Scalable multi-agent local energy trading - Meeting regulatory compliance and validation in the Cardiff grid

Meritxell Vinyals

Scalable multi-agent local energy trading - meeting regulatory compliance and validation in the Cardiff grid

Meritxell Vinyals[a]

[a] *Université Paris-Saclay, CEA, List, F-91120 Palaiseau, France*

Abstract

With the recent approval of energy directives (e.g. the EU Clean Energy Package) that state the rights and obligations of Local Energy Communities (LECs), time had come for smart grid technologies to show that they can comply with the complexity of the new regulatory environment when optimising LEC energy exchanges. This paper meets this challenge by modelling LECs operation by means of a novel energy coordination network that satisfies all stated requirements, such as the right of each community member of selecting their own supplier and the subsequent need of independent metering. The optimisation of the community is decentralised among a set of autonomous computational entities (a.k.a. *agents*), achieving scalability, and orchestrated via an agent interaction protocol based on the well-known Alternative Direction Method of Multipliers (ADMM) that keeps the preferences and cost structures of prosumers private. The approach is validated via extensive simulations using a dataset based on real data related to an existing energy grid in Cardiff (UK). Empirical results show that by optimising the use of prosumers' flexible resources to maximise the share of PV self-generated energy at community level our approach achieves higher community self-consumption ratios (up to 59% increment) and significantly reduces prosumer energy bills (589£/month additional reduction expected on summer months) w.r.t. optimising houses individually.

Keywords: local energy community; agent coordination network; alternative direction method of multipliers; peer-to-peer energy trading; distributed energy resources; regulatory environment

1. Introduction

Local energy communities[1] (LECs) have proven to be an effective and cost-efficient way to integrate the growing share of renewable energy as well as meeting citizen's needs all together. From a citizen perspective, they encourage

Email address: meritxell.vinyals@cea.fr (Meritxell Vinyals)

[1] Here LECs encompasses other terms used in the literature for equivalent structures such as smart grid neighborhoods, citizen energy communities, ...

passive consumers to become active prosumers[2] that not only maximise their individual self-consumption but also trade and share their surplus and flexible energy use within the local community. The benefits of such local energy exchange is two-fold: on one side it improves the security and efficiency of the supply since local energy balancing reduces transmission losses and stress on the infrastructures; and on the other, the direct sharing of electricity among prosumers leads to financial savings with more affordable renewable energy.

Given such benefits it is not surprising that LECs has been the focus of numerous industrial projects and research publications [22, 28, 21, 32, 30, 11]. Equally importantly, its viability has been repeatedly demonstrated through the implementation of actual demonstrators in R&D projects (see [32] for overviews on existing pilots), such as the pioneer Brooklyn microgrid project in New York, the Enerchain Project in Europe or the Valley Housing Project in Australia, among many others. However, the lack of a regulatory framework that legitimise the operation of such local communities have hindered so far its wider adoption at a larger scale [28, 30]. Fortunately, the recent approval of directives that clearly state the rights and obligations of LECs has finally filled this regulatory gap. In particular, the EU commission took a big step in the subject by adopting in June 2019, as part of the Clean Energy Package, a new directive [7] and a new regulation [8] which amend existing legislation on key issues for LEC design. Of particular interest here, [7] legitimize the operation of LECs (to which they refer to as *citizen energy communities*) by stating that all EU consumers shall be entitled to generate electricity for either consume it, store it, sell it back to the market or, very relevant here, share/trade it with other community members (as mentioned above, local energy trading is one of the pillars of LECs).

Therefore, with this legislation framework in place, time had come for smart grid technologies to show that they can comply with this new regulation environment when optimising LEC's operation. For example, one of the requirements stated in [7] is the right of each community member to choose their own energy supplier, in the same way that end-users who are not members of the community. That means that existing works restricted to configurations in which the whole community establishes a single supplier contract does not cover this general case. Notice also that this right also leads to the need to consider independent metering in the design differentiating the energy taken from/feeded to the supplier from those exchanged within the community.

Against this background, this paper models the underlying operation of LECs by means of a novel energy coordination network that integrates the equirements stated by recent directives. The energy coordination network is distributed among different autonomous computational entities (a.k.a. *agents*) and optimised in a decentralised way via the Alternating Direction Method of Multipliers (ADMM) protocol that orchestrates the active information exchange between the community participants required to ensure an effective energy exchange within the LEC. The resulting multi-agent energy trading platform en-

[2]Prosumers are consumers who also produce electricity.

able end-users to trade and share energy with other community members complying with the current regulatory environment. Concretely, this paper makes the following contributions to the state-of-the-art:

- We review the current energy regulatory environment and identify the overriding requirements to be satisfied for the correct operation and development of LECs.
- We model LECs operation by means of an energy coordination network showing how our novel design captures the identified requirements.
- We distribute this coordination network among different agents using ADMM as a coordination mechanism.
- We define for each type of agent which is the local optimisation problem to be solved as to participate in ADMM and define solutions of low computational complexity for each of them. This involved the formulation of a novel faster specialised procedure to solve the projection operator associated to a new type of agent: the XNOR balancing net.
- We validate our approach via extensive simulations using a realistic data set based on consumption and flexible smart appliance loads of 84 users (29 with local PV production and storage) from an existing neighbourhood in Cardiff (UK). We benchmarked our approach at: (i) household level w.r.t. the results obtained by the currently employed approach; and (ii) at community level w.r.t. the household optimised results.

This paper builds on previous work presented as a conference paper in [34] which presented an initial energy coordination network model of LECs and a preliminary evaluation on a synthetic data set. Herein that initial study is consolidated and extended in several important ways by: (i) enhancing the model to consider a larger number of complex devices such as storage units; (ii) showing how the model complies with the new regulatory environment; and (iii) performing extensive simulations with the Cardiff realistic dataset.

The paper is structured as follows. Section 2 reviews the related literature whereas Section 3 summarises the main requirements from recent regulations to be satisfied by LECs. Section 4 mathematically formalises the LEC operation under the identified requirements as an energy coordination network. Section 5 describes the decentralised optimisation protocol that results from applying ADMM on the proposed network and the local optimisation steps for each agent type. Section 6 proposes a payment scheme to distribute the community gains whereas Section 7 discusses the requirements for a real-world deployment of our framework. Section 8 presents and discusses our experimental results. Finally, Section 9 concludes and outlines some lines of future work.

2. Related work

After many analysis [28, 21, 30, 11] weighting the advantages and drawbacks of centralised vs. distributed structures in LECs, main arguments have advocated for the use of distributed schemes as to respect their fundamentally consumer-centric nature and scalability needs as well as the privacy and autonomy concerns of individual domestic households. Thus, the rest of this related work section is devoted to such decentralised approaches.

Following [32], research in this area can be divided into three categories, namely research on: (1) modelling the decision making process of exchanging energy between prosumers by coordinating their DERs with the aim to achieve certain energy related objectives either at individual or at community level [30, 29]; (2) the impact of energy sharing on the physical energy network; and (3) the development of energy sharing platforms that securise the information and financial transactions among prosumers via distributed ledger technologies (e.g. blockchain [23, 17]). Our work falls into the first category, focusing on the objectives of renewable energy usage maximisation and bill reduction.

Several works [16, 10, 19] have proposed to coordinate the local energy trading among prosumers by means of the same mechanisms widely used at grid scale, i.e. through double auction (DA) markets. For example, DAs have been used by [16] to coordinate residential houses via an optimal bidding strategy and by [10] to guarantee that energy exchanges don't violate network constraints. To avoid competitive advantages, [19] studied a DA-based energy market in which buying consumers were *randomly* paired with selling compatible prosumers leading to individual transaction prices. However, all these auction mechanisms that have been so effective at a global scale usually fail to unlock the flexibility that exist behind individual households, typically composed of small loads with a lot of interdependencies. In more detail, under such schemes participants are faced with a complex quotation decision process in which they need to provide accurate estimates of the slots with more local energy exceeding production and of its price to get profitable results. Also, even in the ideal situation in which all households are provided with perfect forecasting models, if all prosumers shift their flexibility and make an offer for the peak of energy production it may perfectly happen that there is not enough energy surplus to cover all the mobilised consumption whereas opportunities in less prominent peaks are missed. Last but not least, these mechanisms fail when providing households with a bidding language that allows them to express the interdependencies (e.g. shiftable loads) and hence this type of flexibility remains blocked.

Game theory is another mathematical framework that has been extensively used to analyze the strategic decision making of LECs intra-energy trading (we refer to [31] for a recent review on this topic). Stackelberg and non-cooperative Nash games have been applied to LECs trading in [1, 35, 18, 36]. However, here we focus on cooperative solutions in which prosumers explicitly communicate in order to coordinate their behaviours. Thus, closer to us, are [17, 14, 29] that model LECs trading as cooperative games in which prosumers are grouped in coalitions, the gains of each coalition being distributed among its members based on revenue distribution schemes that makes them stable. However, it is worth to notice that in many of these games the grand coalition emerges and, more importantly, the value of a coalition in this domain involves computing collaborative solutions in which the behaviours of individual members are coordinated, needing a framework as the one proposed in this paper.

Some works in the literature took the same approach as in this paper, applying distributed optimisation schemes to organise the local energy trading. For example, consensus+innovation and primal-dual gradient algorithms has been respectively used in [27, 12] to orchestrate the operation of a fully decentralised local market. But, by far the most used popular in this landscape has been ADMM. The closer related work with regard to energy coordination networks and ADMM for LEC's optimisation has been presented in [33, 22, 4, 25, 23]. However, such previous studies either model households using abstract models and/or do not consider energy storage and/or lack of any realistic validation and most importantly, all of them, do not model the current legislation environment as we do in this paper. Thus, [33] is based on a very abstract design in which a household is simply modelled using a measure of dissatisfaction w.r.t. each allowed profile and, as a consequence, no legislation requirement it is easy foreseen in that model. [22] focuses on modeling the interface between the community and external parties/markets (i.e. wholesale, balancing, ancillary services,...) and on the evaluation of different communities strategies (i.e. autonomy, peak-shaving, ...) on a small community of 15 prosumers. Interestingly, the objective of [4] is to prove the effectiveness of the approach in a multi-flow setting (considering not only electricity but also heating) with realistic scenarios from a neighborhood in Woerden, Netherlands. However, the design in [4] is restricted to communities with one mutual electricity supplier contract and does not generalise to the case in which each community member may chose its own supplier. The work in [25] tackled the issue of non-convexity of discrete households loads and power equations and proposed several approaches layered on the top of ADMM to handle discrete variables whereas keeping the agent local problems convex. Complementary to us, [23] showed how to secure and fully decentralise the aggregation step of ADMM when optimising the energy exchanges in a microgrid by using a blockchain. We should also mention the work of Lilliu et al. in [15] which used the same Cardiff dataset to evaluate their approach at three levels: the household, the LEC and an third level that considers congestion management requests. However, the similarity with that work ends here since their approach consisted on applying simple heuristic search to solve each of these levels and their model does not take into account any legislation requirement.

Finally, many studies [20, 24, 9] worked on innovative pricing mechanisms that promote the LEC renewable energy sharing, overcoming the well-known limitations of the extensively used feed-in tariffs (i.e. non-coordination of users, generation of new peaks and highly subsidisation). Nevertheless, although efficient LEC interaction needs proper pricing schemes such works do not tackle the problem of coordinating prosumers' flexibility to maximise their profits under such schemes, making them complementary to our approach (as discussed in Section 6 our approach can indeed be potentially combined with these schemes).

3. Requirements

This section reviews different directives and regulations on LECs with the objective of drawing the main requirements to be satisfied by these structures. Table 1 summarises the main identified requirements, classifying them depending if they apply at household or at community level.

Prior to official directives, we should mention the existence of some de facto standards, like the Universal Smart Energy Framework (USEF)[3], that have been extensively used as a reference framework for LECs' design. It is worth to highlight that the approach presented in this paper is USEF compliant.

The first requirement listed in Table 1 is one of the basics from self-consumption: in a given time, a prosumer can either import or export energy from/to the grid but never both (R1). See that this could not be allowed without jeopardizing the whole self-consumption principle: export tariffs are there to allow a prosumer to share their energy surplus, not the energy needed to cover her needs.

Then, we focus on the recently adopted EU Clean Energy Package [7] which sets out the basic principles to regulate LEC's operation, namely: the right of consumers to consume, store and sell their self-generated energy, overriding in this way some of the barriers still present in several countries such as the obligation/prohibition to feed/export self-generated electricity. Notice however that this export applies only to self-generated energy: consumers are not allowed to export energy to the grid that is not coming from self-generation (R2). Importantly, this directive lays the foundation for the regulation of local energy exchange by stating that LECs should not face regulatory restrictions when they apply existing or future technologies to share electricity produced using generation assets within the community among their members. Again, notice that this sharing is restricted to energy self-produced within the community (R5).

Another important requirement from [7] is that the energy community should not impose restrictions on its members when selecting an energy supplier. So community members should be able to choose and change their electricity supplier in a same way as before (R4). This means that even if it is possible for a community to have one mutual energy supplier contract this is not the general case since the directive states that it should be ensured that each customer is free to purchase energy from the supplier of its choice. A direct consequence of considering members of the community with different suppliers (or with the same supplier but different contracts) is that the connection and metering points for such different contracts should remain differentiated (i.e. need for independent metering). In other words, the connection that a prosumer has with the community should remain different from the one it has with the supplier (R3).

Finally, to guarantee community engagement [29], community members should gain economical benefits for sharing their self-produced energy, compared to their conventional energy supplier contract (R6).

[3] `www.usef.energy`

In-home	R1	In a given time, a prosumer can either import energy from the grid or export energy to the grid but never both.
	R2	The only energy that a prosumer can export through an export tariff is the local self-generated energy.
Community	R3	The connection that a prosumer has with the community should remain different from the one it has with the supplier.
	R4	Even being member of a community, each prosumer has the right to choose its own supplier.
	R5	The energy exchanged within a community is restricted to the energy produced within the community (i.e. a community member can not resell to other members the energy imported from the grid).
	R6	Community members should receive a financial compensation for sharing their energy higher than the one offered by traditional energy suppliers.

Table 1: Requirements for a correct LEC operation at in-home and community level.

4. Multi-agent local energy trading model

This section formulates the problem of optimising prosumers' energy exchanges in a LEC as an energy coordination network. It also details how this energy coordination network is distributed among the two types of socio-economical agents that participate in the optimisation, namely: (i) the HEMS (Home Energy Management System) agent, that carries out an in-home optimisation behind the meter with the objective of minimising the prosumer energy bill; and (ii) the AGR (Aggregator) agent, which enables and coordinates an energy exchange zone (i.e. the AGRNet) at community level. Fig. 1 depicts the process flow for the two types of socio-economical agents. Next section reviews the model of energy coordination networks. Then, Section 4.2 details the in-home coordination network that lies in the boundaries of the HEMS agent whereas Section 4.3 does the same for the community network and the AGR agent.

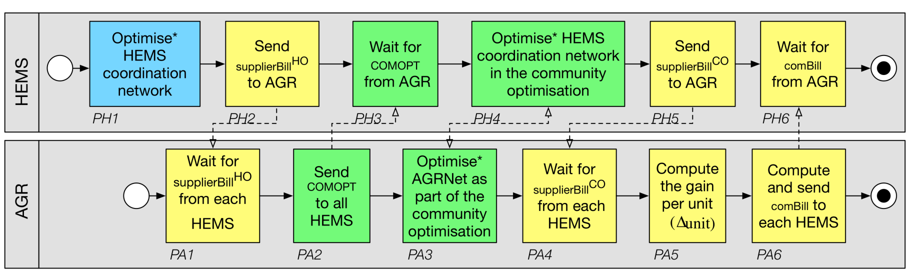

*Optimisation performed by running ADMM on the corresponding coordination network until convergence or maximum time reached.

Figure 1: Process flow of the two types of socio-economical agents (HEMS and AGR). Processes related to the in-home optimisation are depicted in blue, processes related to the community optimisation in green and processes related to the payment mechanism in yellow.

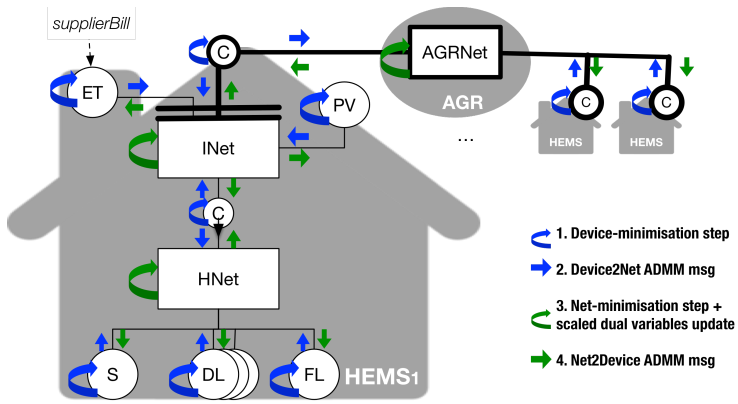

Figure 2: The LEC energy exchange coordination network composed of HEMS coordination networks connected through an AGR. Bold lines highlight the extensions added to perfom community optimisation. Arrows in blue/green stand for ADMM messages/steps.

4.1. Energy coordination networks

Following [13], a coordination network is composed of a set of devices, D, a set of nets, N, and a set of terminals, T. A terminal is a connection between one device and one net. In an *energy* coordination network, a net is an agent that represents a virtual energy exchange zone (i.e. where energy exchange among devices is determined) whereas a device is an agent that owns one or multiple transfer points, each modelled as a terminal. Thus, in an energy coordination network, each terminal $t \in T$ is associated to an energy schedule $(p_t(1), \ldots, p_t(H)) \in \mathbb{R}^H$ over a time horizon $H \in \mathbb{N}^+$. Then, for each time slot $\tau \in [1, H]$, $p_t(\tau)$ is the energy consumed (if $p_t(\tau) > 0$, otherwise produced) by a device through terminal t, during the τ-time slot. Fig. 2 illustrates an energy coordination network modeling an energy community where nets are represented by rectangles, terminals by lines and devices by circles.

For each device $d \in D$, we use d to refer to both the device itself as well as to the set of terminals associated with it, i.e., we say $t \in d$ if terminal t is associated with device d. The set of all energy schedules associated with device d is denoted by $p_d = \{p_t | t \in d\}$, which we can associate with a $|d| \times H$ matrix. In addition to $|d|$ terminals, device d is also associated with an objective function $f_d : \mathbb{R}^{|d| \times H} \to \mathbb{R}$, where $f_d(p_d)$ is the exploitation cost of device d for the energy schedule p_d. Moreover, every device has a set of constraints, denoted as C_d, that p_d should satisfy in order to be feasible.

Analogously to devices, every net $n \in N$ has $|n|$ terminals, an objective function $f_n : \mathbb{R}^{|n| \times H} \to \mathbb{R}$ and a set of constraints C_n that $p_n = \{p_t | t \in n\}$ should satisfy in order to be feasible. Since each net $n \in N$ models an energy exchange zone, the set of constraints C_n should always include the energy balancing condition, i.e. $\sum_{t \in n} p_t(\tau) = 0, \ \forall \tau \in [1, H]$.

4.2. The HEMS agent: the in-home coordination network

The prosumer, represented by the corresponding HEMS agent, is the individual provider of flexibility. To perform the optimisation each HEMS should have a good representation of the technologies providing flexibility, of the prosumer energy behaviour and of the conditions of the energy supplier tariff contracted by the prosumer. As shown in Fig. 2, all of this is modelled through device agents that encode individual operational costs and constraints. Notice that although we restrict our description to those devices present in the validation dataset, our approach can integrate any other device following the same modelling principles.

Thus, each HEMS coordination network contains an External Tie (ET) device which models the connection and the characteristics of the tariff that the prosumer has with its chosen supplier. In particular, the ET device from Fig. 2 is characterised by a local objective function (f_{ET}) that encodes the price of importing (i.e. P^{imp}) and exporting (i.e. P^{exp}) in a given time slot τ as follows:

$$f_{ET}(p_{ET}) = \begin{cases} -P^{imp} \cdot p_{ET} & \text{if } p_{ET} \leq 0 \\ P^{exp} \cdot p_{ET} & \text{if } p_{ET} > 0 \end{cases}$$

Notice that by modeling the energy supplier at home level, our model satisfies the requirement R4, namely even being member of a community, each prosumer has the right to choose its own supplier. For the sake of the optimisation, the network also includes for each prosumer a Fixed Load (FL) that aggregates all the expected[4] consumption of the prosumer that is inelastic, i.e. does not provide any flexibility and can not be controlled. Moreover, as discussed in Section 1, the effectiveness of LECs depends on having some prosumers with local energy surplus to trade. In the Cardiff dataset, self-production comes exclusively from photovoltaic (PV) units. Hence, the in-home coordination network includes PV devices that model[4] this self-production. Finally, the model includes two energy resources that provide flexibility: deferrable loads (DL), those are loads that can be shifted by a certain time, and storage (S), namely batteries that can take in or deliver energy. In more detail, the S device in Fig. 2 constrains that the charge (q_s) is positive and does not exceed its capacity (Q^{max}) in any time slot $0 \leq q_s(\tau) \leq Q^{max}, \quad \tau = 1, \ldots, H$ being $q_s(\tau)$, which incorporates in the same equation the linear charging (η^c) and discharging (η^d) losses and the storage initial charge (Q^{init}), defined as:

$$q_s(\tau) = Q^{init} + \sum_{t=1}^{\tau} \eta^c \cdot p_s^+(\tau) + \sum_{t=1}^{\tau} \frac{1}{\eta^d} \cdot p_s^-(\tau)$$

where: $p_s^+(\tau) \geq 0, p_s^-(\tau) \leq 0, p_s^+(\tau) \cdot p_s^-(\tau) = 0, p_s(\tau) = p_s^+(\tau) + p_s^-(\tau)$.

Now, the most straightforward energy coordination network one can think of the in-home optimisation problem consists in connecting all in-home devices

[4]The model of such devices take as input the corresponding forecasted consumption or production profiles.

(S, DLs, PV, ET, ...) to a home net that guarantees the energy balancing condition. However, a problem with this approach is that, in case of a contracted export tariff, it allows the configuration in which the prosumer exports energy to the grid that it is not coming from the PV, not satisfying R2. To avoid this problem, the HEMS coordination network contains two energy balance nets: the home net (HNet) and the interface net (INet). All the in-home devices with exception of the PV and the ET (i.e. FL, DL, S) will be connected to the HNet. Conversely, PV and ET devices are connected to the INet. Finally, the INet is connected to the HNet through a connector (depicted in Fig. 2 as C with an arrow) that only allows the flow to go from the INet to the HNet (i.e. all the devices connected to HNet only import energy from the INet). More formally, the connector device imposes the following local constraints: $\eta_1 \cdot p_1(\tau) = -p_2(\tau)$ if $0 \leq p_1(\tau) \leq P_1^{max}$ and $p_1(\tau) = -\eta_2 \cdot p_2(\tau)$ if $0 \leq p_2(\tau) \leq P_2^{max}$ where (P_1^{max}, P_2^{max}) are the energy flow limits and (η_1, η_2) are the linear transmission efficiency coefficients respectively in each direction.

This configuration guarantees that only energy produced from the PV is exported through the ET by means of the corresponding export tariff. In more detail under this model, the PV energy can: (i) be used immediately within HNet for local usage; or (ii) stored to be used later within HNet for local usage; or (iii) exported immediately through the export tariff. Notice that the model doesn't consider configurations in which the PV energy is stored and exported to the grid latter. This is because any of these configurations will be clearly suboptimal (i.e. the storage of energy involves charging/discharging/storing losses and its indiscriminate use without need can reduce its life time).

Finally, this model satisfies R1 because since there is a single terminal between ET and INet, the power planned through this connection at a given time can be positive or negative but never both.

As detailed in Section 5, the in-home optimisation protocol run by HEMS (process $PH1$ in Fig. 1) will be the result of applying ADMM over this network.

4.3. The AGR agent: the community coordination network

Fig. 2 highlights in bold the extensions added to the in-home optimisation energy network to perform community optimisation. Thus, we observe that the AGR agent is represented by a net through which it creates a local energy exchange zone for community exchanges. From an optimisation perspective, if the prosumer joins a local community it adds also a connection (depicted in Fig. 2 as a C device) with the AGR that represents her exchanges with the community. Since the connection with the community is mapped through the AGR terminal whereas the connection with the supplier is mapped through the ET device, this model satisfies requirement R3. By means of the AGR terminal the prosumer can: (i) import energy in order to be stored or to be used for immediate consumption of in-home device(s); and (ii) export energy from its local PV production. However, the energy exported through this AGR connection should be restricted to the energy produced within the community (R5). In other words, the model needs to guarantee that prosumers will not sell to other actors (i.e. other prosumers in the community) the energy imported from the grid (i.e. imported by means of their corresponding supplier tariff). To avoid those scenarios and satisfy requirement R5, we need to change the type of INet from in-home optimisation to not only consider the typical balancing constraints but also to check that no transfer of power takes place between the AGRNet and the ET (i.e. to check that for any time slot, both power schedules have the same sign). This new type of net, that we called XNOR energy balance net is represented in Fig. 2 as a bold line on the terminals of INet that should have the same sign, i.e. terminals connecting the onnector to the AGRNet and to the ET. Formally, let $n' \subseteq n$ be the subset of terminals constrained to have the same signal in a XNOR balance net. Then, in addition to the power balance constraint, a XNOR balance net locally imposes the same sign among the n' terminals: $\left(\bigcap_{t \in n'} p_t(\tau) \geq 0\right) \bigcup \left(\bigcap_{t \in n'} p_t(\tau) \leq 0\right),\ \forall \tau = 1, \ldots, H$.

As detailed next in Section 5, the community optimisation protocol will be the result of applying ADMM over this network as a whole, i.e. with HEMS agents being in charge of the computations and messages exchanged within their in-home energy network (process $PH4$ in Fig. 1) and the AGR agent being in charge of the AGRNet (process $PA3$ in Fig. 1).

5. Multi-agent local energy trading algorithm

In our approach, agents compute the most efficient energy exchanges in a decentralised way by executing the ADMM protocol over the respective energy coordination network, i.e. over the HEMS energy coordination network to perform individual in-home optimisation and over the whole community network to perform community optimisation. ADMM [3] is an iterative algorithm that allows solving the underlying social welfare optimisation problem of a coordination network in a distributed way. Formally, an energy coordination network defines the following social welfare optimisation problem:

$$\min_{p \in \mathbb{R}^{|T| \times H}} \sum_{d \in D} f_d(p_d) + \sum_{n \in N} f_n(p_n) \text{ s.t. } \forall d \in D : p_d \in C_d, \forall n \in N : p_n \in C_n \quad (1)$$

that searches for the set of energy schedules that minimizes the sum of the local cost functions of all devices and nets that compose the coordination network whereas satisfying all the constraints in their local constraint's sets.
To form the augmented Lagrangian, nets' objective functions are defined over a duplicated copy of the original variables (i.e. denoted as $\dot{p}$) and the equality constraint ($p = \dot{p}$) is relaxed via a Lagrange multiplier. Then, ADMM consists on applying the following four steps at each iteration $k+1$ (being $\rho > 0$ the augmented Lagrangian penalty parameter and $||\cdot||_2^2$ the square of the euclidean norm, i.e. the sum of squares):
1. [*device-minimization* step] Each device d computes:

$$p_d^{k+1} = \arg \min_{p_d \in C_d} f_d(p_d) + \frac{\rho}{2} \sum_{t \in d} ||p_t - \dot{p}_t^k + u_t^k||_2^2 \quad (2)$$

2. [*device2net* msg] Each device d sends to each neighbouring net $t \in d$ its updated energy schedule p_t^{k+1}.

3. Each net n computes the:
   3.1. [*net-minimization* step]:

$$\dot{p}_n^{k+1} = \arg\min_{\dot{p}_n} f_n(\dot{p}_n) + \frac{\rho}{2} \sum_{t \in n} ||p_t^{k+1} - \dot{p}_t + u_t^k||_2^2 \tag{3}$$

   3.2. [(price) *scaled dual variables update*]:

$$u_n^{k+1} = u_n^k + (p_n^{k+1} - \dot{p}_n^{k+1}) \tag{4}$$

4. [*net2device* msg] Each net n sends to each neighbouring device $t \in n$ its updated device energy schedule $\dot{p}_t^{k+1}$ and the prices for scaled dual variables u_t^k.

These steps are run until convergence of all nets. Following [3], a net converges when the norm of the primal and dual residuals are small being the net primal residual defined as the difference between the energy schedule send by devices and those required by the net, i.e. $p_n^k - \dot{p}_n^k$, and the net dual residual as the difference between the energy required by the net into two consecutive iterations weighted by the scaling parameter, i.e. $\rho \cdot (\dot{p}_n^k - \dot{p}_n^{k-1})$.

Fig. 2 depicts the different ADMM steps for the devices and nets that compose the LEC energy exchange coordination network. We observe that at each iteration, each device agent computes a step (Eq. 2) that minimises its operating cost and constraints (i.e. encoded by f_d and C_d) and a penalty that depends on the messages passed to it through its terminals by its neighbouring nets in the previous iteration (i.e. $\dot{p}_t^k$,u_t^k). Similarly, each net agent computes its minimisation (Eq. 3) and scaled dual variables steps (Eq. 4) with an argument that depends on messages passed to it through its terminals by its neighbouring devices in the previous iteration (p_t^{k+1}). We also observe that in terms of socio-economical agents, the HEMS and the AGR agents exchange one message in each direction at each iteration. In more detail, each HEMS sends to the AGR (in a *device2net* msg) its updated community energy exchange schedule and the AGR respond to each HEMS (in a *net2device* msg) with its updated proposal of community energy exchange schedule along with the price signals. Hence, the size of each exchanged message is linear to the considered time horizon (e.g. 96 values in the experiments considering a day-ahead schedule of 15 min. interval).

Now, notice that the net and the device minimisation steps require the evaluation of a proximal operator whose complexity depends on the structure of the particular local objective function (f_x) and constraints (C_x). In what follows we identify low computational cost procedures for computing the minimization steps of all types of considered agents. Table 2 summarises these procedures as well as their computational complexity. Notice that for balance nets, connector and storage we simply apply low cost procedures already formulated in the literature. Also, for photovoltaic/fixed load devices no optimisation is needed

	Type	Optimisation procedure	Time complexity
Nets	Balance net	From [13]: projection on an hyperplane for each time step	$O(H)$
	XNOR net	Algorithm 1	$O(H)$
Devices	Photovoltaic/ fixed loads	Return the given inelastic schedules	$O(1)$
	External Tie	From [13, 4]: closed form solution in form of arithmetic expression for each of the two cases of the piecewise constraint for each time step	$O(H)$
	Connector	From [4]: closed form solution in form of arithmetic expression result from comparing two projections on the hyperplane for each time step	$O(H)$
	Deferrable load	Evaluating each possible outcome in the deferrable interval	$O(H)$
	Storage	From [4]: connector solver + Dykstra's projection method	Convergence in norm.

Table 2: Summary of the procedures used to solve each net/device local optimisation problem as well as their computational time complexity.

since they provide no flexibility. Hence, in what follows we only detail the local optimisation steps for XNOR balance nets and deferrable loads.

XNOR balance net. In what follows we describe a $O(H \cdot |n'| \cdot log|n'|)$ procedure to compute the net-minimisation step for XNOR nets being $|n'|$ the number of terminals involved in the XNOR constraint (in our case the ARGNet and ET terminals and hence $|n'| = 2$ and complexity reduces to $O(H)$).

Notice that the problem is separable for each time step ($\tau \in [1, H]$) and for each of the two cases, namely for the cases when all XNOR terminals are: (i) positive ($\bigcap_{t \in n'} p_t(\tau) \geq 0$); and (ii) negative ($\bigcap_{t \in n'} p_t(\tau) \leq 0$). In what follows we will detail the positive case. The resolution of the negative case follows in a straightforward manner by changing the corresponding constraint in Eq. 6.

To solve the net minimisation step for the positive case, XNOR balancing nets need to solve the following problem for each time slot $\tau \in [1, H]$:

$$\arg\min_{\dot{p}_n} \frac{\rho}{2} ||\dot{p}_n - v_n||_2^2 \tag{5}$$

$$\text{s.t.} \sum_{t \in n} \dot{p}_t(\tau) = 0,\ \dot{p}_t(\tau) \geq 0, \forall t \in n' \text{ where } v_n = p_n^{k+1} + u_n^k \tag{6}$$

This problem can be casted as a variation of the Euclidean projection onto the simplex in which the inequality of Eq. 6 only applies to a subset of the entries. Therefore, as a basis for our algorithm, whose pseudocode is formally described in Algorithm 1, we use the efficient projection onto the simplex algorithm proposed in [5]. Let $v_{n'}$ be the vector of preferred values at time step τ for the terminals that are in the XNOR constraint and let μ denote the vector obtained by sorting $v_{n'}$ in descending order (this ordering is the costliest operation). At line 2 Algorithm 1 computes γ that is the highest index for which the members of $v_{n'}$ will remain strictly positive after substracting the value to satisfy the balance constraint. Then, line 3 computes θ, the value that will be substracted to each preferred value at time step τ of each terminal of the net n. Finally, the algorithm outputs as the optimal value of each variable its preferred value minus θ and bounded positive in case of XNOR terminals.

Input: Vector $v_n \in \Re^{|n|}$

1 Sort $v_{n'}$ into $\mu : \mu_1 \geq \mu_2 \geq \ldots \geq \mu_{|n'|}$;

2 Find $\gamma = \max\{i \in 1 \ldots |n'| : \mu_i - \frac{1}{i+|n|-|n'|}(\sum_{j=1}^{i} \mu_j + \sum_{j \in n \setminus n'} v_j) > 0\}$;

3 Define $\theta = \frac{1}{\gamma+|n|-|n'|}(\sum_{i=1}^{\gamma} \mu_i + \sum_{j \in n \setminus n'} v_j)$;

Output: $\dot{p}_n = [v_n \setminus v_{n'} - \theta, max\{v_{n'} - \theta, 0\}]$

Algorithm 1: XNOR balance net-minimisation step procedure.

Deferrable load. A deferrable load[5] restricts its required energy profile (P^{req}) to be executed within the interval $[T^{min}, T^{max}]$. The problem is separable and solved independently for each time slot in which the deferrable load can be scheduled (i.e. $T^{min} \ldots T^{max} - |P^{req}|$), being optimal the one with minimal cost in the objective function. Hence time complexity is defined by the length of the interval where the deferrable load can be scheduled (i.e. H in the worst-case).

6. Community payment mechanism

As a result of running ADMM, prosumers find the most cost-effective energy exchanges and configurations but this does not determine how the economic benefits derived from increasing the community self-consumption are shared between them. Although the values of ADMM scaled dual variables on convergence can be interpreted as competitive market prices [3] (see [33, 23] for examples of using this price interpretation in a LEC), here we decided to take a different approach in order to satisfy the directives requirements. Thus, for the sake of completeness, in what follows we propose a budget balanced individual rational payment scheme that distribute the benefits of the community at individual level respecting requirement R6. Nevertheless it is important to notice that our approach is not restricted to this specific payment mechanism and in fact it can be combined with other more complex payment schemes such as the ones proposed in [20] or [24], for example.

To compute the community payments the AGR agent will first compute the gain per unit (Δ_{unit}) of energy exchanged within the community obtained by the community optimisation (CO) with respect to individual home optimisation (HO). This is computed after the AGR agent received from each prosumer $p \in P$ (processes *PA1* and *PA4* in Fig. 1) the amounts that p would pay to her supplier for her exchanges with the grid[6] in: (i) the in-home optimisation solution ($supplierBill_p^{HO}$); and (ii) in the community optimisation solution ($supplierBill_p^{CO}$). Then Δ_{unit} is computed (process *PA5* in Fig. 1) by

[5]This modelling differs from the one in [13] which required the consumption of a minimum amount of energy (not an energy profile) within an interval.

[6]The exchanges with the grid are determined by the terminal with the ET device.

the AGR agent as the difference between the sum of what prosumers (P) pay to their suppliers after in-home optimisation (i.e. $\sum_{p \in P} supplierBill_p^{HO}$) and after community optimisation ($\sum_{p \in P} supplierBill_p^{CO}$) divided by the units of energy exchanged within the community after community optimisation.

Finally, the AGR agent computes (process *PA6* in Fig. 1) the payment of each prosumer p to the community as follows:

$$\begin{aligned} comBill_p = & (supplierBill_p^{HO} - supplierBill_p^{CO}) \\ & - (1-\alpha) \cdot \Delta_{unit} \cdot U_p^{con} - \alpha \cdot \Delta_{unit} \cdot U_p^{prod} \end{aligned} \tag{7}$$

where U_p^{con}, U_p^{prod} stand respectively for the units consumed from and produced for the community by prosumer p and the parameter $\alpha \in [0,1]$ is the fraction of the unit gain that production is rewarded w.r.t. consumption.

Notice that each prosumer ends up paying ($comBill_p + supplierBill_p^{CO}$) the same amount than after individual home optimisation ($supplierBill_p^{HO}$) but with some reduction for every unit consumed/produced from/for the community. Therefore, this payment scheme is individually rational (i.e. prosumers pay equally or less by joining the community than w.r.t. optimising alone their individual households) and hence it satisfies requirement R6. Moreover, the scheme is also budget balanced meaning that the mechanism collects and disburses the same amount of money from/to the agents (i.e. no subsidy is needed).

7. Multi-agent local energy trading framework requirements

The infrastructure required will depend on the particular deployment scheme. The most straightforward deployment[7] of our framework consists of having each HEMS agent running at its home premises and the AGR agent running at community level. In terms of computational infrastructure, this deployment requires an energy box/computer at each household capable of executing the corresponding HEMS agent and a cloud/community computer to execute the AGR agent. As mathematically proven in Section 5 and corroborated with experiments in Section 8, the execution of both, HEMS and AGR agents, is based on scalable procedures and only requires low computational resources.

In terms of communication, each agent needs access to an internet connection to perform the required AGR-HEMS two-way communication. As shown in Fig. 1, each HEMS communicates to the AGR its supplier billing amount after in-home (process $PH2$) and community optimisation (process $PH5$) and as part of the ADMM algorithm when performing community optimisation (process $PH4$). On the other hand, the AGR communicates with each HEMS to start the community optimisation (process $PA2$), during the ADMM protocol for community optimisation (process $PA3$) and to communicate its community billing amount (process $PA6$). As detailed in Section 5, during the ADMM community optimisation, each HEMS and the AGR exchange one linear size message

[7] A deployment with all agents running in a cloud is also a possibility.

in each direction at each iteration, requiring frequent low-bandwith connection. Finally, in terms of physical infrastructure, our approach requires each home to be equipped with smart metering capable of distinguish between the amount exchanged with the grid from the amount exchanged with the community.

8. Results and discussion

This section discusses the results obtained from validating our approach on the Cardiff dataset. Section 8.1 describes the dataset along with other experimental settings. Section 8.2 analyses the results for the individual home optimisation and Section 8.3 for the community optimisation.

8.1. Experimental settings and dataset

The dataset used to conduct our simulations has been created in the context of the MAS^{2}TERING project [6] and it is based on the historical data and characteristics of a low voltage grid in Cardiff, UK.[8] In what follows we provide a brief description of the dataset. For a complete description we refer the reader to the corresponding deliverables of the MAS^{2}TERING project, e.g. [26].

The dataset builds on daily demand and generation data[9] of 84 consumers (29 of which are prosumers with local PV production and storage) from the Cardiff grid. Also, information was made available regarding the supplier tariff chosen by customers. Table. 3 details the four different tariffs present in the dataset. The assignment of a prosumer to a tariff depends on the prosumers' choices and does not vary with the scenario. Regarding export prices for self-generation, all prosumers have the same tariff with 0.0491£/kWh exported.

In [6], this reference data was projected in future smart grid scenarios based on studies carried out by UK National Grid. Our experiments focus on the projected profiles corresponding to the *2030 Green scenario*. Each projection fixes which prosumers have PV generation and can consequently self-produce and store energy (in the dataset all users with PV are also equipped with a home 3.3kWh storage unit with charging/discharging efficiencies set as $\eta^c = \eta^d = 0.99$). Also, each projection assigns to each household a set of smart appliances (dishwashers, electric ovens, electric vehicles, freezers, fridges, tumble dryers and washing machines) that can provide flexibility. The loads of these smart appliances are extracted from the consumer baseline load and added as deferrable loads. The times of the day when each smart appliance can be scheduled are pre-determined but depend on the habits and choices of prosumers.

Experiments were performed for different day types (i.e. Weekdays, Saturdays and Sundays) and seasons. However, for the sake of space, the results discussed in the paper are restricted to Weekdays and to the two most differentiated seasons, i.e. Summer and Winter. Results obtained for other day types do

[8]For full access to the data please send an email to the MAS^{2}TERING coordinator.

[9]Consumption/production values were recorded for each of the 96 time slots, each time slot referring to a 15 min interval.

not present any significant difference whereas Summer and Winter seasons provide a good overview of the benefits of the approach over the year. In the dataset the amount of flexible load in the community is fixed to 395kWh/day and does not vary with the season. However, as analysed in next sections, due to the significantly lower PV production (220kWh/day in Winter vs. 500kWh/day in Summer) and higher demand (2268kWh/day in Winter, around 30% plus than in Summer months due to heating consumption), there is little excess of self-generated energy to trade among prosumers after in-home optimisation in Winter. Therefore, next sections focus on the results obtained for Summer, reporting on results obtained for Winter for comparison.

	00:00-07:00	**07:00-16:00**	**16:00-19:00**	**19:00-24:00**
Off-peak saver 2 bands	0.1149		0.1851	0.1149
Off-peak saver 3 bands	0.0869	0.1267	0.2785	0.0869
Eco 2020 2 bands	0.1244	0.1554		
Economy 7	0.0508	0.1627		

Table 3: TOU tariffs (£/kWh) from the Cardiff dataset.

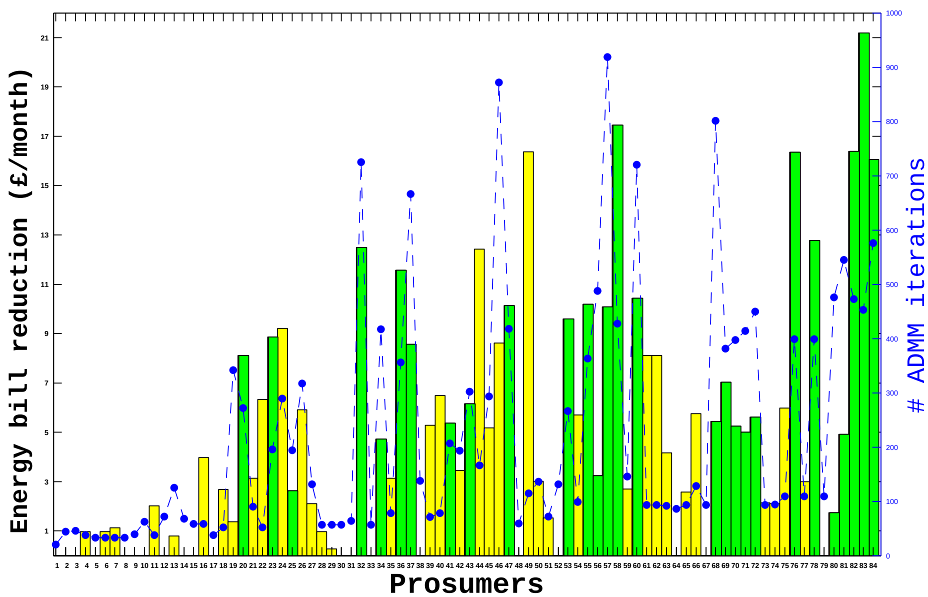

Figure 3: For each prosumer: a) absolute reduction in the monthly energy bill after in-home optimisation for Summer profiles (left y-axis, bar plot) and number of ADMM iterations to perform the optimisation (right y-axis, blue line plot). Prosumers are ordered by their level of flexibility. Green refers to prosumers with storage and local production.

8.2. In-home optimisation results

This section evaluates the results from applying the solution at the level of the individual home (i.e. only in-home optimisation is performed). The results are

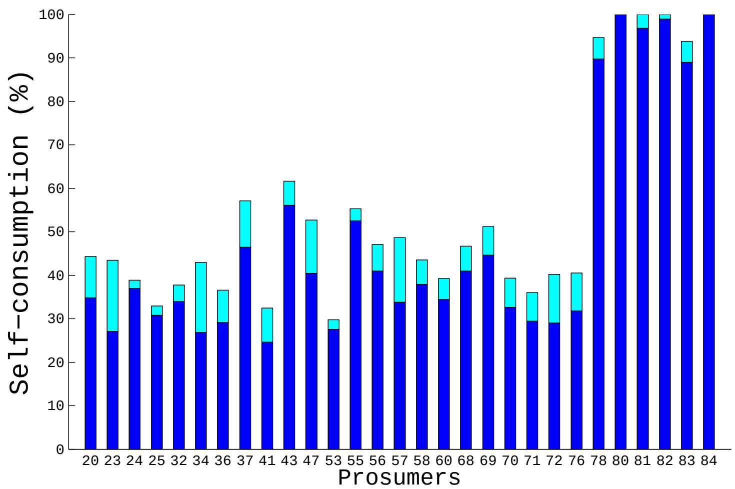

Figure 4: Home level self-consumption of prosumers with local production. Dark blue: current self-consumption. Light blue: additional increment after in-home optimisation.

benchmarked against a non-optimised scenario in which storage units are not used and all deferrable loads are just scheduled to its default value.

Fig. 3 on the left y-axis and bar plot depicts the absolute reduction in the monthly energy bill of each prosumer after performing the in-home optimisation for Summer profiles. Prosumers have been ordered by the amount of their flexible consumption. As expected we observe that all twenty-nine prosumers with PV and storage (depicted in green bars) exhibit important reductions on the energy bill (from 1.74£/month to 21.2£/month) as a result not only of moving flexible loads but more importantly of using the storage to increment its PV self-consumption. The results obtained for Winter months after in-home optimisation were very similar with reductions that went from 1.74£/month to 20.2£/month. In Fig. 3 we observe important reductions on consumers with no self-production (and thereby with no storage) exclusively as a result of shifting their demand from expensive to cheaper hours in their respective energy tariffs. Also, as expected the very first consumer in this list and the only one with no flexibility gets no reduction as a result of the in-home optimisation (as we will see in next section this consumer can still benefit from reductions at community level). Finally, it is worth to notice that even when the level of flexibility of a prosumer plays a role, in Fig. 3 we can identify up to 18 prosumers[10] that even having flexible loads their current flexibility does not allow them to reschedule them when energy is cheaper leading to no reduction on their bills. In the next section we will see how all these prosumers would be able to benefit from community optimisation by exploiting their flexibility to match the exceeding

[10]Prosumers #2,3,5,8-10,14,15,17,30,31,33,38,48,52,64,67,79.

production at community level.

Fig. 3 on the right y-axis and blue line plot depicts for each prosumer the number of ADMM iterations needed to solve the in-home optimisation problem (process PH1 in Fig. 1). We observe a trend in which the algorithm is likely to take more iterations if the prosumer has a storage associated or more flexible loads. Nevertheless, we observe that in all cases the algorithm never takes more than 1000 iterations to optimise a single household (execution times[11] varied between 8 sec. and 4 min. depending on the prosumer, taking each iteration around 0.2 sec. in average). Similar results and trend were obtained when solving Winter profiles, showing no significant difference between seasons.

Fig. 4 shows for the twenty-nine prosumers with local production how self-consumption at home level is increased after the in-home optimisation in Summer. We observe that for most of these prosumers the percentage of self-consumption is larger after this optimisation (increments up to 16% are observed in Fig. 4). We also observe how 4 out of the 5 prosumers with more flexibility reach after the in-home optimisation total self-consumption. Hence, the community optimisation will not be able to further optimise the self-consumption of these prosumers but they can still get cost reductions as a result of shifting their flexible loads to meet the exceeding production at community level. For Winter each prosumer already consumes most of its local PV production after in-home optimisation, leaving little surplus to trade in the community. In more detail, 9/29 prosumers were already reaching a total self-consumption before any optimisation whereas increments up 27% were obtained as a result of home optimisation. Hence, Summer months in the Cardiff dataset present a much more interesting scenario for the community optimisation than not Winter ones.

8.3. Local community optimisation results

This section evaluates the results from applying our approach at the level of the local community. In community optimisation, as a result of the trading between individual houses, prosumers' controllable assets may be shifted to follow the excess of community local production, incentivised by participants economic benefits (every unit of energy produced and consumed at the community has a gain). We took advantage from the schedules obtained for in-home devices after in-home optimisation using them as a warm start point for the ADMM algorithm at community level (this standard trick is known to often reduce the number of iterations needed by the algorithm [13, 3, 25]).

Fig. 5 depicts the amount of daily energy (kWh/day) imported from (in yellow) and exported to (in green) the community by each prosumer for Summer profiles (corresponding to a total of 665 kWh/day at community level). As expected prosumers with self-production and storage are the ones with more units exchanged since in addition of their flexibility they can also exchange their surplus of self-produced energy. In terms of economical benefits, these local exchanges lead to a total reduction of 589£/month at community level and to

[11] All tests were performed on a 2.5-GHz Intel Core i5 with 16GB of RAM.

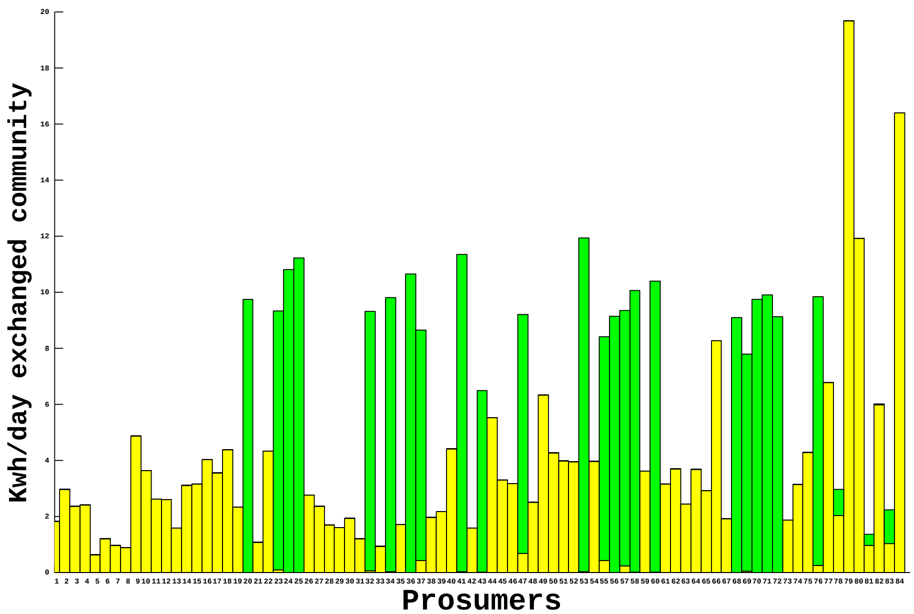

Figure 5: Daily energy imported from (in yellow) and exported to (in green) the community by each prosumer. Prosumers are ordered by their level of flexibility.

a gain per kWh exchanged within the community (Δ_{unit}) of 0.03£/kWh. How this gain per kWh exchanged is shared between consumption and production depends on the particular choice of α. For example, by setting $\alpha = \frac{1}{2}$ the unit consumed from/produced for the community will have the same reduction of 0.015£/kWh whereas by setting $\alpha = \frac{2}{3}$ the unit produced by the community is rewarded twice as much as the unit consumed from the community. Unlike in individual home-optimisation, here all prosumers will get some reduction on the energy bill since all of them consume from and/or produce for the community (reductions on the energy bill vary from 0.56£/month to 21.12£/month setting $\alpha = \frac{2}{3}$). Also, for the majority of participants this gain is much larger than the one obtained from optimising consumption at home level. Instead, for winter profiles the benefits of community optimisation in comparison with the Summer profiles were significantly diminished (a total of 200 kWh/day of units exchanged within the community leading to a total community reduction of 137£/month) as a result of a lower PV production which was mostly already consumed at home level after in-home optimisation.

The total number of ADMM iterations needed to perform this community optimisation was 2787 for Summer profiles corresponding to an execution time of around 3 minutes[12]. For the experiments with Winter profiles results were

[12]The test was performed on a cluster each agent being assigned to a different execution node with a 2.5-GHz Intel Core i5 with 16GB of RAM.

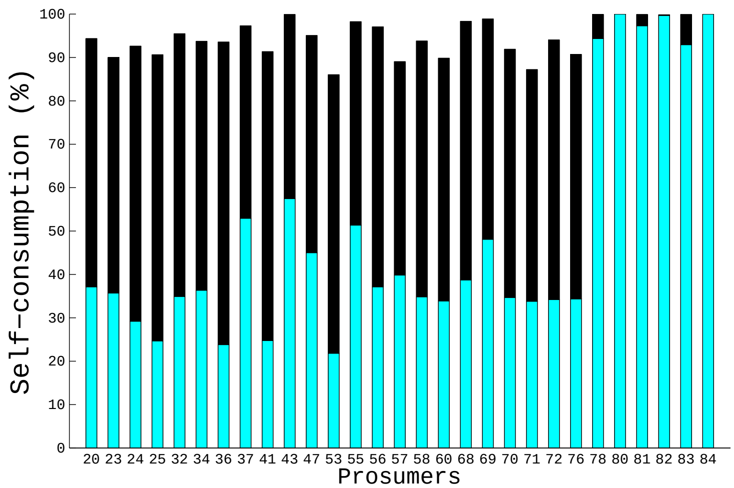

Figure 6: Self-consumption at home level (light blue) and at community level (black) for each of the prosumers with local production after community optimisation for Summer profiles.

similar (i.e. 2720 iterations needed).

Fig. 6 shows for the twenty-nine prosumers with local production the self-consumption level at home level and at community level after community optimisation for Summer profiles. We observe that the self-consumption percentages at home level are, as expected, similar to the ones obtained after the in-home optimisation (even slightly lower since it may be more profitable for prosumers to export to the community and use instead the energy from their storage). Nevertheless, such percentages are much lower than the percentages of self-consumption at community level obtained after the community optimisation. In particular, Fig. 6 shows that the percentages of self-consumption at community level for Summer profiles are extremely high with values that vary between 86 and 100%, depending on the prosumer (this corresponds to increments up to 59% when compared with home self-consumption levels from Fig. 4). The results obtained for Winter profiles showed how the community reached total self-consumption, i.e. all the excess of self-generated PV energy (substantially less than in Summer) was traded at community level.

9. Conclusions and future work

The EU Clean Energy package adopted in 2019 opens way for a major transition of the European energy landscape towards customer empowerment and local energy communities. This paper takes as input this new regulatory environment to elaborate on the cooperative energy network model of LECs. As a result of this analysis, six design requirements were identified (two at household level and four at community level) as for LEC optimising technologies to comply with these new directives. The list of requirements included rights of community members such as the right to trade their self-produced energy (and only their self-produced energy) within the community and the right to choose their own supplier as well as economical requirements that state their minimum level of compensation. We show how our model of cooperative energy network handles all these overriding restrictions for LECs operation.

A multi-agent coordination protocol, based on the well-known ADMM method, has been executed over such cooperative energy network in order to determine the most effective energy exchanges within the LEC as well as the corresponding configuration of flexible and controllable devices in a decentralised way (i.e. each agent solves its own local optimisation problem in a coordinated manner). To efficiently solve the local optimisation agent problems as to participate in the coordination protocol we also formulated low computational cost procedures for all the seven agent types (two net types and five devices types) present in the validation dataset. In particular, to satisfy the requirement that community members should be able to trade their self-produced energy, and only their self-produced energy, within the community our model included a novel type of net agent which in addition to ensure the traditional energy balancing constraints also restricts any transfer of energy between the supplier and the community differentiated connections. To solve the local optimisation problem for this new type of agent, we also proposed a faster computational procedure by casting it as a variation of the Euclidean projection onto the simplex and exploiting the structure of the local constraints. Finally, we proposed a payment mechanism that distributes the gains of the community guaranteeing that each member would be better off by joining the community, compared to optimising its household individually. However, as highlighted in the paper, our model is not restricted to such payment mechanism and indeed seeing the effect on the fairness and level of incentivation when our model is combined with more complex payment mechanisms existing in the literature is an open line of work.

Our approach was evaluated on a dataset based on an existing low-voltage grid in Cardiff (UK) composed of 84 residential customers with smart appliances, 29 of which with photovoltaic generation and residential energy storage. Extensive simulations show that our approach applied at community level achieves higher self-consumption ratios (up to 59% increment w.r.t. self-consumptions ratios obtained when optimising houses individually) and significantly reduces community expenses (a total of 589£/month reduction expected on summer months w.r.t. no community bills). Our results also showed that whereas prosumers in the Cardiff dataset may benefit significantly from performing optimisation at house level throughout the year, the benefits from community optimisation would get significantly reduced during Winter as a result of the lower PV generation that gives little surplus to trade at community level. Hence, we conclude from our validation that if community optimisation should be profitable all over the year the sources of local renewable energy should be diversified and not be exclusively restricted to PV as in the Cardiff dataset.

As part of future work we aim to extend our model to consider a third optimisation level as in [15, 22] in which the community is connected to the wholesale market and/or to other communities, enabling participants to sell their energy surplus in there if they cannot find a buyer from their own community, or offer flexibility services (e.g. manage of the imbalance risk, congestion management,...). Moreover, notice that even when some of the device agents in our model (i.e. the photovoltaic and fixed load devices) use as input consumption/production forecasted values, we do not tackle in this paper how to make our approach resilient to possible deviations in such provided forecast values. Hence another line of future work is to extend our experiments to consider such deviations by applying receding horizon control [13] or/and by updating to robust extensions of the ADMM algorithm [2] resilient to forecast deviations.

Acknowledgements

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

[1] A. Anees, T. Dillon, and Y.-P. P. Chen. A novel decision strategy for a bilateral energy contract. *Applied Energy*, 253:113571, 2019.

[2] A. Attarha, P. Scott, and S. Thibaux. Affinely adjustable robust admm for residential der coordination in distribution networks. *IEEE Transactions on Smart Grid*, 11(2):1620–1629, 2020.

[3] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011.

[4] R. Denysiuk, F. Lilliu, M. Vinyals, and D. R. Recupero. In *Proceedings of the 12th International Conf. on Agents and Artificial Intelligence, ICAART 2020, Volume 1, Valletta, Malta, February 22-24, 2020*, pages 28–39. SCITEPRESS, 2020.

[5] J. C. Duchi, S. Shalev-Shwartz, Y. Singer, and T. Chandra. Efficient projections onto the l_1-ball for learning in high dimensions. In *Proceedings of the 25th Int. Conference on Machine Learning (ICML 2008), Helsinki, Finland, June 5-9, 2008*, pages 272–279, 2008.

[6] EU Project. Mas^{2}tering: Multi-agent systems and secured coupling of telecom and energy grids for next generation smartgrid services, 2014-2017. `http://www.mas2tering.eu/`.

[7] European Commission. *Directive (EU) 2019/944 on common rules for the internal market for electricity*, 2019. Available at `https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32019L0944`.

[8] European Commission. *Regulation on the internal market for electricity (EU) 2019/943*, 2019. Available at `https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32019R0943&from=EN`.

[9] G. Fridgen, M. Kahlen, W. Ketter, A. Rieger, and M. Thimmel. One rate does not fit all: An empirical analysis of electricity tariffs for residential microgrids. *Applied Energy*, 210:800–814, 2018.

[10] J. Guerrero, A. C. Chapman, and G. Verbi. Decentralized p2p energy trading under network constraints in a low-voltage network. *IEEE Transactions on Smart Grid*, 10(5):5163–5173, 2019.

[11] J. Guerrero, D. Gebbran, S. Mhanna, A. C. Chapman, and G. Verbi. Towards a transactive energy system for integration of distributed energy resources: Home energy management, distributed optimal power flow, and peer-to-peer energy trading. *Renewable and Sustainable Energy Reviews*, 132:110000, 2020.

[12] M. Khorasany, Y. Mishra, and G. Ledwich. A decentralized bilateral energy trading system for peer-to-peer electricity markets. *IEEE Transactions on Industrial Electronics*, 67(6):4646–4657, 2020.

[13] M. Kraning, E. Chu, J. Lavaei, and S. P. Boyd. Dynamic network energy management via proximal message passing. *Foundations and Trends in Optimization*, 1(2):73–126, 2014.

[14] W. Lee, L. Xiang, R. Schober, and V. W. S. Wong. Direct electricity trading in smart grid: A coalitional game analysis. *IEEE Journal on Selected Areas in Communications*, 32(7):1398–1411, 2014.

[15] F. Lilliu, A. Loi, D. Reforgiato Recupero, M. Sisinni, and M. Vinyals. An uncertainty-aware optimization approach for flexible loads of smart grid prosumers: A use case on the cardiff energy grid. *Sustainable Energy, Grids and Networks*, 20:100272, 2019.

[16] W. Liu, D. Qi, and F. Wen. Intraday residential demand response scheme based on peer-to-peer energy trading. *IEEE Transactions on Industrial Informatics*, 16(3):1823–1835, 2020.

[17] F. Luo, Z. Y. Dong, G. Liang, J. Murata, and Z. Xu. A distributed electricity trading system in active distribution networks based on multi-agent coalition and blockchain. *IEEE Transactions on Power Systems*, 34(5):4097–4108, 2019.

[18] K. A. Melendez, V. Subramanian, T. K. Das, and C. Kwon. Empowering end-use consumers of electricity to aggregate for demand-side participation. *Applied Energy*, 248:372–382, 2019.

[19] E. Mengelkamp, P. Staudt, J. Garttner, and C. Weinhardt. Trading on local energy markets: A comparison of market designs and bidding strategies. In *14th Int. Conference on the European Energy Market (EEM)*, pages 1–6. IEEE, 2017.

[20] M. Mihaylov, R. Rdulescu, I. Razo-Zapata, S. Jurado, L. Arco, N. Avellana, and A. Now. Comparing stakeholder incentives across state-of-the-art renewable support mechanisms. *Renewable Energy*, 131:689–699, 2019.

[21] D. K. Molzahn, F. Dörfler, H. Sandberg, S. H. Low, S. Chakrabarti, R. Baldick, and J. Lavaei. A survey of distributed optimization and control algorithms for electric power systems. *IEEE Transactions on Smart Grid*, 8(6):2941–2962, 2017.

[22] F. Moret and P. Pinson. Energy collectives: a community and fairness based approach to future electricity markets. *IEEE Transactions on Power Systems*, 34(5):3994–4004, 2018.

[23] E. Munsing, J. Mather, and S. Moura. Blockchains for decentralized optimization of energy resources in microgrid networks. In *2017 IEEE Conference on Control Technology and Applications (CCTA)*, pages 2164–2171, 2017.

[24] P. A. Narbel. Rethinking how to support intermittent renewables. *Energy*, 77:414–421, 2014.

[25] P. Scott and S. Thiébaux. Distributed multi-period optimal power flow for demand response in microgrids. In *Proceedings of the 6th Int. Conf. on Future Energy Systems, e-Energy, Bangalore, India, July 14-17*, pages 17–26. ACM, 2015.

[26] M. Sisinni, I. Grimaldi, C. Ponchel, S. Breton, D. Pires, V. Giordano, T. Messervey, and J. Espeche. D6.1 - detailed use cases scenarios, 2016. `http://www.mas2tering.eu/papers-documents-tools/`.

[27] E. Sorin, L. Bobo, and P. Pinson. Consensus-based approach to peer-to-peer electricity markets with product differentiation. *IEEE Transactions on Power Systems*, 34(2):994–1004, 2019.

[28] T. Sousa, T. Soares, P. Pinson, F. Moret, T. Baroche, and E. Sorin. Peer-to-peer and community-based markets: A comprehensive review. *Renewable and Sustainable Energy Reviews*, 104:367–378, 2019.

[29] W. Tushar, T. K. Saha, C. Yuen, T. Morstyn, M. D. McCulloch, H. V. Poor, and K. L. Wood. A motivational game-theoretic approach for peer-to-peer energy trading in the smart grid. *Applied Energy*, 243:10–20, 2019.

[30] W. Tushar, T. K. Saha, C. Yuen, D. Smith, and H. V. Poor. Peer-to-peer trading in electricity networks: An overview. *IEEE Transactions on Smart Grid*, 11(4):3185–3200, 2020.

[31] W. Tushar, C. Yuen, H. Mohsenian-Rad, T. Saha, H. V. Poor, and K. L. Wood. Transforming energy networks via P2P energy trading: The potential of game-theoretic approaches. *IEEE Signal Processing Magazine*, 35(4):90–111, 2018.

[32] W. Tushar, C. Yuen, T. K. Saha, T. Morstyn, A. C. Chapman, M. J. E. Alam, S. Hanif, and H. V. Poor. P2P energy systems for connected communities: A review of recent advances and emerging challenges. *Applied Energy*, 282:116131, 2021.

[33] R. Verschae, T. Kato, and T. Matsuyama. Energy management in prosumer communities: A coordinated approach. *Energies*, 9(7), 2016.

[34] M. Vinyals, M. Velay, and M. Sisinni. A multi-agent system for energy trading between prosumers. In *Proceedings of the 14th International Conference on Distributed Computing and Artificial Intelligence*, pages 79–86, 2018.

[35] Y. Wang, K. Lai, F. Chen, Z. Li, and C. Hu. Shadow price based coordination methods of microgrids and battery swapping stations. *Applied Energy*, 253:113510, 2019.

[36] C. Zhang, J. Wu, Y. Zhou, M. Cheng, and C. Long. Peer-to-peer energy trading in a microgrid. *Applied Energy*, 220:1–12, 2018.