

Cooperatives for Demand Side Management

Ramachandra Kota¹ Georgios Chalkiadakis² Valentin Robu³ Alex Rogers³ and Nicholas R. Jennings³

Abstract. We propose a new scheme for efficient demand side management for the Smart Grid. Specifically, we envisage and promote the formation of cooperatives of medium-large consumers and equip them (via our proposed mechanisms) with the capability of regularly participating in the existing electricity markets by providing electricity demand reduction services to the Grid. Based on mechanism design principles, we develop a model for such cooperatives by designing methods for estimating suitable reduction amounts, placing bids in the market and redistributing the obtained revenue amongst the member agents. Our mechanism is such that the member agents have no incentive to show artificial reductions with the aim of increasing their revenues.

1 Introduction

In recent years, with environmental and economic concerns regarding energy sustainability becoming increasingly important, research in AI and multiagent systems has, with a growing pace, been taking up the challenge of implementing the vision of the *Smart Grid* [10, 13]—creating robust, intelligent electricity supply and distribution networks to achieve the highest energy efficiency possible.

Virtual Power Plants (VPPs), in particular, are expected to play a crucial role interconnecting and automatically dispatching distributed energy generation, storage, or other demand-side resources, via the use of intelligent software components [2, 7]. On one hand, VPPs have been hailed as a means to achieve the incorporation of the numerous distributed renewable energy generation resources (such as small-to-medium scale wind or solar power generators), into reliable large-scale entities mirroring the operation of conventional power plants [4, 5, 9]. On the other hand, the term VPP has also been widely used, primarily in North America, to denote the amalgamation of consumers acting as “power plants” attempting to counter the effects of peak-time consumption—via participation in “critical peak pricing programs”, or the provision of demand-response consumption reduction services [2]. That is, VPPs of (mainly household) consumers might be rewarded with better consumption rates for reducing their energy demand over some period; or, VPPs of (mainly industrial) consumers, managed by a specialised intermediary company⁴ offering demand-response services, agree, for a cash reward, to step in and contribute to the “trimming down” of the demand curve in the event of an impending critical period [1].

In this paper, we, as well, advocate the use of VPPs of energy consumers to contribute to energy demand reduction. However, we go one step further — rather than dealing with the problem of offering demand reduction services in the event of a critical peak, we

focus on designing mechanisms for enabling the more ambitious *demand management* services [12]. Unlike demand response, demand management refers to consumers providing a regular reduction in demand for some periods (e.g., when electricity generation costs are high). Against this background, we propose the creation of *cooperatives* of consumers, or companies representing consumers, which strive to provide demand management services via participation in the electricity market. In other words, the cooperative acts as the exact analog of a regular power plant selling electricity; however, rather than offering energy, it offers demand reduction services instead, thus extending the electricity markets to include “negawatts” [8].

In our work, the consumers’ cooperative implements a *demand side management scheme (DSMS)*. We term such a cooperative of reducing consumers a *Cooperative for Demand Side Management (CDSM)*. The CDSM effectively recruits suitable electricity consumers as members who agree to participate in the scheme by attempting to reduce their energy consumption when requested. CDSM members can range from large to medium-size consumers (such as factories, commercial buildings, and university campuses), and can be represented by automated *agents* interacting with a central CDSM-operating agent. As the CDSM members are all self-interested agents, there is a need for an effective DSMS within the CDSM ensuring desirable behaviour from the agents. A central contribution of our work lies in designing such a mechanism. In what follows, we use the terms “members” and “agents” interchangeably.

1.1 Motivations for a CDSM Mechanism

The CDSM provides demand management services, in the sense that it has a continuous presence in the electricity market, bidding to provide its reduction services, as it deems profitable, at the market’s regular trading intervals. At the same time, as will be detailed later, the services of a CDSM agents are potentially used only at some subset of those intervals (if any), on any day. Also, the agents’ services are requested a day ahead, unlike demand response services which are requested with only a few minutes notice. In this way, CDSM agents have the potential to *shift* their consumption load while being able to accommodate their business needs or maintain their comfort levels (e.g., they could shift their manufacturing, pumping or cooling activities to different time periods, if they so choose).

Thus, the CDSM is not a demand response service aiming to balance generation and consumption in the event of an emergency; rather, it is a *proactive* demand management scheme, contributing to the flattening of the energy consumption curve for the day ahead. While demand reduction companies have to wait for critical periods to make a big profit, those offering demand management services will aim to prevent those periods from occurring in the first place.

In this context it is also important to note that, from the point of view of the market and the network operator (hereby termed the *Grid*), the CDSM is the equivalent of a regular electricity provider. It

¹ Secure Meters Ltd., Winchester, UK, email: rck05r@ecs.soton.ac.uk

² Electronic and Computer Engineering, Technical University of Crete, Greece, email: gehalk@intelligence.tuc.gr

³ Electronics and Computer Science, University of Southampton, UK, emails: {vr2, acr, nrj}@ecs.soton.ac.uk

⁴ See, e.g., Enernoc: <http://www.enernoc.com/>

supplies the Grid with the equivalent energy of a requested amount of electricity, but achieves this through *reducing* electricity consumption rather than generating more electricity. Thus, to make a profit and maintain its presence in the market, the CDSM has to be a *reliable* provider—if not, it will be suffering “penalties” (imposed on it in the *balancing market*) for not meeting its agreed targets.

In more technical terms, we propose a novel *mechanism for effective demand side management* by allowing electricity consumer cooperatives to participate in the electricity market by offering demand reduction services. Our mechanism is incentive compatible, in the sense that the CDSM members do not, in expectation, gain by inflating their baseline consumption to show an artificial demand reduction. This is achieved via a randomized selection approach for choosing the agents to offer reduction services at particular trading intervals, and the employment of a payment function that encourages agents to restrict their consumption appropriately. At the same time, they have a real monetary incentive to participate in the scheme.

Our approach can also be seen as an energy conservation tool alternative to *dynamic energy pricing* [3]. Though the dynamic pricing of energy consumption has been advocated by economists as a means to avoid market inefficiencies and the “moral hazard problems” generated by the existing demand reduction schemes (i.e., using reduced flat consumption tariffs, or payments upon reaching a reduced consumption target⁵) it has itself been highly controversial, as it advocates the complete liberalization of household energy pricing. It has thus failed to attract much practical support as a demand side management mechanism. Our mechanism cannot be gamed by individuals, because of the structure of its business model—rewarding consumers on a *case-by-case basis* for their *exact* reduction over specific short time intervals, rather than over long periods. Moreover, in contrast to dynamic pricing, it is unlikely to be controversial, because consumers choosing not to participate will not be negatively affected.

The approach we propose follows on a recent line of multiagent systems work demonstrating that mechanisms with certain desirable properties (such as efficiency and incentive compatibility) can be effectively used in the Smart Grid domain [6, 11]. However, it also develops a novel business model which is implementable given the current electricity markets structure — since there are only minor regulatory changes required to allow the operation of companies offering demand management services, alongside the regular producers participating in the electricity markets. Even if rules allowing CDSMs to participate in the market directly were not to be implemented, the creation of consumer cooperatives benefiting from Grid-originated “rebate” offers or “energy credits” for reducing electricity consumption at particular chosen periods (i.e., intervals when electricity price is high), would still be an important leverage to help achieve energy conservation. Either way, the creation of CDSMs offers a powerful tool to combat the instantiation of the “Tragedy of the Commons” threat in this domain, incentivizing consumers to save energy when it is mostly needed, and consuming it when it is cheaper to do so.

In summary, this paper (i) provides an entirely novel model for demand side management through the formation of consumer cooperatives to participate in the electricity markets; (ii) designs an incentive compatible mechanism determining the behaviour of such cooperative in the market and revenue distribution among their members.

2 Background and Notation

As explained in Section 1, the CDSM’s main goal is to profitably operate in the electricity market. Thus, it attempts to maximize pay-

⁵ Note that such schemes can be easily *gamed* by individuals. For instance, a home resident away on vacation might still benefit in cash terms from its perceived “savings” in energy consumption over some period [3].

ments from selling its reduction services in the energy *spot market*, and minimizing any losses from the *balancing market*. In this section, we first briefly describe these two types of markets that are prevalent in most countries with liberalised electricity markets.

2.1 The Energy Markets

The energy spot market is a managed market for trading electricity, while quickly handling the imbalances between supply and demand schedules so that electricity distribution is not affected. It can be run or monitored by a market operator, typically the *independent system operator (ISO)*—which is usually, the national Grid, tasked with running the market and maintaining the whole system in balance [7]. The market determines the price (the “spot”) at which deals are struck through bilateral trading among participants. Several forms of bilateral trading might be in use even within one single country’s spot market, depending on the amount of time available and the energy quantities to be traded. In most countries, a day is divided into 48 half-hour time slots denoting electricity *trading intervals*, and, for each of these, prices are market-determined. In many cases, the ISO requires the provision of certain *ancillary services* contributing to power stabilization and system restoration by (perhaps a subset of) market participants (e.g., certain power generators) [7]. This can lead to the establishment of a secondary ancillary services market.

The spot market, as explained above, determines the price for electricity some time ahead of the actual time it is going to be generated and consumed. However, given the uncertainties surrounding electricity consumption patterns and (increasingly, given the rising penetration of intermittent renewable sources) generation capabilities, the amounts of energy actually delivered can vary substantially from those originally agreed in the spot market. In order to keep the system in balance, certain providers end up (perhaps through ISO intervention) generating excess of energy, while others fall short. The perceived imbalances are settled *a posteriori* in the *electricity balancing market*, through side-payments arranged by the ISO among the over- and under-producing suppliers. The balancing market energy price tends to be different than the spot market price. It depends, to some extent, on the ancillary services and guarantees provided by the participants. In many cases, the perceived “penalties” suffered by under-performers in the balancing market can be quite substantial. At the same time, the price paid to surplus producers in the balancing market may tend to be lower than what they would have received from the spot market. Thus a supplier is best off by meeting its spot market contracts via maintaining high levels of production reliability.

2.2 Basic Notation

Here we introduce the notation used in the rest of the paper. First, we denote the set of CDSM member-agents by S , while S_k denotes the subset of agents chosen to reduce their consumption in the time slot (trading interval) k . Second, $p^s(k)$ is the system-wide CDSM estimate of the energy price at the (day-ahead) spot market’s trading interval k , while $p^{b-}(k)$ is the CDSM’s estimate of the shortfall energy price at the balancing market for that same slot. This is the price of the penalty paid by any provider when it fails to satisfy its actual bid amount with the realised amount. Similarly, $p^{b+}(k)$ is the surplus energy price at the balancing market for the slot (the rate paid to a supplier if it supplies in excess to its bid amount). Typically, p^{b-} tends to be higher than p^s and p^{b+} tends to be lower than p^s .

Next, $y_i^a(k)$ denotes the *baseline consumption* (historical data-calculated average consumption) of agent i at time slot k . The *amount actually consumed* by i at time slot k is denoted by $y_i(k)$.

Agents are also required to provide their *reduction capacity* $\eta_i(k)$, applicable to time slot k of any day; this is provided by i at the time of joining the CDSM, and can be updated by it later if needed. Similarly, an agent provides the CDSM with a *minimum rate* $\pi_i(k)$ sought for its reduction services at k . This is based on the member's business needs. Note that, since the member's primary business is *not* selling reduction services, it should understand that $\pi_i(k)$ cannot be higher than the retail price p_i^c it itself pays for energy consumption, since p_i^c is, in most real-world circumstances, significantly higher than the spot price [7]. We also assume p_i^c is uniform across time periods (i.e., we assume there is no *dynamic energy pricing* used; this is realistic, since, as explained earlier, our model attempts to avoid the main shortcomings of dynamic pricing). Finally, α_i denotes the *performance accuracy* of i , representing how good it is in satisfying its reduction commitments (based on past data). Formally,

$$\alpha_i = \alpha_i^T = \sum_{t=0 \dots T} \gamma^{T-t} \alpha_i^t / \sum_{t=0 \dots T} \gamma^{T-t} \quad (1)$$

where $\alpha_i^t = (y_i^a(t) - y_i(t)) / \eta_i(t)$, and t represents the time slots that i was asked to reduce, with T being the most recent one. Also, $\gamma \leq 1$ is a discount factor, progressively "forgetting" agent i 's past performance accuracy. It is possible that $\alpha_i > 1$, this would mean that, historically, i has been observed to be reducing more than expected. Moreover, though it is expected that $\alpha_i \geq 0$, in the rare occasion that α_i becomes less than 0 (i.e., if it *consistently overconsumes* when asked to reduce), it is reset to 0.

3 The CDSM Mechanism

In this section, we describe our proposed scheme (the DSMS), outlining the algorithms and payment functions of the CDSM. In brief, the CDSM functioning is as follows:

1. Determine the time slots for participation in a given day.
2. For each selected slot k , choose the subset of agents ($S_k \in S$) for reduction. Those *not* selected ($S \setminus S_k$) will be expected to *not* increase their consumption during k .
3. Place bids in the market, carry out reduction and obtain revenue.
4. Distribute the revenue among the members, by paying the agents selected to reduce in the relevant time period according to how well they meet their reduction commitments, and penalizing the rest for any increase in their baseline consumption.

In the rest of this section, we explain each step in detail.

3.1 Bid Determination Process

To participate in the spot market on a given day, the CDSM has to place its bids some time in advance (presumably, one day ahead). To do this, the CDSM first chooses a subset of the trading intervals of the day (i.e., those most profitable in expectation) over which to participate. This is done by determining whether a slot k belongs in the top ξ slots with the highest expected $p^s(k)$ in that day. The number of participatory slots ξ can be determined by the CDSM based on the count and type of its members, and particularly, information regarding its members' reduction potential and impact of such a reduction on their underlying business.

Now, for each of these slots, the CDSM has to choose a subset of agents that would be requested to reduce consumption for that slot. This is achieved through a randomised selection policy. This selection process, along with certain accompanying constraints, helps ensure that it is *not* profitable for agents to fake or alter their consumption baseline in anticipation of better returns through the DSMS. Fol-

lowing the selection of participants, the CDSM determines its bid $\tilde{Q}(k)$ (quantity to reduce)⁶ for k .

As introduced earlier, S_k denotes the *reducers set* for time slot k —the subset of agents (out of S) chosen to reduce their consumption in that time slot. Hence, the agents not present in this subset ($S \setminus S_k$) though not expected to reduce their consumption in the time slot, are, however, expected to not increase it.

Determining S_k : With probability $\rho < 0.5$, each member agent i is sampled and included in the reducers set, conditioned on their minimum rate for that time slot $\pi_i(k)$ being *lower* than the expected spot market price $p^s(k)$ for that period. The process is given in Algorithm. 1. Here, in line 3, \mathbf{x}_i denotes a random variable sampled

```

1  $S_k \leftarrow \emptyset$ ;
2 foreach  $i \in S$  do
3   if  $\pi_i(k) \leq p^s(k)$  AND  $\mathbf{x}_i \leq \rho$  then
4     add  $i$  to  $S_k$ ;
   end
end
```

Algorithm 1: Determining S_k

from the uniform distribution between 0 and 1, to determine whether i should be chosen for participation in this time slot k . Condition $\pi_i(k) \leq p^s(k)$ ensures that no agent that requests a rate higher than the expected spot market price $p^s(k)$ is selected. In Section 3.3, we prove that, given the above selection process and the restriction on ρ to less than 0.5, a member's best strategy is to reveal its genuine consumption baseline, rather than artificially inflate it with the hope of better revenues from the DSMS.

Bid Calculation: After choosing the reducers set S_k , the bid to be placed in the market, for this time slot k , is calculated as:

$$\tilde{Q}(k) = \sum_{\forall i \in S_k} \alpha_i(k) \eta_i(k) \quad (2)$$

Although the potential reduction amount that was declared by a member is $\eta_i(k)$ for time slot k , this value is tempered with the performance factor α_i (from Eq. 1), which is based on the member's historical performance. In this way, the reduction request presented to the member is more realistic than its initially declared capacity. If the bid is accepted in the market, the reduction amount requested from a chosen agent i is, $\tilde{q}_i(k) = \alpha_i(k) \eta_i(k)$.

3.2 Revenue and Revenue Redistribution

The market operator (ISO/Grid) knows the member list of the CDSM (as given by S). Hence, based on the CDSM bid accepted, the operator looks at the consumption of all the CDSM members for the particular time slot k to determine the performance of the CDSM. The ISO is able to calculate this via the use of appropriate smart metering equipment measuring the members' consumption. It can compare any i member's realised consumption $y_i(k)$ in a trading interval k versus its average consumption $y_i^a(k)$ for k . The sum total of all the members' differences from their average consumptions gives the overall reduction achieved by the CDSM. Note that this CDSM-wide reduction depends not just on the reduced consumption of the chosen agents (S_k), but also on any increase/decrease in the consumption of the agents not in the chosen set ($S \setminus S_k$). If any such agent has increased its consumption over its baseline in the time slot, it results in

⁶ To be more precise, the CDSM's bid for a trading interval k would have to be of the form $\langle \tilde{Q}(k), p^s(k) \rangle$, with $p^s(k)$ being an *ask price* set by the CDSM itself—however, for our purposes here it suffices to simply equate this to the anticipated spot market *equilibrium price*.

a lesser overall reduction of the CDSM. Similarly, any agent within the chosen set S_k might also end up increasing its consumption instead of reducing it as per expectation. Therefore, the actual revenue obtained by the CDSM for a given time slot also depends on the imbalance amounts, which in turn depend on the increase/reduction of consumption of all agents within the CDSM. Given that, the goal of the redistribution functions presented here are twofold:

- Agents that are selected to reduce in the particular time period should be incentivised to meet their reported reduction targets.
- Agents that have not been selected to reduce in the particular time period should be incentivised not to increase their consumption.

Now, let the total CDSM revenue obtained for time slot k be $R(k)$. This value includes the spot market payment, as well as penalties paid (or received) from the balancing market. The payment received in the spot market is just based on the *promised* reduction, that is, the accepted bid of the CDSM. Hence, $R(k)$ can be written as:

$$R(k) = p^s(k) \tilde{Q}(k) \begin{cases} -p^{b-}(k) [\tilde{Q}(k) - Q(k)], & \text{if } \tilde{Q}(k) > Q(k) \\ +p^{b+}(k) [Q(k) - \tilde{Q}(k)], & \text{otherwise} \end{cases} \quad (3)$$

where $\tilde{Q}(k)$ is the accepted bid amount (refer Eq. 2), while

$$Q(k) = \sum_{\forall i \in S} \{y_i^a(k) - y_i(k)\} \quad (4)$$

denotes the actual delivered reduction of (all) CDSM members during the trading interval.⁷ Thus, $\tilde{Q}(k) - Q(k)$ represents the net difference between promised and actual reduction, for which the CDSM is potentially penalized in the balancing market. Also if $\tilde{Q}(k) < Q(k)$, then the CDSM actually receives *positive* payments in the balancing market, since it would have contributed more energy (via reduction) than what it had actually promised.

The CDSM then has to distribute this revenue in a fair way amongst its members, based on their performance during the specific trading interval. It is important to realize though, that *only* members that were actually *chosen* to reduce at a given time slot get rewarded for reducing. That is, even if some non-chosen members actually reduced, they are still excluded from (immediate) rewards for that trading interval. In this way, we make sure that members are encouraged to participate during the bid process rather than reducing consumption as and when it suits them and then expecting payment. In contrast, any member that increases its consumption beyond its baseline during the time slot k , in which the CDSM is participating in the market, should be penalised for adversely affecting the revenues of the CDSM. The chosen members themselves, are rewarded according to their contribution towards reduction. For example, if a chosen member had actually increased its consumption beyond its baseline rather than reducing as required, it will be penalised for this increase in addition to the penalty for not achieving the desired reduction.

Therefore, the revenue $R_i(k)$ of agent i at time slot k is divided into two components — a positive component $R_i^+(k)$ which is its payment for participation in reduction, and a negative component $R_i^-(k)$ denoting any penalties imposed on the agent. Thus,

$$R_i(k) = R_i^+(k) - R_i^-(k) \quad (5)$$

The positive component of the revenue is given by:

$$R_i^+(k) = \begin{cases} \bar{\pi}_i(k) \tilde{q}_i(k), & \text{if } i \in S_k \\ 0, & \text{otherwise} \end{cases} \quad (6)$$

⁷ If $Q(k) < 0$, the CDSM actually consumed more energy than normal, resulting in even higher penalties in the balancing market.

where $\bar{\pi}_i(k)$ is the payment rate awarded to the agents, such that $\bar{\pi}_i(k) \leq p^s(k)$ and $\bar{\pi}_i(k) \leq p_i^c$. The first condition ensures that the member does not get a rate that is better than the spot market rate $p^s(k)$. The latter condition is needed to ascertain incentive compatibility (see Section. 3.3). Note also that, in the real world, it is indeed normally the case that $p^s \leq p_i^c$. Thus it is perfectly valid to simply award each member the actual spot market rate p^s instead of some other $\bar{\pi}_i(k)$. Following this, the negative component is given by:

$$R_i^-(k) = \begin{cases} p^{b-}(k) [\tilde{q}_i(k) - q_i(k)], & \text{if } \tilde{q}_i(k) > q_i(k) \\ 0, & \text{otherwise} \end{cases} \quad (7)$$

where $q_i(k) = y_i^a(k) - y_i(k)$. Hence, if a member has consumed more than its baseline (irrespective of whether it was chosen or not), i.e., $y_i(k) > y_i^a(k)$, the value of $q_i(k)$ would be negative, thus leading to a higher value of $R_i^-(k)$, as indeed should be the case.

Finally, given the $R_i(k)$ amount received by each agent, the sum of the revenues paid out by the CDSM is given by $\sum_{j \in S} R_j(k)$. Note that some of the $R_i(k)$ can be negative, meaning that the amount will be paid by the agent to the CDSM for that interval k . Now as the revenue received by the CDSM from the market is $R(k)$; there will remain an amount after the payout: $R^e(k) = R(k) - \sum_{j \in S} R_j(k)$ (note $R^e(k) \geq 0$ because while the CDSM penalises the badly performing agents adequately, it does not additionally reward those who reduce beyond expectation). This excess amount can be managed in ways suiting the CDSM. For instance, it can be distributed among the selected agents, those with positive $R_i^+(k)$. Otherwise, it can be considered as the profit of the CDSM and used for its maintenance or divided amongst the agents in some other fashion—e.g., as a regular payment based on average reduction amount and performance factor.

3.3 Ensuring Incentive Compatibility

Given the DSMS, a member agent might attempt to “game” the scheme to its advantage by trying to generate an artificial consumption baseline with the prospect of making more money from the scheme. That is, it could attempt to unnecessarily over-consume electricity (not actually required for its underlying business) during some interval consistently in order to provide an “artificial” demand reduction later. In order to avoid this problem, we use a solution inspired from randomised mechanism design to show that, in expectation, self-interested agents in our system do not have an incentive to over-consume electricity, in order to exaggerate their *baseline* consumption profile. Specifically, we prove that if: (i) agents are sampled and picked with probability $\rho \leq 0.5$ for the reducers set S_k at each participating time slot, and (ii) the rate offered to agent i for reduction services for any slot is at most equal to the retail consumption price p_i^c it has to pay as part of its regular tariff; then the agent does not have an incentive to artificially inflate its baseline for that slot.

Theorem 1 *If, at any trading slot k , the CDSM samples and includes any agent i in the reducers set S_k for that slot with probability $\rho \leq 0.5$; and, if the rate offered to the agent for its services is $\pi \leq p_i^c$, then it is not profitable in expectation for i to provide an increased baseline consumption at k through “burning” electricity.*

Proof: Consider the case that an agent i intends to unnecessarily increase its baseline consumption for a particular time slot with the hope of gaining more revenue through the DSMS. Let this excess consumption be denoted by δ . That is, the member consumes δ energy more than what it actually needs in order to fake its baseline consumption. The cost of this excess consumption for an occurrence

of the time slot is $p_i^c \delta$. However, the baseline is measured by considering the consumption of the agent over a particular window of days. Let the window length be N . Therefore, in order to maintain an inflated baseline consumption, the agent will have to consume excess δ energy in that time slot for every day throughout the entire period, except for the days when it is asked to reduce.

Let us assume that K is the number of occasions that the agent is called upon to reduce its consumption in that time slot within this N -days window. Hence, the cost to the agent for faking its consumption through “burning” excess energy δ during $N - K$ days is:

$$\text{cost} = (N - K) p_i^c \delta \quad (8)$$

For the days when it is actually called for reduction (whose count is given by K), we can assume that the agent is able to obtain extra revenue for the excess δ that it manages to show in its reduction amount. Assume, without loss of generality, that the rate π the CDSM awards agent i for its reduction services persists throughout the K days under consideration, and is such that $\bar{\pi}_i(k) = \pi \leq p_i^c$. Then, the expected overall revenue gain for i by with the fake baseline, over the N -days window is:

$$\text{gain} = K \pi \delta \quad (9)$$

Now, as the probability of being called upon, $\rho \leq 0.5$, we have $K \leq 0.5N$, and, therefore, $K \leq (N - K)$. In addition, it holds that $\pi \leq p_i^c$. Therefore, clearly, $\text{cost} \geq \text{gain}$. Thus, the agent will *not* gain in expectation by “burning” electricity to fake an inflated baseline. ■

4 Experimental Evaluation

In this section, we describe the simulations conducted and the results obtained for validating our model. For this we used a large data set of 36 small and medium-scale industrial energy consumers (of different types) based in India. These 36 consumers are used to model a CDSM with a corresponding 36 member set. For each of these 36 industries, the data-set included the energy consumption for each of the 48 half-hourly trading intervals of every day over a 6 month period. Based on this data, we estimated their maximum demand elasticity by looking at the ranges of their consumptions, in addition to discussions with domain experts. Specifically, the maximum reduction capacities $\eta_i(k)$ for each agent i (for the 48 time slots) were estimated based on the variance of the demand in the actual data. This is a reasonable model because the variance in their consumption gives an indication of the possible elasticity of their demand. Such indirect modelling was necessitated by the inaccessibility of any data of consumers providing demand management services.

In more detail, if agent i is selected to reduce in time slot k , then its actual reduction is modelled by multiplying $\eta_i(k)$ with a sample from a *beta* distribution $B(\alpha, \beta)$. Beta distribution was chosen because it is somewhat similar to a normal, but has finite support and is non symmetric, giving more weight to the cases where a member reduces less than expected, than to those where reduction is more than required. If the member is not in the reducers set, then its ability to maintain its baseline consumption is computed as being sampled from a normal distribution $N(\mu = 0, \sigma_i)$, where $\sigma_i = 0$ means that there is no variance from the baseline (i.e., the member manages to consume exactly as its baseline). Following that, the price parameters were set as follows: the spot price $p^s = \text{£}0.05$, balancing prices being $p^{b+} = \text{£}0.03$ in case of a surplus, and $p^{b-} = \text{£}0.08$ in case of a shortfall. These values roughly match the long-term averages from the UK electricity market⁸. The retail consumption price for all agents was set at $p_i^c = \text{£}0.10$, similar to the prevalent retail tariffs

⁸ Indian data was unavailable because the electricity market is not liberalised

in the UK. In our simulation, on each day the CDSM only participates in the market for the periods between 8 a.m. and 8 p.m as this is usually the high demand period. For each trading interval during this period, as per the DSMS, each CDSM member has a probability $\rho = 0.49$ of being selected. Given the low member size of the CDSM, we set the probability of selection as high as possible (but less than 0.5). Through the experiments, we sought to study the 3 main aspects governing the economic viability of our mechanism:

1. The expected revenue to a member from joining a CDSM, as a percentage of its general electricity consumption costs.
2. The monetary incentive to join a CDSM, assuming that even single consumers are allowed to participate in the market in a similar way to a CDSM.
3. The efficiency of the learning mechanism for the performance factor α_i of each member.

We discuss the results (shown in Fig. 1) in the following sub-sections. Please note that the error bars are too small to be visible.

4.1 Revenue from participation in a CDSM

In the first set of experiments, we compare the average revenue of an agent from our proposed scheme to the total cost of its electricity bill (for the same period the CDSM is active in the electricity market). Note that the price paid by the member per kWh for its own consumption is the retail price p_i^c , which is twice as high as the spot market price p^s . Moreover, we denote by η_i , the average reduction capacity of agent i (averaged over all the time slots k) and similarly, by y_i , its average consumption per time slot. Using this, we can define an *elasticity index* $E_i = \eta_i / y_i$, that shows how flexible the member is in terms of its reduction services (i.e., denoting reduction capacity as a ratio of its average demand).

The results from these experiments are shown in Fig 1(a). The agents are ordered by their elasticity E_i (denoted on the x-axis), while the y-axis shows the percentage of the cost of their average monthly electricity consumption bills covered by the total revenue made from CDSM participation for a 30-day period. We observe that the revenue gained by participating in CDSM versus the cost of their consumption ranges from 5% to about 25%. In fact a quarter of the members (9 out of 36) achieve revenues of more than 15% of their bill. There is also a nearly linear relationship between the elasticity of the members and their revenues from the DSMS. This confirms the intuition that members that are most flexible with their demand stand to gain the most from CDSM participation. This is because they reduce more in a given period relative to others, and get paid more.

4.2 Revenue from being in a CDSM vs. Singleton

The second set of experiments were to observe specifically how participation in a CDSM generates more revenue to the agents than acting alone in the market (if it were allowed). To this end, we simulated the scenarios in which the 36 industrial consumers participate in the market individually, using the same mechanism as a CDSM, receiving the payments and penalties like a CDSM would. In this context, it is important to note, however, that from a Grid perspective, it is always preferable to interact with CDSMs, because of the *no-increase in consumption* commitments imposed on all the members of the CDSM even when only some of them are actually reducing consumption. If there were no CDSM, and the agents were participating directly in the market, it wouldn't be very useful to obtain a commitment of reduction from some agents, if the other agents (not bound by a CDSM-wide commitment) were free to increase their consumption over their baseline in the same time slot (thus resulting in no overall reduction in demand).

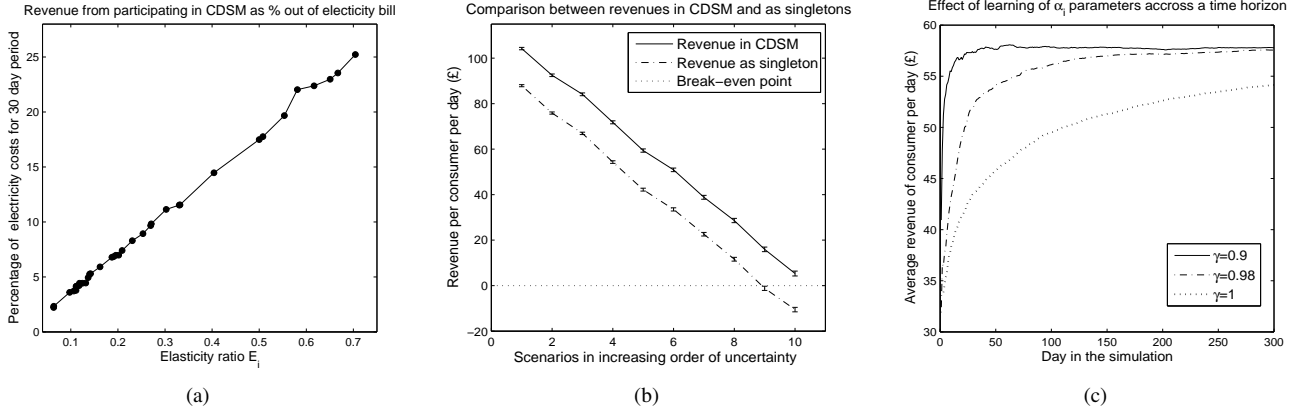


Figure 1. (a) The average revenue from participating in a CDSM, as a percentage of the total electricity bill of the member. (b) Revenue of a member in a CDSM vs. participating directly as singleton in the market, for different uncertainty scenarios; and (c) Effect of learning the performance factor α_i .

In Fig 1(b), we present the average daily revenue for all 36 agents in the 2 settings (as a member of a CDSM and acting as a singleton) for different “uncertainty scenarios”. An uncertainty scenario in this context describes the ease with which the agents stick to their reduction and non-increase targets (if the uncertainty in their business is higher, it will be more difficult for them to respect their commitments). Specifically, an increase in the uncertainty from 1 to 10 denotes a corresponding increase in two parameters: increase (via 0.1 increments) in the α factor of the $B(\alpha, \beta)$ distributions from 2 to 3 (with $\beta = 1$) for agents who are expected to reduce; and increase in the σ (standard deviations) of $N(0, \sigma)$ from 0.1 to 1.1 for agents who are expected to maintain their baseline. Expectedly, results show that the higher the uncertainty, the less profits the agents make, as they are unable to respect their commitments. However, interestingly, being in a CDSM is always more beneficial than acting as a singleton for all the scenarios. This is because the failure to respect one’s commitments can be averaged out more easily between the members of a CDSM. For the same reason, when uncertainty is very high, singleton agents no longer have an incentive to participate in the market (as their revenue becomes negative), unlike those in the CDSM.

4.3 Learning the Performance Factor α_i

The last set of experiments studied the ability of the CDSM to accurately learn the performance factor of its members (see Eq. 1). For this, we start with a setting in which there is an “incorrect” model of the members learnt over a period of 250 time slots (roughly 5 days), in which the α_i parameters of all members was set to zero. Then, we ran simulations over another period of 300 days, in which the CDSM updates α_i of the members in the correct fashion using Eq. 1. Fig. 1(c) shows the average revenue of a member over this period for three different values of the discount factor γ . The results clearly show that, if there is some *forgetting* (i.e., when $\gamma < 1$; here, $\gamma = 0.90$ and $\gamma = 0.98$), then the model quickly learns the “true” value of α_i and the daily revenue converges to the maximum level. However, for the case when there is no such forgetting, (i.e., $\gamma = 1$), the model still converges, but at a much slower rate, as the initial incorrect information persists for longer even in the face of new data.

5 Conclusions

We proposed a novel model for effective demand side management for the emerging Smart Grid. Using principles of multi-agent mechanism design, we presented a demand side management

scheme towards the formation of cooperatives of electricity consumers (CDSM). These cooperatives participate in the existing electricity markets just like the typical energy producers, but by providing *demand reduction* services. We designed the mechanism of the CDSM including methods for utilising its member agents’ services, placing bids in the electricity markets and revenue redistribution amongst the agents. We also evaluated our approach empirically. Simulation results show that participating in such a scheme can help consumers cover up to 25% of their electricity consumption costs.

Acknowledgements: This work is supported by the iDEaS project (www.ideasproject.info).

REFERENCES

- [1] M.H. Albadi and E.F. El-Saadany, ‘A summary of demand response in electricity markets’, *Electric Power Systems Research*, **78**(11), 1989-1996, (2008).
- [2] P. Asmus, ‘Microgrids, virtual power plants and our distributed energy future’, *The Electricity Journal*, **23**(10), 72–82, (2010).
- [3] James Bushnell, Benjamin F. Hobbs, and Frank A. Wolak, ‘When it comes to demand response, is ferc its own worst enemy?’, *The Electricity Journal*, **22**(8), 9–18, (October 2009).
- [4] G. Chalkiadakis, V. Robu, R. Kota, A. Rogers, and N. R. Jennings, ‘Co-operatives of distributed energy resources for efficient virtual power plants’, in *Proc. of the 10th AAMAS*, (2011).
- [5] A. Dimeas and N. Hatzigiorgiouris, ‘Agent based control of virtual power plants’, in *14th Int. Conf. on Intelligent System Applications to Power Systems*, pp. 1–6, (2007).
- [6] E. Gerding, V. Robu, S. Stein, D. Parkes, A. Rogers, and N. R. Jennings, ‘Online mechanism design for electric vehicle charging’, in *The Proc. of the 10th AAMAS*, pp. 811–818, (2011).
- [7] D. S. Kirschen and G. Strbac, *Fundamentals of Power System Economics*, Wiley, 2004.
- [8] A. Lovins, ‘The negawatt revolution- solving the co2 problem’, in *Proc. of the Green Energy Conference*, (1989).
- [9] D. Pudjianto, C. Ramsay, and G. Strbac, ‘Virtual power plant and system integration of distributed energy resources’, *IET Renewable Power Generation*, **1**(1), 10–16, (2007).
- [10] S. 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*, (2012).
- [11] H. Rose, A. Rogers, and E. Gerding, ‘A scoring rule-based mechanism for aggregate demand prediction in the smart grid’, in *The Proc. of the 11th AAMAS*, (2012).
- [12] G. Strbac, ‘Demand side management: Benefits and challenges’, *Energy Policy*, **36**(12), 4419 – 4426, (2008).
- [13] U.S. Department of Energy. Grid 2030: A national vision for electricity’s second 100 years, 2003.