arXiv:2009.02166v1 [cs.MA] 4 Sep 2020

Collaboratively Optimizing Power Scheduling and Mitigating Congestion using Local Pricing in a Receding Horizon Market

Coen van Leeuwen
TNO
Groningen, Netherlands
coen.vanleeuwen@tno.nl

Joost Stam
University of Twente
Enschede, Netherlands
j.stam-2@student.utwente.nl

Arun Subramanian
TNO
The Hague, Netherlands
arun.subramanian@tno.nl

Koen Kok
TNO
Groningen, Netherlands
koen.kok@tno.nl

Abstract—A distributed, hierarchical, market based approach is introduced to solve the economic dispatch problem. The approach requires only a minimal amount of information to be shared between a central market operator and the end-users. Price signals from the market operator are sent down to end-user device agents, which in turn respond with power schedules. Intermediate congestion agents make sure that local power constraints are satisfied and any potential congestion is avoided by adding local pricing differences. Our results show that in 20% of the evaluated scenarios the solutions are identical to the global optimum when perfect knowledge is available. In the other 80% the results are not significantly worse, while providing a higher level of scalability and increasing the consumer's privacy.

Index Terms—Smart Grid, Multi-Agent Systems, Market, Self-Organization

I. INTRODUCTION

Three trends set a challenge for future power grids. Firstly, the transition towards sustainable energy sources leads to more renewable energy, but also to a larger fraction of unpredictable and intermittent production. Secondly, the electrification of various systems such as transport (electric vehicles), heating (heat pumps), and in general an increase of electricity-consuming devices leads to a huge growth of power consumption. And thirdly, the distribution of energy generation leads to a very different pattern in the load of the power transmission grid, than it was designed for.

The control of a vast number of small power units, both consuming and producing, is extremely difficult to do completely top-down, so a centralized control strategy cannot be used [1]. At the same time the power infrastructure is aging, and was not built for the emerging pattern of distributed *prosumers* [2]. This is why we need self-organizing control algorithms to schedule the use of electric devices while taking into account constraints of the power distribution infrastructure. This aids *distribution system operators* (DSO) and *transport system operators* (TSO) to maintain power balance and make sure there is no congestion, i.e. the grid capacity is not overloaded.

II. PROBLEM STATEMENT

The problem at hand is a variation of the economic power dispatch problem [3], [4], where the operation of a set of generators is optimized, such that power is provided to consumers in the most cost-effective manner. Traditionally this problem would include controllable generators (power generation plants), constraining transmission resources (transformers, stations, cables), and end-consumers having a static load. Currently, with distributed energy resources, and demand response—the possibility to control the load of consumers using flexibility of smart devices—the problem changes significantly.

A formal definition of the problem is as follows:

$$\min_{\mathbf{x}} \quad \sum_{i \in \mathcal{N}} J_i(\mathbf{x}_i), \tag{1a}$$

$$\text{s.t.} \quad \sum_{i \in \mathcal{N}} \mathbf{x}_i = \Theta, \tag{1b}$$

$$\mathbf{A}\mathbf{x} \leq \mathbf{b}, \tag{1c}$$

$$x_i^{\min} \leq \mathbf{x}_i \leq x_i^{\max}, \quad \forall i \in \mathcal{N}, \tag{1d}$$

$$\mathcal{C}_i(\mathbf{x}_i) = 0, \quad \forall i \in \mathcal{N}. \tag{1e}$$

With all parameters defined in Table I, the objective function as defined in (1) is the same as for the traditional economic dispatch, which can be summarized as: to find a set of electrical powers $\mathbf{x}_i$ for every device $i \in \mathcal{N}$ in the grid, that minimizes the sum of all costs $J_i(\mathbf{x}_i)$. In the original dispatch problem defined, $\mathcal{N}$ defines only the producers. However, in our problem formulation (1) $\mathcal{N}$ denotes the full set of devices in the grid, including *producers and consumers*. This means that contrarily to the traditional economic dispatch problem, not only the generators are controllable, but also the flexible loads of end consumers.

Note that the costs in (1a) do not necessarily refer to the financial cost of a dispatch. Rather, this generic model only optimizes some *social welfare*, which defines a desirable outcome for all participants involved. Depending on the situation at hand, one could minimize the amount of greenhouse gasses emitted, or maximize the amount of renewable energy used. However, for the rest of this paper the costs are defined as the energy losses of the devices. By minimizing the amount of energy losses we attempt to find a dispatch that is both economic and sustainable.

TABLE I
USED NOTATIONS AND PARAMETERS IN THIS PAPER.

Symbol	Description	Typical value	Unit
$\mathbf{A}$	topology matrix of child relations		
α	forecast accuracy	0.9	
$\mathbf{b}$	vector of congestion thresholds	$[2, 3, \ldots, 3.5] \times 10^4$	W
β_i	congestion threshold of i	3×10^4	W
$\mathcal{C}_i$	constraint function of i		
$\mathcal{C}_i^{Z}$	constraint function of i of type Z*		
γ_i	PV operation cost of i	0.2	
$\mathbf{e}_i$	$1 \times T$ vector of energy of i	$[0.0, 1.1, \ldots, -0.3] \times 10^3$	Wh
e_{it}	energy of i at t	200	Wh
$e_i^{\min}$	minimum energy of i	0.0	Wh
$e_i^{\max}$	maximum energy of i	2×10^3	Wh
ϵ	$1 \times T$ vector of errors	$[0, 11, \ldots, -225]$	W
$\epsilon^{\max}$	upper bound on the error	10^{-3}	W
η_i	storage efficiency of i	0.9	
Θ	target power of the cluster	$[1.9, -1, \ldots, 2.6] \times 10^4$	W
i	an agent index	$0, 1, \ldots, n$	
J_i	cost function of i		
J_i^{Z}	cost of i of type Z*		
λ_i	storage leakage of i	360	W
$\mathcal{M}_i$	set of children of i		
$\mathcal{N}$	set of all agents		
ρ	$1 \times T$ vector of prices	$[0.5, 0.3, \ldots, 0.8]$	
T	time horizon	24	
t	time slot index	$0, 1, \ldots, T$	
τ	duration of a time slot	1	h
$\mathbf{x}_i$	$1 \times T$ vector of scheduled powers of i	$[1.1, -0.5, \ldots, 0.4] \times 10^3$	W
$\bar{\mathbf{x}}$	$1 \times T$ vector of aggregated powers of child nodes	$[24.0, -18.6, \ldots, 8.7] \times 10^3$	W
$\mathbf{x}_i'$	$1 \times T$ vector of expected powers of i	$[1.0, -0.4, \ldots, 0.8] \times 10^3$	W
$\hat{\mathbf{x}}$	$1 \times T$ vector of average historic powers	$[0.8, 0.6, \ldots, -0.5] \times 10^3$	W
x_{it}	scheduled power of i at t	-1.7×10^3	W
$x_i^{\min}$	minimum power for i	-2×10^3	W
$x_i^{\max}$	maximum power for i	4×10^3	W
ψ_i	heat pump coefficient of performance (COP) of i	4.0	

Note that powers are indicated from the grid perspective. That means positive powers indicate power going from the grid to the device, and negative powers are from the device back to the grid.

*Z indicates a type of agent, which can either be MO for market operator, CO for congestion, LOAD for a static load agent, PV for an agent with photovoltaic solar panels, or ST for a storage agent.

Let us discuss the constraints of (1), i.e. (1b–1e). The constraint (1b) states that the sum of all powers in the system has to meet a specific target Θ. In a balanced (island) grid this target has to be zero for every time slot, but in a connected grid this target is the contracted load with the transmission operator. Put differently, Θ is the net power input of the grid to the rest of the world.

Constraint (1c) defines the limitations of the physical power grid; a topology matrix A specifies which device is connected to which grid components, and b is the rating of those components—the maximum amount of power that it can safely transmit.

Constraints (1d) and (1e) represent the constraints of the devices in the grid. Specifically, constraint (1d) states that a device cannot produce or consume more power than it physically can, which is represented by lower and upper bound, $x_i^{\min}$ and $x_i^{\max}$, respectively. However, (1e) also takes into account another *private* constraint $\mathcal{C}$. We refer to this constraint as private, as it only concerns the state of a single device, and its state may concern information that we should not require to share with other parties. E.g. an end-user should not need to share its intent to run his or her washing machine with the rest of the neighborhood. This constraint differs per

device, and specifies device limitations such as a battery that cannot hold more than a certain amount of energy, and cannot discharge when already empty. Flexible loads also may have constraints considering the time at which they can turn on or off based on the consumer's settings. We will elaborate on these constraints, as well as the cost functions of the device agents in Section III-B.

A. Related Work

In existing studies, different strategies are used for energy management. Four different main categories are defined in [1] based on whether there is a distributed aspect of decision-making and whether there is one- or two-way communication. One of the strategies defined in [1] uses *Transactive Control*, where distributed systems decide locally on their device management, using two-way communication in a market-based control scheme. The authors compare this approach with traditional top-down switching, price-reaction and centralized optimization strategies and show that transactive control is capable of using the full flexibility potential of the smart grid devices, while maintaining the end-user privacy. This claim is consistent with earlier studies [5], [6] that have shown that using a market-based control in a multi-agent system

can provide equally optimal results as a centrally optimized system, under certain conditions.

Multi-agent based methods are already getting attention in the smart grid domain due to many desirable properties: robustness, user-friendliness, attack resistance and scalability [7], [8]. The economic dispatch problem is very well suited to be represented using multi-agent systems [9]. There are many studies that use a multi-agent based approach to model and solve the problem, such as the two-way message passing agents using consensus algorithms to find an allocation that is optimal [10]. Other methods use a completely decentralized method involving reinforcement [11], utility maximization [12] or model predictive control [13].

In [14], a strategy is proposed to schedule the power consumption and production of generators and loads based on a method called *Negotiated Predictive Dispatch*. In this approach wind and conventional generators, as well as static and flexible loads, are controlled on the transmission level. Agents propose power schedules, which are aggregated by a central market operator, which then updates a price program using gradient descent, meaning the price of congested timeslots increase, and that of underused time slots decrease. In doing so the authors show that they are able to balance power production and consumption, while satisfying power grid constraints. However, this approach focuses on the transmission level, whereas we consider the distribution level power grid to be at a much more imminent risk of congestion. At the transmission level, the scale is much larger than at the distribution level, both geographically and considering the power levels involved. Moreover, the power grid has a very different topology at the transmission level. A major drawback of this approach when applied at the distribution level, is that using a global power price for congestion mitigation is not only unfair for agents in non-congested areas, but it would also be converging to a solution much too slowly. A similar approach to the negotiated dispatch is provided by [15] in which a gradient descent on the pricing is used, followed by a local optimization of agents.

Another approach is given by [16], in which power programs are proposed by device agents that are able to satisfy cluster constraints to reach a target energy consumption, while also staying within the limits of the grid. In their approach, called *Profile Steering*, agents are only motivated to accomplish the cluster goal, which is to consume or produce energy at a specific target amount, and will propose alternative power programs whenever constraints are violated. Proposals that reduce the constraint violations the most, are then selected as the new candidate programs. In our view, the problem of this approach lies in the lack of motivation for the participant to sacrifice private rewards for running an alternative program. The only objective is the cluster goal, which means that the price of having limited resources is paid by a select set of individuals that offer the most flexibility, or conversely, the profit is reaped by those who are able to maximally make use of remaining capacity.

The authors of [17] propose an approach based on DCOPs (Distributed Constraint Optimization Problem) to solve the economic dispatch problem, as well as the real-time demand-response balancing problem. Their algorithm is able to find optimal solutions for a system of controllable generators and (predicted) loads, taking into account transmission network constraints. They show that finding the optimal solution using Dynamic DCOPs is possible, but their solution does not scale very well. A relaxed version of the problem, in which soft constraints may be (temporarily) violated, scales better, but still for relatively small time horizons, even considering their implemented solution on a GPU.

We note that related work on the topic of this chapter listed here is not exhaustive. For further detailed background, a comprehensive overview of different mechanisms for solving optimization problems in smart grids, particularly where demand response is involved, we refer to [18], [19], [20].

III. Self-Organizing Economic Dispatch

Our contribution is a self-organizing method for coordinating the power scheduling of smart grid devices on the distribution level energy grid. This is a multi-agent based heuristic approach for find a solution to the economic dispatch problem as defined in (1). This multi-agent based approach allows for great scalability, while ensuring privacy and final control of the end-user.

At the distribution grid level, devices are typically consumer devices, such as PV-panels, household batteries, heat pumps, ventilation or air-conditioning units. The controllability, or flexibility of such devices is often limited, and bound by the device limitations and the user preferences. With increasing numbers of such devices, using a strictly "top-down" approach is intractable, since the search space grows exponentially with each added device. For this reason we use a hierarchical approach, in which the market operator delegates its task to intermediate "congestion" agents which then independently solve subproblems—which is to make sure that the total power throughput at their allocated point does not exceed a certain threshold. Making such subdivisions is justified, since in the power grid, transformers are effectively branches of the topology, and two nodes under different transformers are independent of one another; hence, transformers are the logical point to place congestion agents in the control hierarchy.

Our approach, denoted as Local Pricing Receding Horizon (LP-RH) is based on the economic incentive that end-user devices should have, to provide its owner with a service, in the most affordable way. Put differently, the system will use price differences to stimulate agents to schedule their power consumption or production in a balanced way, such that any grid constraints are satisfied. The overall scheme is simply to gather expected power programs from the connected devices, and then iteratively adjust energy prices to steer the agents into a certain power program. The method is explained in more detail in the following section, and is similar to the Negotiated Price Dispatch proposed by [14].

A. Local Pricing Receding Horizon

In order to create a power planning taking into account forecasts and/or predicted power programs of clients, we use a Receding Horizon (RH) approach, depicted in Fig. 1 which means that at any point we only take a fixed horizon of T
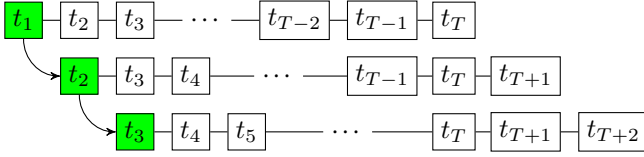
Fig. 1. The principle of a receding horizon market is that every time step the time horizon for optimization shifts by one. Planned power programs (white) are turned into fixed contracts for the current time step (green), which determines the outcome of the algorithm. This figure is redrawn from [14].

Algorithm 1 LP-RH Algorithm

$\quad$ $\mathbf{x}_i = \text{createPowerprogram}(i, \rho)$
1: **if** i is a Device agent **then**
2: $\quad$ $\mathbf{x}_i = \arg\min_{\mathbf{x}} C_i(\mathbf{x}, \rho)$ {Run local optimization}
3: **else**
4: $\quad$ **repeat**
5: $\quad\quad$ **for all** $j \in \mathcal{M}_i$ **do**
6: $\quad\quad\quad$ $\mathbf{x}_j = \text{createPowerprogram}(j, \rho)$
7: $\quad\quad$ **end for**
8: $\quad\quad$ $\mathbf{x}_i = \sum_{j \in \mathcal{M}_i} \mathbf{x}_j$
9: $\quad\quad$ $\epsilon = C_i(\mathbf{x}_i)$
10: $\quad\quad$ $\rho = \text{adjustPrices}(\epsilon, \rho)$ {Using gradient descent}
11: $\quad$ **until** $\epsilon \leq \epsilon^{\max}$
12: **end if**
13: **return** $\mathbf{x}_i$

---

program time units (PTU) into account. A PTU is typically 15 minutes or one hour, throughout this paper we will use PTU length of one hour so $T = 24$. However, none of the mentioned approaches are limited to this convention. When the algorithm has converged and a power program is found for the next T PTUs that satisfies all constraints, the first PTU becomes the "current" situation, and the projected power program becomes a contract. The time horizon now shifts by one, and the entire system starts again.

The Local Pricing Receding Horizon algorithm is shown as pseudocode in Algorithm 1; an example graph on which the algorithm could run is shown in Fig. 2. The algorithm describes how any agent finds a new power program $\mathbf{x}_i$ to satisfy its local constraint C_i. Every device agent in the grid does this by solving a local optimization problem (line 2), taking into account local constraints as explained in Section III-B.

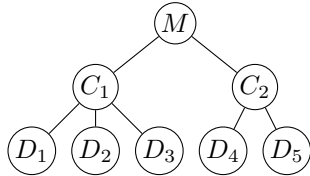Prices ρ are updated by the market operator agent and

Fig. 2. Simple example topology of a graph which could be running the LP-RH algorithm. One market operator (M) is connected to two congestion agents (C_i), which are connected to a total of five device agents (D_i). The device agents can represent PV panels, consumer loads, storage agents, or any other leaf nodes in the power grid. Topologies where congestion agents are connected to more congestion agents are also possible.

the congestion agents in order to get a power program $\mathbf{x}$ satisfying the grid constraints. When the market operator or any congestion agent runs the LP-RH algorithm, all $j \in \mathcal{M}_i$, where $\mathcal{M}_i$ are the immediate children of agent i, are requested to propose a power program based on the price ρ (line 6). Here the function is called recursively until all agents have determined a new power program. Note, that if any of the children are congestion agents, they determine their power program by forwarding the prices to *their* children and returning the sum of all received programs. Many agents will not know exactly how they will function in the future, as contexts might change, a user might behave differently than anticipated, or weather conditions may act up; this is why receding horizon method updates the time iteratively. With the sum of all power programs, the market operator can determine the error ϵ, which is the sum of all local constraint violations C for every $t \in T$.

In line 9 the constraint C_i is used to compute the local constraint violation and stored as an error variable ϵ. The value of ϵ is taken into account to adjust the prices in line 10. In the "adjustPrices" function, a gradient descent approach is used, which linearly interpolates the last two errors as a function of the price. We then choose the price at which the error is projected to reach zero. Note that this means we assume the reaction of the devices linearly depends on the prices, which will only hold under very specific circumstances. To overcome this issue, we iteratively repeat the process until $\epsilon \leq \epsilon^{\max}$ and assume that a non-linear response can be described as a series of linear pieces. The parameter $\epsilon^{\max}$ then represents an upper bound on the error, which we can use as a convergence criterion.

B. Agent Behavior

Agents in the system are characterized by the device that they have to assign a power program for. In the distributed case, every agent locally optimizes a local cost function J_i, which is part of the global optimization problem (1). Also, local constraints C_i have to be taken into account which are represented by constraint (1e). The behavior of the agents can be defined by three functions for the costs, the local constraint and the power program; an overview of this is shown in Table II. In our model we consider the following types of agents:

1) Market Operator: This is the root node of the tree (M in Fig. 4), as far as the local distribution grid concerns. In the physical grid, it corresponds to the transformer that connects the local low voltage (LV) grid to the medium voltage (MV) grid. We assume that there is no energy loss at the market operator, so

$$J^{\text{MO}} = 0. \tag{2}$$

Its goal is to find a solution to (1), and its private cost would be the same as the target constraint in (1b). The market operator runs Algorithm 1, using a deviation of the target profile Θ as an error, which is minimized by the algorithm. Hence, its local constraint is defined as

$$C^{\text{MO}} = \bar{\mathbf{x}} - \Theta, \tag{3}$$

Agent type	Cost	Constraint	Power
Market Operator	$J^{\mathrm{MO}} = 0$	$\mathcal{C}^{\mathrm{MO}} = \bar{\mathbf{x}} - \Theta$	$\bar{\mathbf{x}} = \sum_{i \in \mathcal{N}} \mathbf{x}_i$
Congestion agent	$J_i^{\mathrm{CO}} = 0$	$\mathcal{C}^{\mathrm{CO}} = \bar{\mathbf{x}}_i - \beta_i$ iff $|\bar{\mathbf{x}}_i| < \beta_i$	$\bar{\mathbf{x}}_i = \sum_{j \in \mathcal{M}_i} \mathbf{x}_j$
Load agent	$J_i^{\mathrm{LOAD}} = 0$	$\mathcal{C}_i^{\mathrm{LOAD}} = 0$	$\mathbf{x}_i$
PV agent	$J_i^{\mathrm{PV}} = \mathbf{x}_i - \mathbf{x}'_i$	$\mathcal{C}_i^{\mathrm{PV}} = 0$ iff $\tau \mathbf{x}_i \rho \geq \gamma_i$	$\mathbf{x}_i = \mathbf{x}'_i$ or 0
Storage agent	$J_i^{\mathrm{ST}} = \mathbf{x}_i(1 - \eta'_i)$	$\mathcal{C}_i^{\mathrm{ST}} = 0$ iff $e_i^{\min} < \mathbf{e}_i < e_i^{\max}$	$\mathbf{x}_i = f(\rho)$

where $\bar{\mathbf{x}}$ denotes the power program of the market operator. Since the market operator is not a device in the grid itself, its power is the sum of the powers of its children, which in case of the market operator are all nodes in the cluster:

$$\bar{\mathbf{x}} = \sum_{i \in \mathcal{N}} \mathbf{x}_i. \qquad (4)$$

The value of (3) can either be negative or positive, which respectively means either the total power production or the total consumption is too high. The market operator sets an initial price of $\rho = 0.5$ for all timeslots, and then uses Algorithm 1 to minimize the constraint value until it reaches zero, in order to satisfy the global constraint (1b).

2) Congestion Agent: This is an intermediate node on the grid tree (C in Fig. 4) connected to a parent node who is either the market operator or another congestion agent. It corresponds to a component in the grid where congestion might potentially occur, such as a transformer. Equivalently to the market operator, we assume that no energy is lost here, so:

$$J_i^{\mathrm{CO}} = 0. \qquad (5)$$

This agent has a constraint that aims to limit the power usage of that part of the grid, this corresponds to the constraint (1c). The congestion agent also uses Algorithm 1 to minimize the error of its local constraint, which is defined as

$$\mathcal{C}_i^{\mathrm{CO}} = \begin{cases} \bar{\mathbf{x}}_i - \beta_i, & \text{if } |\bar{\mathbf{x}}_i| \geq \beta_i, \\ 0, & \text{otherwise.} \end{cases} \qquad (6)$$

Here β_i is the power limit of the agent i, which means it is the maximum power throughput of the grid at the point i represents. Again $\bar{\mathbf{x}}_i$ denotes the power profile of the congestion agent, but is now equal to the sum of all devices under the current node:

$$\bar{\mathbf{x}}_i = \sum_{j \in \mathcal{M}_i} \mathbf{x}_j, \qquad (7)$$

Similar to the constraint of the market operator in (3) the constraint value can become negative or positive, and is used to compute the error ϵ in Algorithm 1.

3) Load Agent: This is an agent responsible for an uncontrollable load in the grid (represented by any device agent D in Fig. 4). This could be a consumer household, an office, street lighting, or any other non-flexible load, and thus only participates in the problem as part of constraint (1b). In the distributed system however, it is responsible for making a forecast of the power usage, and this forecast will be updated with more accurate information as the time horizon shifts.

The load agent has no attached cost in the global problem, and no need to locally compute any optimal behavior. This is equivalent to stating that its cost and constraint correspond to

$$J_i^{\mathrm{LOAD}} = 0, \qquad (8)$$

$$\mathcal{C}_i^{\mathrm{LOAD}} = 0. \qquad (9)$$

Its corresponding load profile $\mathbf{x}_i$ is fixed to some profile that constrains the global problem (1). In our experiments its values are taken from real households as described in Section IV-B.

4) PV Agent: The PV agent (any D in Fig. 4) has some flexibility to offer to the optimization function by allowing curtailment in reference to the expected generation. We assume that curtailment is binary, in that either the PV generates power as normal, or it is switched off and produces no power at all. Curtailing means that there is potential energy lost, and hence the cost function of a PV agent is defined as

$$J_i^{\mathrm{PV}} = \mathbf{x}_i - \mathbf{x}'_i, \qquad (10)$$

where $\mathbf{x}'_i$ indicates the *expected* power, when not curtailing. This expected power is taken from the scenario, which will be detailed in the Section IV.

When reacting to prices in the distributed system, a decision is made in order to decide whether to curtail based on the price profile. If the operational running costs γ_i of the PV is *more* than the power that would be generated by it, there is no point in running the generator (from an economic point of view). Hence, the local constraint and its corresponding decision rule of a PV agent can be written as

$$\mathcal{C}_i^{\mathrm{PV}} = \begin{cases} 1 & \text{if } \tau \mathbf{x}_i \rho < \gamma_i, \\ 0 & \text{otherwise,} \end{cases} \qquad (11)$$

$$\mathbf{x}_i = \begin{cases} 0 & \text{if } \tau \mathbf{x}'_i \rho < \gamma_i, \\ \mathbf{x}'_i & \text{otherwise.} \end{cases} \qquad (12)$$

For the PV agent it also holds that the corresponding expected power profile $\mathbf{x}'_i$ determines the global problem (1). Its values in our experiments are taken from real PV panels as described in Section IV-B.

5) Storage Agent: A storage agent (again a leaf node D in Fig. 4) provides flexibility by allowing to store some energy in a local storage like a battery or a heat buffer. There are limits to the amount of energy that can be stored, either because of the physical limitations of the storage device, or because of the end-user settings. Moreover, a storage agent has some efficiency, which defines energy loss when energy is put into it, or out from it. This means that we can define the cost

function of the storage agent as the energy lost during charging or discharging

$$J_i^{\text{ST}} = \mathbf{x}_i(1 - \eta_i'), \tag{13}$$

where

$$\eta_i' = \begin{cases} \eta_i & \text{if } \mathbf{x}_i \geq 0, \\ \eta_i^{-1} & \text{otherwise.} \end{cases} \tag{14}$$

This difference makes sure that the loss is correlated to the internal power of the battery when charging or discharging. I.e. $\mathbf{x}_i$ defines the power at the grid side of the storage, when charging a lower power effectively charges the battery, and conversely when discharging a higher power is required to provide some power level to the grid.

In order to define the constraints of the storage agent we must define the update function for the amount of energy $\mathbf{e}_i$ stored as the cumulative sum of the powers

$$e_{it} = e_{i0} + \tau \sum_{t'=t_0}^{t'=t} \eta_i' \mathbf{x}_{it'} - \lambda_i. \tag{15}$$

Here λ_i represents the leakage or self-discharge rate of the storage agent i, and e_{i0} is the stored energy at the start of the experiment. Let us now define the following constraints on the storage agent

$$\mathcal{C}_i^{\text{ST}} = \begin{cases} 0, & \text{if } e_i^{\min} < \mathbf{e}_i < e_i^{\max}, \\ 1, & \text{otherwise.} \end{cases} \tag{16}$$

The minimum and maximum energy values are defined by $e_i^{\min}$ and $e_i^{\max}$, respectively.

For a storage device, the power limits denote the maximum charge and discharge rates. They are defined by $x_i^{\max}$ and $x_i^{\min}$, respectively, in (1d). A special case of a storage device is a heat pump, which (electrically powered) heats a house in order to keep the temperature within comfortable levels. The heat pump only allows charging and only discharges through leakage, hence for a heat pump $x_i^{\min} = 0$ and $\lambda_i > 0$.

When a storage agent has to update its expected power profile in line 2 of Algorithm 1, some response function ($f(\rho)$ in Table II) is required, which is "economically sane" and has the following characteristics:

- return $x_i^{\max}$ for low prices and $x_i^{\min}$ for high prices,
- be a monotonically decreasing for increasing prices,
- have a "plateau" of zero power response for some intermediate price ($\rho = 0.5$), which is wider for less efficient devices.

The final characteristic allows more efficient devices to respond to subtle price change, and have less efficient devices respond to more extreme prices. This way devices with a higher efficiency are used first when flexibility is needed, and (when parameterized correctly) less efficient devices will only be used when required. In our implementation we chose the simplest response, where the agent linearly decreases its power from $x_i^{\max}$ to zero at $\rho = \eta_i/2$. It then stays at zero symmetrically around $\rho = 0.5$ until $\rho = 0.5/\eta_i$, and then decreases its power response linearly to $x_i^{\min}$. This power/price relation is depicted in Fig. 3.
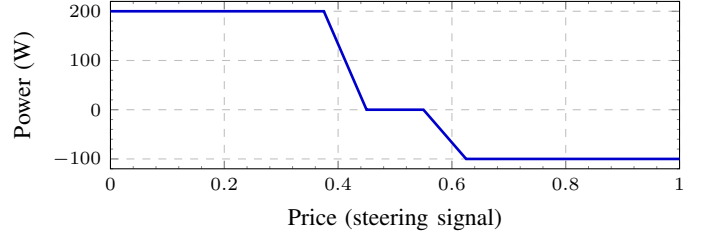

Fig. 3. The strategy of the storage agent with $x_i^{\max} = 200\,\text{W}$, $x_i^{\min} = -100\,\text{W}$ and $\eta_i = 0.9$ shows the power response for an increasing price. The *plateau* at $0\,\text{W}$ extends from $\rho = \eta_i/2$ to $\rho = 0.5/\eta_i$.

An alternative strategy for the storage agent would be to make use out of any differences in the price, charging and discharging as soon as the price differences are large enough to overcome its efficiency. From a strictly economic perspective, this is the optimal strategy to maximize its own benefit. However, this leads to very "binary" behavior with minimal and maximal charging rates [12] and thus, little room for optimizing from the market operator and congestion agent. Therefore, a linear strategy is implemented as depicted in Fig. 3, allowing to solve the overall optimization problem (1).

IV. EXPERIMENT SETUP

The LP-RH algorithm was empirically evaluated by running a simulation of an LV grid with a set of realistic household load profiles and PV production profiles for a series of 24 PTUs.

A. Distribution Network Topology

For the topology of the network we use the European Low Voltage Test Feeder [21] network. This dataset is used to benchmark power and energy algorithms on realistic European distribution networks. In this paper we superimposed six points on the topology, where we monitor and mitigate any potential congestion. These points are strategically chosen to separate the problem into independent subproblems. The resulting topology with the congestion points are shown in Fig. 4.

B. Household Load and PV Profiles

The household consumption and production profiles are taken from a pilot study [22], in which 92 residential consumers were monitored over the course of a year (from March through November 2018). The data was preprocessed such that we have separated information on the consumption of houses, and of the PV installations. Data is anonymized and randomized per month, so that we can select data from any specific month for a base load of a household, or a residential PV installation.

In the experiment the 54 households from Fig. 4 were assigned a random instance of the load *and* PV profiles from the same month (i.e. all households were equipped with PV panels). The daily load consumption varied between 4.98 kWh and 29.39 kWh, and the total PV production varied between 826 Wh and 18.8 kWh. Furthermore, 16 randomly selected
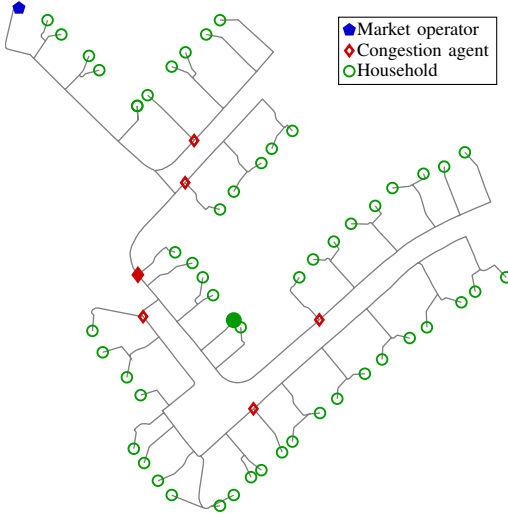
Fig. 4. The topology of the IEEE LV feeder network used for the simulation. The figure depicts the connections between household consumers, the intermediate congestion agents, and the root market operator. The "filled" markers represent the agents from Section VI.

households were assigned a household battery, and again 16 were chosen to have a heat pump installation.

Every household and PV installation in the simulation would select a random profile from the dataset, which was used as its power profile. The objective Θ was set to a total net consumption of the LV grid using the mean total power consumption of the houses including PV production. Choosing the target profile Θ in this way corresponds to a situation in which the energy provider of the simulated neighborhood would agree to a contract for the average behavior of the households, and subsequently attempts to use flexibility to account for any deviations from the normal.

The batteries were all dimensioned with storage capacities of $e^{\max} = 10.8\,\text{kWh}$, and maximum charge and discharge rates of $x^{\max} = 4000\,\text{W}$ and $x^{\min} = -4000\,\text{W}$, respectively. The batteries charging efficiencies were all set at $\eta = 0.9$. The heat pumps were estimated to have a working energy capacity of $e^{\max} = 2\,\text{kWh}$, this means the difference between the thermal energy capacity of the house at the minimum and maximum comfortable user temperature is $2\psi\,\text{kWh}$, where $\psi = 4.0$ is the coefficient of performance of the heat pump. Then, the heat pumps have a $x^{\max} = 1600\,\text{W}$ and $x^{\min} = 0\,\text{W}$, an efficiency of $\eta = 1$ and a constant leakage rate of $\lambda = 360\,\text{W}$. Finally, to ensure convergence, the maximal error value is set to $\epsilon^{\max} = 10^{-3}$ in the experiments for this paper

C. Forecast Uncertainty

The predicted load and production power profiles $\hat{x}$ of the load and PV agents were generated by taking average profiles of the complete dataset. These average profiles are considered as taken from historic data and hence, provide a ground for predicting the power program of future PTUs. When agent i determines its power prediction $\mathbf{x}'_i$, it will compute a weighted average between its assigned power profile $\mathbf{x}_i$ and the average

profile $\hat{\mathbf{x}}$, such that:

$$\mathbf{x}_{it} = (1-\alpha)\mathbf{x}'_{it} + \alpha\hat{\mathbf{x}}_t, \tag{17}$$

$$\alpha = \sqrt{\frac{t-1}{T-1}}, \tag{18}$$

such that at $t = 1$ the prediction equals the selected profile $\mathbf{x}'_{it}$. At $t = T$, the prediction is simply the average power $\hat{\mathbf{x}}_t$.

V. CENTRALIZED SOLVER

In order to address the performance of the LP-RH algorithm, a centralized optimization approach is also introduced to provide lower bounds to its results. A mixed integer linear program (MILP) solver was used to find these bounds for the problem stated in (1). The centralized optimization approach considers the exact same scenario that was solved by the decentralized algorithm; the energy loss is minimized and the same set of constraints apply. A fundamental difference lies in the availability of information. Whereas detailed information is only shared locally in the LP-RH algorithm, the centralized optimization approach assumes complete knowledge of the current system state for the decision-maker; i.e., no limits are imposed on the spacial flow of information within the network. With these features in mind, two versions of the MILP were formulated: Receding Horizon Centralized Solver and Perfect Information Centralized Solver.

A. Receding Horizon Centralized Solver

The receding horizon centralized solver (RHCS) is the most similar to the LP-RH algorithm. It uses a receding horizon approach (as depicted in Fig. 1) to find an ideal dispatch solving the consecutive sub-problems. Uncertainty about future device states is again simulated by (17). Because this solver has perfect information at the *current* state for each iteration of the receding horizon, its solution represents a lower bound on the solution found by the LP-RH algorithm.

B. Perfect Information Centralized Solver

In the perfect information centralized solver (PICS), a single centralized optimization problem is solved for the complete dispatch. In addition to perfect spacial information, this version of the MILP also has perfect *temporal* information about the complete system state; this means that no uncertainty about future states is simulated. Referring back to (17), this is equivalent to setting $\alpha = 1$, which results in perfect predictions for each profile. The solution of this MILP represents an absolute lower bound for problem (1).

VI. RESULTS

In this section the results of the described experiments are shown. Before comparing the performance of the different solvers, which is done in Section VI-A, the behavior of the LP-RH algorithm is shown in this section. The graphs in this section show examples of single simulation runs, demonstrating the behavior of the different agent types. Powers are shown as power consumption, this means a net consumption is shown as positive power, and conversely negative powers indicates a
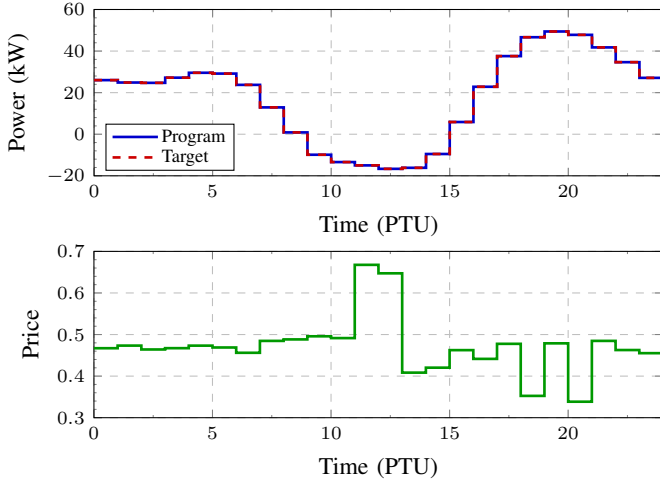
Fig. 5. The power program of the market operator and the price profile as the outcome of Algorithm 1, show the results of a 24-hour simulation of the problem with the LV feeder network and 92 random households.
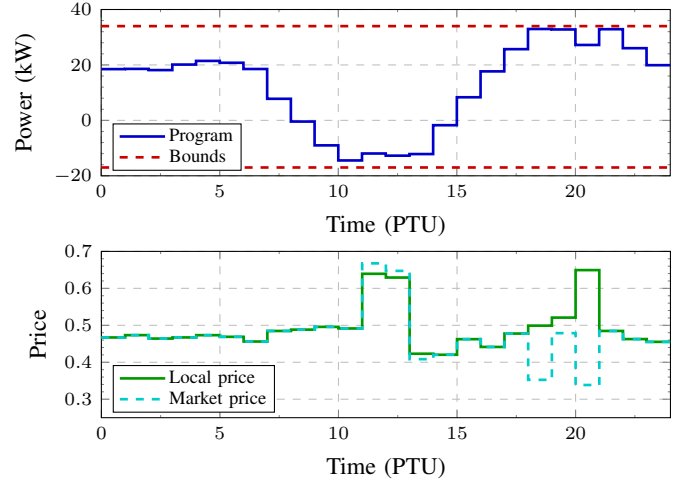

Fig. 6. The power program of a congestion agent shows some periods of congestion at the production and consumption bounds (red dashed lines), and the corresponding changes in price profiles (from Algorithm 1) relative to the global price of the market operator.

net power production. In Fig. 5 the final result of the market operator is shown, where the market operator has found a price profile such that the target profile is met exactly.

Fig. 6 shows the power profile of one of the three congestion agents that is directly connected to the market operator. Its power congestion limits are set such that in the peak moments of the day there is some congestion expected. This results in a price difference shown in the power program, as around the peak PV production ($t = 12, 13, 14$) the local price is slightly lower than the market price, leading to a lower net power production. Similarly, at the end of the day ($t = 19, 20, 21, 22$) a positive power congestion was mitigated by increasing the price.

The power program of a heat pump agent is shown in Fig. 7. This heat pump is connected directly to the congestion agent from Fig. 6, hence its price profile should be identical. What is most obvious in this graph is that the high price at $t = 12$ leads to a zero power consumption, and quickly after that, the power consumption rises in order to maintain comfortable temperatures. Again, around $t = 19, 20, 21$ the relatively high prices lead to a power consumption of zero, which was apparently feasible because of the high power consumption leading up to it.

Finally, Fig. 8 shows the total power programs of all devices of the four classes in this experiment summed up. In this figure, the power profiles of the loads and the PVs are the direct result of the chosen profiles, and are the input for problem (1). We can see that for the majority of the experiment, all used flexibility is from the heat pumps with the high efficiency. Only at times of the congestion will the less efficient battery agents be used.

A. Comparison with Central Solvers

In a validation experiment 105 random problem instances (15 permutations for seven months) were created, and solved by the three different algorithms. In Fig. 9 the results are shown for the *feasible* instances. A solution is considered
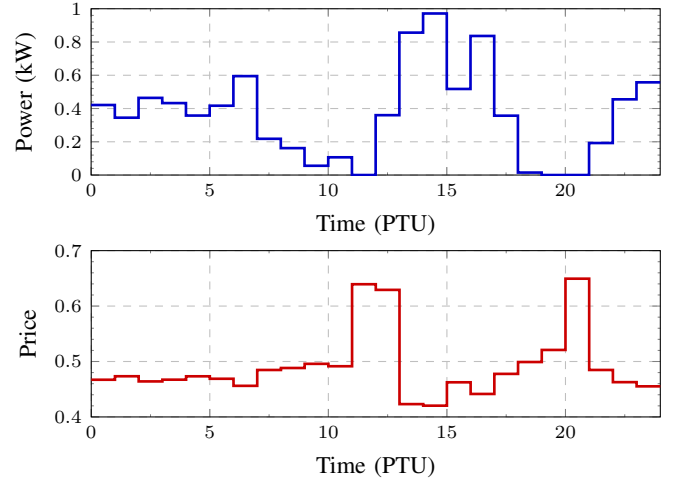

Fig. 7. The power program of a heat pump agent shows the power being mostly used at moments where the price is low, relieving the need to charge when the price is high; or considering the view of the grid, when a lower power consumption is required.

feasible if all three solvers were able to find a solution that satisfies all constraints. The LP-RH algorithm did not find a correct solution for 26 problem instances, 11 of which were found to be overconstrained according to RHCS. For the feasible instances the LP-RH algorithm found solutions that were not significantly worse than the PICS, and in 21 instances found a solution with the exact same cost. In 27 instances LP-RH found a solution that was equally well as the RHCS or even slightly better—this seems to contradict the initial statement that RHCS acts as a lower bound for LP-RH; however, this is due to the way both solvers deal with uncertainty. The LP-RH algorithm allows the congestion agents to violate power constraints for future PTUs, as long as eventually they are not actually congested due to imperfect forecasts according to (17). The RHCS algorithm does not allow power constraints to be violated at any point in the future,
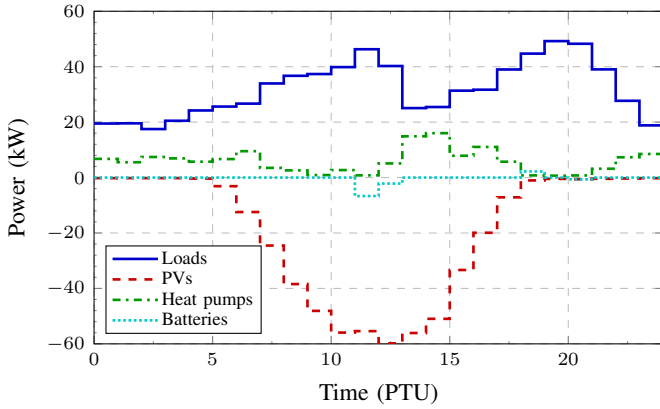
Fig. 8. The total power consumption per device type show that the batteries are used far less than the heat pumps, since they have a lower efficiency. The PV panels did not have to be curtailed in this run.
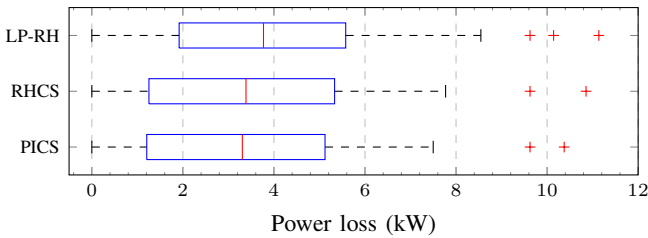

Fig. 9. Compared with centralized optimization solvers, the LP-RH algorithm performs equally well. In this box plot the median is shown as a line, the boxes indicate the 25% and 75% percentiles, and the tails are capped at a maximum length of the box width; outliers are drawn separately.

and hence might react too strict when a future congestion is mistakenly predicted.

In the set of *infeasible* problems, i.e. problems that do not have a valid solution satisfying all constraints, LP-RH does find solutions that minimize the power loss. However, since these solutions either temporarily overload congestion agents (where $\mathbf{x}_a gen > \beta_i$), or does not match the target profile Θ exactly, they are not a fair comparison, since they do not strictly solve (1). In a separate run the problems were relaxed, by increasing the power rating of the congestion agents. This resulted in the LP-RH not being able to find a feasible solution in only 14 problem instances, but the results were otherwise very similar to the ones reported here in Fig. 9.

VII. CONCLUSIONS

We have introduced an algorithm for self-organizing smart grids, by solving the economic dispatch problem using a decentralized market based approach with local pricing. The LP-RH algorithm using a hierarchical approach was shown to be able to solve the problem using a fairly simple interaction scheme in which pricing information is sent down the hierarchy tree, and planned or forecasted power programs are sent back up. Using a gradient descent approach, the market operator is capable of tuning the pricing to find feasible solutions to minimize the power losses in the grid.

In our experiments we found that in 20% of the problems LP-RH did not perform any worse than a perfect-information centralized solver. In the other 80% our algorithm did not perform significantly worse. The benefit of LP-RH over a centralized solver are in the robustness and scalability of the solution, as well as in the preserved privacy of the end-consumers.

In the implementation of the response of the storage agent, we intentionally did not choose to respond with an optimal power program given the price signal. Particularly, when a high price is expected in the future, the agent will not "proactively" charge to avoid having to charge later, or vice versa. This behavior could be implemented at the agent quite easily using a dynamic programming approach, but it would lead to very extreme behavior, e.g. very binary behavior of charging or not-charging at full capacity even for small price differences. This binary behavior is hard to deal with in the rest of the hierarchical tree, and does not lead to any problems per se, but might be improved upon in a future continuation of this work.

Other variations of the problem may include other device types, for instance time-shiftable devices such as washing machines or dishwashers. Also, using electric vehicles (EV) as an additional type of agent, providing energy flexibility is a very interesting extension, which will undoubtedly lead to other complications because of their high power ratings. Finally, an integration with a real time balancing algorithm such as [23] would be very fruitful to complete the needs of the future smart grid.

REFERENCES

[1] K. Kok and S. Widergren, "A society of devices: Integrating intelligent distributed resources with transactive energy," *IEEE Power Energy Mag.*, vol. 14, no. 3, pp. 34–45, 2016.

[2] Q. Gemine, D. Ernst, and B. Cornélusse, "Active network management for electrical distribution systems: Problem formulation, benchmark, and approximate solution," *Optimization and Engineering*, vol. 18, no. 3, pp. 587–629, Sep. 2017.

[3] D. W. Ross and S. Kim, "Dynamic economic dispatch of generation," *IEEE Trans. Power App. Syst.*, vol. PAS-99, no. 6, pp. 2060–2068, 1980.

[4] M. A. Abido, "Environmental/economic power dispatch using multiobjective evolutionary algorithms," in *Power Engineering Society General Meeting*. Toronto, Canada: IEEE, Jul. 13–17 2003, pp. 920–925.

[5] F. Ygge and H. Akkermans, "Decentralized markets versus central control: A comparative study," *Journal of Artificial Intelligence*, vol. 11, no. 1, pp. 301–333, 1999.

[6] H. Akkermans, J. Schreinemakers, and K. Kok, "Emergence of control in a large-scale society of economic physical agents," in *Proc. AAMAS*, New York, NY, USA, Jul. 19–23 2004, pp. 1232–1234.

[7] M. Pipattanasomporn, H. Feroze, and S. Rahman, "Multi-agent systems in a distributed smart grid: Design and implementation," in *Power Systems Conference and Exposition*. Seattle, WA, USA: IEEE, Mar. 15–18 2009, pp. 1–8.

[8] C. M. Colson and M. H. Nehrir, "Comprehensive real-time microgrid power management and control with distributed agents," *IEEE Trans. Smart Grid*, vol. 4, no. 1, pp. 617–627, 2013.

[9] N. Cai, N. T. T. Nga, and J. Mitra, "Economic dispatch in microgrids using multi-agent system," in *North American Power Symposium*. Champaign, IL, USA: IEEE, Sep. 9–11 2012, pp. 1–5.

[10] M. Hu, J.-W. Xiao, S.-C. Cui, and Y.-W. Wang, "Distributed real-time demand response for energy management scheduling in smart grid," *International Journal of Electrical Power & Energy Systems*, vol. 99, pp. 233–245, 2018.

[11] W. Liu, P. Zhuang, H. Liang, J. Peng, and Z. Huang, "Distributed economic dispatch in microgrids based on cooperative reinforcement learning," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 29, no. 6, pp. 2192–2203, 2018.

[12] N. Li, L. Chen, and S. H. Low, "Optimal demand response based on utility maximization in power networks," in *Power Engineering Society General Meeting*. Detroit, MI, USA: IEEE, Jul. 24–28 2011, pp. 1–8.

[13] M. A. Velasquez, J. Barreiro-Gomez, N. Quijano, A. I. Cadena, and M. Shahidehpour, "Distributed model predictive control for economic dispatch of power systems with high penetration of renewable energy resources," *International Journal of Electrical Power & Energy Systems*, vol. 113, pp. 607–617, 2019.

[14] J. Warrington, S. Mariéthoz, and M. Morari, "Negotiated predictive dispatch: Receding horizon nodal electricity pricing for wind integration,"in *International Conference on the European Energy Market*. IEEE, May 25–27 2011, pp. 407–412.

[15] G. Binetti, A. Davoudi, F. L. Lewis, D. Naso, and B. Turchiano, "Distributed consensus-based economic dispatch with transmission losses," *IEEE Trans. Power Syst.*, vol. 29, no. 4, pp. 1711–1720, 2014.

[16] T. van der Klauw, M. E. T. Gerards, G. Hoogsteen, G. J. M. Smit, and J. L. Hurink, "Considering grid limitations in profile steering," in *International Energy Conference*. IEEE, Apr. 4–8 2016, pp. 1–6.

[17] F. Fioretto, W. Yeoh, E. Pontelli, Y. Ma, and S. Ranade, "A distributed constraint optimization (DCOP) approach to the economic dispatch with demand response," in *Proc. AAMAS*, São Paulo, Brazil, May 8–12 2017, pp. 999–1007.

[18] J. S. Vardakas, N. Zorba, and C. V. Verikoukis, "A survey on demand response programs in smart grids: Pricing methods and optimization algorithms," *IEEE Commun. Surveys Tuts.*, vol. 17, no. 1, pp. 152–178, 2014.

[19] S. Parhizi, H. Lotfi, A. Khodaei, and S. Bahramirad, "State of the art in research on microgrids: A review," *IEEE Access*, vol. 3, pp. 890–925, 2015.

[20] A. R. Jordehi, "Optimisation of demand response in electric power systems, a review," *Renewable and Sustainable Energy Reviews*, vol. 103, pp. 308–319, 2019.

[21] IEEE, "European low voltage test feeder network (v2)," Available online: http://sites.ieee.org/pes-testfeeders/files/2017/08/European_LV_Test_Feeder_v2.zip, 2015, accessed: Nov. 08, 2018.

[22] A. Rassa, C. van Leeuwen, R. Spaans, and K. Kok, "Developing local energy markets: a holistic system approach," *IEEE Power and Energy Magazine*, vol. 17, no. 5, pp. 59–70, 9 2019.

[23] K. Kok and A. Subramanian, "Fast locational marginal pricing for congestion management in a distribution network with multiple aggregators," in *International Conference and Exhibition on Electricity Distribution*, Madrid, Spain, Jun. 3–6 2019.