

Multiagent Coordination for Energy Consumption Scheduling in Consumer Cooperatives

Andreas Veit

Karlsruhe Institute of Technology
Karlsruhe, Germany
andreas.veil@student.kit.edu

Ying Xu, Ronghuo Zheng, Nilanjan Chakraborty, Katia Sycara

Carnegie Mellon University
Pittsburgh, PA 15213

{yingxl, ronghuoz}@andrew.cmu.edu, {nilanjan, katia}@cs.cmu.edu

Abstract

A key challenge to create a sustainable and energy-efficient society is in making consumer demand adaptive to energy supply, especially renewable supply. In this paper, we propose a partially-centralized organization of consumers, namely, a consumer cooperative for purchasing electricity from the market. We propose a novel multiagent coordination algorithm to shape the energy consumption of the cooperative. In the cooperative, a central coordinator buys the electricity for the whole group and consumers make their own consumption decisions based on their private consumption constraints and preferences. To coordinate individual consumers under incomplete information, we propose an iterative algorithm in which a virtual price signal is sent by the coordinator to induce consumers to shift demand. We prove that our algorithm converges to the central optimal solution. Additionally we analyze the convergence rate of the algorithm via simulations on randomly generated instances. The results indicate scalability with respect to the number of agents and consumption slots.

Introduction

A key challenge in creating an energy-efficient society is to adapt electricity demand to supply conditions, e.g., by reducing peak electricity demand. The energy demand management becomes more critical when the energy supply uncertainty rises as is the case with increasing penetration of renewable energy in the electricity market (Medina, Muller, and Roytelman 2010a). A straightforward way to manage consumer demand is via direct load control (DLC), in which utility companies directly control the power consumption of consumers' appliances by switching them on/off. In small scale pilot studies, DLC has been successful in reducing peak energy consumption by better matching of supply and demand. However, the biggest drawback of DLC is that consumers may not be comfortable with utility companies having direct control over their appliances (Rahimi and Ipakchi

2010; Medina, Muller, and Roytelman 2010b). An indirect method of controlling the overall demand is to use tariffs (such as time-of-use (TOU) pricing) to incentivize consumers to shift peak time energy use. Recent technological advances in smart meters and smart appliances, have enabled direct and real time participation (RTP) of an individual consumer in the energy market through the use of software agents. A key feature of RTP programs is that each customer communicates with the utility companies individually which may lead to undesirable effects like herding (Ramchurn et al. 2012).

In (Mohsenian-Rad et al. 2010) it is argued that a good demand side management program should focus on controlling the aggregate load (which is also important for economic load dispatching (Wood and Wollenberg 1996)) of a group of consumers instead of individual consumers. Therefore, in this paper we introduce and study the problem of coordinating a group of consumers called *consumer cooperatives*. A consumer cooperative allows *partial centralization* of consumers represented by a group coordinator (mediator) agent, who purchases electricity from utilities or the market on their behalf. Such consumer configurations can potentially increase energy efficiency via aggregation of demand to reduce peak power consumption, and direct participation in the energy markets. The coordinator is not a market maker or a traditional demand response aggregator (Jellings and Chamberlin 1993) since it neither sets energy prices nor aims to incur profits by selling to the market. Rather, its role is akin to a social planner's in that it manages the demand of its associated consumer group for cost effective electricity allocation. It has to ensure that the demand goals and constraints of group members (consumers) are fulfilled, while also helping flatten out peak demands for the group. The consumers espouse the goals of the group, however they are not willing to totally disclose their demand goals and constraints to either other firms or the coordinator. Moreover, the members autonomously decide how to shift their loads to help the group flatten peak demands. Real world consumer groups coordinated in the above manner can be formed naturally in many application scenarios, especially when they are

geographically co-located, e.g., industrial parks/technology parks, commercial estates, large residential complexes.

The partially centralized coordination model offers several advantages to the stakeholders. For individual consumers, participation in such energy groups allows them to retain control of their own appliances. In addition, the consumers can obtain electricity at better prices than they would have obtained if they bought individually. The price advantage is due to three reasons. First, because of its size, the mediated participation of the consumer group in the market allows the group to enter into more flexible purchase contracts. This has the result that the price paid by the consumers reflect more accurately the actual cost of production (which is not the case in current long term fixed contract structures (Kirschen 2003)). Second, by buying as a collective, the group can benefit from volume discounts. The situation here is analogous to group insurance programs in companies. Third, in negotiated electricity contracts, the price usually consists of two components, one coming from the actual energy production cost and the other as a premium against volatility in the energy demand and/or supply. Buying as a group helps in reducing the premium against volatility since the coordinated demand management can achieve higher stability of demand and reduce demand peaks. From the utility's perspective, the consumer groups are large enough to be useful in demand response programs and have more predictable demand shifts compared to individual consumers.

The technical challenge in designing a consumer cooperative is that the central coordinator does not know the constraints of the individual consumers and thus cannot compute the optimal demand schedule on its own. Furthermore, the actual cost of electricity consumption will depend on the aggregate consumption profile of all agents and the agents may not want to share their consumption patterns or consumption constraints with other agents. Therefore, in this paper, we design an algorithm that enables the central agent to coordinate the consumers to achieve the optimal centralized consumption load where the agents have private knowledge about their consumption constraints.

We design an iterative algorithm where, in each iteration, the coordinator sends a *virtual price signal* to the consumers and the agents compute their consumption profiles based on this price signal and send it back to the coordinator. We prove that this iterative algorithm ensures that the agents converge to the optimal centralized schedule. This provably optimal demand scheduling algorithm for consumer cooperatives is the primary contribution of this paper.

Related Work

Demand response programs for managing consumer side demand have been studied from a variety of directions ranging from direct load control to indirect incentive based control (please see (Medina, Muller, and Roytelman 2010a) and references therein for an overview). Here, we will restrict our discussion to demand management using variable price signal as this is most relevant to our current work.

In the extant literature on indirect demand shaping it is (implicitly or explicitly) assumed that the utility companies

can send a price signal to software agents at the smart meters that respond to this price and schedule the appliances for the future (examples include (Philpott and Pettersen 2006; Chu and Jong 2008; Parvania and Fotuhi-Firuzabad 2010; Pedrasa, Spooner, and MacGill 2010; Conejo, Morales, and Baringo 2010; Kim and Poor 2011; Tanaka et al. 2011; Dietrich et al. 2012; Roozbehani et al. 2012; Mohsenian-Rad and Leon-Garcia 2010)). However, with such price signals, in theory, there is a possibility of instabilities like *herding phenomenon*, whereby agents move their consumptions towards the low price times and thus cause a spike in demand, thereby increasing the energy cost. Various heuristics have been proposed, but there is no algorithm with provable guarantees to solve this problem (Ramchurn et al. 2012; Voice et al. 2011; Ramchurn et al. 2011). A key feature of this work is that the agents communicate directly with the utilities and the focus is on controlling the overall load by interacting directly with each consumer. There is no interaction among the consumers.

In contrast, (Mohsenian-Rad et al. 2010) have focused on a utility or a generator controlling the load of a group of consumers. The authors accomplish this by allowing the individual consumers to interact with one another. The problem is formulated in a game-theoretic framework and the consumers coordinate in an iterative manner and exchange their demand profiles (but not their consumption constraints with each other). In this paper, we consider consumers interacting as a group with the utility, but our consumer architecture is different. We assume the consumers to form a cooperative with a central coordinator and our agents do not share their consumption profiles with other agents.

Problem Formulation

We consider a central coordinator purchasing electricity from the electricity market to support demands of a consumer group. We assume that the price is known over the whole planning horizon. This is true when the agent group has a long term electricity contract (say yearly) and the agents planning horizon is shorter (say 1 day). The contract is not a flat rate contract since in this case there would be no economic incentive for the agents to shift their demands.

We consider N members in a group with the planning period divided into M discrete time slots. The number of discrete time slots we consider is dependent on the market price structure, which can be different in practice based on the utility companies. For example, $M = 2$ for time-of-use pricing with different prices during day and night, whereas $M = 24$ in an hourly pricing scheme. Let $\mathbf{R}$ be an $N \times M$ matrix where each row of the matrix, $\mathbf{r}_i$ is the electricity demand of the agent i , $i \in \{1, 2, \dots, N\}$. We call $\mathbf{r}_i$ the *demand profile* of agent i . Each entry r_{ij} is the electricity demand of agent i for time slot j . The total aggregated demand in time slot j is $\rho_j = \sum_{i=1}^N r_{ij}$. The average market price of a unit electricity at time slot j is defined as $p_j(\rho_j)$.

We assume a typical market price function where the prices are different in each time slot and has a threshold structure. For each time slot, every unit of electricity consumed below a specified threshold is charged at a lower




Figure 1: Two-level Threshold Pricing Rates by BC Hydro.

price, while any additional unit exceeding that threshold is charged at a higher price. Thus, the marginal electricity price in a time slot, denoted by $p_j^M(\rho_j)$, is a non-decreasing function of the total demand. The marginal price at a given demand is payment increment (decrement) for adding (reducing) one unit of electricity. Figure (1) shows an example for a two level threshold price function adopted from BC Hydro.¹ The marginal price of a two-level threshold structure can

formally be written as follows: $p_j^M(\rho_j) = \begin{cases} p_j^H & \rho_j > h_j \\ p_j^L & \rho_j \leq h_j \end{cases}$

with $p_j^H > p_j^L$, where h_j is the threshold for consumption in time slot j . We further assume that the high price in any time slot is greater than the low price in any other time slot, i.e., $p_j^H > p_k^L, \forall j, k$. Let $x^+ = \max\{0, x\}$ and $x^- = \min\{0, x\}$. The energy cost for time slot j is thus:

$$p_j(\rho_j) \rho_j = p_j^H (\rho_j' - h_j)^+ + p_j^L (\rho_j' - h_j)^- + p_j^L h_j \quad (1)$$

The demand profile of each agent $\mathbf{r}_i$ must satisfy their individual constraints. We assume that the total demand of each agent during the whole planning period is fixed, i.e., $\sum_{j=1}^M r_{ij} = \tau_i$, where τ_i is the total demand for agent i . The overall demand can come from two types of loads, shiftable loads and non-shiftable loads. We will consider loads where the consumption constraints are given by a constraint set $\mathcal{X}_i$ which is private knowledge of the agent i . An agent does not share this constraint set $\mathcal{X}_i$, neither with other firms nor with the coordinator. Unless otherwise specified we will assume that $\mathcal{X}_i$ can be expressed by constraints of the form $r_{ij} \in [\underline{r}_{ij}, \bar{r}_{ij}]$ specifying the minimum and maximum consumption of agent i in time slot j , where $\underline{r}_{ij} \geq 0$. Note that these constraints are linear and $\mathcal{X}_i$ is a convex polytope. In some application scenarios, when an agent determines its energy consumption profile, it has to consider additional cost associated with the consumption schedule. For example in a factory the energy is usually for production. Thus changing the energy consumption schedule may mean changing the production process, thereby, the production cost. For agent

i , we denote this cost by $g_i(\mathbf{r}_i)$. We assume this cost function to be convex. The overall cost function of each agent is then $\sum_{j=1}^M p_j(\rho_j) r_{ij} + g_i(\mathbf{r}_i)$. With the objective to minimize the sum of all agents costs, the *central energy allocation problem* can be written as:

$$\begin{aligned} \min C(\mathbf{R}) := & \sum_{i=1}^N \sum_{j=1}^M p_j(\rho_j) r_{ij} + \sum_{i=1}^N g_i(\mathbf{r}_i) \\ \text{s.t. } & \mathbf{r}_i \in \mathcal{X}_i, \sum_{j=1}^M r_{ij} = \tau_i. \end{aligned} \quad (2)$$

where the energy allocations r_{ij} are the optimization variables. Note that the above problem is defined on a convex set $\mathcal{X}_i$. Although the objective function is non-linear it is convex because of the following. First, $\sum_{i=1}^N \sum_{j=1}^M p_j(\rho_j) r_{ij} = \sum_{j=1}^M p_j(\rho_j) \rho_j$ is convex and non-decreasing in ρ_j as indicated by Equation (1). Together with $\rho_j = \sum_{i=1}^N r_{ij}$, we can conclude that $\sum_{j=1}^M p_j(\rho_j) \rho_j$ is convex in $r_{ij}, \forall i, j$ (Boyd and Vandenberghe 2004). Since $g_i(\mathbf{r}_i)$ is also convex, the total cost function $C(\mathbf{R})$ is a summation of convex functions and so also convex. Thus Problem (2) is a convex minimization problem.

Solution Approach

Although the problem in (2) is a convex optimization problem, since the constraints and preferences of the agents are private knowledge, the consumption profiles cannot be computed directly by the central coordinator. Since the constraints in Problem (2) are agent-specific, they are naturally separable. The objective function, although a sum of the individual costs of each agent, is coupled, because the price of electricity in any time slot, j , depends on the aggregated consumption of all agents ρ_j . Therefore, we use a primal decomposition approach to solve the problem in which the sub-problems correspond to each agent optimizing its own energy cost subject to their individual constraints. The central coordinator has to compute the appropriate information to be sent to the agents so as to guide the consumption pattern towards time slots with lower prices (this corresponds to the master problem in primal decomposition methods).

An intuitive solution approach would be for the coordinator to tell the agents the aggregated demand in each time slot. Knowing the market price, the agents could then solve their individual optimization problems. This approach has problems because the agents don't know the constraints and consumption preferences of the other agents in the group, while their costs strongly depend on the consumption of the other agents. For example all agents knowing the market price and the current aggregated demand could shift as much demand as possible to a supposedly cheap time slot. This would lead to a herding phenomenon, where all agents would move to the cheap time slot resulting in a rise in the demand and thus increasing the total cost. This effect is shown in Figure (2) where part (a) illustrates an initial setting before the shift and part (b) shows a demand profile resulting from the herding phenomenon. Thus, the key challenge is to design the price signal which the coordinator sends to the agents.

We propose a novel virtual price signal that the coordinator uses to guide the agents' demand profiles. A virtual price

¹ BC Hydro is a Canadian utility company. This price is obtained from: www.bchydro.com.




Figure 2: A comparison of demand profiles in Initial, Herding and Coordinated scenarios.

signal is not the final price the agents have to pay, but information about what they would have to pay, given the current aggregated demand. The goal of designing the virtual price signal is to enable the agents to foresee the possible price increment/reduction caused by their demand shifting. Therefore, the virtual price signal, $s_{ij}(r_{ij})$, is a function of the agents demand in each time slot. To design the virtual price signal the coordinator first computes the amount of demand that should be ideally shifted in each time slot. As shown in Figure (2a), this amount, denoted by Δ_j , $j = 1, 2, 3$, is the difference between the total demand and the threshold in each time slot. To avoid herding the amount Δ_j needs to be divided amongst the agents and a threshold price signal needs to be designed for each agent, so that the price below the threshold is lower than the price above the threshold. This serves to penalize the total demand in a time slot going above the threshold. Thus, the agents know how much demand they can shift at what prices and can solve their individual optimization problem. The exact calculation of the price signal $s_{ij}(r_{ij})$ is shown in the following section. Given the price signal, the optimization problem each agent solves is

$$\min C_i(\mathbf{r}_i) := \min \sum_{j=1}^M s_{ij}(r_{ij}) r_{ij} + g_i(\mathbf{r}_i) \quad (3)$$

$$\text{s.t. } \mathbf{r}_i \in \mathcal{X}_i, \sum_{j=1}^M r_{ij} = \tau_i.$$

Note that the above problem, like the central problem, is a convex optimization problem and is thus solvable.

Coordination Algorithm

We will now present the algorithm for energy allocation and the details of virtual price signal design. Note that because of their individual constraints and cost functions some agents might not be able to shift as much demand as was assigned to them by means of the virtual price signal. This implies that the aggregated demand shift can be less than the amount that could have been achieved. Figure (2c) shows this scenario where the total demand in the second time slot remains above the threshold. To shift the remaining amount, another price signal dividing the excess amount would be necessary. This motivates us to design an iterative algorithm for the coordinator to update the virtual price signal based on the consumer's feedback and thus gradually adjust the individual demands to the central optimal solution. Each iteration consists of two steps: First, the central coordinator aggregates

the demand submitted by the agents and computes virtual price signals for each agent. Second, the individual agents use the price signal to solve their individual cost optimization problem and report their new demand profile to the coordinator.

Overview of algorithm

Recall that $\mathbf{r}_i$ denotes the demand profile of agent i and that $\mathbf{R}$ is the matrix of the demand profiles of all agents. Let $\mathbf{r}'_i$ be the updated demand profile of agent i after an iteration and $\mathbf{R}'$ be the new demand profile of all agents.

Initialization: All agents compute an initial energy consumption profile $\mathbf{r}_i$ by solving Problem (3) based on the market prices and send it to the coordinator.

1. The coordinator adds up the individual demands to determine the aggregated demand ρ_j and then calculates the amount of demand to be shifted in each time slot Δ_j . Finally the coordinator divides that demand amongst all agents and computes the virtual price signals $s_{ij}(r_{ij})$.
2. The coordinator sends the virtual price signals to all agents.
3. After receiving the virtual price signal, all agents individually calculate their new demand profiles $\mathbf{r}'_i$ according to their optimization Problem (3).
4. The agents send their new demand profiles back to the coordinator.
5. The coordinator compares the new demand profiles to the old profiles. If no agent changed its demand profile, i.e., $\mathbf{R} = \mathbf{R}'$, the coordinator stops. Otherwise, it sets $\mathbf{R} = \mathbf{R}'$ and goes to step (1).

Coordination with virtual price signal

In this section we explain how the central coordinator coordinates the firms in demand shifting via virtual price signals. The key idea is to use the marginal electricity cost in each time slot to create the virtual price signals for the agents. Shifting demand from time slots with high marginal cost to those with low marginal cost can be beneficial provided an appropriate amount is shifted. An appropriate shift implies that the resulting aggregated demand profile does not lead to higher marginal price in the former cheap time slots. As an example, in Figure (2a) a demand shift from time slot

2 to time slot 1 could be beneficial, as long as the shifted amount does not exceed Δ_1 . In order to limit the agents demand shifts they have to be able to foresee the price changes caused by their demand shifting.

The virtual price signal: The virtual price signal is a threshold price function providing marginal prices and also the demand levels at which the prices apply. The virtual price signal is structured as follows

$$s_{ij}^M(r_{ij}) = \begin{cases} p_j^H & r_{ij} > h_{ij} \\ p_j^L & r_{ij} \leq h_{ij} \end{cases} \quad (4)$$

in which h_{ij} is the threshold for agent i . Although the prices p_j^H and p_j^L are the market prices, we call this a virtual price signal, because the threshold h_{ij} can change if the agents change their demand profiles. The coordinator chooses $h_{ij} < r_{ij}$ to induce the agents to reduce or $h_{ij} > r_{ij}$ to increase demand in one time slot. With Δ_{ij} as the amount the coordinator wants agent i to change demand in time slot j , the threshold h_{ij} is updated based on the demand submitted in the last iteration:

$$h_{ij} = r_{ij} + \Delta_{ij} \quad (5)$$

Thus, the agents know that at the current market price they can at most change their demand in time slot j by Δ_{ij} . For the demand exceeding h_{ij} , they need to pay a higher price.

The demand, Δ_j , the coordinator wants to change in time slot j is calculated as the difference between the current aggregated demand and the threshold of the market price:

$$\Delta_j = h_j - \rho_j \quad (6)$$

Since the coordinator doesn't want the agents to change their demand more than Δ_j the coordinator has to ensure that $\sum_i \Delta_{ij} \leq \Delta_j$. The coordinator calculates the shift attributed to each agent proportional to that agents share of demand in that time slot: $\Delta_{ij} = \frac{\Delta_j r_{ij}}{\sum_i r_{ij}}$.

The agent's response to the virtual price signal: Having received the virtual price signal the firms will independently optimize their demand profile in order to minimize their cost according to Problem (3). Together with the virtual price signal the agents' objective function $C_i(\mathbf{r}'_i) = \sum_{j=1}^M s_{ij}(r'_{ij}) r'_{ij} + g_i(\mathbf{r}'_i)$ can be written as:

$$\sum_{j=1}^M \left[p_j^H (r'_{ij} - h_{ij})^+ + p_j^L (r'_{ij} - h_{ij})^- + p_j^L h_{ij} \right] + g_i(\mathbf{r}'_i) \quad (7)$$

Since the virtual price signal divides the amounts to be shifted amongst the agents so that the agents have to pay the high price p_j^H for demand exceeding their individual threshold, no agent will shift too much demand, based on a false impression of possible cost reduction.

Convergence of the algorithm

In this section we prove that the above introduced iterative algorithm converges to the optimal solution. First we will show that the algorithm converges and then we will show

that the converged solution is an optimal solution under certain conditions. In Lemma 1, we will show that the algorithm strictly reduces cost in every iteration. We will use this fact in Theorem 1 to show that the algorithm always converges. Then in Theorem 2 we will show that, when $g_i(\cdot) = 0$, the converged solution is an optimal solution.

Lemma 1. *The algorithm strictly reduces the total cost in every iteration: $C(\mathbf{R}') < C(\mathbf{R})$.*

Proof. In the following we give a sketch of the proof, but omit some algebraic steps because of space constraints. In the beginning of each iteration the total cost for the consumer group based on market prices as of Problem (2) is equal to the cost of the sum of the agents' Problems (3) based on the virtual price signals $\sum_{i=1}^N C_i(\mathbf{r}_i) = C(\mathbf{R})$:

$$\begin{aligned} & \sum_{i=1}^N \sum_{j=1}^M s_{ij}(r_{ij}) r_{ij} + \sum_{i=1}^N g_i(\mathbf{r}_i) \\ &= \sum_{i=1}^N \sum_{j=1}^M \left[p_j^H \left(r_{ij} - \left(r_{ij} + \frac{(h_j - \rho_j) r_{ij}}{\sum_i r_{ij}} \right) \right)^+ \right. \\ & \quad \left. + p_j^L \left(r_{ij} - \left(r_{ij} + \frac{(h_j - \rho_j) r_{ij}}{\sum_i r_{ij}} \right) \right)^- \right. \\ & \quad \left. + p_j^L \left(r_{ij} + \frac{(h_j - \rho_j) r_{ij}}{\sum_i r_{ij}} \right) \right] + \sum_{i=1}^N g_i(\mathbf{r}_i) \\ &= \sum_{j=1}^M \left[p_j^H (\rho_j - h_j)^+ + p_j^L (\rho_j - h_j)^- + p_j^L h_j \right] + \sum_{i=1}^N g_i(\mathbf{r}_i) \\ &= \sum_{i=1}^N \sum_{j=1}^M p_j(\rho_j) r_{ij} + \sum_{i=1}^N g_i(\mathbf{r}_i) \end{aligned}$$

If the algorithm has not stopped, at least one agent has changed its demand profile, i.e., $\exists i$ with $\mathbf{r}'_i \neq \mathbf{r}_i$. From Problem (3) agents only change their demand profile, if that reduces their cost. Thus, given the virtual price signal for agent i the cost of the new demand profile $\mathbf{r}'_i$ is strictly lower than of its previous demand profile $\mathbf{r}_i$: $C_i(\mathbf{r}'_i) < C_i(\mathbf{r}_i)$.

After all agents have submitted their new demand profile the new aggregated demand is computed as: $\rho'_j = \sum_{i=1}^N r'_{ij}$. Now we show that for every time slot j the total cost given the new aggregated demand and market prices is lower or equal to the sum of the agents' individual cost given their new demand profiles and the virtual price signals. With Equations (1) and (7) we get the difference of the two costs:

$$\begin{aligned} & \forall j \sum_{i=1}^N p_j(\rho'_j) r'_{ij} - \sum_{i=1}^N s_{ij}(r'_{ij}) r'_{ij} \\ &= \begin{cases} \left[\sum_{i: r'_{ij} \leq h_{ij}} (p_j^H - p_j^L) (r'_{ij} - h_{ij}) \right] \leq 0 & \rho'_j > h_j \\ \left[\sum_{i: r'_{ij} > h_{ij}} (p_j^L - p_j^H) (r'_{ij} - h_{ij}) \right] \leq 0 & \rho'_j \leq h_j \end{cases} \end{aligned}$$

Since $\sum_{i=1}^N p_j(\rho'_j) r'_{ij} \leq \sum_{i=1}^N s_{ij}(r'_{ij}) r'_{ij} \forall j$, it also holds for the sum of all time slots: $C(\mathbf{R}') \leq \sum_{i=1}^N C_i(\mathbf{r}'_i)$. Thus, the total cost is strictly reduced in each iteration:

$$C(\mathbf{R}') \leq \sum_{i=1}^N C_i(\mathbf{r}'_i) < \sum_{i=1}^N C_i(\mathbf{r}_i) = C(\mathbf{R}) \quad \square$$

Theorem 1. *The iterative algorithm always converges.*

Proof. From the definition we have that problem (2) is convex and a lower bound on the total cost can be obtained by the sum of the individual initial demand profile costs at market prices. From Lemma 1 we have that the algorithm reduces the total cost in each iteration. Thus, it can be concluded that the algorithm converges. $\square$

Theorem 2. *Assuming $g_i(\cdot) = 0$, the converged solution $\mathbf{R}$ is optimal.*

Proof. We now prove by contradiction that when the algorithm has converged to the solution $\mathbf{R}$, then no other solution $\mathbf{R}'$ exists with lower cost respect to the central problem (2). Assume there exists a solution $\mathbf{R}'$ with $C(\mathbf{R}') < C(\mathbf{R})$, then there exist 2 time slots $\{j, k\}$ st. $\rho'_j < \rho_j$ and $\rho'_k > \rho_k$ and $p_j^M(\rho_j) > p_k^M(\rho_k)$. Otherwise the cost would not be lower. We will show that for time slots $\{j, k\}$ where $\rho_j > \sum_i r_{ij}$ and $\rho_k < \sum_i r_{ik}$ it always holds that $p_j^M(\rho_j) < p_k^M(\rho_k)$, thus our assumption that $p_j^M(\rho_j) > p_k^M(\rho_k)$ is contradicted.

If $\rho_j \neq h_j$, $\rho_k \neq h_k$ we get that all time slots not at the threshold with lower marginal cost than time slot j are at their maximum constraint, because otherwise a shift from j to that time slot would be beneficial for at least one agent and the algorithm would not have stopped. Time slot k is not at the maximum constraint. It follows $p_j^M(\rho_j) < p_k^M(\rho_k)$.

If $\rho_j = h_j$, $\rho_k = h_k$ then $p_j^M(\rho_j) = p_j^L$ (as ρ_j decreases) and $p_k^M(\rho_k) = p_k^H$ (as ρ_k increases). From the the market price structure we have $p_k^H > p_j^L$. It follows $p_j^M(\rho_j) < p_k^M(\rho_k)$.

If $\rho_j \neq h_j$, $\rho_k = h_k$ and if $\rho_j < h_j$ then $p_j^M(\rho_j) = p_j^L$ and $p_k^M(\rho_k) = p_k^H$ thus $p_j^M(\rho_j) < p_k^M(\rho_k)$. If $\rho_j > h_j$ then $p_j^M(\rho_j) = p_j^H$ and $p_k^M(\rho_k) = p_k^H$. We have $p_k^H > p_j^H$, because otherwise a shift from j to k would be beneficial for at least one agent and the algorithm would not have stopped. It follows $p_j^M(\rho_j) < p_k^M(\rho_k)$. The case of $\rho_j = h_j$, $\rho_k \neq h_k$ works similarly.

It follows that $\mathbf{R}$ is the optimal solution to the central problem, as no solution with lower cost exists. Thus, the proposed iterative algorithm converges to the optimal solution. Since the problem is convex that solution is also the global optimal solution (Boyd and Vandenberghe 2004). $\square$

When $g_i(\cdot) \neq 0$ the algorithm can get stuck in a suboptimal solution, for which we have counter examples. However if in the converged solution all aggregated demands are not at the threshold, the solution is still optimal. When some thresholds are hit we need an additional phase, where the coordinator queries the agents for their individual valuation of an additional unit in those time slots. The coordinator uses this information to adjust the price signals.

Simulation Results

In the previous section we proved the convergence of the introduced algorithm to the global optimal solution. Since we do not have a bound on the number of iterations needed for convergence, we conducted simulations on randomized




Figure 3: Reduction of total cost over the course of the algorithm for different agent populations sizes.

data to get an indication of the convergence rate and scalability. In this section we present the results of our experiments. In our experiments we considered populations of up to 200 agents and varied the number of time slots up to 48. For the analysis we had to specify the market prices and the agents' individual constraints. We based the prices on the pricing model of BC Hydro as in Figure (1) and generated the values for the prices uniformly distributed $p_j^L \in [4, 8]$, $p_j^H \in [8, 16]$ and $h_j = 10 \cdot N$, where N is the number of agents. We generated the constraints for the lower and upper bounds and total consumption uniformly distributed $\underline{r}_{ij} \in [6.5, 11.5]$, $\bar{r}_{ij} \in [r_{ij}, 2r_{ij}]$ and $\tau_i \in [\sum_{j=1}^M r_{ij}, 0.4 \sum_{j=1}^M r_{ij} + 0.6 \sum_{j=1}^M \bar{r}_{ij}]$. Figure (3) illustrates the results of our experiments, showing how the total cost is reduced over the course of the algorithm. Since the agents have no individual virtual price signals in the first iteration, they optimize their cost according to the market prices. The graph shows the cost reduction for the different agent populations. Each data point is the average over all variations of numbers of time slots. The main finding is that over all experiments we see that 90% of the cost reduction is reached within 3 iterations. The average cost reduction is 5.3% of the initial cost. Since the simulation was conducted only on randomized data, this cost reduction is just a first indication of possible gains. We plan to evaluate our approach on real consumption and pricing data in future work.

Conclusion

In this paper, we presented an iterative coordination algorithm to minimize the energy cost of a consumer cooperative given private information about the demand constraints. We designed a virtual price signal that coordinates the consumers' demand shifts such that the total cost is reduced in each iteration. Then we proved the convergence of the algorithm to the global optimal solution. Finally our simulations indicate that the algorithm scales with the number of agents and consumption slots.

Future work will first check the individual rationality of an agent to join the cooperative and also the incentive com-

patibility of an agent to report his consumption in each iteration truthfully. We will also look at scenarios with the presence of storage and generation. The consumer group may have centralized and/or decentralized generation and/or storage facilities. We also aim to consider uncertainty of prices. The electricity price is uncertain, when the planning horizon is longer than the horizon for which prices are known. Moreover we may consider the setting of the coordinator being a profit making entity.

References

- Boyd, S., and Vandenberghe, L. 2004. *Convex optimization*. Cambridge university press.
- Chu, C., and Jong, T. 2008. A novel direct air-conditioning load control method. *Power Systems, IEEE Transactions on* 23(3):1356–1363.
- Conejo, A.; Morales, J.; and Baringo, L. 2010. Real-time demand response model. *Smart Grid, IEEE Transactions on* 1(3):236–242.
- Dietrich, K.; Latorre, J.; Olmos, L.; and Ramos, A. 2012. Demand response in an isolated system with high wind integration. *Power Systems, IEEE Transactions on* 27(1):20–29.
- Jellings, C. W., and Chamberlin, J. H. 1993. *Demand Side Management: Concepts and Methods*. PennWell Books, Tulsa, OK.
- Kim, T., and Poor, H. 2011. Scheduling power consumption with price uncertainty. *Smart Grid, IEEE Transactions on* 2(3):519–527.
- Kirschen, D. 2003. Demand-side view of electricity markets. *Power Systems, IEEE Transactions on* 18(2):520–527.
- Medina, J.; Muller, N.; and Roytelman, I. 2010a. Demand response and distribution grid operations: Opportunities and challenges. *Smart Grid, IEEE Transactions on* 1(2):193–198.
- Medina, J.; Muller, N.; and Roytelman, I. 2010b. Demand response and distribution grid operations: Opportunities and challenges. *Smart Grid, IEEE Transactions on* 1(2):193–198.
- Mohsenian-Rad, A., 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, A.; 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.
- Parvania, M., and Fotuhi-Firuzabad, M. 2010. Demand response scheduling by stochastic scuc. *Smart Grid, IEEE Transactions on* 1(1):89–98.
- Pedrasa, M.; Spooner, T.; and MacGill, I. 2010. Coordinated scheduling of residential distributed energy resources to optimize smart home energy services. *Smart Grid, IEEE Transactions on* 1(2):134–143.
- Philpott, A., and Pettersen, E. 2006. Optimizing demand-side bids in day-ahead electricity markets. *Power Systems, IEEE Transactions on* 21(2):488–498.
- Rahimi, F., and Ipakchi, A. 2010. Demand response as a market resource under the smart grid paradigm. *IEEE Transactions on Smart Grid* 1(1):82–88.
- Ramchurn, S.; Vytelingum, P.; Rogers, A.; and Jennings, N. 2011. Agent-based control for decentralised demand side management in the smart grid.
- Ramchurn, S.; Vytelingum, P.; Rogers, A.; and Jennings, N. 2012. Putting the 'smarts' into the smart grid: A grand challenge for artificial intelligence. *Communications of the ACM* 55(4):86–97.
- Roosbehani, M.; Ohannessian, M.; Materassi, D.; and Dahleh, M. 2012. Load-shifting under perfect and partial information: Models, robust policies, and economic value.
- Tanaka, K.; Uchida, K.; Ogimi, K.; Goya, T.; Yona, A.; Senjy, T.; Funabashi, T.; and Kim, C. 2011. Optimal operation by controllable loads based on smart grid topology considering insolation forecasted error. *Smart Grid, IEEE Transactions on* 2(3):438–444.
- Voice, T.; Vytelingum, P.; Ramchurn, S.; Rogers, A.; and Jennings, N. 2011. Decentralised control of micro-storage in the smart grid. In *AAAI Conference on Artificial Intelligence*.
- Wood, A., and Wollenberg, B. 1996. *Power generation operation and control*, John Wiley & sons. New York.