Coordinated Energy Management for Inter-Community Imbalance Minimization

Rodrigo Verschae, Hiroaki Kawashima, Takekazu Kato, and Takashi Matsuyama

Graduate School of Informatics, Kyoto University, Yoshida-Honmachi, Sakyo-ku, Kyoto 606-8501, Japan

Abstract

With the increase in demand flexibility and the rapid introduction of uncontrollable renewable power sources, effective schemes for managing the power usage of end-users are required. In this line, we propose a energy management framework for coordinating the power usage of communities of networked agents. More specifically, communities (groups of loads) coordinate to minimize their aggregated power imbalance, while taking into account each community's objectives and constraints, as well as the preferred power usage pattern of each end-user. For having a robust coordination that can work under unexpected events, we propose to assign the agents to communities using a measure of the flexibility of sets of agents. The coordination framework builds on the alternating directions method of multipliers (ADMM), algorithm that is used to implement a distributed coordination using a hierarchical architecture. While the distributed coordination allows to manage the power usage of each end-user, the hierarchical architecture enables the integration, in a single framework, of energy management problems that would be otherwise handled independently. We illustrate and analyze the coordination framework using three simulated scenarios.

Keywords: Coordinated energy management, imbalance minimization, power balancing, demand response, optimization, distributed coordination.

1. Introduction

While the balancing of generation and consumption is crucial for the operation of electric power systems, shaping the demand can help reducing operational, capacity and investment costs [1]. The introduction of information technologies in the management of power systems [2] is making the demand more flexible (thanks to the ability of real-time sensing, controlling, and scheduling the power usage, and the introduction of mobile loads and local storage), which together with the rapid introduction of uncontrollable renewable generation sources, is requiring new balancing schemes to effectively manage the increasing flexibility of the demand.

*Corresponding Author: Rodrigo Verschae. Email: rodrigo@vision.kuee.kyoto-u.ac.jp.
The final published version of the paper is available at: http://dx.doi.org/10.1016/j.renene.2015.07.039

The traditional solution to balance the supply and demand has been to control the generation (by dispatching of generation units), with an (aggregated) demand that is not controlled. Also, and seeking to reduce costs (e.g. associated to high consumption peaks), the control of the power consumption through *demand response* (DR) programs has been introduced [3] [4].

DR programs can be roughly classified in two categories: *i*) *price-based*, such as time-of-use (TOU), critical-peak-pricing (CPP), and real-time-pricing (RTP), and *ii*) *event (or incentive)-based*, such as direct load control (DLC), curtailment, and demand reduction bid programs. While event-based programs are designed to be seldom used and therefore is not clear how they can deal with uncontrollable power sources, price-based programs (e.g. RTP) may cause peak rebounds because all agents receive the same price signal and may decide to use power at a similar time (see [5] for a deeper discussion on these issues).

Seeking to address these issues, a new category of methods for controlling the power usage of a group of users has been recently proposed (see e.g. [6] [7] [8] [9] [10] [11]). We refer to this new category as *coordinated energy management*[1]. Take for example [11], where a coordinated home energy management system (CoHEMS) was introduced to balance the supply and demand of homes by a coordinated scheduling of deferrable appliances. The key difference between *demand response* and *coordinated energy management* is that in the latter the agents coordinate their power usage seeking the benefit for the whole community, while in existing methods each agent acts independently, usually after receiving a top-down signal sent by an operator. Additional benefits are that: *i*) it can more effectively manage the aggregated power usage and avoid generating peak rebounds (thanks to the coordination), and *ii*) it can consider each user's quality of life and privacy (thanks to the inherent distributed architecture).

In power systems and energy markets, there are constraints and objectives associated to different entities. For example, the distribution network has a radial structure with constraints associated to subparts of its tree structure (e.g. constraints associated to transmission substations). Similarly, in energy markets (e.g. wholesale, balancing and retail), various mechanisms are used to minimize energy cost and to ensure the operational stability at the different levels of the power grid, and these markets have objectives associated to subparts of the network. Moreover, these markets usually work independently, but their integration in a single framework could be beneficial (see e.g. [12] [13] [14]). Therefore, a coordination architecture able to consider objectives and constraints in an integrated way is helpful in many scenarios.

In this context, we propose a framework for coordinating the power usage of networked agents (appliances, factories, etc.), with the main goal being to minimize the imbalance among communities, while including objectives and constraints for each community and taking into account each user's quality of life / activities. Although the framework does not need to respond to external requests, it can provide demand response as a service, and it can be used to manage both, power usage and power generation. Also of importance is the robustness of the coordination architecture, therefore we propose an optimization

[1] This kind of approach has been refereed to as coordinated DR [5]. However, we believe that *coordinated energy management* is a more suitable name, as it applies to more general cases: the coordination does not need to respond to external requests, and each user can be a prosumer (a producer and a consumer).

framework for building communities such that each community has flexibility to handle unexpected events.

The present article builds on our previous work [15], giving a complete description of it, and including new experimental results. The main contribution is two fold: *i*) first, a framework for coordinated energy management for inter-community imbalance minimization is introduced, and *ii*) second, a framework for building these communities, taking into account each agent's flexibility, is proposed. In order to illustrate the coordination framework, we use coordinated electric vehicle (EV) charging, but the coordination framework can be used in other applications. Three scenarios are considered: *i*) *Imbalance minimization with community constraints*, *ii*) *Imbalance minimization for demand response as a service*, and *iii*) *Imbalance minimization robustness under unexpected events*, scenarios where the framework's robustness against different conditions is also analyzed.

The article is structured as follows. We first formulate the imbalance minimization problem (Section 2.1), we continue by presenting the coordination framework (Section 2.2) and by describing the framework for assigning agents to communities (Section 2.3). Finally we present the example scenarios (Section 3) and conclude (Section 4).

2. Coordination Framework

Our overall goal is to minimize the power usage imbalance among communities, while considering each community's goals (e.g. consumption targets) and constraints (e.g. local consumption and generation), and also taking into account each user's preferred power usage pattern (which is related to the user's quality of life / activity).

Communities. There are many ways to define a community. The agents that are part of a community can be predefined depending on the users (agents) goals, values (e.g. to consume renewable power sources) or by considering some physical closeness (e.g. the users' in neighborhood). Some users may have objectives or constraints for its agents (either physical or contractual), and this will require that those agents are in the same community during the coordination. For example, one could consider the radial power distribution network, where some nodes (e.g. a sub-transmission substation, or load bus) of this tree structure could define the communities.

A relevant case corresponds to each community being managed by a *power retailer*. Then, the inter-community imbalance minimization corresponds to an imbalance minimization among power retailers, and this can be seen as an imbalance market, where the power retailers negotiate the power usage and prices.

2.1. Problem formulation

We consider N agents, $i \in \mathcal{N} = \{1, \ldots, N\}$, each one having a power usage described by a vector $x_i \in \mathbb{R}^T$, with T the number of time-slots for a given time period (e.g. $T = 144$ can represent a 24 hour period with 10 min time-slots). The components of the power profile x_i are positive at times power is consumed, and negative when power is generated. Agent i can correspond to a single power using device (e.g. an appliance), or to a larger entity (e.g a household, a factory, etc). Although we focus on power using agents, considering generation and storage is straight-forward.

We assume each agent can sense (measure), plan and schedule/manage its power usage. Sensing allows the agent to build a model of the power usage uncertainty and flexibility, while planning allows to define its future consumption, and scheduling implements the control. Note that in the current paper we do not address the problem of real-time control, which must consider the stochastic nature of appliance usage and of renewable generation (see [16] for an example of an agent that could implement this real time control).

The agents are organized in J non-overlapping communities[2], with the agents in community j represented by the index set $\mathcal{N}_j$. We note $N_j = |\mathcal{N}_j|$ the cardinality of community j. $\{\mathcal{N}_j\}_{j\in\mathcal{J}}$ corresponds to a partition of $\mathcal{N}$ ($\mathcal{N} = \cup_{j=1}^{J}\mathcal{N}_j$ and $\mathcal{N}_j \cap \mathcal{N}_k = \emptyset$ for $j \neq k$), thus $N = \sum_{j=1}^{J} N_j$ is the total number of agents.

We formulate the inter-community imbalance problem as the following optimization problem:

$$\underset{(x_i)_{i\in\mathcal{N}}}{\text{minimize}} \sum_{i\in\mathcal{N}} f_i^t(x_i) + \sum_{j\in\mathcal{J}} h_j^t(\sum_{i\in\mathcal{N}_j} x_i) + g^t(\sum_{i\in\mathcal{N}} x_i), \tag{P1}$$

where there are objectives functions for three levels: agent, community, and global (with the function g^t measuring the inter-community imbalance); and where the t superscript is used to indicate that these functions may change over time. This optimization is run at each time t, but in the following we omit the t superscript for clarity.

We call $f_i : \mathbb{R}^T \to \mathbb{R}$ the function that measures the local cost (disutility) of agent i to achieve the profile x_i, while $h_j : \mathbb{R}^T \to \mathbb{R}$ is a function that measures the cost/constraints associated to the aggregated profile $\sum_{i\in\mathcal{N}_j} x_i$ of the community j, while the function $g : \mathbb{R}^T \to \mathbb{R}$ measures the imbalance cost associated to the global aggregated profile $\sum_{i\in\mathcal{N}} x_i$. Thus, the goal is to minimize this imbalance cost taking into account community and agent objectives.

Global objective. The function $g(v)$ measures the imbalance cost associated to the aggregated power usage $v = \sum_{j\in\mathcal{J}} w_j$, with $w_j = \sum_{i\in\mathcal{N}_j} x_i$, the aggregate power usage of community j. Many options are possible, but we use the squared Euclidean norm of the deviation from a reference profile: $g_r(v) \propto ||v - r||^2$, where $r \in \mathbb{R}^T$ can represent an aggregated planned usage profile (e.g. determined in a day-ahead market) or an expected generation power profile (e.g. a PV forecast). When $r = 0$, the goal is to flatten the power usage, and in general $r = \sum_{j\in\mathcal{J}} r_j \in \mathbb{R}^T$, with $r_j \in \mathbb{R}^T$ the target power usage of the community j.

Community objective. The objective $h_j(\sum_{i\in\mathcal{N}_j} x_i)$ for each community $j \in \mathcal{J}$ can encode a shared cost, a set of constraints (e.g. $h_j(x) = 0$ for $x \in \mathcal{X}_j$, and $+\infty$ otherwise), or could be zero for all x. The constraints can encode physical restrictions (power usage limits), energy cost-related bounds (e.g. budgets), consumption targets, among others. We consider two cases:

- Usage bounds: $h_j(x) = -\sum_{\tau=0}^{T} \log 1_{[x_{(\tau)} \leq v^j_{(\tau)}]}$, with $1_{[\circ]}$ the indicator function, $x_{(\tau)}$ the τ^{th} component of $x \in \mathbb{R}^T$, and $v^j_{(\tau)} \in \mathbb{R}$ a power usage limit at time τ.
- Minimize deviation from plan: $h_j(x) \propto ||x - r_j||^2$, with r_j a power usage target.

[2]**Notation:** we use the index j to refer to communities and the index i to agents.

Agent objective. For modeling and controlling each agent's power usage, we use the family of objective functions proposed in [9]. While the coordination is not restricted to this particular family, we use it as it can take into account user dissatisfaction and uncertainty measures in the appliances' control, and it comprises a probabilistic model that can be learned from real daily sensing data through statistical or machine learning techniques. The cost of agent i is defined as:

$$f_i(x_i) = \min_{u_i \in \mathcal{U}_i} \left[-\log P(x_i, u_i) \right], \tag{1}$$

which implements a probabilistic generative model $P(x_i, u_i) = P(u_i)P(x_i|u_i)$, with the probabilities $P(u_i)$ and $P(x_i|u_i)$ being used to measure how natural the control $u_i \in \mathcal{U}_i$ is, and the certainty of achieving the profile x_i given the control signal u_i, respectively. We further define the corresponding disutility functions, $f_i^{x|u}(x_i, u_i) = -\log P(x_i|u_i)$, and $f_i^u(u_i) = -\log P(u_i)$, thus $f_i(x_i) = \min_{u_i} \left[f_i^u(u_i) + f_i^{x|u}(x_i, u_i) \right]$. Thus, in case $P(x_i, u_i)$ becomes zero, $f_i(x_i)$ becomes $+\infty$, indicating that the associated profile cannot be achieved. A simple example consists of $P(x_i|u_i)$ equal to the delta function $\delta(x_i - \chi_i(u_i))$, and $\chi_i : \mathcal{U}_i \to \mathbb{R}^T$ being a mapping from control signals u_i to profiles x_i. More complex mappings can consider possible uncertainties in the mapping from control signals to power signals (see [9] for more details).

The control u_i is closely related to the user's QoL, while the power profile x_i is not necessarily so. Thus, the function $P(u_i)$ can be related to this QoL while also considering soft/hard constraints for the control u_i. A key aspect in power balancing is to schedule the power usage. Thus, for modeling the involved probabilities a hidden semi-Markov model (HSMM) [17] [18] is used, method that explicitly models time-varying signals as a sequence of time intervals (or modes) and is popular in the speech recognition community. Basically, a sequence of discrete states $s_{i,t} \in \mathcal{Q}_i = \{q_{i,m}\}_{m=1,\ldots,M_i}$, is used to describe the profile x_i, where each state represents the mode of appliance i at time t, and the control variables u_i are these modes at each time $t = 1, \ldots, T$.

During the coordination, the agents solve a optimization problem using functions as defined in Eq. (1) when using a HSMM model. The optimization problem is presented in the next paragraphs (specifically Eq. (2) which can be solved efficiently using dynamic programming [9]).

2.2. Coordination

To implement the coordination among agents, we introduce a *coordinator* for each community and a *global coordinator*. The basic coordination scheme is presented in Fig. 1, where the global coordinator exchanges messages with the community coordinators, while each community coordinator exchanges messages with the agents in its community. Put in other way, the agents coordinate with the help and via the coordinators.

Proximal operator. For an agent to take part of the coordination, the agent (with cost function f) has to be able to solve the following optimization problem:

$$\text{prox}_{f,\rho}(v) = \arg\min_x f(x) + \frac{\rho}{2}||x - v||^2, \tag{2}$$

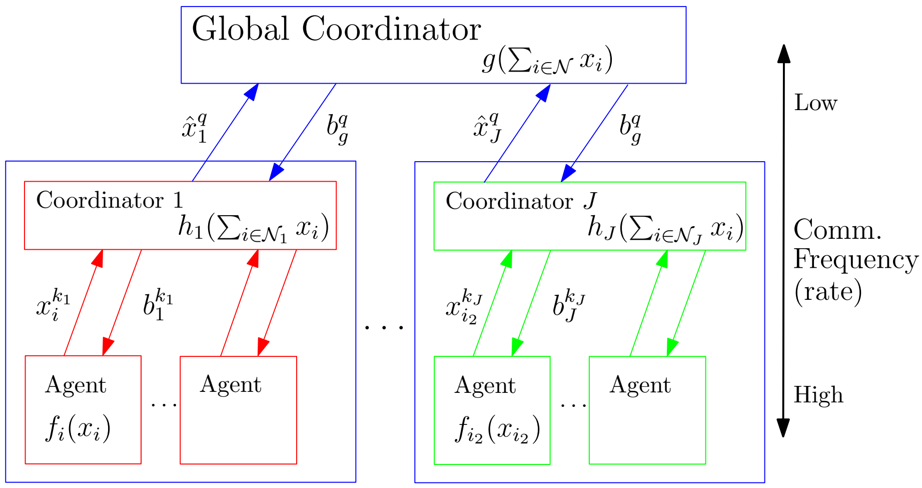

Figure 1: Distributed hierarchical architecture. The agents, with the help of community and global coordinators, cooperate to minimize $\sum_{i\in\mathcal{N}} f_i(x_i) + \sum_j h_j(\sum_{i\in\mathcal{N}_j} x_i) + g(\sum_{i\in\mathcal{N}} x_i)$, namely they seek to minimize the inter-community imbalance while taking into account local, and community objectives (see text in Section 2 for details).

which means that it implements a proximal operator[3] [19] [20], with the value of v depending on a message sent by a coordinator. A (global/community) coordinator implements a proximal operator, plus some extra steps. The details of this coordination, and the relation between the parameters, the exchanged messages, and each entity's objective function, is specified in the following.

The Alternating Direction Method of Multipliers (ADMM). ADMM [21] is used to implement the coordination. In its basic form, ADMM solves problems of the form:

$$\underset{x}{\text{minimize}}\ f(x) + g(x). \tag{P2}$$

By introducing a duplicate variable $z = x \in \mathbb{R}^T$ and a Lagrange multiplier u (associated to the constraint $z = x$), ADMM consist on applying the iterative algorithm:

$$x^{k+1} := \text{prox}_{f,\rho}(z^k - u^k) \tag{3a}$$
$$z^{k+1} := \text{prox}_{g,\rho}(x^{k+1} - u^k) \tag{3b}$$
$$u^{k+1} := u^k + x^{k+1} - z^{k+1}, \tag{3c}$$

with ρ a penalty parameter, and k the iteration index.

The convergence of ADMM has been widely studied [21]. It is known to converge when the objective functions (f and g) are closed, proper and convex, and the augmented Lagrangian has a saddle point. Under these assumptions, the ADMM iterates satisfy

[3] For clarity we have written $\rho \in \mathbb{R}$ as a parameter of the operator. This is not standard notation; usually $\text{prox}_{f/\rho}(v)$ is used instead.

residual, objective and dual variable convergence. ADMM is known to converge for biconvex problems and under non-convex constraints [21], and its convergence has been studied for non-convex [22] and time varying functions [23].

We note that the x-step in Eq. (3a) and the z-step in Eq. (3b) are solved separately, with the x-step requiring only access to f and the y-step having only access to g, thus allowing a distributed implementation. We will use the ADMM algorithm to implement a distributed coordination at two levels, inter- and intra-community coordination.

2.2.1. Inter-community coordination

First, we re-write the problem (P1) by grouping the agents according to the given communities $\mathcal{N}_j$:

$$\min_{(x_i)_{i\in\mathcal{N}}} \sum_{j\in\mathcal{J}} \left[\sum_{i\in\mathcal{N}_j} f_i(x_i) + h_j(\sum_{i\in\mathcal{N}_j} x_i) \right] + g(\sum_{j\in\mathcal{N}} \sum_{i\in\mathcal{N}_j} x_i). \tag{P3}$$

This problem corresponds to a *sharing problem* [21] among communities, with $g(\sum_{j\in\mathcal{J}} \sum_{i\in\mathcal{N}_j} x_i)$ the cost shared among J communities. To solve it we use ADMM, and similarly to what we did with problem (P2) we introduce duplicate variables. In this case the duplicate variables of aggregated profile of each community $w_j = \sum_{i\in\mathcal{N}_j} x_i$ are introduced, leading to the optimization problem:

$$\begin{aligned} \min_{(x_i)_{i\in\mathcal{N}},(w_j)_{j\in\mathcal{J}}} & \sum_{j\in\mathcal{J}} \left[\sum_{i\in\mathcal{N}_j} f_i(x_i) + h_j(\sum_{i\in\mathcal{N}_j} x_i) \right] + g(\sum_{j\in\mathcal{J}} w_j) \\ \text{s.t. } & w_j - \sum_{i\in\mathcal{N}_j} x_i = 0 \ \forall j \in \mathcal{J}. \end{aligned} \tag{P4}$$

By applying ADMM to problem (P4) we obtain the following iterative procedure[4], with iteration index q:

$$(\tilde{x}_i^{q+1})_{i\in\mathcal{N}_j} := \arg\min_{(x_i)_{i\in\mathcal{N}_j}} \sum_{i\in\mathcal{N}_j} f_i(x_i) + g_j^q(\sum_{i\in\mathcal{N}_j} x_i), \forall j \tag{4a}$$

$$\bar{w}^{q+1} := (1/J)\ \mathrm{prox}_{g,\rho_g/J}\left(J(\bar{x}_g^{q+1} + \nu^q)\right) \tag{4b}$$

$$\nu^{q+1} := \nu^q + \bar{x}_g^{q+1} - \bar{w}^{q+1}, \tag{4c}$$

with $\bar{w}_g^q = \frac{1}{J}\sum_{j\in\mathcal{J}} w_j^q$, $\bar{x}_g^q = \frac{1}{J}\sum_{j\in\mathcal{J}} \hat{x}_j^q$, and $\nu \in \mathbb{R}^T$ the Lagrange multiplier associated to each constraint $w_j = \sum_{i\in\mathcal{N}_j} x_i$ (note that the Lagrange multiplier is equal for all j ($\nu = \nu_j\ \forall j \in \mathcal{J}$) [21]). Here, $\hat{x}_j^q = \sum_{i\in\mathcal{N}_j} \tilde{x}_i^q$ is the group proposed profile, and $\tilde{x}_i^q$ is the profile of agent i at iteration q. The parameter ρ_g is a global penalty term.

Note that in Eq. (4a) we have introduced the function:

$$g_j^q(v) = h_j(v) + \frac{\rho_g}{2}||v - \hat{x}_j^q + b_g^q||^2, \tag{5}$$

[4]The derivations are omitted for space reasons; see [9] [15] [21] for similar derivations. Note that we have used the property $\mathrm{prox}_{f,\rho}(Jx) = (1/J)\ \mathrm{prox}_{f,\rho/J^2}(x)$.

which depends on $\hat{x}_g^q$, the previous proposed profile by the community (known by the coordinator j), and on $b_g^q = \bar{x}_g^q - \bar{w}_g^q + \nu_g^q$, a signal determined by the global coordinator.

2.2.2. Intra-community coordination

At each iteration q of the algorithm in Eq. (4), each community solves its corresponding problem in Eq. (4a):

$$\underset{(x_i)_{i\in\mathcal{N}_j}}{\text{minimize}} \sum_{i\in\mathcal{N}_j} f_i(x_i) + g_j^q\left(\sum_{i\in\mathcal{N}_j} x_i\right) \tag{P5}$$

with $g_j^q(v)$ managed by the community *coordinator*. The function g_j^q depends on b_g^q (Eq. (5)), a signal defined by the global coordinator at the iteration q (at the global level).

It is not difficult to see that problem (P5) is a simpler case of problem (P3), also corresponding to a sharing problem. Then, similarly to problem (P3) we introduce duplicate variables, in this case for each agent power usage $z_i = x_i \; \forall i \in \mathcal{N}_j$, and solve it using ADMM, obtaining the iterative algorithm (with iteration index k):

$$x_i^{k+1} := \text{prox}_{f_i,\rho_j}(x_i^k - b_j^k) \; \forall i \in \mathcal{N}_j, \tag{6a}$$

$$\bar{z}_j^{k+1} :=(1/N_j) \; \text{prox}_{g_j^q,\rho_j/N_j}\left(N_j(\bar{x}_j^{k+1} + \nu_j^k)\right), \tag{6b}$$

$$\nu_j^{k+1} := \nu_j^k + \bar{x}_j^{k+1} - \bar{z}_j^{k+1}, \tag{6c}$$

with $b_j^k = \bar{x}_j^k - \bar{z}_j^k + \nu_j^k$ a broadcast signal determined by the community coordinator. Here $\nu_j^k \in \mathbb{R}^T$ corresponds to the scaled Lagrange multipliers [21], ρ_j is a penalty parameter, and we note $\bar{a}$ to represent the average of a set of variables $\{a_i\}_{i\in\mathcal{N}}$, (i.e. $\bar{a} = \frac{1}{N}\sum_{i\in\mathcal{N}} a_i$).

Case of no community objective functions. When community j has no objective function (i.e. $h_j(x) = 0, \; \forall x$), Eq. (6b) can be solved analytically (with $\alpha_{\rho,j} = \frac{\rho_j}{\rho_g N_j}$ and $\hat{r}_j^q = \frac{1}{N_j}(b_g^q - \hat{x}_j^q)$):

$$\bar{z}_j^{k+1} = \frac{1}{1+\alpha_{\rho_j}}\left[\alpha_{\rho_j}\left(\bar{x}_j^{k+1} + \nu_j^k\right) - \hat{r}_j^q\right], \tag{7}$$

meaning that $\bar{z}_j^k$ and ν_j^k are linear in all the required variables (the same holds for the residuals used in the algorithm's stopping criterion; see below). Thus, in this case the state of the coordination can be calculated by any entity having access to $\hat{r}_j^q$ and to the sequence $(\hat{x}_j^k)_{k=0,\ldots,K_j}$ (with K_j the number of iterations).

2.2.3. Implementation

Distributed coordination. The coordination is implemented using the architecture in Fig. 1, which corresponds to a two-level distributed coordination: an inter-community coordination (in blue) and an intra-community coordination (in red and green). Take the case of the coordination of community j (the inter-community coordination is analogous): the first step (x-step in Eq. (6a)) is calculated concurrently by each agent of the community j after receiving the broadcast signal b_j^k from the coordinator, while the second and third steps (z- and ν-steps in Eq. (6b) and Eq. (6c)) are calculated by the community coordinator, which first aggregates the profiles $(x_i^{k+1})_{i\in\mathcal{N}_j}$, calculates $\bar{x}_j^{k+1}$, $\bar{z}_j^{k+1}$ and ν^{k+1}, and then broadcasts b_j^{k+1} to all agents.

There are two advantages of the two-level distributed formulation that it is worth highlighting: *i*) the problems solved by two communities are de-coupled (agents at different communities do not work synchronously), and *ii*) the communication rate at the upper (inter-community) level can be lower than the one within each community.

Starting conditions. There are no requirements for the initial values of the optimization variables. We use the following values for the community level (the inter-community level is analogous): the coordination starts with a broadcast signal $b_j^0 = 0$ (zero vector of dimension T), and $b_g^0 = 0$, the initial value of the profiles are also always set to zero $x_i^0 = 0$, thus agent i sends the solution of $\text{prox}_{f_i,\rho_j}(0)$ at the first iteration. Given that the intra-community coordination is nested in the inter-community coordination, the intra-community coordination always starts with the observed values (for x_i^k and $\bar{z}^k$) in the previous coordination.

Stopping criterion. For problem (P1) solved in Eq. (3) (ADMM), the stopping criterion proposed in [21] can be used. It requires the evaluation of $s_p^k = x^k - z^k$ and $s_d^k = \rho(z^{k+1} - z^k)$, the primal and dual residuals respectively. The termination criterion are: $||s_p^k|| \leq \epsilon^{pri}$ and $||s_d^k|| \leq \epsilon^{dual}$, with $\epsilon^{pri} > 0$ and $\epsilon^{dual} > 0$ feasibility tolerances.

Communication link failures. When an agent or community is not able to communicate with its upper level, the following behavior is used. Assume community j and the global coordinator cannot communicate at iteration q. In such case the global coordinator uses the last received message from community j, lets say $\tilde{x}_j^{(p)}$ (with $p \leq q$), as the profile that community j plans to follow, while community j modifies its objective function such that temporarily its goal is to have a profile as close as possible to the last communicated one. Once the communication is restored, the community goes back to use its original objective function, allowing the coordination to stay close to its expected behavior, and to resume the coordination once the communication is restored. This fall back protocol is defined in terms of the last communicated profiles (instead of broadcast signals), so that the global coordinator can continue coordinating with the remaining communities.

Data package losses and early stopping as inexact solutions of the proximal operator. Distributed algorithms can suffer from problems occurring during the data transmission (e.g. package loses, delays, errors, etc.) that could affect the convergence of the algorithm. Also, in some cases it may be required to limit the number of iterations, e.g. for reducing the communication requirements within a community. One of the advantages of using ADMM is that it is known to converge even if the solutions of the proximal operator are inexact (under some suboptimality measures [21] [24]). This means that ADMM can converge even if the intra-community coordination is solved approximately, e.g. when:

- Some data packages are lost during the communication (a lost profile can be approximated using previously received packages), which increases the robustness of the algorithm and can simplify the communication protocol.

- The number of iterations for a community is limited (the x-step is solved approximately for a community), which allows reducing the communication requirements within a community.

Intra-day coordination. In order to deal with the stochastic nature of part of the power consumption (associated to the users' unpredictable living activities) or with the uncertainty in the generation (associated to external conditions, such as weather in the case of photovoltaics), we use the receding horizon optimization presented in [25]. This *intra-day coordination* consists of running the proposed coordination several times during the day, taking into account possible changes in the agents' objective function and the observed power usage. This is done considering changes in the objective function which reflect changes in user preference, forecast of renewable generation, or energy cost (e.g. associated to a peak price signal). This means that the proposed framework does not handle power usage in real-time: it assumes that the real-time control is implemented by each agent (e.g. using [16]), and that it can incorporate new information as it is available through the day. In other words the coordination is used for "planning", and the agent is expected to follow the plan, but the whole community will collaborate when an agent cannot follow its planned power usage or new information is available.

2.2.4. Exchange market interpretation

As we saw, the (scaled) Lagrange multipliers are the same for all agents within a community ($\nu_j = \nu_i \ \forall i \in \mathcal{N}_j$, i.e., they are in consensus [21]). Also, it is not difficult to see that once the algorithm converges, the broadcast signal b_j^q converges to the Lagrange multiplier ν_j. The Lagrange multipliers can be interpreted as clearing prices of an exchange market [21], and the coordination determines both, the optimal power usage and the clearing prices. In that sense, the coordination works as a Tâtonnement process [26] [13], with the role of each coordinator being similar to the one of secretary of market.

In case there is no community objective function for each of the communities, and if $\rho_j = \rho_g \forall j$, the Lagrange multipliers at the inter- and intra-community levels converge to the same values. Therefore all agents participate in single a hierarchical exchange market that converge to the same clearing prices at the different levels of the coordination.

2.3. Robust communities

In cases when the agents can be assigned to any community (e.g. when the communities have no objective function), we can do the assignment seeking that each community is able to work independently and to handle unplanned situations during the coordination; thus obtaining a robust system. For this we propose a framework for assigning agents to communities and a measure of the agents' flexibility to be used in the assignment.

Agent assignment. We formulate the assignment of the agents in $\mathcal{N} = \{1, \ldots, N\}$ into J communities as a welfare maximization problem:

$$\underset{(\mathcal{N}_1,\ldots,\mathcal{N}_J)\in\mathcal{P}(\mathcal{N})}{\arg\max} \sum_{j=1}^{J} W_j(\mathcal{N}_j), \tag{P6}$$

where we seek to maximize the welfare of J communities over all possible elements of the partition $\mathcal{P}(\mathcal{N})$. The set function $W_j : 2^{\mathcal{N}} \to \mathbb{R}_+$, evaluated in a set $S \subseteq \mathcal{N}$, gives the welfare for community j when being assigned the agents in S. We will use the function W_j to measure the flexibility of community j, but other (or additional) criteria can also be used.

The assignment problem in (P6) is NP-hard [27], so when the number of agents is large, finding the optimal is not feasible. However, in case the functions W_j are monotone[5] and submodular[6], approximation algorithms exist [27] [28]. These algorithms are γ-approximations, meaning that in expectation their solution is at least γOPT, with OPT the optimal assignment. We consider the off-line setting [28], case for which an optimal approximation is obtained (with the best achievable bound for an approximation algorithm), with $\gamma = (1 - 1/e) \approx 0.632$.

The optimal approximation [28] is obtained by allocating each agent at random using a particular probability distribution, as shown in Algorithm 1 (WSAgentAssignment). This algorithm is ran a few times and the best selection kept. In general the probability distribution (line 2 in Algorithm 1) needs to be estimated once using Algorithm 2 (EstimateDistribution) (for a given set of functions $\{W_j\}_j$ and agents), while in the particular case when the welfare function is the same for all communities, $W(S) = W_j(S)\ \forall j$, the optimal approximation is obtained by allocating the agents uniformly at random.

In addition to having a bound that ensures a solution that is not too far from the optimal one, this algorithm has two more advantages: it does not require knowing the analytic form of W_j (just requires monotonicity and submodularity), and it can be used when some agents have been already assigned to communities (to build the communities incrementally).

Algorithm 1 WSAgentAssignment $\left(\mathcal{N}, J, (W_j)_{j=1}^{J}\right)$

Input: $\mathcal{N}$: set of input agents; J: number of communities
Input: $(W_j)_{j=1}^{J}$: input community welfare functions
1: $\mathcal{N}_j \leftarrow \emptyset\ \forall \in 1, \ldots, J$ // Initialize sets
2: $\omega \leftarrow$ EstimateDistribution $\left(\mathcal{N}, J, (W_j)_{j=1}^{J}\right)$
3: **for** each $i \in \mathcal{N}$ **do**
4: $\quad j \leftarrow$ sample ω_i
5: $\quad \mathcal{N}_j \leftarrow \mathcal{N}_j \cup i$ // Assign agent i to set $\mathcal{N}_j$
6: **end for**
Output: $(\mathcal{N}_j)_{j=1}^{J}$

Flexibility measure. Each community should be flexible such that it can adjust under unplanned events. From the point of view of the coordination, this can be seen as the capability of a community to have various power profiles, while considering that not all achievable profiles are equally preferred. To model this flexibility, we take the profiles that can be achieved by a set of agents and model them using a probability distribution

[5] A function $f : 2^{\mathcal{N}} \rightarrow \mathbb{R}$ is *monotone* if $f(S) \leq f(P)\ \forall S \subseteq P$.
[6] A function f is *submodular*, if for all S_1, S_2 and s such that $S_1 \subseteq S_2$ and all $s \notin S_2$, we have that $f(S_1 \cup s) - f(S_1) \geq f(S_2 \cup s) - f(S_2)$.

Algorithm 2 EstimateDistribution $\left(\mathcal{N}, J, (W_j)_{j=1}^{J}\right)$

Input: $\mathcal{N}$: input agents, and J: number of communities
Input: $(W_j)_{j=1}^{J}$: input community welfare functions

$\delta \leftarrow \frac{1}{(JN)^2}$, $\Delta \leftarrow 0$, $y_{ji}(\Delta) \leftarrow 0 \forall i \in \mathcal{N}, j \in \{1, \ldots, J\}$
while $\Delta < 1$ **do**
    Let $R_j(\Delta)$ be a random set containing each agent i independently with probability $y_{ji}(\Delta)$
    for all i, j **do**
        $\omega_{ji}(\Delta) \leftarrow \mathbf{E}\left[W_j(R_j(\Delta) \cup i) - W_j(R_j(\Delta))\right]$
        //expected marginal gain of community j from agent i
    end for
    for each i **do**
        $j_i(\Delta) \leftarrow \arg\max_j \omega_{ji}(\Delta)$ //break ties arbitrarily
        for each j **do**
            $y_{ji}(\Delta + \delta) \leftarrow y_{ji}(\Delta)$
            if $j == j_i(\Delta)$ **then**
                $y_{ji}(\Delta + \delta) \leftarrow y_{ji}(\Delta) + \delta$ // preferred group
            end if
        end for
    end for
    $\Delta \leftarrow \Delta + \delta$
end while
Output: ω

of their aggregated profile. Then, we seek that this distribution is as "spread" as possible, and measure this using the *entropy* of the community's aggregated profile[7].

Let us call $X_i \in \mathbb{R}^T$ the multivariate random variable associated with the profile of agent $i \in \mathcal{N}_j$, and $Y_j = \sum_{i \in \mathcal{N}_j} X_i$ the random variable associated to the aggregated profile of the community j. We want the entropy of the aggregated profile $H(Y_j)$ to be large, so that the community j is flexible. Given that the entropy of a sum of multivariate random variables evaluates the *entropy of sums*[8] (known to be submodular and monotone [30]), we can use Algorithm 1.

We consider a particular case for the entropy as a flexibility measure. We assume the random variables X_i, $i = 1, \ldots, N$ are independent multivariate normals (MVN). Thus Y_j, is also a MVN. If we call $\mathcal{K}_{Y_j}$ the covariance matrix of Y_j, then the entropy of Y_j can be written as $H(Y_j) \propto \log |\mathcal{K}_{Y_j}|$, where an irrelevant additive factor depending on the dimension T has not been considered. We define the set function measuring the flexibility as:

$$W^H(\mathcal{S}) = \log \left| \sum_{i \in S} \mathcal{K}_i \right|, \tag{8}$$

with $\mathcal{K}_i$ the covariance matrix of agent i. There are various ways to obtain the covariance matrices $\{\mathcal{K}_i\}_i$: *i*) every agent i could estimate its matrix $\mathcal{K}_i$ and send it to the coordinator, or *ii*) the profiles $\{x_i^k\}_k$, obtained by a coordinator during past negotiations can be used as samples to approximate the covariance matrix: $\mathcal{K}_i = \frac{1}{K-1} \sum_{k=1}^{K} (x_i^k - \bar{m}_i)(x_i^k - \bar{m}_i)^T$, with $\bar{m}_i = \frac{1}{K} \sum_{k=1}^{K} x_i^k$ (this is used later in the example scenarios). To avoid numerical issues when calculating the log determinant, we use the sum of the non-zero singular values (SVD) of the covariance matrix.

3. Example Scenarios

We illustrate the proposed coordination for inter-community imbalance minimization in three scenarios: *i*) *coordination under community constraints*, *ii*) *coordination for demand response as a service*, and *iii*) *coordination under unexpected events*. These scenarios were selected to illustrate what can be achieved by the coordination, and to analyze the algorithm and its robustness. We consider the problem of coordinated Electric Vehicles (EVs) charging, where the charging is controlled by an EMS (the coordinating agent) at the households[9]. The coordination should help avoiding producing large consumption peaks due many EVs being charged at a similar time (e.g. during the evening after the EVs are plugged).

To focus on the coordination, we only considerer EVs, assume their charging cannot be interrupted, and neglect the power usage of other devices – although any controllable appliance could be considered in the coordination: e.g. A/C (level adjustment), laundry (usage timing), and batteries.

[7]An agent that can generate many profiles with equal probability is flexible (high entropy). If it has only one profile, it is not flexible (low entropy).

[8]Note that this is not the same as the entropy of a joint probability distribution, which is more common and also submodular [29].

[9]In the following we use *EV* and *EMS* as interchangeable terms.

General setup. The power usage is coordinated over a day using profiles of dimension $T = 144$ (10-minute time slots). We assume the EVs must be charged (3kWh) continuously, and the charging requires $1000[W]$. Thus charging one EV takes about 18 time-slots (3 hrs). Each EV has its preferred starting time and allows some variability around it. We consider a day-ahead coordination, followed by an intra-day coordination required due to changes in objective functions (that represent unforeseeable changes in user preference or weather conditions). Unless stated, we use $\rho_g = \rho_j = 0.1 \times 10^{-6}\ \forall j$.

Agent model. We use the probabilistic model proposed in [9] and briefly described in Section 2.1. Each EV is represented by three modes ($M = 3$): $\mathcal{Q} = \{q_1, q_2, q_3\}$ indicating the period before, during, and after charging respectively, and each mode is defined by its output and duration probability distributions. We assume the $1000[W]$ required charging happens during mode q_2, and this is modeled using an output distribution given by $P(x_{i,t}|s_t) = 1_{[x_{i,t}=\mathcal{X}(s_t)]}$, with $1_{[\circ]}$ the indicator function, $x_{i,t}$ the profile x_i at time t, and $\mathcal{X}(s_t) = 0, 1000, 0[W]$ for $s_t = q_1, q_2, q_3$ respectively. Thus $f_i^{x|u}(x_i, u_i) = -\sum_{t=1}^{T} \log 1_{[x_{i,t}=\mathcal{X}(s_t)]}$, with $u_i \triangleq s_{1:T} \in \{q_1, q_2, q_3\}^T$, and x_i uniquely determined by u_i. We assume the duration of mode m for EV i follows a Gaussian distribution[10] $\mathcal{N}(\mu_{i,m}, \sigma_{i,m}^2)$.

For the coordination in Scenarios 1 and 2, we consider two types of users:
A) $\mu_{i,1}$ uniformly distributed in $[30, 35]$; $\sigma_{i,1} = 9$, and *B*) $\mu_{i,1}$ uniformly distributed in $[100, 105]$; $\sigma_{i,1} = 3$.
The parameters of the other two modes are given by: $\mu_{i,2} = 18$ (3 hrs); $\sigma_{i,2} = 1$, and $\mu_{i,3} = T - (\mu_{i,1} + \mu_{i,2})$; $\sigma_{i,3} = 10$, for both types of users. Thus, agents of *type A* are more flexible in their start charging time (larger sigma), than the *type B* ones. Agents of *type A* prefer to charge the EVs in the early morning, while agents of *type B* prefer to charge in the evening.

Day-ahead coordination. In the three scenarios, a day-ahead coordination is performed, where the agents, together with the global coordinator, perform a coordination and determine their planned aggregated power profile. More precisely, the coordinator seeks to balance the aggregated power profile using a global cost of the form $g^0(\sum_{i\in\mathcal{N}} x_i) = \alpha||\sum_{i\in\mathcal{N}} x_i - r^0||^2$, with $\alpha = 0.1 \times 10^{-6}/N$, and using the algorithm in Eq. (6). The profile $r^0 \in \mathbb{R}^T$ is a predefined power usage target depending on the scenario. The obtained day-ahead aggregated profile $r_j = \sum_{i\in\mathcal{N}_j} r_{(i)}$, with $r_{(i)}$ the day-ahead coordinated profile of agent i, can be used as a reference during a intra-day coordination.

Intra-day coordination. We use the distributed architecture (as in Fig. 1) with $J = 4$ communities, and the agents are assigned to communities using the proposed assignment framework. An intra-day coordination can take place due to: (i) agents not following their scheduled power usage, or (ii) agents (or coordinators) having changed their objective functions. The used global and community objectives are specific to each scenario.

From the formulation and implementation point of view, we briefly mention two key issues. The objective functions may change over time [23], thus we replace $f_i(x_i)$ by $f_i^t(x_i)$ and $g(\sum_i x_i)$ by $g^t(\sum_i x_i)$, with the index t representing the current time. Second, when

[10]The PDF can be learned from sensed data, but we have set it manually.

solving Eq. (2) not all components of the variable x_i are optimization variables: if we write $x_i = (x_{i,1}, \ldots, x_{i,t-1}, x_{i,t}, \ldots, x_{i,T})$, the components $(x_{i,1}, \ldots, x_{i,t-1})$ are observations (have fix values), while $(x_{i,t}, \ldots, x_{i,T})$ are optimization variables. This changes the way the generative model is solved [25].

3.1. Scenario 1: Imbalance minimization with community constraints

This scenario illustrates the coordination under community constraints and shows the influence of the parameters ρ_g and ρ_j in the convergence of the algorithm. Here $N = 400$ electric vehicles (EVs) coordinate their power usage seeking to match the local PV generation. A day-ahead PV forecast indicates a generation that provides, through the day, about 2/3 the energy required to charge the EVs. The EVs coordinate to balance the remaining required power (r^0 is the PV forecast).

Fig. 2 shows the evolution of the coordinated profile for the given day-ahead PV forecast (sunny, shown in green). The obtained coordinated profile (after 150 iterations) is shown in red. Recall that some EVs prefer to charge in the morning, while others prefer the evening, but most PV generation occurs during the afternoon, thus the EVs must modify their charging time. The dashed (cyan) curve shows the profile obtained without coordination. Fig. 2 (bottom) indicates a fast convergence of the global coordinator's cost.

Now we assume the coordinators receive an updated PV generation forecast in the early morning (at $t = 30$), indicating a cloudy period during the time range $[57, 72]$. Thus, the 4 communities coordinate to match the updated PV forecast. Fig. 3(a) shows the evolution of the coordinated profile. Note that just 40 iterations are needed for the upper level of the hierarchy, as the intra-day coordination uses day-ahead profiles as initial conditions.

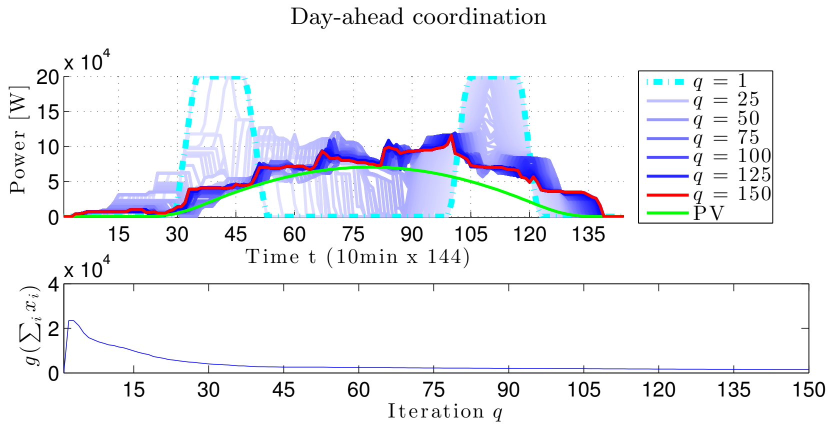

Figure 2: Scenario 1. Power balancing with PV generation forecast. Subfigures. *Top*: Coordinated profile evolution (dashed cyan: first iteration; light-blue/dark-blue earlier/later iterations) and PV forecast (green). *Bottom*: Global objective evolution.

We consider the case where each community (e.g. a multi-dwelling) has the constraint that *all local PV generation of a community should be used inside the community* (in some countries –e.g. Spain– there have been high penalties for local generation sold back to the

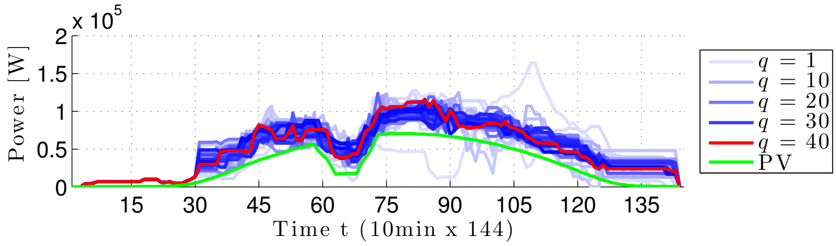

(a) Without community constraints

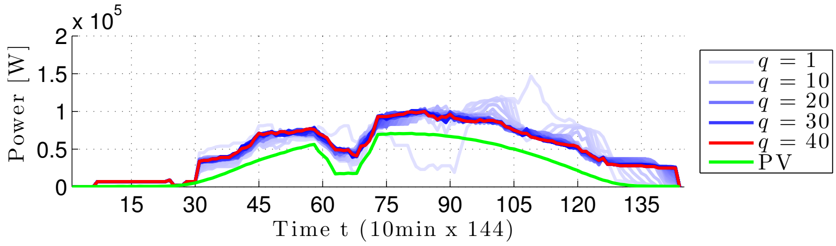

(b) With community constraints

Figure 3: Scenario 1. (Intra-day). Aggregated power profile. Scenario: An updated PV forecast indicates a reduction in generation in the time range $[57, 72]$. In the case of community constraints (b), each of the communities must consume more than the local PV.

grid). Formally, the objective of community j at time t is $g_j^t(x) = -\sum_{\tau=t}^{T} \log 1_{[x_{(\tau)} \geq v_{(\tau)}^{j,t}]}$, with $v_{(\tau)}^{j,t} \in \mathbb{R}^T$ the PV generation forecast of time-slot τ (at time t), and $x_{(\tau)}$ the τ^{th} component of $x \in \mathbb{R}^T$; i.e., the consumption of community j should be larger than its generation at all times. Fig. 3 shows the obtained coordinated profile, and Fig. 4 the obtained communities' coordinated profiles, with and without community constraints for different values of $\rho = \rho_g = \rho_j$. For the case without constraints (in Fig. 4(a)) we can observe that for some time periods the coordinated profiles has larger generation than consumption (see for example the time range $[90, 110]$ for community 1 (top-left plot)), while in the case with community constraints (Fig. 4(b)) all four communities fulfilled the constraint at all times. We can also observe that the convergence of the global cost does not depend too much on value of ρ and it is faster in the case with community constraints. From the coordinated profiles (Fig. 4(b)) we can see that smaller values of ρ_j are preferable.

3.2. Scenario 2: Imbalance minimization for demand response as a service

This scenario illustrates how the EVs can coordinate to balance their power usage and provide a DR service, and we also compare this to the use of price-based demand response. The global coordinator seeks to flatten the power usage, but it may receive a signal indicating a critical-peak-pricing (CPP) event, thus an updated coordination is needed. In this scenario we also briefly study two features of the coordination: (a) its robustness under communication problems, and (b) the use of constraints in the number of iterations for the intra-community coordination.

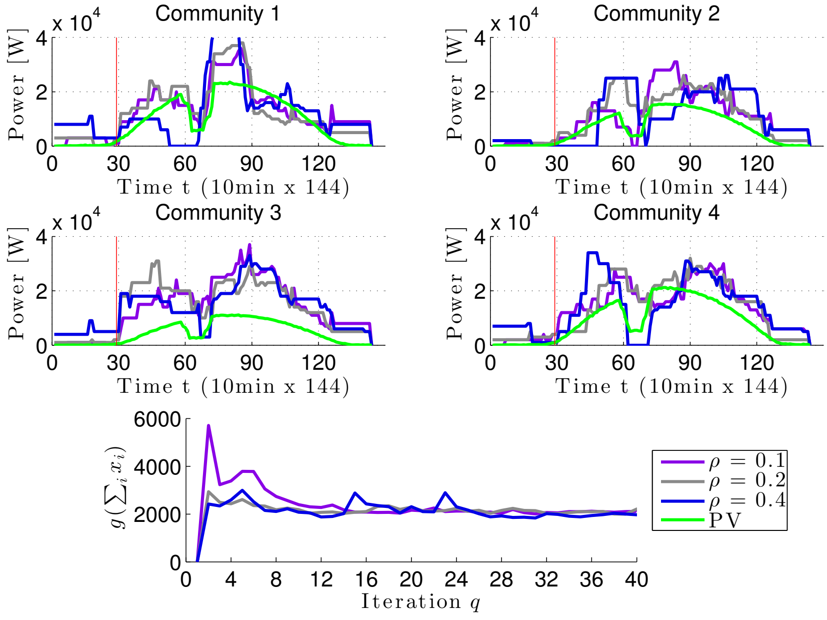

(a) Without community constraints

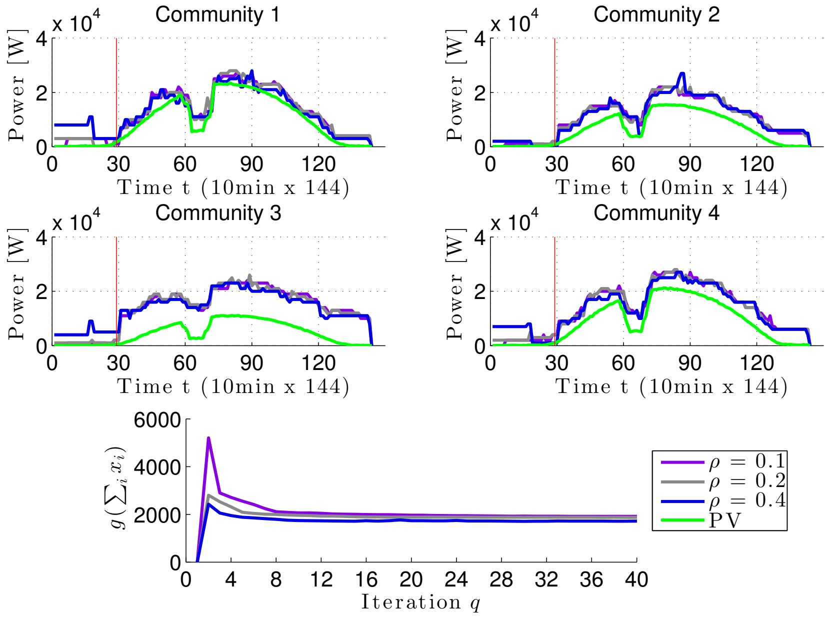

(b) With community constraints

Figure 4: Scenario 1 (Intra-day). Coordinated profiles (top) and global objective convergence (bottom) with and without constraints. Different values of $\rho = \rho_g = \rho_j$ are used. Note that the PV generation differs among the communities. In (b), the constraint (demand larger than supply for each community) is fulfilled at all times.

In the day-ahead coordination, the $N = 400$ agents balance their power usage, obtaining the coordinated profile shown in Fig. 5 (black). Note that the power is less than $0.5 \times 10^5 [W]$ in the time range $[0, 85]$, slightly larger for $[86, 139]$, while no power is used in the time range $[140, 144]$. As in Scenario 1, this is due to the distinct flexibility of the two agents' types.

Now we assume that at time t_{cpp} there is a signal indicating a CPP event for the time range $[40, 50]$. This price signal, received by the global coordinator, is intended to cause a reduction in the power usage during that period and requires an intra-day coordination. To integrate this CPP signal, the global coordinator's objective at time τ is: $g^{(\tau)}(x) = \alpha ||x||^2_{W_\tau(\beta)}$, with $W_\tau(\beta) \in \mathbb{R}^{T \times T}$ and $\alpha = 0.1 \times 10^{-6}/N$. Normally $W_\tau(\beta)$ is the identity matrix (the goal is to flatten the power usage), but it can be used to indicate a time-varying price. We use[11]:

$$W_\tau(\beta) = \begin{cases} I & \text{if } \tau < t_{cpp} \\ diag(b_{cpp}) & \text{otherwise,} \end{cases} \tag{9}$$

$$\text{with } b_{cpp} = (\underbrace{1, \ldots, 1}_{T_1}, \underbrace{\beta, \ldots, \beta}_{T_2}, \underbrace{1, \ldots, 1}_{T_3}),$$

where $t_{cpp} = 30$ (the time when the CPP event is informed), with $T = T_1 + T_2 + T_3$, $T_1 = 39$ (the duration before the CPP event), $T_2 = 11$ (the duration of the CPP event), and $T_3 = 94$. Here we use a value of $\beta = 10$, representing a 10-fold increase in the quadratic penalty cost for the range $[T_1 + 1, T_1 + T_2]$.

Fig. 5 (red) shows the intra-day coordination results under this CPP signal. The obtained profile (red) reduces the power usage to more than half compared to the day-head one for the time range $[40, 50]$, without producing any peak rebound. Recall that the agents are not interruptible, thus the EVs that started charging before $\tau = 30$ might have not finished charging, causing the low (but not zero) consumption in the first part of the period $[40, 50]$.

Fig. 6 presents the results obtained by a *price-based demand response* program. This is formulated as each agent solving independently (no coordination) the problem $\arg\min_x f_i^{(\tau)}(x) + \alpha ||x||^2_{W_\tau(\beta)}$, with α and W_τ as defined above. Here we considered different values of $\beta \in 1, 2, \ldots 6$. We can observe that the obtained aggregated profile produce a large peak rebound for large values of β, while small values of β produce high consumption peak during the CPP period. On the contrary, by coordinating (Fig. 5), a more balanced power usage pattern is obtained, with a low power usage during the critical peak pricing period and without causing peak rebounds.

Also, when one community cannot communicate with the global coordinator (is shown in Fig. 5, blue)) the remaining communities and the global coordinator still coordinate, achieving a profile very similar to the one obtained with full communication (Fig. 5, red), and far from the ones obtained by demand response (Fig. 6).

Inexact solutions of the proximal operator. ADMM is known to converge even if the x- and z- steps are inexact (under some suboptimality measures [21] [24]), which allows solving the x-step approximately for a community. We analyze two uses of this ability.

[11] We are assuming a quadratic cost of the power usage. This simplification follows from the fact that part of fuel cost of a thermal generation unit can be approximated as a quadratic function of the electric power output [31].

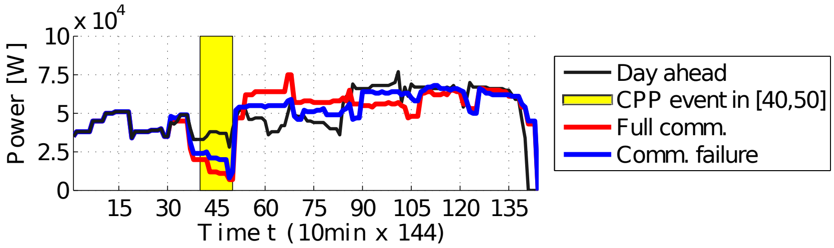

Figure 5: Scenario 2: Coordinated DR as a service. After the day-ahead coordination (black), there is a CPP event for the period [40, 50] (yellow). Intra-day coordinated profiles under full communication (red) and communication failures (blue) are similar.

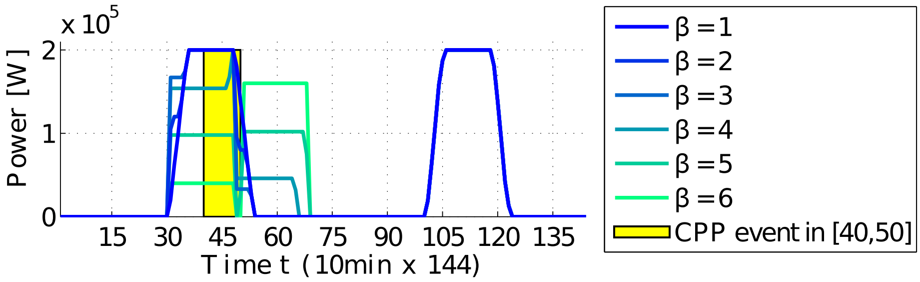

Figure 6: Scenario 2 (intra-day). Aggregated profile for a *price-based demand response* program (no coordination). β indicates a critical pick price in [40, 50].

- *Number of iterations*: Given that the coordination within a community is nested in the coordination at the upper levels of the hierarchy, we may want to restrict the time used in the intra-community coordination (e.g. to reduce the processing power in the agents or the required communication bandwidth). We do this by limiting the number of iterations in the community's coordination. Fig. 7 shows the obtained results when we limit k, the number of iterations within each community. We can see that using 10 or more iterations gives good results, with 20 and 25 iterations giving very similar profiles (Fig. 7(a)) and convergence speeds (Fig. 7(b)). However, using k equal to 1 might require a large number of iterations at the global level, which should be avoided.

- *Data package losses*: The coordination is capable of handling data package losses by approximating lost packages (by using the last available solution to the coordinator). Fig. 8 shows the result of approximating the solution of the x-step for a percentage (in $\{0, 20, 40, 60, 80\}$) of lost packages. It can be observed that the larger the number of iterations, the smaller the deviation from the optimal solution (Fig. 8, bottom) and the smaller the norm of the power profile (Fig. 8, top). Moreover, in terms of deviation from the optimal solution (see Fig. 8, bottom), for each number of iterations larger of equal than 10, the obtained results are very similar for all loss rates. Thus, data package losses have a negative effect only for small number of iterations (k less than 10).

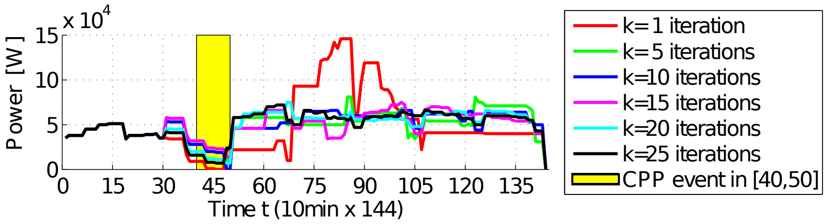

(a) Aggregated coordinated profiles

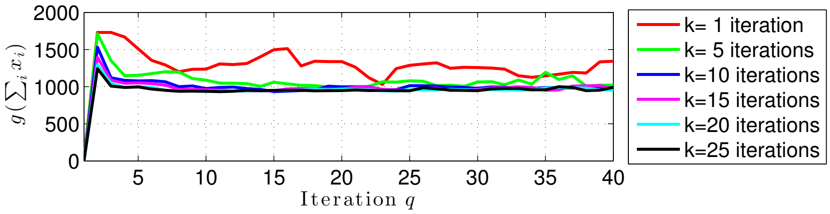

(b) Global objective evolution

Figure 7: Scenario 2 (Intra-day). Number of iterations, k, for the groups. (a) Aggregated coordinated profiles for all EVs (for $q = 40$ iterations at the global level), and (b) Global objective.

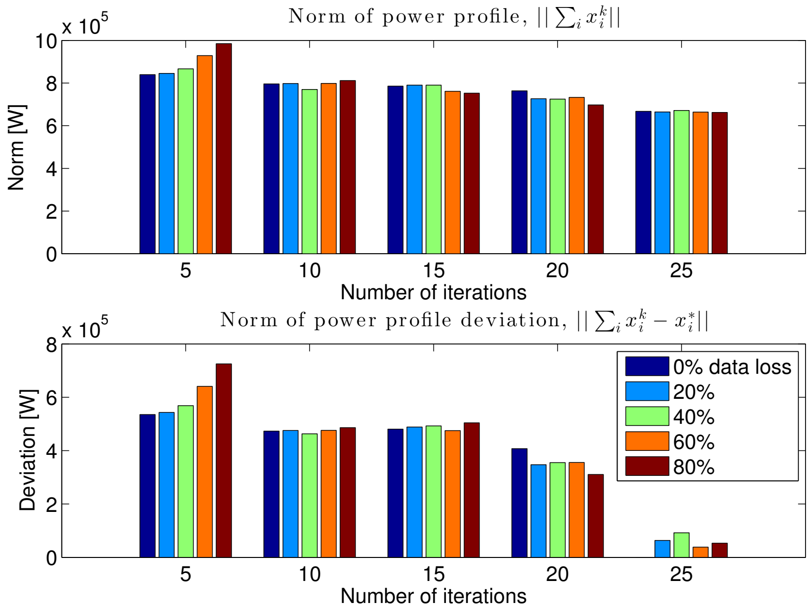

Figure 8: Scenario 2 (Intra-day). Number of iterations (k) versus data package losses. Top: Norm of the aggregated power profile. Bottom: Norm of deviation $||\sum_i x_i - \sum_i x_i^{25}||$ (with x_i^{25} the profile obtained after 25 iterations without package losses). Note that data package losses have a negative effect only for few iterations (k less than 10).

3.3. Scenario 3: Imbalance minimization robustness under unexpected events

Having a robust coordination is very important, in particular in case of emergencies and unexpected events. In the current scenario we perform a more exhaustive analysis of the coordination robustness, taking into account two factors: the number of communities that cannot communicate, and the percentage of agents that do not follow their planned power usage.

Communication link failures. We consider two types of users, with a percentage of them changing their behavior during the day. We use the same model as in the previous scenarios, but with different parameters. We let all EVs set a preferred starting time $\mu_{i,1}$ uniformly distributed in $[50, 70]$, and $\sigma_{i,1}$ take values in $\{3, 9\}$ (evenly distributed) among the users, i.e. about half the users is more flexible in their start charging time than the other half. The remaining modes' duration probabilities are defined as in the previous scenarios ($\mu_{i,2} = 18$ (3 hrs); $\sigma_{i,2} = 1$, and $\mu_{i,3} = T - (\mu_{i,1} + \mu_{i,2})$; $\sigma_{i,3} = 10$).

The agents are assigned to $J = 4$ communities, seeking to maximize the flexibility of each community, using Algorithm 1 (WSAgentAssignment) with 10^4 samples. The covariance matrices of the agents, $\mathcal{K}_i$, are estimated by the global coordinator using the profiles observed during the day-ahead coordination $\{x_i^k\}_{k=1}^K\ \forall i$.

We consider that some agents deviate from their planned charging time, while also contributing less to the coordination. This information is available at time $\tau = 10$, and is modeled as the agents in $\hat{\mathcal{N}} \subseteq \mathcal{N}$ changing their preferred starting time and reducing their starting time flexibility. In order to analyze the robustness of the coordination, we modify the proportion of agents that change their behavior: $p = |\hat{\mathcal{N}}|/|\mathcal{N}|$ with $p \in \{0, 0.2, 0.4, 0.6, 0.8\}$. At time τ, the agents in $\hat{\mathcal{N}}$ update their parameters as follows: $\sigma_{i,1} \leftarrow 1$ and $\mu_{i,1} \leftarrow \mu_{i,1} + \delta_i$, with δ_i following a uniform distribution in $[-10, 10]$. In other words, besides the change in preferred starting time, the flexibility diminishes considerably ($\sigma_{i,1}$ goes from 3 or 9 to 1). Other parameters do not change (except $\mu_{i,3} = T - (\mu_{i,1} + \mu_{i,2})$).

The global objectives in the day-ahead and intra-day coordination are different. While in the former case the goal is to balance the aggregated power profile, in the later case (at each time τ) the goal is to minimize its deviation from a reference r. Here we consider this reference to be aggregated power usage $r = \sum_i r_i$ obtained in the day-ahead coordination. Thus, the global cost at time t is $g^t(\sum_{i\in\mathcal{N}} x_i) = \alpha||\sum_{i\in\mathcal{N}} x_i - r||^2$, with $\alpha = 10^{-6}/N$. In the cases where one or more communities cannot communicate with the global coordinator, the fall back procedure described in Section 2.2.3 is used.

Fig. 9 summarizes the results for different values of p and for different numbers of communities not being able to communicate with the global coordinator. It can be observed that when 20% or 40% of the agents change their planned pattern usage, the coordination obtains a profile similar to day-ahead plan, even when there is only communication within the communities. Moreover, when 80% (p=0.8) of the agents change their behavior and become inflexible, and there is only coordination within the communities (Fig. 9, red curve), the aggregated profile is far from the *no-coordination* case. These results show the robustness of the proposed coordination: the coordination still achieves an aggregated profile that is not too far from the planned one, even if many agents change their planned behavior and some communities cannot communicate.

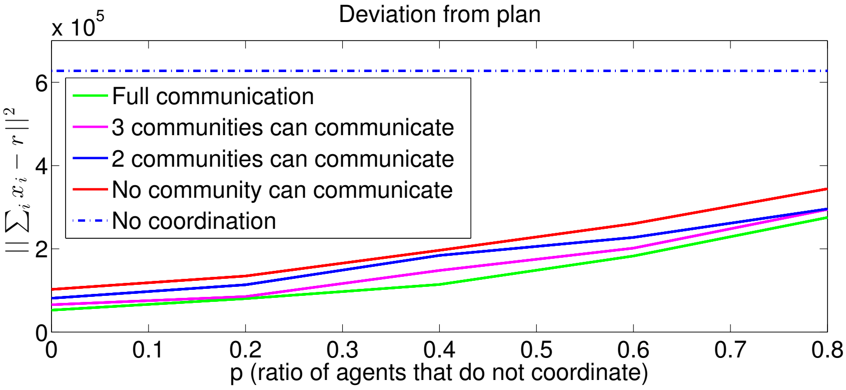

Figure 9: Deviation from the reference when a set of groups cannot communicate with the global coordinator and when $p \times 100\%$ of the agents do not follow the plan. The *no coordination* case (blue dotted line) is shown for comparison.

Recurrent deviation from plan. The proposed framework does not explicitly model the stochastic nature of power consumption or supply, and in its basic form is designed to work with controllable devices. It assumes the agents do their best effort to follow the day-ahead planned profiles, but in real-world scenarios this may not always apply. To handle unexpected events of stochastic nature, the *intra-day coordination* is used. In the current scenario we analyze the case several agents change their cost function different times of the day.

Similarly to previous experiments, we let all EVs set a preferred starting time $\mu_{i,1}$ uniformly distributed in $[50, 55]$, and $\sigma_{i,1}$ evenly distributed in $\{3, 9\}$. The remaining modes' duration are defined as in the previous scenarios ($\mu_{i,2} = 18$; $\sigma_{i,2} = 1$, and $\mu_{i,3} = T - (\mu_{i,1} + \mu_{i,2})$; $\sigma_{i,3} = 10$). These are the parameters used in the day-ahead coordination (by all agents), and the parameters used in the intra-day coordination by all agents that do not change their cost function.

We consider that a percentage p of the agents (with $p \in \{80, 60, 40\}$) change their cost function. Each of these agents change its parameters at time τ_i as follows: $\sigma_{i,1} \leftarrow 1$ and $\mu_{i,1} \leftarrow \mu_{i,1} + \delta_i$, with δ_i uniformly distributed in $[0, 10]$ (there is a change in preferred starting time and the flexibility highly decreases). Other parameters do not change (except $\mu_{i,3} = T - (\mu_{i,1} + \mu_{i,2})$). The value of τ_i follows a Gaussian distribution $\mathcal{N}(\mu_{i,1} - \Delta_i, 5)$, with $\Delta_i \in \{0, 10, 20, 30\}$. A negative value of τ_i indicates that the agent changes its plan before the day-ahead coordination, a positive one indicates the time at which the agent will update its cost function, while when τ_i is larger than μ_i, the agent just start using power usage the new cost function is never used in the coordination).

Fig. 10 presents the obtained results (average over 10 runs with $N = 256$ agents). In the top row we can observe how the power is balanced by the community, with and without the intra-day coordination, in terms of the l_2-norm of the aggregated power $||\sum_i \breve{x}_i||^2/||\sum_i x_i||^2$ (relative to $\breve{x}_i$, the profile without coordination). In the bottom row we can observe results in terms of the deviation from the day-ahead plan: $||\sum_i x_i - x_i^*||^2/||\sum_i x_i^*||^2$ (with x_i^* the day-ahead plan). In both figures, smaller values are better. Note that the intra-day coordination helps in balancing the power usage and in achieving profiles closer to the planned one, even if large number of agents do not follow the plan and this information is only know short time before the power used.

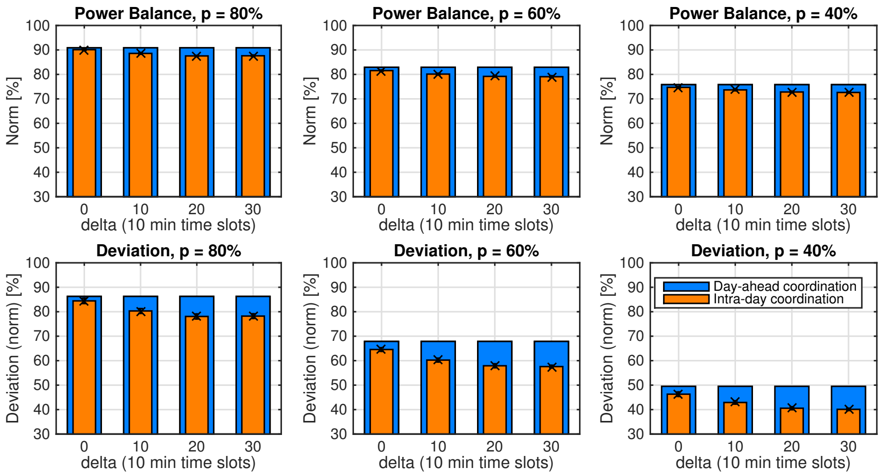

Figure 10: Normalized power balance and deviation from day-ahead plan when a percentage $p \in \{80, 60, 40\}$ of the agents do not follow the plan and change their cost function during the day at different times. Results are the average over 10 runs. See text for details.

4. Conclusions

We have proposed a coordinated energy management framework for inter-community imbalance minimization. The end-users in each community coordinate and plan the aggregated power usage pattern, allowing to balance the aggregated energy usage. The coordination is done via (global and community) coordinators in a hierarchical distributed architecture, that can take into account global, community and end-users' objectives and constraints. Seeking to have a robust system that can better respond to unplanned changes and communication problems, we have also proposed a framework for assigning the agents to communities by taking into account a measure of the flexibility of each agent. The proposed framework was evaluated in three simulated example scenarios, where we demonstrate that the proposed framework can be used to coordinate groups of agents and minimize inter-community imbalance, while taking into account the agent and the groups objectives. In addition, it was observed that the framework *i*) is robust under communication problems (failure in communication links and data package losses), *ii*) has relatively low communication requirements, *iii*) can handle the case when agents do not follow their specified plan, and *iv*) is not strongly dependent on the algorithm's parameters (number of iterations and penalty parameter ρ). Also of importance is that the coordination is more effective that traditional price-based method (e.g. critical peak pricing), and that the communities can coordinate to provide demand respond as a service. Possible future directions include: the design of an incentive model for the end-users, to explicitly consider the stochastic nature of generation and consumption, to manage distributed uncontrollable generation (e.g. photo-voltaics) and distributed storage, and to extend the coordination framework to manage the community's power usage in real-time.

Acknowledgment

This work was supported by CREST from the Japan Science and Technology.

References

[1] P. Siano, Demand response and smart grids – a survey, Renewable and Sustainable Energy Reviews 30 (0) (2014) 461 – 478. doi:http://dx.doi.org/10.1016/j.rser.2013.10.022.

[2] P. Palensky, D. Dietrich, Demand side management: Demand response, intelligent energy systems, and smart loads, IEEE Transactions on Industrial Informatics 7 (3) (2011) 381–388. doi:10.1109/TII.2011.2158841.

[3] A. Barbato, A. Capone, Optimization models and methods for demand-side management of residential users: A survey, Energies 7 (9) (2014) 5787–5824. doi:10.3390/en7095787.

[4] L. Gelazanskas, K. A. Gamage, Demand side management in smart grid: A review and proposals for future direction, Sustainable Cities and Society 11 (0) (2014) 22 – 30. doi:http://dx.doi.org/10.1016/j.scs.2013.11.001.

[5] M. Alizadeh, X. Li, Z. Wang, A. Scaglione, R. Melton, Demand-side management in the smart grid: Information processing for the power switch, IEEE Signal Processing Magazine 29 (5) (2012) 55–67. doi:10.1109/MSP.2012.2192951.

[6] S. Caron, G. Kesidis, Incentive-Based energy consumption scheduling algorithms for the smart grid, in: IEEE International Conference on Smart Grid Communications (SmartGridComm), 2010, pp. 391–396. doi:10.1109/SMARTGRID.2010.5622073.
URL http://dx.doi.org/10.1109/SMARTGRID.2010.5622073

[7] M. Alizadeh, T.-H. Chang, A. Scaglione, Grid integration of distributed renewables through coordinated demand response, in: IEEE 51st Conference on Decision and Control (CDC), 2012, pp. 3666–3671. doi:10.1109/CDC.2012.6426122.

[8] O. Ardakanian, C. Rosenberg, S. Keshav, Distributed control of electric vehicle charging, in: 4th International Conference on Future Energy Systems (e-Energy '13), ACM, New York, NY, USA, 2013, pp. 101–112. doi:10.1145/2487166.2487178.

[9] H. Kawashima, T. Kato, T. Matsuyama, Distributed mode scheduling for coordinated power balancing, in: IEEE International Conference on Smart Grid Communications (SmartGridComm), 2013, pp. 19–24. doi:10.1109/SmartGridComm.2013.6687927.

[10] S. Kishore, L. Snyder, Control mechanisms for residential electricity demand in smartgrids, in: IEEE International Conference on Smart Grid Communications (SmartGridComm), 2010, pp. 443–448. doi:10.1109/SMARTGRID.2010.5622084.

[11] T.-H. Chang, M. Alizadeh, A. Scaglione, Real-time power balancing via decentralized coordinated home energy scheduling, IEEE Transactions on Smart Grid 4 (3) (2013) 1490–1504. doi:10.1109/TSG.2013.2250532.

[12] L. Gkatzikis, I. Koutsopoulos, T. Salonidis, The role of aggregators in smart grid demand response markets., IEEE Journal on Selected Areas in Communications 31 (7) (2013) 1247–1257.

[13] H. Kim, M. Thottan, A two-stage market model for microgrid power transactions via aggregators, Bell Labs Technical Journal 16 (3) (2011) 101–107. doi:10.1002/bltj.20524.

[14] M. Ali, A. Alahivl, F. Malik, M. Humayun, A. Safdarian, M. Lehtonen, A market-oriented hierarchical framework for residential demand response, International Journal of Electrical Power & Energy Systems 69 (0) (2015) 257 – 263. doi:http://dx.doi.org/10.1016/j.ijepes.2015.01.020.

[15] R. Verschae, H. Kawashima, T. Kato, T. Matsuyama, A distributed hierarchical architecture for community-based power balancing, in: IEEE International Conference on Smart Grid Communications (SmartGridComm), 2014, pp. 169–175.

[16] T. Kato, K. Yuasa, T. Matsuyama, Energy on demand: Efficient and versatile energy control system for home energy management, in: Smart Grid Communications (SmartGridComm), 2011 IEEE International Conference on, 2011, pp. 392–397. doi:10.1109/SmartGridComm.2011.6102354.

[17] M. Ostendorf, V. Digalakis, O. A. Kimball, From hmms to segment models: A unified view of stochastic modeling for speech recognition, IEEE Transactions on Speech and Audio Processing 4 (1995) 360–378.

[18] S. Levinson, Continuously variable duration hidden markov models for automatic speech recognition, Computer Speech and Language 1 (1) (1986) 29 – 45. doi:http://dx.doi.org/10.1016/S0885-2308(86)80009-2.

[19] P. Combettes, J.-C. Pesquet, Proximal splitting methods in signal processing, in: H. H. Bauschke, R. S. Burachik, P. L. Combettes, V. Elser, D. R. Luke, H. Wolkowicz (Eds.), Fixed-Point Algorithms for Inverse Problems in Science and Engineering, Springer Optimization and Its Applications, Springer New York, 2011, pp. 185–212. doi:10.1007/978-1-4419-9569-8_10.

[20] N. Parikh, S. Boyd, Proximal algorithms, Foundations and Trends in Optimization 1 (3) (2014) 127–239. doi:10.1561/2400000003.

[21] S. Boyd, N. Parikh, E. Chu, B. Peleato, J. Eckstein, Distributed optimization and statistical learning via the alternating direction method of multipliers, Foundations and Trends in Machine Learning 3 (1) (2011) 1–122. doi:10.1561/2200000016.

[22] S. Magnússon, P. Chathuranga Weeraddana, M. G. Rabbat, C. Fischione, On the Convergence of Alternating Direction Lagrangian Methods for Nonconvex Structured Optimization Problems, ArXiv e-printsarXiv:1409.8033.

[23] Q. Ling, A. Ribeiro, Decentralized dynamic optimization through the alternating direction method of multipliers, IEEE Transactions on Signal Processing 62 (5) (2014) 1185–1197. doi:10.1109/TSP.2013.2295055.

[24] J. Eckstein, D. P. Bertsekas, On the douglas-rachford splitting method and the proximal point algorithm for maximal monotone operators, Mathematical Programming 55 (1-3) (1992) 293–318.

[25] R. Verschae, H. Kawashima, T. Kato, T. Matsuyama, A distributed coordination framework for on-line scheduling and power demand balancing of households communities, in: European Control Conference (ECC), 2014, pp. 1655–1662. doi:10.1109/ECC.2014.6862394.

[26] D. A. Walker, Walras's theories of tatonnement, Journal of Political Economy 95 (4) (1987) pp. 758–774.

[27] A. Krause, D. Golovin, Submodular function maximization, in: L. Bordeaux, Y. Hamadi, P. Kohli (Eds.), Tractability: Practical Approaches to Hard Problems, Cambridge Univ. Press, 2014, pp. 71 – 104.

[28] J. Vondrak, Optimal approximation for the submodular welfare problem in the value oracle model, in: Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing, STOC '08, ACM, New York, NY, USA, 2008, pp. 67–74. doi:10.1145/1374376.1374389.

[29] T. M. Cover, J. A. Thomas, Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing), 99th Edition, Wiley-Interscience, 1991.

[30] M. Madiman, On the entropy of sums, in: Proceedings of the 2008 IEEE Information Theory Workshop, 2008, pp. 303–307.

[31] B. Ummels, M. Gibescu, E. Pelgrum, W. Kling, A. Brand, Impacts of wind power on thermal generation unit commitment and dispatch, IEEE Transactions on Energy Conversion 22 (1) (2007) 44–51. doi:10.1109/TEC.2006.889616.