainder of this paper is organized as follows: In Section II we discuss our consumption model, the pricing mechanism, and our performance metrics. In Section III, we describe our approach for finding a near-optimal participation-rate distribution setting for consumers. In Section IV, we present our experimental results. We discuss the relevance of our work in Section V, and finally, we deliver our conclusions in Section VI.

II. CONCEPTUAL FRAMEWORK

This section describes our consumption and load profile model. We also introduce the dynamic pricing mechanism and define participation rate. As others have proposed previously, automated intelligent agents can be used to represent consumers in the scheduling, making the entire process seamless from the consumers' perspective [5], [9], [13]. Finally, we define performance metrics used for evaluation. Table I summarizes the key notation used in this and subsequent sections.

A. Consumption Profile

Our proposed consumption model aims to realistically express and distinguish the different times of the day when *usage groups* use appliances as well as the right time to shift

TABLE I. NOTATIONS SUMMARY

Notations	Description
\mathcal{N}	Set of consumers.
\mathcal{A}	Set of all appliances.
A_i	Set of appliances owned by consumer i .
\mathcal{T}	Set of time slots during a day.
$\mathcal{U}^{i,a}$	Consumer i 's consumption profile for appliance $a \in \mathcal{A}_i$.
$u_j^{i,a}$	The j^{th} usage group of consumer i for appliance a .
$S_j^{i,a}$	Set of possible starting times of consumer i 's appliance a for load shifting.
$s_j^{i,a}$	The normal starting time of appliance a preferred by consumer i .
$d_j^{i,a}$	The usage duration of consumer i 's appliance a .
$\mathbf{I}_j^{i,a}$	The realization load of consumer i for appliance a in each usage group.
$L(t)$	The total load of all consumers at time slot $t \in \mathcal{T}$.
$L^i(t)$	Consumer i 's load at time slot $t \in \mathcal{T}$.
$c^i(t)$	The cost of consumer i for time slot t .
$c(i)$	The daily cost of consumer i .
ρ_i	The participation rate of consumer i .

that use. Failure to distinguish such usage groups leads to meaningless schedules.

Example 1: Consider a consumer who needs to schedule x morning time slots (usage group 1) and y evening time slots (usage group 2) for appliance a . She declares $x + y$ time slots available for load shifting in the morning and $x + y$ time slots in the evening. A model that fails to distinguish between different usage groups requirements might schedule $x + y$ time slots in the morning (because the morning electricity price is cheaper), which is unintended.

In addition, we need to schedule an appliance in a continuous range, because turning an appliance off during its usage time inconveniences the user. However, this does not reduce the expressiveness of our model. When a consumer intends discontinuous usage of an appliance, we can use different usage groups.

Let \mathcal{N} be the set of all consumers and \mathcal{A} be the set of all appliances. We denote the set of appliances owned by user $i \in \mathcal{N}$ by $A_i \subseteq \mathcal{A}$. Power consumption (in Watts) for each appliance $a \in \mathcal{A}$, is defined as the function $w : a \rightarrow \mathbb{R}$. The *consumption profile* of a consumer i is defined as a set of usage groups, $\mathcal{U}^{i,a}$, for $a \in A_i$. Let $\mathcal{T} = \{t_1, \dots, t_n\}$ denote a set of *time slots* during a day. We hence define each usage group $u_j^{i,a} \in \mathcal{U}^{i,a}$, for $1 \leq j \leq |\mathcal{U}^{i,a}|$, as a tuple $(S_j^{i,a}, s_j^{i,a}, d_j^{i,a})$, where $S_j^{i,a} \subseteq \mathcal{T}$ is the set of *possible starting times* for load-shifting, $s_j^{i,a} \in S_j^{i,a}$ is the consumer's preferred *normal starting time*, and $d_j^{i,a} \in \mathbb{R}$ is the *usage duration* of appliance a (in hours). Appliance a of consumer i is:

- (i) *shiftable* within usage group $u_j^{i,a}$, if $\{s_j^{i,a}\} \subset S_j^{i,a}$, and
- (ii) *non-shiftable* within usage group $u_j^{i,a}$, if $\{s_j^{i,a}\} = S_j^{i,a}$.

B. Load

The *realized load* of consumer i for appliance a in each usage group $u_j^{i,a} \in \mathcal{U}^{i,a}$, denoted as $\mathbf{I}_j^{i,a} = (l_{j,1}^{i,a}, \dots, l_{j,|\mathcal{T}|}^{i,a})$, is her actual energy consumption in watt-hours (Wh). Function $\pi : \mathcal{T} \rightarrow \mathbb{R}$ maps time slots to hours. In addition, we require the realized load to satisfy:

- *Starting-time correctness:* Let the starting time of $\mathbf{I}_j^{i,a}$

be denoted by $l_{j,I}^{i,a}$, where $\forall x, 0 \leq x < I \leq |\mathcal{T}|$, $l_{j,x}^{i,a} = 0$. We require $l_{j,I}^{i,a} \in S_j^{i,a}$.

- *Duration correctness*: $(\sum_{y=1}^{|\mathcal{T}|} l_{j,y}^{i,a}) = d_j^{i,a} \cdot w(a)$.
- *Load continuity*: $\forall x \leq y$, if $l_{j,x}^{i,a} \neq 0 \wedge l_{j,y}^{i,a} \neq 0$, then $l_{j,z}^{i,a} \neq 0$ where $x < z < y$.
- *Load maximality*: $l_{j,t}^{i,a} \leq w(a) \cdot \pi(t)$ for all $t \in \mathcal{T}$.

We define the load at time slot $t \in \mathcal{T}$ for consumer i as:

$$L^i(t) = \sum_{a \in \mathcal{A}_i} \left(\sum_{u_j^{i,a} \in \mathcal{U}^{i,a}} l_{j,t}^{i,a} \right). \quad (1)$$

Then, the total load at time slot $t \in \mathcal{T}$ over all consumers is $L(t) = \sum_{i \in \mathcal{N}} L^i(t)$.

C. Pricing Mechanism

We require our price function to be *a) increasing* and *b) strictly convex* on the total load for each time slot. Assuming that consumers prefer lower prices, price functions that satisfy the requirements above encourage consumers to shift their consumption from peak time slots to off-peak time slots. An example of a real energy cost function that satisfies both the above assumptions is the quadratic cost function for thermal generators [3], [14], [15]:

$$p(L(t)) = c_1 L(t)^2 + c_2 L(t) + c_3. \quad (2)$$

However, any function satisfying the two requirements above can also be used.

The cost to consumer i for time slot t is calculated as proportional to her load over total load at a time slot t :

$$c^i(t) = \frac{L^i(t)}{L(t)} p(L(t)). \quad (3)$$

The daily cost to consumer i is the sum of her costs over all time slots:

$$c(i) = \sum_{t \in \mathcal{T}} c^i(t). \quad (4)$$

Since time slots with higher loads will be charged at a higher price, the consumer whose objective is to minimize cost will be encouraged to shift consumption to the time slot with the lowest load. Note that a consumer does not require a detailed knowledge of other consumers' consumption. The only information needed by consumer i to minimize $c(i)$, is the price function and the aggregate load of all consumers for each time slot, $L(t)$, which the energy company communicates to consumers.

D. Participation Rate

Given the cost function $c(i)$ in Eq. 4, a rational consumer would aim to minimize costs, i.e., choose the realized load $l_{j,I}^{i,a}$, for all $u_j^{i,a} \in \mathcal{U}^{i,a}$ and for all $a \in \mathcal{A}_i$, such that $c(i)$ is minimized. However, if a majority of consumers alter their consumption at the same time, moving to less loaded time slots, these time slots might become the new peaks. To avoid this peak-shifting problem, we introduce the notion of *participation rate*, which is the probability that a consumer will alter her consumption schedule. For each consumer i , we denote $\rho_i \in \mathbb{R}$ as her participation rate, where $0 \leq \rho_i \leq 1$.

E. Performance Metrics

We evaluate the DSM system's performance using two metrics, which are also used below in the experimental section of this paper.

1) *Peak to average ratio*: We measure DSM system's performance by how much we can reduce the peak load. One widely used measurement is *peak to average ratio* (PAR): the day's peak load divided by the average load on the day. Using our model, we define:

$$PAR = \frac{|\mathcal{T}|}{\sum_{t \in \mathcal{T}} L(t)} \max_{t \in \mathcal{T}} L(t) \quad (5)$$

In general, a high peak load (relative to the average) results in a high PAR. Hence, a lower PAR is preferable. Another commonly used metric, called *load factor*, is $\frac{1}{PAR}$. However, in the sequel, we use PAR instead. We denote $PAR_x(k)$ as the PAR of the DSM system x at day k . Even though PAR can be highly effective measure DSM system performance on a particular day, there is difficulty in deciding which system is better (lower PAR) than the other when their daily PARs overlap.

Example 2 (Overlapping PAR): Consider two DSM systems, A and B . Assume that $PAR_A(k) = 2$, if k is odd, and $PAR_A(k) = 4$, if k is even. Furthermore, $PAR_B(k) = 5$, if k is odd, and $PAR_B(k) = 3$, if k is even. A snapshot of the system on day 1 suggests A is better than B , while snapshot on day 2 suggests the opposite.

2) *Area under PAR curve*: The *area under PAR curve* (AUP) metric is more appropriate when dealing with multiple DSM systems exhibiting the overlapping PARs problem above.

Given a set of measurements of system x on a set of days D , we define the AUP of x as the the sum of its PARs, i.e.

$$AUP_x(D) = \sum_{i \in D} PAR_x(i) \quad (6)$$

A DSM system x is better than y relative to a set of days D if $AUP_x(D) < AUP_y(D)$.

Example 3: Consider again DSM systems A and B from Example 2. Given a set of days $D = \{1, \dots, 100\}$, we know that A is better than B for the period D , because $AUP_A(D) = 300$ and $AUP_B(D) = 400$.

III. FINDING NEAR-OPTIMAL STRATEGY DISTRIBUTION

Consumers' strategy distribution is defined as a distribution of participation rates adopted by the consumers. It is *optimal* if it brings the system to its optimum performance (based on a predefined metric), given consumers' consumption profiles. A key question we then need to answer is, *how can one find the best strategy distribution according to a predefined objective?*

When we consider the possibility of assigning a particular participation rate – which is a real number – to each consumer, we face a problem with an infinite space. In order to be able to explore the solution space, we make the participation rate discrete. We use a Monte Carlo Tree Search (MCTS) with Upper Confidence-bounds applied to Trees (UCT) to find a near-optimal solution [16]. The method is widely known through

its application in multi-armed bandit problem [16] and general game playing [17]. While the basic idea of MCTS is to traverse the representation tree in a random fashion, the role of UCT is to guide these random traversals. Instead of a completely random exploration of the tree, we gradually explore good branches (exploitation). After a number of exploitations on the good branches, the algorithm switches to explore bad branches (in case there was an unexplored good branch there).

Problem definition: Given a set of possible participation rates $P = \{r_i \mid r_i \in \mathbb{R}, 0 \leq r_i \leq 1\}$, *assignment granularity* of $g \in \mathbb{N}$, and *assignment* $0 \leq v_i \leq g$, we must compute the fraction $f_i = \frac{v_i}{g}$ of consumers for each $r_i \in P$ such that assigning f_i of consumers' satisfies the predefined objective (e.g., minimum PAR).

A. Strategy-Tree

We cast the problem of finding the best possible strategy distribution into a tree state-space search problem. The key idea is the gradual assignment of a fraction of consumers to a particular participation rate until all consumers have been assigned to a particular participation rate. As we explore the next level of the tree, the chosen child represents a proportion of consumers assigned to a particular participation rate, i.e., the choice made at level k defines the fraction f_i of consumers who will be assigned participation rate r_k . When we arrive at the leaf, it means that all consumers have been assigned to their respective participation rates. The quality of the solution found is a function of the path that we choose to traverse the tree from the root to the leaf.

Definition 4 (Strategy-Tree): Let the set of possible participation rates be $P = \{r_i \mid r_i \in \mathbb{R}, 0 \leq r_i \leq 1\}$, and the *assignment granularity* be $g \in \mathbb{N}$. This determines the fine-grained fraction of consumers assigned to participation rates.

- *Node.* A node α in the tree is a set $\{v_1^\alpha, \dots, v_{|P|}^\alpha \mid v_i^\alpha \in \mathbb{N}, 0 \leq v_i^\alpha \leq g, 1 \leq i \leq |P|\}$, where v_i^α denotes a fraction $\frac{v_i^\alpha}{g}$ of the consumers assigned to r_i .
- *Root.* The root r of the tree is a node in level 0 and has zero assignment for all elements, i.e. $v_i^r = 0$, for $1 \leq i \leq |P|$.
- *Node sum.* For a node α , we denote the sum of its assignment as $\sigma(\alpha) = \sum_{i=1}^{|P|} v_i^\alpha \leq g$.
- *Leaf.* A node α in the tree is a leaf, if $\sigma(\alpha) = g$.
- *Children.* Each non-leaf node α at level k has $g - \sigma(\alpha) + 1$ children, i.e. $\gamma_0(\alpha), \dots, \gamma_{(g-\sigma(\alpha))}(\alpha)$. If $\gamma_j(\alpha)$ is a non-leaf node, then $\gamma_j(\alpha) = \{v_1^\alpha, \dots, v_{k+1}^\alpha + j, \dots, v_{|P|}^\alpha\}$. Intuitively, choosing $\gamma_j(\alpha)$ means that we assign $\frac{j}{g}$ of consumers to the participation rate $r_k \in P$, i.e., $f_k = \frac{j}{g}$. Hence, the depth of the tree is bounded by $|P|$.

Example 5: Assume that we have $P = \{0.0, 0.3, 0.7, 1.0\}$, and $g = 10$. Then, we have $r = (0, 0, 0, 0)$ as the root. The root r has 11 children, i.e. $\gamma_0(r), \dots, \gamma_{10}(r)$, where $v_1^{\gamma_j(r)} = j$, and $v_i^{\gamma_j(r)} = 0$, for $2 \leq i \leq 4$ and $0 \leq j \leq 10$.

B. Node Value

For a non-leaf node α , its valuation $Q(\alpha)$ is defined as the average of the value of its children. For a leaf node, its value is determined by evaluating the participation-rate distribution encoded in the node with a DSM simulation for a specific duration (e.g., our experiments simulate a 365-days period).

Let ρ_i be the participation rate of consumer i . At each simulation step, consumer i with probability ρ_i adjusts her consumption schedule, i.e., realized load (defined in Section II-B) to minimize her cost $c(i)$ as in Eq. 4. Once the simulation is finished, the system performance (such as PAR or AUP) is set as the value of this leaf node.

Another key strength of UCT is that, since the evaluation here can be based on any predefined metrics, we can tailor it to match any objective we have in mind. As we will show below (see Section IV), we focus on PAR (and AUP) since it is widely used in the literature. We can also use other metrics if needed, such as utilitarian social welfare (total cost), or egalitarian social welfare (maximum cost of a consumer).

C. Tree Traversal

For a node α , we define $\Gamma(\alpha)$ as the set of all children of α . Given that $\gamma \in \Gamma(\alpha)$, we define:

- $Q(\gamma)$ as the valuation of γ ,
- $N(\alpha)$ as the number of times we visit α ,
- $N(\alpha, \gamma)$ as the number of times we visit γ from α .

From α , the next child to be visited, γ^* , is chosen by:

$$\gamma^* = \operatorname{argmax}_{\gamma \in \Gamma(\alpha)} \left\{ Q(\gamma) + \theta \sqrt{\frac{\ln(N(\alpha))}{N(\alpha, \gamma)}} \right\} \quad (7)$$

The above equation captures the balance of exploration and exploitation of UCT. On one side, the first term, $Q(\gamma)$, captures the quality of the next child (the higher the value, the better the node), and increases the chance that it will be visited (exploitation of good nodes). On the other side, the fewer visits to a child, the larger the value of the second term. This increases its chance to be visited (exploration of under-visited nodes). Thus, θ is the *balancing parameter* to adjust how aggressive the exploration is.

D. Regularization

1) *Laziness:* Lower participation rates means less frequent consumption schedule changes. Adopting a low participation rate strategy while maintaining high benefit (low energy cost) is attractive to consumers. Hence, a solution that assigns more consumers to low participation rates (*laziness property*) is desirable.

Let ρ_i be a participation rate adopted by consumer i . For a DSM system x , and its set of consumers N_x , its *laziness score* is defined as the average of the participation rates adopted, i.e.

$$\lambda(x) = \frac{\sum_{i \in N_x} \rho_i}{|N_x|}. \quad (8)$$

Given two systems x and y , x is *lazier* than y if $\lambda(x) < \lambda(y)$.




Fig. 3. UCT performances under different balancing parameter, θ , with 95% confidence interval. Measured by AUP (the lower the better). Inset: detailed view when UCT starts to flatten out.

2) *Regularization*: In order to have a lazier system, we introduce a *regularization* parameter to our UCT approach by slightly modifying the selection mechanism of the next child to be visited. The lazier the system encoded in a node, the lower the penalty incurred (see the second term of Eq. 9). Given α as the current node, P as the set of possible participation rates, and $\gamma = \{v_1^\gamma, \dots, v_{|P|}^\gamma\} \in \Gamma(\alpha)$ as a child of α , we replace $Q(\gamma)$ in Eq. 7 with $Q^\lambda(\gamma)$, where

$$Q^\lambda(\gamma) = Q(\gamma) - \frac{\sum_{i=1}^{|P|} (v_i^\gamma \cdot r_i)}{\sum_{i=1}^{|P|} v_i^\gamma}. \quad (9)$$

The first term is the valuation of γ , The second term is the laziness score calculation of the system encoded in γ .

IV. EXPERIMENTS

Our experiments use hourly time slots throughout the day, i.e. $\mathcal{T} = \{1, \dots, 24\}$. Unless otherwise stated, we set the number of consumers as $|\mathcal{N}| = 100$. The dataset we use to plot Figure 1 does not have a more detailed customer load profile, such as appliance-level measurement. Thus, we define the consumers' consumption profiles based on [18], [19]. See Table II for example.¹ In addition, as background knowledge, we add several usage groups and load shifting constraints. For the price function, we set $c_1 = 0.01$, $c_2 = 2$, $c_3 = 100$ (see Eq. 2). However, as we mention in Section II, any increasing and strictly convex function of the total load for each time slot can also be used. The simulation is set to 365 days.

A. Parameter Tuning

1) *UCT configuration*: This experiment investigates which value of θ (see Section III-C about Tree Traversal) will give us the best result. We measure UCT performance based on AUP using different θ , i.e., $\theta = 1$, $\theta = 10$, $\theta = 100$, $\theta = 500$, and $\theta = 1000$. We do not try $\theta > 1000$ because our experiments show that the maximum AUP reached by the system is around 850; hence, giving $\theta = 1000$ is more than sufficient for the exploration. We use a set of possible participation rates $P = \{0.0, 0.1, \dots, 1.0\}$ and we run the experiments up to 1,000 tree traversals for each θ setting.

Figure 3 shows that the AUPs flatten out after 80 traversals and remain stable up to 1,000 traversals. For the rest of the experiments, we use $\theta = 1000$. Up to 1,000 traversals, $\theta =$




Fig. 4. PAR measurement of the best solution found by UCT with 1,000 traversals and balancing parameter $\theta = 500$ under different participation rate granularities.




Fig. 5. Cumulative distribution of participation rate assigned by UCT under different possible participation rate granularities. There is no assignment for participation rate larger than 0.5.

100, $\theta = 500$, or $\theta = 1000$ give us equal performance, and are better than $\theta = 1$, or $\theta = 10$. For the rest of the experiments, we use $\theta = 500$.

2) *Set of possible participation rates*: In this experiments, we investigate the impact of the granularity of the set of possible participation rates. We set three different granularities, i.e., $P_1 = \{0.0, 0.1, \dots, 1.0\}$, $P_2 = \{0.0, 0.05, \dots, 1.0\}$, and $P_3 = \{0.0, 0.02, \dots, 1.0\}$. Figure 4 shows that using P_1 the PAR decreases rapidly in the beginning, but P_2 and P_3 outperform it in the long run.

In order to have a better understanding of the distribution of participation rate among consumers, we plot its cumulative distribution. Figure 5 shows that in P_1 , P_2 , and P_3 more than 90% of consumers are assigned to a participation rate of 0.1 or less, and we assign no consumers to high participation rates (> 0.5). This means that our solution has the desirable property of assigning the majority of consumers to low participation rates (infrequent schedule changes).




Fig. 6. Price level per kWh paid on average (95% confidence interval) by 1,000 consumers grouped by their participation rate. Participation rate distribution used is the result from P_3 as shown in Figure 5.

¹Our source code is available at <https://github.com/tritritri/effective-dsm>

TABLE II. CONSUMER CONSUMPTION PROFILE EXAMPLE

Appliance	Usage groups	Possible starting time	Normal starting time	Duration
Microwave	usage group 1	07:00 - 09:00	07:00	15 minutes
	usage group 2	18:00 - 21:00	19:00	20 minutes
Television	usage group 1	19:00 - 22:00	20:00	1 hour 30 minutes
Washing machine	usage group 1	16:00 - 22:00	19:00	1 hour 25 minutes
Light (non-shiftable)	usage group 1	18:00	18:00	7 hours
Dishwasher	usage group 1	02:00 - 16:00, 21:00 - 23:00	21:00	40 minutes




Fig. 7. PAR of three different approaches (UCT, uniform 0.05, and turn-taking), and PAR due to the herding effect (all consumers responded).

B. Price Fairness

The participation rate is a new concept in efforts to help reduce PAR and avoid peak-shifting due to the herding effect. However, one might suspect that different participation rates might introduce price unfairness among consumers. Figure 6 shows the price paid by consumers grouped by participation rates. We take the distribution of participation rates obtained for P_3 (see Figure 5) for 1,000 consumers on 365 days. At the end of day 365, we group the consumers having the same participation rate and compute the average price paid per kWh (kilowatt-hour) for each group.

The average price paid per kWh by different participation rate groups was 0.1510 with a standard deviation 0.0007. This fairness (especially for consumers with a low participation rate) is brought about because as long as there are consumers who shift from the peak time slot, the price for that slot will decrease (the price for the other slots will slightly increase). As a result, the price difference between slots also decrease.

C. The Herding Effect and Other Approaches

Figure 7 shows the PAR results due to the herding effect (denoted as *all response*). When we have the herding effect, instead of converging to the minimum, the PAR of the system oscillates. Figure 7 also shows the result found by UCT, and other approaches, i.e., assigning a participation rate 0.05 to all consumers [13] (denoted as *uniform 0.05*), and allowing only one consumer per day to change her consumption schedule [3] (denoted as *turn-taking*). We run the experiments with 1,000 consumers over 365 days.

The performance of the solution found by UCT and *uniform 0.05* outperform *turn-taking*. This is because the response of one consumer per day has too little effect on the system. The PARs shown by both, UCT and *uniform 0.05*, are very competitive (with AUP measured 418.2 and 419.18 respectively). In the solution found by UCT, more than 87% of the population has participation rate 0.06 or less. While Ramchurn




Fig. 8. PAR measured on static and dynamic environment with 20% and 70% usage deviation. Eventually, their PAR reach about the same level. We have more oscillation on the environment with higher usage deviation.

et al. shows empirically that participation rates should be reasonably small [13], we validate it here algorithmically. In addition, our proposal to introduce non-uniform participation rates enables us to accommodate a population with different levels of willingness to participate (which is more realistic) in a DSM program.

D. Dynamic Environments

These experiments evaluate the participation rate distribution found by UCT using P_3 in dynamic environments (see Figure 5), where appliance usage durations vary from day to day (as in the real world). Here, we define $x\%$ usage deviation as a range $[-\frac{x}{100}d_j^{i,a}, +\frac{x}{100}d_j^{i,a}]$ that consumer i can add to her normal duration, $d_j^{i,a}$, of using appliance a within usage group $u_j^{i,a}$ (defined in Section II-A).

Figure 8 shows that although all settings eventually reach roughly the same PAR, systems with higher deviation have more unstable PAR. On one hand, this shows that our solution also works to lower PAR in dynamic environments. On the other hand, this is a warning to the energy companies: assuming that the deployment of a DSM tool is successful and we enjoy low PAR, back-up generators are still needed to satisfy the demand levels of any unexpected peak-loads that may occur. The performance of UCT compared to other approaches in this environment exhibits the same pattern as in Figure 7, where UCT and *uniform 0.05* have similar performance, but they both outperform *turn-taking*.

E. Regularization and Laziness

The main reason for having regularization and laziness in the system is because lower participation rates are preferable (less response needed). To analyze the system further, under different regularization parameters, we run experiments with $\beta = 10$, $\beta = 100$ and without regularization, i.e., $\beta = 0$. Figure 9 shows that by applying $\beta = 100$, almost half of




Fig. 9. Cumulative distribution of participation rate assigned by UCT with different regularization parameter β . There is no assignment for participation rate larger than 0.5.

consumers are assigned to participation rate 0.02 (almost half of consumers need to respond once every 50 days).

In the regularized version of UCT, we introduce a penalty term when evaluating the leaf nodes, which might not allow the optimal solution selection. Thus, we should observe the AUP of the systems before making any further conclusion. At the end of day 365, AUPs of the system with $\beta = 0$, $\beta = 10$, $\beta = 100$ are 421.84, 422.63, and 459.54 respectively. We observe an almost 10% difference between the system with $\beta = 0$ and $\beta = 100$. Depending on how DSM stakeholders value laziness, one might adjust the β higher or lower as a trade-off between an optimal (in terms of PAR or AUP) or a lazy system.

V. RELATED WORK

Palensky and Dietrich [20] and Strbac [21] gave an overview of DSM. Mohsenian-Rad et al. proposed a game theoretic model for DSM, and proved the systems' convergence [3]. In order to deal with the herding effect, they proposed turn-taking mechanisms, which require a centralized control. Li, Chen, and Low extended this model by taking into account consumer utility and inconvenience [2]. In our model, a consumer can define her flexibility to shift consumption in time. Within these shiftability constraints, no deterioration of utility is assumed. Ramchurn et al. proposed to set all consumers' participation rate to 0.05 and showed empirically that the herding effect can be handled using a small enough participation rate [13]. In this study, we extend their work by solving the herding effect algorithmically and offering a framework for DSM designers to adjust their systems towards either consumer convenience or system performance.

Previous works have proposed using additional devices for DSM. Vytelingum et al. [5] and Mishra et al. [22] presented a DSM mechanism which assumed the existence of energy storage on consumers' premises. Ganu et al. developed a device called *nPlug*, which can perform scheduling and voltage sensing to identify peak and off-peak periods on the grid [23].

The large battery capacity of future fleets of electric vehicles (EVs) has been considered as a significant extra load that will have to be satisfied. Charging one EV battery can consume 32kWh (comparable to one household's daily consumption) in just a few hours [24], [25]. Too many EVs charging their batteries simultaneously could create a new peak load. Several approaches have been proposed to mitigate this [26], [27], [28]. However, instead of seeing EVs as one of the future's key

challenges, the Vehicle-to-Grid (V2G) concept sees them as an opportunity: their battery capacity could be leveraged for DSM by discharging power back into the grid when supply is short [29], [30], [31].

VI. CONCLUSION

We used Monte Carlo Tree Search with UCT to find near-optimal (non-uniform) participation rate distributions to address the problem of peak-shifting caused by the herding effect. In addition, we exploited the laziness property of a DSM system to have more consumers assigned to low participation rates. A low participation rate means less frequent schedule changes. Thus, the lazier the system, the more convenient it is for consumers. However, laziness comes at a price – it reduces system performance. For this purpose, we provided a regularization parameter which can be used to align the system towards performance or laziness.

With respect to previous approaches in the literature, our method combines the best of both worlds: it enables a fast convergence solution (as in *uniform 0.05*), and admits less frequent schedule changes (as in *turn-taking*). In addition, we validated algorithmically the result found empirically by [13]. We also demonstrated that our solution ensures the price fairness property. In the future, we plan to (i) provide a mechanism to maintain the participation rate distribution in the population, (ii) obtain the desirable properties using a distributed or evolutionary mechanism, and (iii) convert the load data into symbolic representation to reduce data numerosity and granularity (which also gives additional privacy protection for consumers) [32].

VII. ACKNOWLEDGMENTS

We would like to thank partners in Wattalyst project (www.wattalyst.org), Sandeep Kalra, and the anonymous reviewers for their helpful comments. The research leading to these results has received funding from the European Union's Seventh Framework Programme (FP7/2007-2013) under grant agreement no: 288322 (Wattalyst).

REFERENCES

- [1] EG1, "Expert group 1: Functionalities of smart grids and smart meters," Final Deliverable. EU Commission Task Force for Smart Grids, 2010.
- [2] N. Li, L. Chen, and S. H. Low, "Optimal demand response based on utility maximization in power networks," in *Proceedings of the 2011 IEEE Power & Energy Society General Meeting*, July 2011.
- [3] A. 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," *IEEE Transactions on Smart Grid*, vol. 1, no. 3, pp. 320–331, Dec. 2010.
- [4] P. Samadi, R. Schober, and V. Wong, "Optimal energy consumption scheduling using mechanism design for the future smart grid," in *IEEE International Conference on Smart Grid Communications (SmartGridComm)*, Oct. 2011, pp. 369–374.
- [5] P. Vytelingum, T. Voice, S. D. Ramchurn, A. Rogers, and N. R. Jennings, "Theoretical and practical foundations of large-scale agent-based micro-storage in the smart grid," *J. Artif. Intell. Res. (JAIR)*, vol. 42, pp. 765–813, 2011.
- [6] X. Vallve, A. Graillot, S. Gual, and H. Colin, "Micro storage and demand side management in distributed PV grid-connected installations," in *9th International Conference on Electrical Power Quality and Utilisation, 2007. EPQU 2007.*, Oct. 2007, pp. 1–6.

- [7] J.-Y. Le Boudec and D.-C. Tomozei, "Demand response using service curves," in *IEEE PES Innovative Smart Grid Technologies*, 2011.
- [8] T. K. Wijaya, K. Larson, and K. Aberer, "Matching demand with supply in the smart grid using agent-based multiunit auction," in *2013 Fifth International Conference on Communication Systems and Networks (COMSNETS)*, 2013, pp. 1–6.
- [9] H. S. Oh and R. Thomas, "Demand-side bidding agents: Modeling and simulation," *IEEE Transactions on Power Systems*, vol. 23, no. 3, pp. 1050–1056, Aug. 2008.
- [10] M. Pedrasa, T. Spooner, and I. MacGill, "Coordinated scheduling of residential distributed energy resources to optimize smart home energy services," *IEEE Transactions on Smart Grid*, vol. 1, no. 2, pp. 134–143, Sep. 2010.
- [11] T. K. Wijaya, D. Banerjee, T. Ganu, D. Chakraborty, S. Battacharya, T. Papaioannou, D. P. Seetharam, and K. Aberer, "DRSim: A Cyber Physical Simulator for Demand Response Systems," in *IEEE International Conference on Smart Grid Communications (SmartGridComm)*, Oct. 2013.
- [12] CER, "Smart metering trial data publication," 2012.
- [13] S. D. Ramchurn, P. Vytelingum, A. Rogers, and N. Jennings, "Agent-based control for decentralised demand side management in the smart grid," in *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 1*, ser. AAMAS '11. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, 2011, pp. 5–12.
- [14] M. Gonzalez Chapa and J. Vega Galaz, "An economic dispatch algorithm for cogeneration systems," in *IEEE Power Engineering Society General Meeting*, 2004, pp. 989–994 Vol.1.
- [15] P. Rao and S. Nagendra, "Combined heat and power economic dispatch: A direct solution," *Electric Power Components and Systems*, vol. 34, no. 9, pp. 1043–1056, 2006.
- [16] L. Kocsis and C. Szepesvári, "Bandit based monte-carlo planning," in *ECML*, ser. Lecture Notes in Computer Science, J. Fürnkranz, T. Scheffer, and M. Spiliopoulou, Eds., vol. 4212. Springer, 2006, pp. 282–293.
- [17] H. Finnsson and Y. Björnsson, "Simulation-based approach to general game playing," in *Proceedings of the 23rd national conference on Artificial intelligence - Volume 1*, ser. AAAI'08. AAAI Press, 2008, pp. 259–264.
- [18] S. Gottwalt, W. Ketter, C. Block, J. Collins, and C. Weinhardt, "Demand side management – A simulation of household behavior under variable prices," *Energy Policy*, vol. 39, no. 12, pp. 8163 – 8174, 2011.
- [19] N. R. C. Office of Energy Efficiency, "Energy consumption of major household appliances shipped in Canada," Dec. 2005.
- [20] P. Palensky and D. Dietrich, "Demand side management: Demand response, intelligent energy systems, and smart loads," *IEEE Transactions on Industrial Informatics*, vol. 7, no. 3, pp. 381–388, Aug. 2011.
- [21] G. Strbac, "Demand side management: Benefits and challenges," *Energy Policy*, vol. 36, no. 12, pp. 4419 – 4426, 2008.
- [22] A. Mishra, D. Irwin, P. Shenoy, and T. Zhu, "Scaling distributed energy storage for grid peak reduction," in *Proceedings of the 4th International Conference on Future Energy Systems*, ser. e-Energy '13. New York, NY, USA: ACM, 2013, pp. 3–14.
- [23] T. Ganu, D. P. Seetharam, V. Arya, R. Kunnath, J. Hazra, S. A. Husain, L. C. De Silva, and S. Kalyanaraman, "Nplug: A smart plug for alleviating peak loads," in *Proceedings of the 3rd International Conference on Future Energy Systems*, ser. e-Energy '12. New York, NY, USA: ACM, 2012, pp. 30:1–30:10.
- [24] S. D. Ramchurn, P. Vytelingum, A. Rogers, and N. R. Jennings, "Putting the 'smarts' into the smart grid: A grand challenge for artificial intelligence," *Communications of the ACM*, vol. 55, no. 4, pp. 86–97, Apr. 2012.
- [25] A. Hess, F. Malandrino, M. B. Reinhardt, C. Casetti, K. A. Hummel, and J. M. Barceló-Ordinas, "Optimal deployment of charging stations for electric vehicular networks," in *Proceedings of the first workshop on Urban networking*, ser. UrbaNe '12. New York, NY, USA: ACM, 2012, pp. 1–6.
- [26] O. Ardakanian, C. Rosenberg, and S. Keshav, "Distributed control of electric vehicle charging," in *Proceedings of the 4th international conference on Future energy systems*, ser. e-Energy '13. New York, NY, USA: ACM, 2013, pp. 101–112.
- [27] Z. Ma, D. Callaway, and I. Hiskens, "Decentralized charging control of large populations of plug-in electric vehicles," *IEEE Transactions on Control Systems Technology*, vol. 21, no. 1, pp. 67–78, 2013.
- [28] E. Sortomme, M. Hindi, S. MacPherson, and S. Venkata, "Coordinated charging of plug-in hybrid electric vehicles to minimize distribution system losses," *IEEE Transactions on Smart Grid*, vol. 2, no. 1, pp. 198–205, 2011.
- [29] C. Guille and G. Gross, "A conceptual framework for the vehicle-to-grid (V2G) implementation," *Energy Policy*, vol. 37, no. 11, pp. 4379 – 4390, 2009.
- [30] S. Kamboj, W. Kempton, and K. S. Decker, "Deploying power grid-integrated electric vehicles as a multi-agent system," in *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 1*, ser. AAMAS '11. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, 2011, pp. 13–20.
- [31] S. Vandael, T. Holvoet, and G. Deconinck, "A multi-agent system for managing plug-in hybrid vehicles as primary reserve capacity in a smart grid," in *Proceedings of the Second International Workshop on Agent Technologies for Energy Systems (ATES 2011)*, May 2011.
- [32] T. K. Wijaya, J. Eberle, and K. Aberer, "Symbolic representation of smart meter data," in *Proceedings of the Joint EDBT/ICDT 2013 Workshops*, ser. EDBT '13. New York, NY, USA: ACM, 2013, pp. 242–248.