

Agent Cooperatives for Effective Power Consumption Shifting

Charilaos Akasiadis and Georgios Chalkiadakis

Electronic and Computer Engineering

Technical University of Crete

Kounoupidianna Campus, Chania, Greece, GR73100

{akasiadi, gehalk}@intelligence.tuc.gr

Abstract

In this paper, we present a directly applicable scheme for electricity consumption shifting and effective demand curve flattening. The scheme can employ the services of either individual or cooperating consumer agents alike. Agents participating in the scheme, however, are motivated to form cooperatives, in order to reduce their electricity bills via lower group prices granted for sizable consumption shifting from high to low demand time intervals. The scheme takes into account individual costs, and uses a strictly proper scoring rule to reward contributors according to efficiency. Cooperative members, in particular, can attain variable reduced electricity price rates, given their different load shifting capabilities. This allows even agents with initially forbidding shifting costs to participate in the scheme, and is achieved by a weakly budget-balanced, truthful reward sharing mechanism. We provide four variants of this approach, and evaluate it experimentally.

Introduction

The expanding application of Smart Grid and “clean” energy production technologies necessitates the adoption of novel, “intelligent” techniques to better coordinate and run the future power production and distribution process [Fang et al., 2011; Ramchurn et al., 2012]. As technology evolves and electricity demand rises, the task to keep it precisely balanced with supply at all times becomes especially challenging [MIT authors, 2011]. Maintaining demand curve stability, in particular, can alleviate the risk of disastrous electricity network collapses, and leads to financial and environmental benefits—as then some generators can be run on idle, or even be shut down completely [MIT authors, 2011].

To this end, several load control programs have been proposed, where electricity consumers are encouraged to abridge their consuming activities, or shift them to off-peak hours in order to reduce *peak-to-average ratio* (PAR) [MIT authors, 2011; Davito, Tai, and Uhlaner, 2010]. Apart from industrial and other large-scale consumers, the participation of residential customers is also possible, provided that smart meters or similar electricity management systems are available [Sastry et al. 2010, van Dam et al. 2010]. Typically,

such schemes either involve an intermediary company which manages consumers who agree to contribute to the “trimming down” of the demand curve in the event of an impending critical period; or provide a reduced “flat” electricity consumption rate or rebate to consumers for lowering their consumption over a prolonged time period; or require the use of dynamic, real-time pricing (RTP) [Albadi and El-Saadany, 2008; Asmus, 2010; Bushnell et al. 2009]. Most of the “flat-rate” based schemes can be easily manipulated by individuals, however [Bushnell et al. 2009]; while RTP has been strongly criticized for promoting the complete liberalization of household energy pricing. In addition, due to increased levels of consumer uncertainty regarding imminent price fluctuations, RTP can require user manual response or the continuous monitoring of smart meters, leading to difficulties in application. Recent work also shows that RTP mechanisms do not necessarily lead to PAR reduction, because large portions of load may be shifted from a typical peak hour to a typical non-peak hour [Mohsenian-Rad and Leon-Garcia, 2010]. Other approaches aim to optimize consumption schedules via searching for Nash equilibria in specific game settings [Ibars et al. 2010, Mohsenian-Rad et al. 2010], but require players to retain fixed strategies, which is unrealistic for large, open environments.

At the same time, recent work in the multiagent systems community has put forward the notion of Virtual Power Plants (VPPs). These correspond to coalitions of electricity *producers* or *consumers*, who cooperate in order to meet market demands, mimic the reliability characteristics of traditional power plants, and deal efficiently with the issues that accrue [Asmus, 2010; Ramchurn et al., 2012]. In particular, the work of [Kota et al., 2012] has proposed the formation of consumer *Cooperatives for Demand Side Management* (CDSMs), with the aim of “selling back” to the electricity market the amount of load that was not consumed due to proactive reduction measures. In their scheme, which comes complete with certain incentive compatibility guarantees, consumers form cooperatives with the purpose to participate in the wholesale electricity markets *as if they were producers*, essentially selling energy *negawatts* in the form of reduction services. Though visionary, their approach requires a legislature change in order to be applied in real life. Moreover, no guidelines whatsoever as to where to shift consumption to are provided by that model, and deals agreed

there involve reduction promises only.

In contrast, here we propose a simple power consumption shifting scheme which can be applied directly, including residential, commercial and industrial customers. Our scheme motivates self-interested business units, represented by autonomous agents that potentially form coalitions, to shift power consumption from peak intervals to lower demand ones, in order to receive lower electricity price rates for their contribution. In more detail, the independent system operator (ISO), usually the national Grid, gives information a day ahead for the time intervals that consumption needs to be reduced at, and those that it is best to shift consumption to. Consumption during these preferred non-peak intervals is granted a better price. Then, consumers weigh own costs and potential profits, and choose to participate in a shifting operation or not. The employment of a *strictly proper scoring rule*, CRPS, incentivizes agents to report their predicted shifting capabilities as *truthfully and accurately* as possible.

Now, it is conceivable that the Grid would be willing to promise significantly lower electricity rates for considerable shifting efforts only, which cannot normally be undertaken by any consumer alone (due to small reduction capacity or high shifting costs). As a result, agents are motivated to join forces in a cooperative, to coordinate their actions so as to reach the expected reduction levels and make their participation in the scheme worthwhile. This is similar to *group buying* in e-marketplaces, where some members can obtain items that cost more than they are able to pay for alone, but due to group internal price fluctuations set by corresponding mechanisms, the purchase finally becomes advantageous to all [Li et al., 2003; Yamamoto and Sycara, 2001]. Inspired by work in that domain, we devise an *individually rational, incentive compatible, and budget-balanced* reward sharing mechanism which determines variable reduced electricity prices for coalescing agents via internal money transfers, and incentivizes them to participate in the consumption shifting scheme. We propose four variants of this mechanism, and evaluate their effectiveness.

To the best of our knowledge, this is the first specific protocol and mechanism that achieves *large-scale electricity demand shifting* under uncertainty, imposing the necessary constraints to avoid the event of new peaks arising; it does so without the use of “intermediary parties” or real-time pricing; and provides further testimony to the benefits arising from the formation of agent cooperatives in the Smart Grid.

Model and Shifting Scheme

The electricity pricing scheme used in many countries consists of two different rates, a daily and a nightly one. In our model, we also assume that there exist exactly two different price levels $p_{high} > p_{low}$, that characterize *each specific time interval* t , based on a demand threshold τ under which electricity generation costs are lower. The high demand intervals with p_{high} price are considered to be *peak* ones, at which demand needs to be reduced. We denote these as t_h in a T_H set, as opposed to low price intervals $t_l \in T_L$.

Now, given the daily consumption pattern known to the Grid, it would ideally like consumption to drop to a *safety limit* sl that is placed below τ . Consuming at the safety limit

would ensure that some low cost generated load is available in case of an emergency, minimizing the risk that high-cost generators would have to be turned on. That is, the Grid ideally wants to curtail consumption by $Q^{t_h} \geq q_\tau^{t_h}$, where (i) Q^{t_h} is the load normally consumed over the safety limit at t_h , and (ii) $q_\tau^{t_h}$ is the minimum amount of load whose potential removal can, under the Grid’s estimations, allow for a better electricity price to be offered to contributing reducers. Intuitively, $q_\tau^{t_h}$ is a sizable load quantity that makes it cost-effective for the Grid to grant a very low electricity rate, in anticipation of reaching a demand level that is close to the safety limit. We denote the load reduced by some agent i at a t_h as $q_i^{t_h}$, and that shifted by i to some t_l as $q_i^{t_l}$. The following must then hold: **(1)** $\sum_i q_i^{t_h} \geq q_\tau^{t_h}$ that is, the amount of load reduced must be higher than the minimum needed at t_h ; **(2)** $\sum_{t_l} q_i^{t_l} \leq \sum_{t_h} q_i^{t_h}, \forall i$ meaning that every reducer shifts to a subset of non-peak intervals an aggregate load amount of at most the load reduced (over the t_h intervals he participates in); **(3)** $\sum_i \sum_{t_l} q_i^{t_l} \leq Q^{t_h}, \forall t_h \in T_H$, i.e., the sum of all reducing agents shifted load to all non-peak intervals must be at most equal to Q^{t_h} , assuming that the Grid has no interest in further reducing consumption, once it has reached sl ; and **(4)** $\sum_i q_i^{t_l} \leq q_{sl}^{t_l}, \forall t_l \in T_L$, namely, the total load shifted to each t_l must not exceed the $q_{sl}^{t_l}$ quantity originally available under sl at t_l . The objective is to keep demand close to sl in as many intervals as possible.

Our scheme allows sizable load consumption from peak to non-peak intervals where an *even lower price* $p_{group} < p_{low}$ is granted, and which is a *function of the actual load reduction* q in a way that for larger load portions, the price becomes better. We term this price as p_{group} because such reduction will likely be possible only by groups of agents. This price is awarded if *the actual quantity of the load shifted* from t_h exceeds some minimum value $q_{min}^{t_h}$, set by the Grid given its knowledge of $q_\tau^{t_h}$ (e.g., it could be $q_{min}^{t_h} = q_\tau^{t_h}$).

An Efficient Consumption Shifting Scheme An agent i that wishes to participate in the consumption shifting scheme, is characterized by (a) its reduction capacity r_i^t , namely the amount of load that it is willing to curtail (e.g., by shifting) at a time interval t , and (b) its shifting cost $c_i^{t_h \rightarrow t_l}$, that is the cost that occurs if consumption of a unit of energy is shifted from (a peak interval) t_h to (a non-peak) t_l . Given the above, the exact shifting protocol is as follows.

Every day, the Grid announces the forecasted (according to its own information/uncertainty) peak intervals T_H and the most preferable non-peak intervals T_L ; and also announces the (quantity-depended) price rates it awards for consumption in T_L (we elaborate below). It then waits for shifting proposals by business units. Each business unit (consisting of a single consumer or more), can interact with the Grid and state its overall load reduction capability during announced T_H intervals, and a number of intervals $t_l \in T_L$ to which it is willing to shift consumption to. This procedure is called *bidding*. Note that for each $t \in \{T_H \cup T_L\}$ there can be more than one bidders. Furthermore, bidders can pledge to shift some load from one high consumption t_h interval to several low consumption t_l ones. Moreover, each bidder has

to report its confidence regarding its ability to shift r_i^t from interval t , in the form of a normal distribution describing its *expected relative error* regarding its reduction forecast.

Now, to promote efficiency in load shifting and avoid Grid interaction with unreliable participants, the agents need to be motivated to precisely report their actual reduction capacity. To achieve this, we employ a *strictly proper scoring rule*, the *continuous ranked probability score (CRPS)* [Gneiting and Raftery, 2007], which has also been recently used in [Robu et al., 2012] to incentivize renewable energy-dependent electricity *producers* to accurately state their estimated output when participating in a cooperative. A scoring rule $S(\hat{P}, x)$ is a real valued function that assesses the accuracy of probabilistic forecasts, where \hat{P} is the reported prediction in the form of a probability distribution over the occurrence of a future event, and x the actual occurrence itself. The rule is *strictly proper* if it incentivizes forecasters to state their true beliefs P only, and it does so by maximizing expected reward only when $\hat{P} = P$. Use of CRPS allows us to directly evaluate probabilistic forecasts, and the score is given by:

$$CRPS(\mathcal{N}(\mu, \sigma^2), x) = \sigma \left[\frac{1}{\sqrt{\pi}} - 2\phi\left(\frac{x-\mu}{\sigma}\right) - \frac{x-\mu}{\sigma} \left(2\Phi\left(\frac{x-\mu}{\sigma}\right) - 1 \right) \right] \quad (1)$$

In our setting, $\mathcal{N}(\mu, \sigma^2)$ is the uncertainty over the *expected relative error* regarding the reduction capacity, as reported by an agent (and estimated given its private knowledge of consumption requirements and business needs); while x is the actually observed error, ϕ the PDF and Φ the CDF of a standard Gaussian variable. A CRPS value of zero signifies a precise forecast, while a positive value shows the distance between prediction and occurrence. For convenience, we normalize CRPS values to $[0, 1]$, with 0 assigned when we have exact forecast, and 1 assigned when the forecast gets far from the occurrence. To improve readability, we also henceforth note $CRPS(\mathcal{N}(\mu, \sigma^2), x)$ as $CRPS$, and write $CRPS_i$ to denote the CRPS rule applied to i 's performance.

Summarizing, the Grid announces peak intervals t_h that need consumption reduction, and t_l intervals to which shifting is acceptable. The Grid determines and announces a better price rate $p_{group}(q)$ to offer for the consumption of load q at (any) t_l instead of t_h . This price is awarded if *the quantity of the load shifted* from t_h exceeds $q_{min}^{t_h}$. Consumers then make their bid collectively or alone, and state their expected reduction capacity $r_i^{t_h}$ and corresponding uncertainty, along with the intervals that they are willing to shift to. Bids are accepted after consideration of the constraints (1)-(4) stated above. Finally, an agent i whose bid was accepted, receives a reduced electricity bill B_i given its *actual* contributed reduction $\bar{r}_i^{t_h}$ at t_h , and its final consumption at t_l , $q_i^{t_l}$, as follows:

$$B_i = (1 + CRPS_i) q_i^{t_l} p_{group}(\bar{r}_i^{t_h}) \quad (2)$$

Strict propriety is maintained in Eq. 2, since the only factor depending on agent forecasts is the $1 + CRPS_i$ one.

Agent Incentives and Decisions The participation of each agent in the scheme obviously depends on its individual

costs and potential gains. Suppose that an agent i ponders the possibility of altering its baseload consumption pattern¹ by shifting some electricity consumption q from an interval t_h to t_l . This shifting is associated with a cost $c_i^{t_h \rightarrow t_l}$ for the agent. The *gain* that an agent would have for shifting q to t_l given t_l 's lower price p_{low} , would be equal to $gain(i|p_{low}) = q(p_{high} - p_{low} - c_i^{t_h \rightarrow t_l})$ since the agent would be able to consume q at t_l for a lower rate. However, under normal circumstances this gain is *negative* for the agent; if not, then the agent would have already been able to make that shift (and its baseload pattern would have been different). Now, if an *even lower* rate p_{group} is granted for consumption of $q \geq q_{min}^{t_h}$ at t_l s.t. $p_{group} + c_i^{t_h \rightarrow t_l} \leq p_{high}$, then the agent is incentivized to perform the shift.

Agent Cooperatives

In the general case, it is very rare even for large industrial consumers to have reduction capacity $\geq q_{min}^{t_h}$. Therefore, the agents need to organize into cooperatives in order to coordinate their actions and achieve the better rates promised by the Grid for effective consumption shifting. At every given time interval t_h earmarked for potential consumption reduction, only a subset \mathbb{C}^{t_h} of cooperative members might be available for shifting services. We assume that every member agent announces its availability to a cooperative manager agent, along with its reduction capacity $r_i^{t_h}$; its confidence $\mathcal{N}(\mu_i, \sigma_i^2)$ on actually reducing that amount at t_h ; and the set of t_l intervals that it pledges to move consumption to.

Even so, more often than not, it is impossible for all agents in \mathbb{C}^{t_h} to participate in the cooperative effort. This is because their shifting costs of some of them might be so high that do not allow their inclusion in any profitable cooperative bid. Therefore, only a subset C of \mathbb{C}^{t_h} will be selected for participation in the bid. Any such shifting bid is composed by four parts: t_h , the high cost interval to reduce consumption from; r_C , the amount C pledges to reduce at t_h ; a pair (T_l, Q_l) that determines the set of low cost intervals t_l to move consumption to, along the set of corresponding quantities that will be moved to each t_l ; and an estimate of its $\mathcal{N}(\mu_C, \sigma_{C,t_h}^2)$ joint relative error on predicted r_C . The bid is determined so that the collective expected gain from the shifting operation is non-negative (we provide the details of how this is ensured below). Assuming that C was selected and reduced by $\bar{r}_C^{t_h}$, the bill B_C charged to the cooperative for consuming $q_C^{t_l}$ is given by Eq. 2 (substituting C for i): $B_C = (1 + CRPS_C) q_C^{t_l} p_{group}(\bar{r}_C)$.

Now, even if the *collective* expected gain from the bid is positive, it is not certain that all individuals in C have a positive expected gain as well. Nevertheless, with positive collective expected gain, the possibility of internal *gain transfers* is raised, allowing non-negative (expected) gain for all participants. These transfers have to be performed in such a way so that the budget-balancedness of any cooperative bid is ensured, at least in the weak sense. We now describe the process by which the cooperative determines its bid at t_h .

¹Note that this is constantly monitored by the Grid given agent consumption over long time periods, and not "stated" by the agent.

Cooperative Bid Determination Since the Grid-awarded group rate depends on quantity reduced, we (originally) assume that the cooperative attempts to select a subset C with maximal reduction capacity (we modify this assumption in subsequent algorithm variants). We now present an algorithm that achieves this, while ensuring that C and each one of its members has a non-negative gain, and that budget-balancedness is ensured. In what follows, we drop time indices where these are clearly implied.

To begin, let $\hat{p}_i = (p_{\text{high}} - c_i)$ be agent i 's (implicitly stated) *reservation price*, that is, the highest price that i is willing to pay for moving from t_h to t_l (in order to not suffer a loss). The algorithm then proceeds as follows.

First, for every i , we check whether $\hat{p}_i \leq 0$. If that holds for all i , we stop; the problem is infeasible (as *all* agents need to be paid with a rate equal at least \hat{p}_i in order to participate). If that is not the case, then there exist some agents in \mathbb{C}^{t_h} for which there is a price they can accept to pay so as to move some of their consumption to t_l without suffering a loss.

The algorithm then sets $r_i := r_i - \sigma_i r_i$ for all agents in \mathbb{C}^{t_h} , that is, the cooperative makes a *pessimistic* estimate of an agent's expected performance, given its stated uncertainty. The algorithm then ranks the agents by $r_i \hat{p}_i$ in decreasing order. Then, starting from the agent with the highest $r_i \hat{p}_i$ value, we sum these values up in decreasing order, and add the respective agents in a group C . Intuitively, *the algorithm attempts to add in the coalition members with high "potential" to contribute to reduction*—that is, members with potentially high q_i to contribute, while being able to accept a relatively high (though reduced) energy price \hat{p}_i . This process continues until both of the following conditions are met for the *maximum possible* group of agents C :

(i) $\sum_{i \in C} r_i \hat{p}_i \geq r_C p_C$; and (ii) $r_C \geq q_{\min}$ where q_{\min} is the minimum quantity admitting a "group price", $r_C = \sum_{i \in C} r_i$, and $p_C = p_{\text{group}}(r_C)$ is the price rate offered by the Grid for reduction r_C .

To provide further intuition, note that the expected gain of every agent in some group C given p_C is $\text{gain}(j|p_C) = r_j(\hat{p}_j - p_C)$. If we were simply given a C for which this *gain* was positive for every member, then each agent would have been able to just pay p_C and enjoy the corresponding *gain*. However, the reducing set C and individual effective price of its members have to be dynamically determined by the cooperative, so that individual rationality is ensured.

Now, if *all* agents in \mathbb{C}^{t_h} are inserted in C and r_C is still lower than q_{\min} , the problem is infeasible and we stop. Likewise, if all agents are in C and $\sum_{i \in C} r_i \hat{p}_i - r_C p_C < 0$, the problem is again infeasible and we have to stop.

Assume that this has not happened, and *both* conditions have been met for *maximal* C .² This means that there is *at least* one agent j in C with positive expected gain, given p_C . That is, $\text{gain}(j|p_C) = r_j(p_{\text{high}} - c_j - p_C) = r_j \hat{p}_j - r_j p_C > 0$; if not, then no agent has a positive *gain*, and thus $\sum_{i \in C} r_i \hat{p}_i - r_i p_C \leq 0$, leading to $\sum_{i \in C} r_i \hat{p}_i \leq r_C p_C$, contradicting condition (i) above. This also means that agents in

²That is, after a subset C has met $r_C \geq q_{\min}$, we kept adding agents to C until by adding some k we constructed a C' for which $\sum_{i \in C'} r_i \hat{p}_i - r_{C'} p_{C'} < 0$, in which case k is removed from C' .

C are *collectively* willing to pay a total amount for moving their r_i consumptions to t_l , which is greater than what their group will be asked to pay for, given the offer p_C for r_C .

Thus we have ended up with the maximal C so that (i) and (ii) hold, and which contains some agents with positive and some with negative gain given p_C , and which we can now use to implement a gain transfer scheme so that all individual agents in C end up with non-negative gain themselves.

Setting variable effective prices At this point the cooperative pre-assigns different *effective price* rates p_i^{eff} to each contributor, producing bills that must sum up *at least* to B_C . This is done with the understanding that a member's final effective price will eventually be weighted according to its individual contribution, given also that C will receive an actual price rate that will be dependent on its CRPS score.

Thus, the cooperative initially sets $p_i^{\text{eff}} = p_C$, $\forall i \in C$, given the price p_C expected, and proceeds to rank in decreasing order agents in C according to their *expected gain*, that is $\text{gain}(i|p_i^{\text{eff}} = p_C) = r_i(\hat{p}_i - p_C)$. If all agents already have non-negative *gain*, then everyone pays p_C and expects $\text{gain}(i|p_C)$ without need of balancing. If negativities exist, then we must rearrange p_i^{eff} s.t. agents with the highest gain provide some of their surplus to those with negative, to make their participation individually rational. The first step is to count the total negative gain existing and assign negative gain agents a reduced p_i^{eff} s.t. their gain becomes exactly zero. Then, we increase p_i^{eff} of the top agent until its gain is equal to the g_j gain of the $j = i + 1$ agent below (as long as $g_j \geq 0$). Then we do the same for the second top agent, until its gain reaches that of the third. We continue this way until all requested gain is transferred, or one's gain reaches zero. If the latter happens, we move to the top again and repeat.

The p_i^{eff} prices thus determined represent internally pre-agreed prices set ahead of the actual shifting operations. The actual bill b_i that an agent $i \in C$ will be called to pay, however, is determined after the actual shifting operations have taken place, and depends on its actual performance wrt. the performance of other agents also, as follows:

$$b_i = \frac{(1 + \text{CRPS}_i) p_i^{\text{eff}} q_i}{(\sum_{j \in C \setminus \{i\}} (1 + \text{CRPS}_j) p_j^{\text{eff}} q_j) + p_i^{\text{eff}} q_i} B_C \quad (3)$$

Strict propriety is ensured by this rule, as it is an affine transformation of a member's CRPS_i score; and the sum of the b_i bills is always at least as much as the overall bill B_C charged to C , making the mechanism *weakly budget balanced*, and generating some small cooperative *surplus*.

Properties and Algorithm Variants The reward transfer scheme and the overall cooperative bid determination algorithm presented above have several desirable properties. Apart from (weak) budget balancedness, *individual rationality* is ascertained for all agents in C , as they all have non-negative expected gain from participation.

Moreover, the transfer scheme presented is *truthful*. Of course, since the agents operate in a *large, dynamic*, and

open environment, one cannot determine an incentive compatible mechanism in the Bayes-Nash sense, since analysing Bayes-Nash equilibria properties is computationally infeasible in this setting. Indeed, it is next to impossible for a member agent to reason on the unknown capabilities or availability of thousands of other agents, and *no common prior* determining such properties can be reasonably assumed. Given this uncertainty, the best that an agent can do is *to be truthful* regarding its shifting costs, capacity, and corresponding confidence: If the agent states inflated shifting costs, it runs the danger of not being selected for C . Similarly, if the agent states shifting costs lower than its real ones, then it risks suffering a high reduction in expected gain (since the lower these costs are, the higher its p_i^{eff} effective price). In addition, the sheer size and dynamic nature of the problem makes it improbable that a rational consumer would be willing to utilize, on a daily basis, the resources necessary to estimate “potentially beneficial” fake shifting costs, in order to game the scheme. Also, in practice the cooperative could use estimates of industry-dependent shifting cost limits, to fend off any such attempts. Finally, an agent has to be as accurate as possible regarding shifting capacity and corresponding uncertainty, or will suffer a gain loss due to a bad CRPS score.

Last but not least, the computational cost of the overall bid determination process (which is in any case actually performed *offline* and a day in advance) is quite reasonable. Specifically, it is proportional to the cost of sorting at most $|C^{t_h}|$ agents in every t_h twice (once when ranked according to $r_i \hat{p}_i$, and once when ranked according to perceived *gain*).

Now, the scheme presented aims to achieve the lowest possible group price, through the addition of as many agents as possible into the reducing set of agents in any given t_h , as long as budget balancedness and individual rationality are respected. Though this is clearly efficient for the Grid (since it apparently promotes the maximum possible reduction at any t_h), it is not necessarily efficient for the reducing coalition at t_h . That is, it does not necessarily maximize the sum of the members expected *gains*: since agents with potentially high costs keep being added until it is possible for the coalition to sustain them through “gain transfers”, there might exist different reducer sets with higher overall gain.

The bid determination mechanism proposed above can thus be summarized as *Method 1: Rank agents by potential and maximize expected capacity*. We now proceed to consider 3 variants of that approach. We note that all variants include a reward transfer phase, and that, unless stated otherwise, retain all the good properties of our original approach.

Method 2: Rank by potential, meet minimum capacity requirement. This method is exactly the same as the original one, with the difference that we stop adding agents in C the moment when the q_{min} requirement is met.

Method 3: Rank by potential, maximize expected capacity, exclude agents with negative expected gain. This method is exactly the same as the original one, but once q_{min} is met, an agent in the ranked list is added in the coalition only if its expected gain is non-negative wrt. p_{group} at the moment of its entry (otherwise the search for contributors continues).

Method 4: Rank by expected gain, maximize expected capacity. This method ranks prospective contributors by their

expected gain wrt. p_{group} offered at the moment they are checked for entering C . However, whenever an agent enters C , p_{group} changes, and so does the expected gain of every C member. Thus, it has to be recalculated with every new entry, which increases computational complexity.

Experimental Evaluation

In our experiments, we use simulated consumption patterns for 4968 agents, generated from distributions derived after a statistical analysis of 36 small and medium scale actual *industrial consumers* from India³ for a number of (simulated) days. Each simulation day is divided in 48 half-hour intervals. The τ threshold is fixed to 96.5% of the maximum demand across all time intervals. The safety limit is set to 99% of τ , while $q_{min}^{t_h}$ to 1% of the total load at t_h ; and the p_{high} & p_{low} values are set to the day-night prices specified by the public utility company of a given country. The p_{group} rate (in €/KWh) ranges from $p_{group}^{max} = 0.05625$ to $p_{group}^{min} = 0.0214$, depending on reduction size q :

$$p_{group}(q) = \frac{p_{group}^{min} - p_{group}^{max}}{Q^{t_h} - q_{min}^{t_h}} \cdot (q - q_{min}^{t_h}) + p_{group}^{max} \quad (4)$$

with q ranging from $q_{min}^{t_h}$ to a maximum (t_h -specific) Q^{t_h} .

Individual shifting preferences are generated as follows. A beta distribution ($a = 1, b = 43.444$) is sampled twice per agent, giving the means (a higher and a lower) of 2 Gaussians (with $\sigma=0.01$) which are then sampled for each interval, resulting to the actual agent *shifting cost*. The higher-mean Gaussian is used for shifts from intervals with baseline consumption above the agent’s daily average to ones below that average; and the lower-mean one is used for all other shifting operations. Average shifting costs classify the agents into 3 cost levels: *high / medium / low*, numbering 811.86 / 2809 / 1347.14 agents respectively on an average run. *Reduction capacities* are estimated based on the variance of each agent baseline consumption, as this is a good indicator for its demand elasticity. Agent uncertainty stated for bidding, is provided by sampling a beta distribution, with $a = 1, b = 5$ (i.e., the great mass of the agent population has low to average uncertainty) and actual agent shifting actions are provided by sampling another beta, with $a = 4, b = 2$ (modeling the realistic case that at best the agents deliver what they promised, but often fail to do so).

In our experiments, we first assume that *all 4968 agents participate in the cooperative*. Note that running one simulation day involving all agents takes on average only 1.5 sec on a 3.3 GHz PC (with a further 110 sec for demand curves and shifting costs initialization period before simulation starts).

Methods Comparison First, we simulated 100 days and compared the aforementioned methods, in order to choose one for further evaluation. Results are shown in Table 1.

Method 2 clearly ranks lower than all others both in terms of cooperative gains and trimming effectiveness. This is

³The exact same consumers as in the experiments of [Kota et al., 2012]; we thank that paper’s authors for providing the dataset.

	Method 1	Method 2	Method 3	Method 4
Expected Cooperative Gain (€/day)	2094.72	415.73	2099.40	2139.38
Actual Cooperative Gain (€/day)	895.56 ($\sigma=348.03$)	-273.89 ($\sigma=77.30$)	899.34 ($\sigma=349.36$)	933.24 ($\sigma=361.66$)
Cooperative “Surplus” (€/day)	6.20	14.81	6.20	6.13
Expected Reduction (KWh)	32856.45	12868.40	32861.60	32919.06
Final Reduction (KWh)	24454.32 ($\sigma=8083.3$)	9471.053 ($\sigma=2556.5$)	24461.93 ($\sigma=8086.1$)	24539.96 ($\sigma=8058.2$)
Peak (Demand $\geq \tau$) Trimmed (%)	98.616 ($\sigma=0.75$)	42.511 ($\sigma=5.86$)	98.619 ($\sigma=0.73$)	98.639 ($\sigma=0.75$)
Avg. Reducing Coalition Size	47.70598	15.49350	47.82045	49.06499
High Cost Participants (%)	5.609709	4.396607	5.384088	4.151038
Medium Cost Participants (%)	57.78608	56.61016	57.91386	55.56619
Low Cost Participants (%)	36.60421	38.99323	36.70205	40.28277

Table 1: Average Results (100 days simulation); σ denotes standard deviation from average values.

because it forms coalitions of a “minimum” size, capable *in expectation* to shift just q_{\min} . Thus, due to uncertainty governing actual agent behaviour, profits suffer when agent promises fail to materialize (and q_{\min} is not reached).

In contrast, *Methods 1, 3 and 4* all trimmed more than 98% of peak load, and have similar performance. When adopting *Method 3* though, it is possible for an agent that could be favoured by the final p_{group} price to be excluded from the coalition, if its contribution potential was checked early-on in the process, when the p_{group} price awarded at that point happened to grant negative expected gain to that agent. *Method 4* results to the highest cooperative gain, and highest consumption reduction, but is the most expensive computationally. Moreover, from the Grid’s point of view, it is probably not worth it to hand over an additional €37.28/day, 4.1% of the amount “paid” to *Method 1*, for a mere 0.35% increase in consumption reduction. Therefore, *Method 1* appears to be the most appropriate for our purposes, as it is comparable to the rest both in terms of cooperative gains and trimming ability; is cheaper for the Grid to use; and allows even agents with initially negative expected gains to participate in the scheme. For those reasons, we chose *Method 1* for further evaluation in this paper.

Coalition Size vs. Group Price Range Next, we examine the average reducing coalition size formed at each t_h given different p_{group} prices granted for collective consumption shifting. More specifically, we simultaneously increase the p_{group}^{\max} and p_{group}^{\min} values produced by Eq. 4 up to +0.05 of their initial values, and observe the average number of agents in reducing coalitions for each peak interval. Figure 1 shows this concept, where average coalition sizes over 100 simulation days are plotted against group price range variations. It is obvious that as p_{group} increases to get closer to p_{low} , fewer agents decide to contribute—and, subsequently, less consumption is finally shifted. Thus, in order for shifting to take place, the Grid must grant a p_{group} range that provides enough gain to the agents, given individual shifting costs.

Other Observations and Further Insights We measured that an average number of agents participating into each reducing coalition at some t_h is 47.7 individuals. The 36.6% of these are *low shifting cost* agents, whereas 57.8% are *medium* and 5.6% *high cost*. When the reward transfer

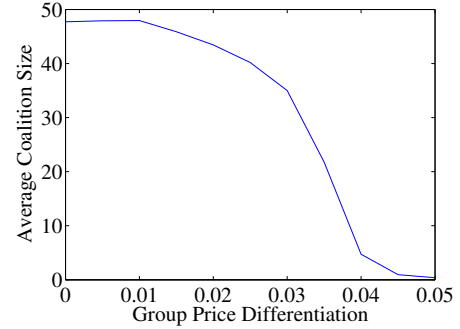


Figure 1: Avg. reducing coalition size vs. p_{group} increase.

scheme is enforced, the actual amount transferred on average during a gain balancing operation at a given t_h is negligible, in the order of €10⁻⁵, and is granted by either *low cost* or *medium cost* consumers.

It is worth noting that the Grid each day grants back to consumers an average of *only* €895.56, from its average daily income of €354064.3. Note that we cannot account for the Grid profits emerging due to reduced generation costs from the evasion of peak intervals, because such an analysis would require information that is typically not disseminated by the Grid operators. However, since we observe that in an average simulation run an 98.616% of the peak load is “safely” shifted, we can infer that the Grid stands to gain from the shifting operations. Another positive side-effect is, of course, that power outages (and resulting costs) become more distant possibilities as the demand curve flattens out.

We also verified experimentally that CRPS incentivizes agent accuracy and efficiency. Specifically, we progressively increased the relative error of reducing agents every 10 simulation days with a step of 0.1 (over a period of 110 days). This leads to an almost linearly increasing (worsening) CRPS score for the agents and their reducing coalitions, with a corresponding reduction in their profits.

In a final set of experiments, we considered settings with considerably fewer consumers participating in the cooperative. With 30% of the agents participating, it is still possible to shift 98.52% of peak load, while 10% of the population manages to shift 93.82% of peak load. With 7% of agents participating, 75% of the total peak load is shifted;

while 4% and 3% of all agents shift 51.86% and 12.66% of peak load respectively. Finally, 2.5% of all agents shift only 0.8% of the peak load. Thus, *membership clearly has to reach a “critical mass”* for the cooperative to be effective.

References

- Albadi, M. H., and El-Saadany, E. F. 2008. A summary of demand response in electricity markets. *Electric Power Systems Research* 78(11):1989 – 1996.
- Asmus, P. 2010. Microgrids, virtual power plants and our distributed energy future. *The Electricity Journal* 23(10):72 – 82.
- Bushnell, J.; Hobbs, B. F.; and Wolak, F. A. 2009. When it comes to demand response, is ferc its own worst enemy? *The Electricity Journal* 22(8):9–18.
- Davito, B.; Tai, H.; and Uhlaner, R. 2010. The smart grid and the promise of demand-side management. *McKinsey*.
- Fang, X.; Misra, S.; Xue, G.; and Yang, D. 2011. Smart grid - the new and improved power grid: A survey. *Communications Surveys Tutorials, IEEE PP*(99):1 –37.
- Gneiting, T., and Raftery, A. E. 2007. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association* 102(477):359–378.
- Ibars, C.; Navarro, M.; and Giupponi, L. 2010. Distributed demand management in smart grid with a congestion game. In *SmartGridComm 2010*, 495 –500.
- Kota, R.; Chalkiadakis, G.; Robu, V.; Rogers, A.; and Jennings, N. R. 2012. Cooperatives for demand side management. In *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI-2012)*, 969–974.
- Li, C.; Rajan, U.; Chawla, S.; and Sycara, K. 2003. Mechanisms for coalition formation and cost sharing in an electronic marketplace. In *ICEC '03*, 68–77. New York, NY, USA: ACM.
- MIT authors. 2011. Engaging electricity demand. In *MIT Interdisciplinary Study on the future of the electric grid*. MIT.
- Mohsenian-Rad, H., and Leon-Garcia, A. 2010. Optimal residential load control with price prediction in real-time electricity pricing environments. *Smart Grid, IEEE Transactions on* 1(2):120 –133.
- Mohsenian-Rad, H.; Wong, V.; Jatskevich, J.; Schober, R.; and Leon-Garcia, A. 2010. Autonomous demand-side management based on game-theoretic energy consumption scheduling for the future smart grid. *Smart Grid, IEEE Transactions on* 1(3):320 –331.
- Ramchurn, S.; Vytelingum, P.; Rogers, A.; and Jennings, N. R. 2012. Putting the “smarts” into the smart grid: A grand challenge for artificial intelligence. *Communications of the ACM*.
- Robu, V.; Kota, R.; Chalkiadakis, G.; Rogers, A.; and Jennings, N. R. 2012. Cooperative virtual power plant formation using scoring rules. In *Proceedings of the 26th Conference on Artificial Intelligence (AAAI-12)*, 370–376.
- Sastry, C. R.; Srivastava, V.; Pratt, R.; and Li, S. 2010. Use of residential smart appliances for peak-load shifting and spinning reserves - cost/benefit analysis. *U.S. Department of Energy*.
- van Dam, S. S.; Bakker, C. A.; and van Hal, J. D. M. 2010. Home energy monitors: impact over the medium-term. *Building Research & Information* 38(5):458–469.
- Yamamoto, J., and Sycara, K. 2001. A stable and efficient buyer coalition formation scheme for e-marketplaces. In *Proceedings of the 5th International Conference on Autonomous Agents (AGENTS '01)*, 576–583. New York, NY, USA: ACM.