Modelling the formation of peer-to-peer trading coalitions and prosumer participation incentives in transactive energy communities

Ying Zhang [a,b], Valentin Robu [a,c,f,*], Sho Cremers [a,d], Sonam Norbu [e], Benoit Couraud [e], Merlinda Andoni [e], David Flynn [e], H. Vincent Poor [f]

[a] *Intelligent and Autonomous Systems Group, CWI (National Research Centre for Mathematics and Computer Science), Amsterdam, The Netherlands*
[b] *Algorithmics Group, Delft University of Technology, Delft, The Netherlands*
[c] *Electrical Engineering Department, Eindhoven University of Technology, Eindhoven, The Netherlands*
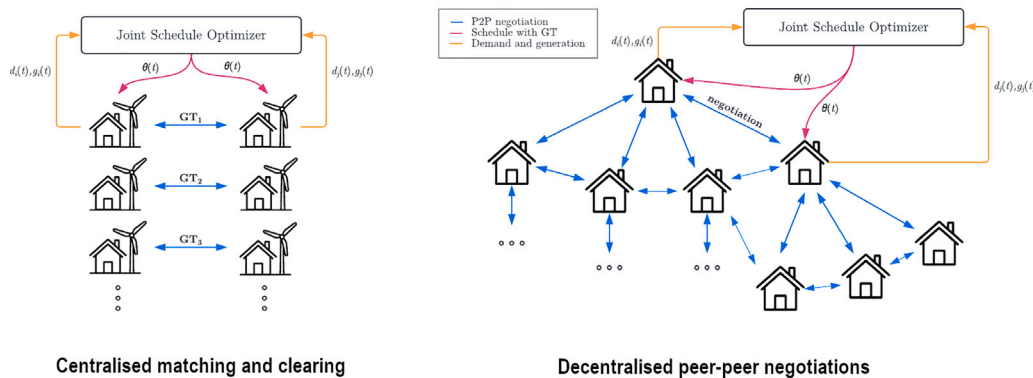[d] *Intelligent Electrical Power Grids Group, Delft University of Technology, Delft, The Netherlands*
[e] *James Watt School of Engineering, University of Glasgow, Glasgow, UK*
[f] *Electrical and Computer Engineering, Princeton University, Princeton, NJ, USA*

GRAPHICAL ABSTRACT

Market Models for Transactive Energy Communities

Centralised matching and clearing

Decentralised peer-peer negotiations

ARTICLE INFO

ABSTRACT

Peer-to-peer (P2P) energy trading and energy communities have garnered much attention over in recent years due to increasing investments in local energy generation and storage assets. Much research has been performed on the mechanisms and methodologies behind their implementation and realisation. However, the efficiency to be gained from P2P trading, and the structure of local energy markets raise many important challenges. To analyse the efficiency of P2P energy markets, in this work, we consider two different popular approaches to peer-to-peer trading: centralised (through a central market maker/clearing entity) vs. fully decentralised (P2P), and explore the comparative economic benefits of these models. We focus on the metric of Gains from Trade (GT), given optimal P2P trading schedule computed by a schedule optimiser. In both local market models, benefits from trading are realised mainly due to the diversity in consumption behaviour and renewable energy generation between prosumers in an energy community. Both market models will lead to the most promising P2P contracts (the ones with the highest Gains from Trade) to be established first. Yet, we find diversity

decreases quickly as more peer-to-peer energy contracts are established and more prosumers join the market, leading to significantly diminishing returns. In this work, we aim to quantify this effect using real-world data from two large-scale smart energy trials in the UK, i.e. the Low Carbon London project and the Thames Valley Vision project. Our experimental study shows that, for both market models, only a small number of P2P contracts i.e. less than 10% of the possible P2P contracts are required to achieve the majority of the maximal potential Gains from Trade. Similarly, only a fraction of prosumers are required to participate in energy trading to realise significant GT; namely we found that 60% of the maximal GT can be realised with only 30% of prosumers' participation, with the percentage of maximal GT reaching 80% when participation increases to 50% of prosumers. Finally, we study the effect that diversity in consumption profiles has on overall trading potential and dynamics in an energy community. We show that in a community with a DF(load diversity factor) = 1, 80% of potential maximal GT can be achieved by 10% of prosumers engaging in P2P trading, while in a community with DF = 1.5, it is beneficial for 40% of the prosumers to trade.

1. Introduction

Energy systems around the world are experiencing rapid changes. Until recently, energy systems were primarily centralised networks managed by large utility companies. However, currently a paradigm shift is taking place in energy systems across the world [1–5]. Recent years have seen many initiatives for creating local energy communities, typically low voltage networks consisting of 50–200 households [6]. Some notable examples of these initiatives include: the Brooklyn Microgrid project,[1] the REflex smart energy demonstrator on the Orkney Islands,[2] the many local projects managed by Community Energy Scotland,[3] the SchoonSchip Community in the north of Amsterdam (location of the GridFriends Project),[4] the Thames Valley Vision project,[5] the Low Carbon London demonstration project[6] etc. However, there are many open challenges and knowledge gaps in setting up such local markets, starting with the best choices for market organisation and establishing contracts. There are currently two main approaches to this problem [4,7].

The first approach involves centralised market clearing, in which a coordinator entity/market maker selects the most promising pairs of prosumers in the community to trade, and establishes contracts in order of decreasing gains from trade. The second approach involves purely decentralised peer-to-peer (P2P) systems where prosumers trade energy directly with one another — the selection of who to trade with is up to individual prosumers, and may be influenced by personal preferences, not just efficiency (e.g. trading energy with a neighbour may be preferable to trading with someone else in the community, who is not a close acquaintance).

The centralised approach of establishing contracts benefits from performing coordinated optimisation and aggregation, leading to efficient use of resources and optimal social allocation for the community. However, centralised approaches may experience scalability issues especially when the number of participating prosumers increases. On the other hand, peer-to-peer systems benefit from their decentralised nature. They can often be privacy preserving and secure, and benefit from the development of new supporting technologies such as decentralised energy contracts and blockchain systems. Furthermore, they allow for a more individual choice and strategy, where prosumers are allowed complete freedom of choice who to trade with, in order to optimise their own benefits.

While decentralisation provides many tangible benefits, it raises many open questions. While with enough trades peer-to-peer systems may approach similar community benefits from a social welfare perspective as the centralised market approach, each trade has some

associated cost, which may vary depending on various factors, e.g. the establishment and deployment of the decentralised platform, development, IT, security and administrative costs to name a few. This cost can make some trades with marginal benefits inefficient and not rational to set up. Hence, an important question is: How many P2P trading contracts are required to realise most of the potential benefits from trade? Another important question is that of prosumer participation. Recall that a community may be composed of up to several hundred prosumers. But are all of them required, and would they actually benefit from participation in P2P energy trading?

This issue is often unaddressed in the literature, leaving the reader to think that every member in a community will benefit from being involved in peer-to-peer trading [4,7,8]. However, the picture is more nuanced, after considering the different energy consumption profiles, physical constraints and the availability of home energy storage. In some cases, even if there are some marginal gains to be had, the potential costs of participating in a P2P trading scheme may outweigh the benefits. For example, consider 2 prosumers, each with their own micro-generation (e.g. solar panels on the roof) and battery. In theory, peer-to-peer trading should definitely provide benefits for these consumers, and this is certainly true at least some of the time. However, consider a setting in which their energy demand profiles are very closely aligned, e.g. both households have similar work patterns, that involve not much consumption during the morning/day, and a consumption peak in the evening. In such a case, P2P energy trading might deliver rather insignificant benefits for the two prosumers, because they will likely have excess energy to trade and residual demand at (mostly) the same time periods throughout each day.

Then, a question would be whether it would be possible to achieve a high percentage of the possible benefits (in our case, measured as fraction of maximal potential Gains from Trade) with only the most promising fraction of prosumers participating in peer-to-peer trading? Can we determine which prosumers would benefit the most, and which other prosumers should they trade with?

Finally, the value of bilateral energy contracts is dependent on the prosumers that establish them. Prosumers value energy quantities differently depending on their own consumption and generation profile with respect to others. The diversity of a community might affect the potential gains achieved from peer-to-peer trading. Thus, the effect of diversity factors in demand profiles across the community needs to be studied. While other authors have discussed these issues, to our knowledge, this paper is the first to examine these questions using a actual real data, from two large-scale UK trials, i.e. Thames Valley Vision and Low Carbon London.

To answer these questions, we consider two popular approaches to P2P markets. The first is a centralised approach to P2P that uses a central matching and clearing mechanism. The second is a decentralised approach using automated negotiation to establish P2P trading contracts. In both models, once a P2P trading contract is established, prosumers form a coalition, which produces a Gains from Trade for both parties (i.e. a financial benefit, in this case a reduction in bill from trading and operating assets jointly, compared to the bill when each agent uses only its own individual assets), and the agents split this

---

[1] https://www.brooklyn.energy/
[2] https://www.reflexorkney.co.uk/
[3] https://communityenergyscotland.org.uk/
[4] https://amsterdamsmartcity.com/updates/project/grid-friends
[5] https://eatechnology.com/resources/projects/new-thames-valley-vision-tvv/
[6] https://innovation.ukpowernetworks.co.uk/projects/low-carbon-london/

Nomenclature

Subscripts and Sets

K	[Sub]set of agents in the coalition ($K \subseteq N$)
N	set of all agents in the community considered in the model
i	for agents
j	for agents

Parameters

E_{bat_i}	Battery capacity installed at agent i premises
P_{PV_i}	Solar PV Power installed at agent i premises
T	Time horizon
λ	Lifetime of the generator
$\tau^b(t)$	Buying price (i.e. import tariff) at t [pence/kWh]
$\tau^s(t)$	Selling price (i.e. export tariff) at t [pence/kWh]
c^g_{kW}	Cost of the generator [pence/kW]
c^{bat}_{kWh}	Cost of the battery [pence/kWh]
dl	The deadline of negotiation
k	number of peers in the negotiation framework

Variables

Ω^a_i	Energy contracts accepted by agent i
Ω_i	Set of all possible energy contracts for agent i
Ω	Set of all possible energy contracts
Θ	Peer-to-peer oracle
ω	Energy contract tuple (i, j, θ)
$\theta(t)$	Energy amount traded at time t decided by Θ
θ	Energy amounts traded decided by Θ
$\tilde{b}_i(T)$	Accumulated bill for agent i at time T after trading
$\tilde{e}^b_i(t)$	Energy bought by agent i after peer-to-peer trading
$\tilde{e}^s_i(t)$	Energy sold by agent i after peer-to-peer trading
$\tilde{e}_i(t)$	Net demand for agent i after peer-to-peer trading
$b_i(T)$	Accumulated bill for agent i at time T
$b_N(T)$	Accumulated bill for community N at time T
$b_K(T)$	Accumulated bill for P2P coalition K at time T
$b_K(T)$	Accumulated bill for coalition K at time T
$c^g_i(T)$	Generator costs for agent i over T
$c^g_K(T)$	Generator costs for K over T
$c^{bat}_i(T)$	Battery degradation cost for agent i over T
$c^{bat}_K(T)$	Battery degradation cost for K over T

$d_i(t)$	Demand for agent i
$d_K(t)$	Aggregated demand for K
$e^b_i(t)$	Energy imported by agent i
$e^b_K(t)$	Energy imported by K
$e^s_i(t)$	Energy exported by agent i
$e^s_K(t)$	Energy exported by K
$e^{p2p}_i(t)$	Energy traded by agent i
$e_i(t)$	Net demand for agent i
$g_i(t)$	Generation for agent i
$g_K(t)$	Aggregated generation for K
$o_i(r)$	The offer that agent i proposes in round r
$p^{bat}(t)$	Battery power at time t
rv_i	The reservation value of agent i, i.e. the minimum utility the agent will accept.
r	The round number of negotiation
t	time step
u^{max}_i	The maximum utility that can be gained by agent i from trading
$v_i(\omega)$	Value of the contract ω gained by agent i

Acronyms

DF	Depreciation Factor
GT	Gains from Trade
$P2P$	Peer-to-peer
RES	Renewable Energy Source
SoC	State of Charge of a battery in kWh

benefit. The exact way energy is exchanged for each clearing period (in our case, we use half hourly clearing periods) throughout the day will be determined by a joint schedule optimiser, that takes into account the demand, generation and assets of both parties.

Our methodology allows modelling realistic peer-to-peer trading using large-scale data sets. Due to computational complexity, prior works often consider small communities with a few prosumers over short time windows, and without the use of large-scale, real world datasets. We argue a more detailed and data-driven approach is needed to provide insights in the effectiveness of peer-to-peer trading in real-world settings. Thus, in our framework, the benefit of energy exchanges are computed using a joint schedule optimiser. This optimiser computes the trading schedule with the maximal Gains from Trade for both parties. The prosumer agents themselves then decide whether to trade, and the redistribution of these benefits. In summary, our paper makes a number of the key new contributions:

- It performs a systematic comparison of two market models: centralised clearing (double auction-type format) vs. decentralised negotiations, in terms of dynamics and trading coalitions that form over time. In order to enable this comparison, we introduce a new negotiation framework, that incorporates a joint optimiser that computes the most profitable energy exchanges between any set (coalition) of prosumers in terms of the notion of Gains from Trade (i.e. joins that can be obtained by reaching a P2P contract and forming a coalition, compared to trading alone). This framework can be used in both types of markets, allowing us to model separately the market organisation layer (centralised clearing vs. negotiation) from the asset control and optimisation layer.

- In most energy community and micro-grid energy schemes, not all consumers are required to participate in peer–peer trading to reach market efficiency (in terms of potential Gains from Trade). In order to quantify this effect, the paper investigates whether the majority of the benefits can be achieved when only a subset of

the most promising prosumers participate in peer-to-peer trading. Moreover, it includes a study of how many peer-to-peer contracts are needed to realise the benefits of energy trading in a real P2P market. Our analysis also includes an evaluation of diversity in demand profiles affects the potential benefits and the marginal value of the contracts.

- Finally, the developed simulation framework can use large-scale real data, and allows for a more realistic negotiation simulation with large communities over longer time periods, with high time granularity than possible in other prior works. Our results and conclusions are based on real demand and generation data sets from two large-scale energy trials in the UK.

The remainder of the paper is organised as follows: In Section 2, relevant literature is presented. Section 3 outlines the models used for the prosumers, energy communities and $P2P$ trading. Two approaches to peer-to-peer trading are outlined in the following two sections, with Section 4 describing the central market clearing mechanism, and Section 5 discussing the negotiation framework. In Section 6, we perform an empirical analysis of the two methods compared to the community setting using realistic energy communities. Finally, Section 7 concludes the paper and highlights some topics for future work.

2. Related work

Energy communities form an increasingly important area of research, with many earlier works studying different aspects and challenges. In recent years a number of projects and initiatives have published systematic reviews of the field. Tushar et al. [1] review the state of the art in connected energy communities, provide extensive background on different aspects of P2P sharing, and describe a number of relevant pilot projects across the globe. Capper et al. [4] provide a systematic review and classification of market design and energy trading models based on 139 recent peer-reviewed journal articles, in research sponsored by the global observatory on peer-to-peer, community self-consumption and transactive energy models.[7] Sousa et al. [7] provide a classification of P2P energy trading markets based on the degree of decentralisation and P2P topology into: full P2P markets, community-based markets and hybrid markets (combination of the two designs, each operating at different layers). Schwidtal et al. [5] discuss the feasibility of different business models for profitable energy community and transactive energy projects.

In this paper, following prior literature, we consider two prominent market models. First, we model markets that use centralised clearing, through a market maker or other centralised entity, that selects the contracts with the optimal Gains from Trade. Double auction is a common mechanism used within this approach, widely used in many applications ranging from wholesale energy markets to the stock market. Being widely understood, it is also the most studied mechanism for local energy markets and has seen many practical applications. Sioshansi [9] discusses double auction mechanisms and their use in electricity markets, while Wang et al. [10] and Capper et al. [4] provide examples of double auction use in P2P energy markets and energy communities. In double auctions, buyers and sellers relay their preferences and utility to a central entity or broker. Then, this central entity will decide the matching of parties, price and will clear the market, which guarantees market efficiency. Wang et al. [10] show that such an auction mechanism can provide these benefits while preserving privacy using a two-level transactional model. Our centralised mechanism matches, essentially, a double auction, where contracts are sorted and cleared in order of decreasing Gains from Trade (GT).

The second model of market organisation we consider is a decentralised peer-to-peer energy market, that allows direct trading and

establishing of contract between prosumers, forgoing the need for an intermediary. As opposed to markets relying on a clearing entity/market maker, peer-to-peer markets allow prosumers to select who they trade with — for example they could make or request offers with several other prosumers of their own choice (e.g. neighbours, friends in the community etc.), and select the most favourable one.

The most widely used mechanism to automate decentralised P2P markets is automated negotiation. Pinto et al. [11] proposed a multi-agent simulation platform to support negotiations of small players in transactive electricity markets, including local energy markets, bilateral contracts and participation in wholesale energy markets. Another example of automated negotiations is provided by Saxena and Abhyankar [12] focusing in electricity markets for distribution grid management. In automated negotiations, prosumers are represented by autonomous agents that negotiate over several topics, e.g. price per unit, energy quantities, etc. Often, the main topic negotiation is concerned about is the price per unit, such as in the works by Guo et al. [13] and Imran et al. [14]. Etukudor et al. [15] discuss a framework where prosumers negotiate the prices for a single day ahead together with the energy quantities traded for multiple periods within the day. Another approach is discussed by Chakraborty et al. [16], who explore an energy lending scheme, where prosumers negotiate over energy quantities that they lend and the time the borrowed energy is returned.

Since peer-to-peer negotiation only considers bilateral agreements, as opposed to an auction mechanism where a market-wide price is set, an important process is also peer selection, i.e. how a prosumer decides with whom to negotiate. Concurrent negotiation approaches have been explored, where sellers negotiate with multiple sellers concurrently and vice versa [17]. However, due to computational complexity, systems favour a selection of peers before the negotiation process. A popular approach is the facilitation of peer selection by a central institution (often the platform provider) [18]. Khorasany et al. [19] propose a single peer selection approach, where prosumers select peers greedily based on maximum expected profit. Other recent works focus on optimising economic objectives when taking physical constraints and storage into account [20], and the fair redistribution of benefits in energy coalitions. Long et al. [21] provide a P2P trading scheme for a community microgrid that focused on the optimality and fairness between participating prosumers. Norbu et al. [22] provide a model for the fair redistribution of benefits achieved by community-owned assets to individual prosumers. Robu et al. [23] look at the fair allocation of gains achieved by virtual power plant formation, while Kota et al. [24] study this issue for demand-side management cooperatives. The concept of Shapley value has been used to ensure fairness, such as in the works of Cremers et al. [25] and Robu et al. [26]. Finally, other works take a social science perspective, and examine the perception of consumers to P2P energy trading schemes [27].

However, although great progress has been made by related works, there are still key remaining questions. Firstly, many peer-to-peer approaches, especially negotiation, suffer from high complexity, and generalising them to large communities has been difficult [28]. Thus, the main contribution of earlier works often lies in their methodologies, since their experiments consider only a small simulation-based setup, and often do not rely on real prosumer demand data. Further research into how these models generalise to larger settings is required. Secondly, while peer-to-peer markets will approach the efficiency of a central market when enough contracts have been established, each contract has some associated costs. Hence, how many P2P energy trading contracts need to be established to achieve the majority of the maximal potential Gains from Trade is still an open question. Thirdly, these works focus on the total economic or technical benefits for either the community or an individual prosumer. However, within communities, especially those based on locality, some prosumers benefit to different degrees (i.e. asymmetrically) from peer-to-peer trading. Efforts have been made to ensure Pareto optimality, i.e. that no prosumer is worse

---

[7] https://userstcp.org/task/peer-to-peer-energy-trading/

off in favour of another prosumer. But the benefits earned by each prosumer are not equal, since they depend on their size, and crucially on the energy consumption and generation behaviour of the prosumer compared to those of the community, e.g. a prosumer that consumes with a different pattern is likely will benefit more. Thus, there is a need to investigate the effects of prosumer participation and diversity on the benefits of trading. This paper aims to model these settings – and thus address these challenges – using a large-scale real-world data set.

3. Modelling prosumers in an energy community

In this section, we first introduce a formal model of a prosumer, by modelling their respective demand and generation profile. Prosumers are also assumed to have a source of energy storage (i.e. a battery).

3.1. Single prosumer model

A prosumer $i \in N$, where N are the prosumers connected to the local energy grid, has access to the battery and some renewable energy resource (RES) (e.g. wind turbine). The generation profile of this prosumer i is represented as $g_i(t), \forall t \in [1, T]$ for the power generated at time step t, where T corresponds to the time horizon for the operation window. Similarly, the demand profile is represented as $d_i(t), \forall t \in [0, T]$. Furthermore, for each prosumer, we compute the optimal battery usage $p^{bat}(t)$, where a positive value indicates charging and a negative value discharging. Finally, a prosumer also has access to the power grid and is able to export and import energy, represented as $e_i(t)$, where it is positive when importing from the grid and it is negative when exporting to the grid.

The energy demand of the prosumer should always be satisfied, either due to self-consumption or by buying power from the grid. Furthermore, excess energy should be either sold or stored. So, the following constraint should always hold:

$$e_i(t) = d_i(t) - g_i(t) + p^{bat}(t) \qquad (1)$$

3.2. Battery control algorithm

Battery usage at each time step was controlled using a heuristic-based algorithm adapted from Norbu et al. [22]. Battery charging and discharging follow directly from the energy generation and consumption. If at some time, the generation exceeds the consumption the remaining energy is used for charging the battery. However, if the battery has reached its maximum capacity, excess energy is instead sold to the grid.

Instead, if consumption exceeds the generation, we instead discharge the battery. If at any time the battery cannot provide sufficient energy, energy is bought from the central power grid.

This battery control algorithm provides a tractable method to approximate an optimal control strategy. Our heuristic-based strategy performs well (essentially close to optimal) under the assumption that flat import and export tariffs are used and that using energy storage is more cost-effective than importing and exporting (see [29] for a discussion). In the current energy market, these assumptions are realistic, since import tariffs are high, while export tariffs are being phased out [30]. Given this aspect is technically complex, and not the main focus of this paper, we provide the full details of the battery control algorithm we use in our experiments in Appendix A.

3.3. Cost computation

The costs for a single prosumer can be calculated as the sum of the costs of importing energy from the grid, the revenue from selling excess energy to the grid, and the depreciation costs of the private energy

assets, i.e. the battery and the generator. It can be computed as follows:

$$b_i(T) = \sum_{t=1}^{T} e_i^b(t)\tau^b(t) - \sum_{t=1}^{T} e_i^s(t)\tau^s(t) + c_i^{bat}(T) + c_i^g(T) \qquad (2)$$

where the $e_i^b(t)$ and $e_i^s(t)$ are the energy bought from and sold to the grid at time t respectively. $\tau^b(t)$ and $\tau^s(t)$ are the import and export tariffs, and $c_i^{bat}(T)$ and $c_i^g(T)$ are the depreciation costs for the battery and RES respectively. Note that, in our model, both individual prosumers and communities are modelled as price-takers w.r.t. the central grid or power distribution company, i.e. they are too small to influence the wholesale market prices. Hence, the import/export tariffs from the grid are inputs for our pricing model.

The depreciation cost of the RES is expressed as follows:

$$c_i^g(T) = \frac{P_{PV_i} \cdot c_{kW}^g \cdot T}{\lambda} \qquad (3)$$

where c_{kW}^g is the cost per rated kW and λ is the lifetime of the generator.

The depreciation cost of the battery can be computed as follows:

$$c_i^{bat}(T) = \frac{E_{bat_i} \cdot c_{kWh}^{bat} \cdot T}{\max\left(\frac{1}{DF}, \lambda\right)} \qquad (4)$$

where SoC_i^{max} is the battery capacity, c_{kWh}^{bat} is the price per kWh of capacity, and DF is the depreciation factor of the battery based on the usage, calculated using "rain-flow" cycle counting [22,31]. We include the depreciation factor in this calculation since the battery's lifetime can be influenced by the number of cycles and depth of discharge. See for full details on the computation of DF.

The cost computation in this section refers primarily to the use of energy assets and the resulting energy bill. In practical implementations, prosumers engaging in trading with peers in an energy community also face other costs, such as the development and maintenance of the energy trading platform, costs of control systems of energy assets, IT and administrative costs for contract deployment, etc. Furthermore, these costs depend on several factors, such as the technological solution selected for the trading platform, the complexity of contracts deployed, the security level and general market conditions. Due to this complexity, in this work we focus on the costs that are directly related to the energy bill and use of energy assets deployed.

3.4. Coalition model

The above energy community model corresponds to a coalitional model where N prosumers are players that can share their energy assets, resulting in a joint cost reduction. The joint cost function of this coalition is defined by the joint bill that prosumers that form the coalition will have to pay. Formally, we consider a coalition $K \subseteq N$ where each prosumer i has their own energy assets. However, once two agents agree to trade (enter a P2P trading contract), they form a coalition, where each prosumer can use the other members' assets if excess generation/storage capacity exists, and they are not being used by the party owning them. Thus, in this model, we can aggregate the generation, demand and storage to form the community's generation, demand and storage, i.e.

$$d_K(t) = \sum_{i \in K} d_i(t) \qquad \forall t \in [0, T] \qquad (5)$$

$$g_K(t) = \sum_{i \in K} g_i(t) \qquad \forall t \in [0, T] \qquad (6)$$

$$SoC_K^{max} = \sum_{i \in K} SoC_i^{max} \qquad (7)$$

Similarly, we can compute the bill for the energy community given these aggregate profiles as follows:

$$b_K(T) = \sum_{t=1}^{T} e_K^b(t)\tau^b(t) - \sum_{t=1}^{T} e_K^s(t)\tau^s(t) + c_K^{bat}(T) + c_K^g(T) \qquad (8)$$

Hence the term "coalition" in this paper denotes a set of agents that agree to exchange energy according to a schedule that optimises the use of their joint energy assets (battery, generation) and joint demand. The term is borrowed by cooperative or coalitional game theory, where coalitions are cooperative groups formed by agents, in order to increase the joint benefits achieved (or in this case, reduce the joint costs or bill). The combined battery and generation assets are then controlled jointly in an optimal way to satisfy their joint demand. This minimises their imports from the grid (or utility company) resulting in a mutual saving, we call the Gains from Trade (c.f. Section 3.6).

3.5. The peer-to-peer trading model

We consider a peer-to-peer model where a prosumer i trades energy with another prosumer j using a joint profile optimiser Θ, which computes the optimal energy exchange at each time step. The contents of the computation of these trades are dependent on the net demand of the prosumers. So we record all trades in *energy contracts* $\omega = (i, j, \theta)$, where i is the energy-receiving prosumer and j is the energy-providing prosumer participating in the trade, and θ are the energy trades found by the profile optimiser Θ for each time step $0..T$. We denote the set of trades by prosumer i as Ω_i^a. Note that a coalition can be formed through multiple contracts, as agents iteratively join and trade with existing prosumers inside the coalition.

The total energy exchanged at time t by a prosumer i is given by

$$e_i^{p2p}(t) = \sum_{(i,j,\theta) \in \Omega_i^a} \theta(t) - \sum_{(j,i,\theta) \in \Omega_i^a} \theta(t), \qquad \forall j \in K \tag{9}$$

The exchanged energy $e_i^{p2p}(t)$ may be positive or negative, in which case a positive value indicates prosumer i received more energy than sent to other prosumers at time t, while a negative value depicts i giving more energy than receiving.

The net demand of the prosumer after trading is given by

$$\widetilde{e_i}(t) = e_i(t) - e_i^{p2p}(t) \tag{10}$$

Similarly to Eq. (1), we can identify two regimes where $\widetilde{e_i^b}(t) = \widetilde{e_i}(t) > 0$, i.e. residual demand and where $\widetilde{e_i^s}(t) = \widetilde{e_i}(t) < 0$, i.e. excess generation.

The bill after trading can be computed as follows:

$$\widetilde{b_i}(T) = \sum_{t=1}^{T} \widetilde{e_i^b}(t)\tau^b(t) - \sum_{t=1}^{T} \widetilde{e_i^s}(t)\tau^s(t) + c_i^{bat}(T) + c_i^g(T) - \sum_{\omega \in \Omega_i^a} v_i(\omega) \tag{11}$$

where v_i is the function that gives the financial value gained by prosumer i from contract ω. This value is decided based on the contract mechanism. In this paper, we investigate two mechanisms, i.e. a centralised market, and bilateral, concurrent negotiations.

3.6. Gains from trade in energy trading

Due to the diversity of demand and energy profiles, aggregation of profiles and assets leads to each coalition requiring less energy import and exports from the central power grid. Furthermore, considering all the distributed residential batteries as a combined battery capacity that could be centrally controlled to optimise the gains of the community leads to more flexibility, further decreasing reliance on the grid, and reducing the joint bill of the community for the residual demand over some time period $b_N(T)$. Considering all distributed storage assets of the community as a single combined battery capacity controlled centrally in both the centralised community model and the peer-to-peer model leads to a cost reduction for the prosumers (as compared to the case where each prosumer uses and controls its own individual assets such as individual battery independently, and thus pays all the imported electricity to his/her main electricity supplier). Using terminology from market design/economics, we call this reduction in costs the 'Gains from Trade' (gains from trading energy between prosumers in the community), as opposed to using only one's own individual assets). Formally, the Gains from Trade (also abbreviated GT) for a community of N prosumers over a time horizon T is defined as:

$$GT_N(T) = \sum_{i \in N} b_i(T) - b_N(T) \tag{12}$$

where, for all agents $i \in N$, $b_i(T)$ denotes the bill of each agent, when controlling its assets individually, and $\widetilde{b_i}(T)$ denotes its bill after trading based on the schedule recommended by the joint profile optimiser. Intuitively, GT is the difference between the sum of the prosumers' bills to the utility company (for residual demand not covered by their own renewable generation) that agents would have to pay jointly, minus the bill for the whole community if they share and control their energy assets together (i.e. appearing a single prosumer to the utility company). Note that here $b_N(T) = \sum_{i \in N} \widetilde{b_i}(T)$, i.e. the sum of the bills of the N prosumers trading inside the community is necessarily the total bill of the community. It is always the case that $b_N(T) \leq \sum_{i \in N} b_i(T)$, hence the $GT \geq 0$ value is greater than zero, since in the worst case, the best control is for everyone to exclusively use their individual assets. However, in most realistic cases, there will be times (clearing periods) when pooling of assets is better.

The above definition is provided for a community of N agents. But actually, in peer-to-peer trading, we can apply this concept the same way for any subset (or coalition) of prosumers $K \subseteq N$ in this community trading over period T (i.e. any subset greater than $|K| \geq 2$ agents). The GT for any coalition K is

$$GT_K(T) = \sum_{i \in K} b_i(T) - b_K(T) \tag{13}$$

Note that unlike conventional game-theoretic approaches, the Gains from Trades that underpin the coalition formation are derived by the results of the optimisation process and energy assets utilisation, rather than a clearly mathematically defined characteristic function. This allows for a more realistic representation and pragmatic approach of the model developed. Furthermore, note that $GT_K(T) \geq 0$, because no P2P contract would be established, and no trading would occur if the agents are worse off by pooling their assets together in a coalition. Similarly, here $b_K(T) = \sum_{i \in K} \widetilde{b_i}(T)$, i.e. the sum of the bills of the K prosumers inside the coalition is the total bill of the coalition. Also important to note that the Gains from Trade for a subset (sub-coalition) of agents in the community will always be less or equal to that of the whole community, i.e. $GT_K(T) \leq GT_N(T), \forall K \subseteq N$). However, the growth of GT is definitely not linear or proportional to the number of prosumers joining the coalition K, in fact quite the opposite: for some prosumers/contracts, their effect on the gains from trade growth is high, while for some it is very marginal (for example, small consumers or with a non-diverse diverse demand profile, who have little scope to trade), In most of the experiments in this paper, it is a more general methodology to report the relative ratio:

$$\frac{GT_K(T)}{GT_N(T)} \text{ for possible coalitions } K \subseteq N$$

in percentage terms, rather than the raw GT numbers, that is often specific to the trial, location in the UK, or energy tariff levels in time period when the trial data was collected.

3.6.1. Joint profile optimiser

The goal of peer-to-peer trading is to maximise the Gain from Trade. However, since every prosumer has different demand and generation profiles, it is not trivial to find the best exchange that minimises the joint total cost. Furthermore, even if we find the optimal (maximal) Gains from Trade, we need to distribute these gains over the prosumers in the contract.

To address the optimisation of the energy exchange, we propose a method using a joint profile optimiser that can compute the energy that should be exchanged between two prosumers. Recall that in the setting considered in this paper, once consumers decide to trade, either

(a) Prosumer A profile after P2P trading

(b) Prosumer B profile after P2P trading
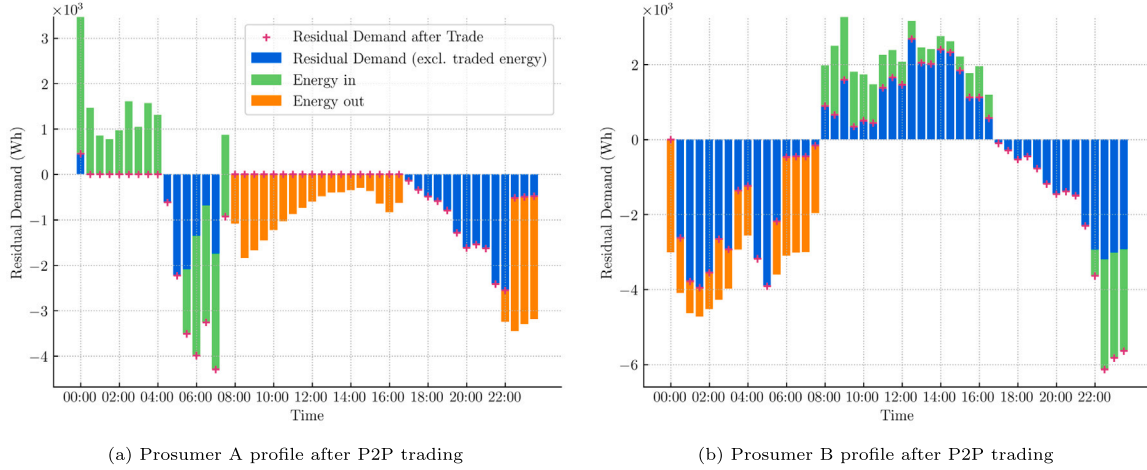
Fig. 1. Example of trading profiles between 2 prosumers: Prosumer A and Prosumer B, who decide to trade and form a coalition, over one example day (48 half-hourly time slots). Note that the imports (green sections) of one prosumer correspond to the orange sections (exports) of the other prosumer. In this case, it appears Prosumer A is a "night owl" type prosumer (i.e. a prosumer who is working and using energy during the night) - we note this is the real [anonymised] profile of a consumer from the Low Carbon London trial data. Important to note that, in the case of these two specific profiles substantial trading occurs, but for most profile pairs in the dataset there can be much less (or even no) P2P trading occurring, depending on the fit between the specific demand profiles.

through centralised or automated negotiation approach, they agree to share their assets, i.e. their flexible assets will be controlled by the Joint Profile Optimiser. Therefore, in both scenarios, consumers allow the Joint Profile Optimiser to use their excess energy generation/storage by other prosumers, at times they do not need it themselves. Hence, we can find the optimal energy contract, minimising the combined bill of prosumers in a coalition, using the aggregated demand, generation and battery of the prosumers in that coalition.

However, we still need to calculate the demand, generation and battery usage of the individual prosumers in the coalition, that would guarantee them the maximum Gains from Trade. This calculation requires satisfying a number of constraints.

First, the net demand at each time step given by

$$e(t) = d(t) - g(t) + p^{bat}(t) \tag{14}$$

for the coalition of two agents is equal to the sum of net demands of the individual agents after trading, i.e.

$$e_{ij}(t) = e_i(t) + e_j(t) \tag{15}$$

Since this would mean that the total energy between the two consumers individually is the same as the expected energy in the coalition.

Second, the state of charge (SoC) of the battery should be equal for the coalition and the sum of individuals, i.e.

$$SoC_{ij}(t) = SoC_i(t) + SoC_j(t) \tag{16}$$

These constraints combined ensure that the system of two individuals behave as a coalition, while preserving their individual profiles. To ensure that these constraints are satisfied we can use our battery control algorithm that balances the demand and generation of each prosumer at each time step.

The basic intuition of the profile optimiser is that, when two prosumer agents decide to trade, their optimal joint generation and demand (load) profile is the same as if they were a larger agent, controlling the assets of both parties. Then, this joint agent would perform optimal control w.r.t. the joint assets, and the aggregated load/generation of both prosumers. In Fig. 1 we provide an example illustration of trade that optimises the joint Gains from Trade between two prosumer agents, labelled Prosumer A and Prosumer B (note these are two real, anonymised prosumers, picked from the Low Carbon London dataset used in this paper). While the joint profile optimiser can operate on two individual prosumer agents (as in this example), it can also operate on the coalition of the two prosumers, which is the case considered in this

work when 2 agents agree to trade, either through centralised matching or through negotiation. In this case, the generation/demand profile used for the coalition is the aggregate of the generations/demands of all prosumers in that coalition. Thus, the profile optimiser can be applied to iteratively build larger coalitions, scalable to any number of agents.

4. Centralised matching and clearing

The centralised market mechanism tries to clear the market by accepting the contracts in order of most Gains from Trade that can be achieved by the agents trading (as computed by the joint schedule optimiser). For a visual overview of the process, see Fig. 2.

To find out which contracts are available, we compute all contracts in the contract space by computing all pairs of prosumers and the resulting trades θ produced by the joint profile optimiser Θ. The contract with the highest gains from trade is accepted and the involved prosumers trade their energy if the power flow resulting from the energy quantities in the contract do not violate any local grid constraint. This verification is done either directly by the DSO or by the Joint Schedule Optimiser. The gains of the accepted contract are split equally between the two prosumers. However, this results in the net demand profiles of these prosumers changing, i.e. the contract space changes. This requires a re-computation of the contract space before a new contract is accepted.

Hence, we can define a single round of accepting a contract as follows:

1. Compute the contract space given the prosumers net demands profiles $\widetilde{e}_i(t)$ ($e_i(t)$ for the initial round)
2. Sort the contracts space in descending order by potential Gains from Trade (GT)
3. Run a powerflow on the local network with the quantities from the trade with the highest potential GT
4. If the local grid constraints are met, accept the contract with the highest GT, otherwise, go the next contract
5. Split the gains equally between the two prosumers

This process continues until no more trades can be made, or the gains of trade drop below a specified threshold. Section 6 provides an overview of the gains of trades evolution for different numbers of rounds. Agents who establish trades with their peers form a coalition.

The proposed method attempts to minimise the number of trades (contracts) between individual prosumers. However, it can be extended to instead minimise the number of prosumers participating in trading.
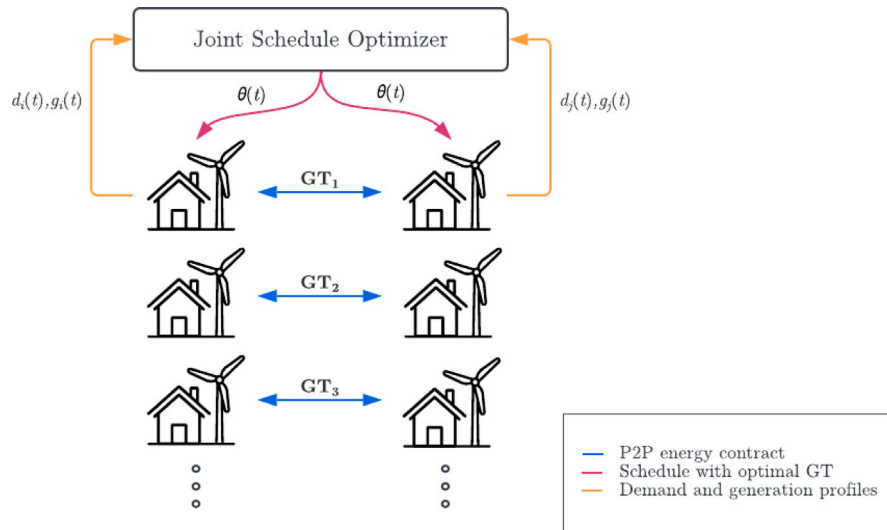
Fig. 2. Centralised matching and clearing; note that contracts are matched in order of decreasing Gains from Trade, i.e. $GT_1 \geq GT_2 \geq GT_3 \geq \cdots$.
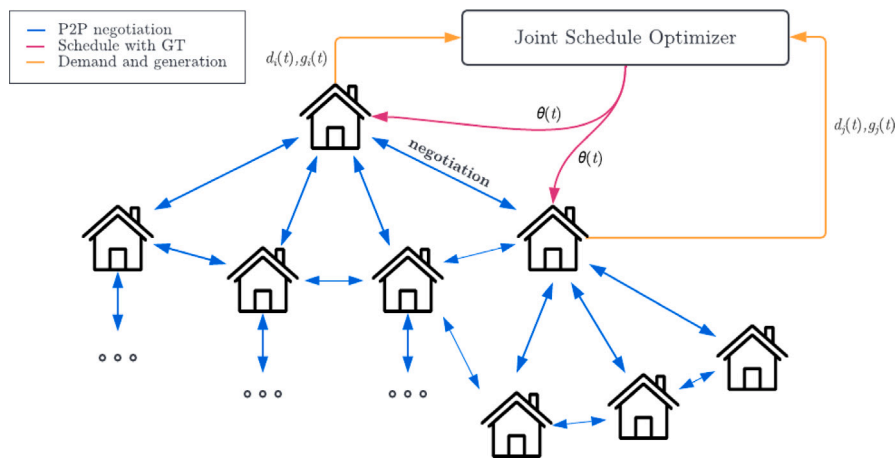

Fig. 3. Decentralised P2P negotiation between several prosumers, note that interaction between the joint profile optimiser and prosumers happen for each bilateral negotiation thread. In the figure, one example is shown.

We can model the trading coalitions that form around the agents that trade with each other. Note that, within these trading coalitions, not every member trades with each other, but only those who have established a trade contract. We can look at the gains those coalitions could realistically produce on their own and model them as a single prosumer for future contracts. This process will minimise the number of agents participating in trade since gains are maximised for groups first.

5. Peer-to-peer negotiation framework

Modelling automated P2P energy negotiations among multiple agents has proven complex [15,16,19], with earlier works often limiting the negotiation domain to a very small number of time steps and quantities. Our framework consists of two main steps, i.e. peer selection and negotiation. During the peer selection step, we select a number of prosumers based on the highest possible savings (Gains from Trade). Finally, during negotiation, the distribution of the gains is decided. These steps are then iterated over until no more new contracts can be established or a specified deadline has been reached. While earlier works have considered multi-issue negotiation over the amount of energy to be exchanged at every time point — resulting in an extremely large negotiation space, we instead consider a negotiation mechanism where agents can use the joint profile optimiser, that computes their

joint optimal energy exchange schedule. Prosumers, therefore, do not need to negotiate over how much energy needs to be exchanged at each time step, as in the proposed setting, the community embraces the best strategy proposed by the profile optimiser considering all batteries aggregated. A visual overview of the process is provided in Fig. 3).

The profile optimiser returns an exchange schedule and its expected joint Gains from Trade (compared to the sum of individual optimisations), but will not specify *how the financial gains will be divided* for each prosumer. Thus, prosumers will need to negotiate only over the *redistribution* of the expected gains from trade inside the coalition being formed. The introduction of a joint profile optimiser enables a reduction of complexity, and modelling larger domains. It allows us to consider trades over very long periods, with many time steps — and hence be able to build large data-driven simulations. While agents have to share their generation/consumption with the optimiser, they do not need to share their full generation/demand information with other agents directly, which preserves privacy.

5.1. Negotiation

The negotiation protocol is based on the concurrent alternating offers protocol. It consists of two phases: An offer phase and a commit phase. In the offer phase, agents keep exchanging offers until at least

one is accepted. While in the commit phase, a single most promising offer is accepted in the market.

5.1.1. Offer phase

The bidding phase is described in Algorithm 1. This phase consists of multiple rounds. During each round, every agent provides a price offer to multiple different agents along with the energy quantities defined by the joint profile optimiser. Note that the price offered by an agent to its partners is the same for all in a single negotiation round. The receiving agent then can accept or reject the offer. In case the offer is accepted, this is made public to other market participants, and a new coalition is formed by the agents reaching the deal. Otherwise, each bidding agent will create another offer during the next round. Note that, if a deal is accepted, the newly formed coalition can continue to trade in the market as a single player, to reach P2P deals with other prosumer agents and/or other coalitions, and the schedule optimiser will compute the Gains from Trade of the new coalition compared to the one already formed.

Algorithm 1 Offer exchange phase

Input: Contract space Ω, sets of potential trading partners P_i for each agent $i \in N$, agent strategy s_i for each agent i
Output: Set of accepted offers and contracts (can be more than one)
  $r \leftarrow 0$
  $\Omega^a \leftarrow \emptyset$
  $O^a \leftarrow \emptyset$
  while $r < deadline$ AND $\Omega^a = \emptyset$ **do**
    for all $i \in N$ **do**
      for all $j \in P_i$ **do**
        select $\omega = (i, j, \theta) \in \Omega$ ▷ Precomputed by the profile optimiser
        make-offer$_{i \rightarrow j}(r, v_i(\omega), s_i)$
      end for
    end for
    for all $o \in$ receive-offers$_i(r)$ **do**
      if accept-offer$_i(r, o, s_i)$ **then**
        select $\omega = (i, j, \theta) \in \Omega$ ▷ Find the contract belonging to the offer
        $\Omega^a \leftarrow \Omega^a \cup \{\omega\}$
        $O^a \leftarrow O^a \cup \{o\}$
      end if
    end for
    $r \leftarrow r + 1$
  end while
  return Ω^a, O^a

5.1.2. Commit phase

In the commit phase, shown in Algorithm 2, the market looks at all accepted offers and approves the one with the highest Gains of Trade after validation through a power flow analysis that the local grid constraints are not exceeded (voltage and lines thermal limits). If the offer with the highest gains results in the violation of grid constraints, the next highest gains' offer is selected and assessed through power flow study, until an offer is compatible with the grid state or until no offer remains. All other offers in that specific round are then rejected, to prevent large regret, as some players in the market have changed. Specifically, as a new coalition has been formed (by the agents with the highest GT exchange having reached a contract, and now participating as a coalition), this new coalition may offer a better deal to some other agent. Hence the acceptance process is iterative, with exactly one coalition being formed in each round. After acceptance, the payments

are settled and energy is exchanged. After the energy exchange the contract space is updated.

Algorithm 2 Commit phase

Input: Contract space Ω, all prosumers $i \in N$, set of all accepted contracts Ω^a and offers O^a in the offer phase
Output: Accepted contract ω^a and new contract space Ω'
  Sort $(\omega, o) \in \Omega^a \times O^a$ on $v_i(\omega)$ in descending order
  for $i \in N$ **do**
    $\omega^{a,i}, o^{a,i} \leftarrow \Omega^a \times O^a$.pop() ▷ The i^{th} contract with the most value
    if $\omega^{a,i}$ meets local grid constraints **then**
      $(i, j, \theta) = \omega^{a,i}$
      $\hat{i} \leftarrow$ update-profiles(i, θ)
      $\hat{j} \leftarrow$ update-profiles(j, θ)
      exchange-payment$_{i \rightarrow j}(o^{a,i})$
      $\Omega' \leftarrow \Omega$ **break**
    end if
  end for
  for all $\omega = (i, k, \theta) \in \Omega$ **do**   ▷ Contracts including agents in ω^a should be updated
    Compute ω' using the profile optimiser
    $\Omega' \leftarrow \Omega' \setminus \{\omega\} \cup \{\omega'\}$
  end for
  return ω^a, Ω'

5.2. Peer selection

During the peer selection process, the mechanism selects for each prosumer a subset of prosumers to trade with. This procedure is shown in Algorithm 3.

In this model, we assume that each prosumer sends offers to a limited number of k promising prosumers it negotiates with (in most simulations in this paper $k = 5$, which basically means that each prosumer negotiates simultaneously with 5 prosumers that are most promising in terms of potential Gains from Trade). The negotiation partners of a prosumer are sorted based on the potential gains as computed by Eq. (12), and the k peers with the highest potential gain are selected as partners to send offers to. The prosumer is able to negotiate with each partner concurrently. Note the negotiating power of different prosumers may be asymmetric: considering some prosumer, it is not guaranteed that for each partner in its top k preferred partners, the considered prosumer is also in those partners' top k. This is natural, as some prosumers have more power to trade, more flexibility to offer, or a more attractive profile than others. On the other hand, if there are a lot of prosumers with similar generation/demand profiles, their negotiating power is limited, as there are a lot of alternative choices.

Algorithm 3 Peer selection algorithm for agent i

Input: Set of all possible energy contracts Ω_i, number of desired trading partners k
Output: List of the top-k trading partners
  Sort $\omega \in \Omega_i$ on $v_i(\omega)$ in descending order
  return $\Omega_i[1..k]$

5.3. Agent strategies

The negotiation strategy dictates the offer and acceptance behaviour of an agent. Different negotiation strategies affect how the Gains from Trade inside a trading coalition is redistributed. But, it would not significantly affect the overall benefits of peer-to-peer trading. For this work, we model agents as having a linear concession strategy for dividing the expected gains during negotiations, but other strategies could also be employed.
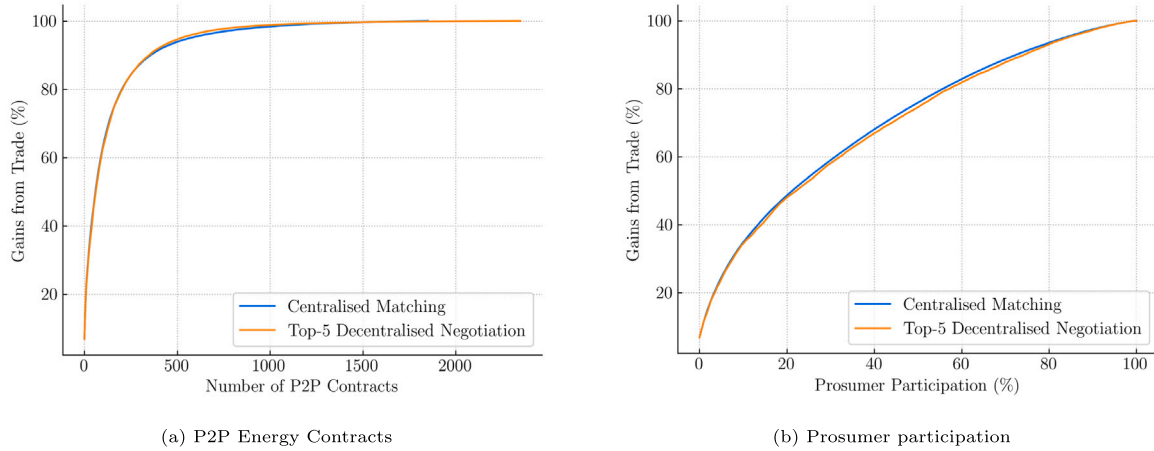
(a) P2P Energy Contracts

(b) Prosumer participation

Fig. 4. Convergence of the Gains from Trade (as a percentage of the maximum possible Gains from Trade).

In more detail, in the first round the agent proposes an offer that is equal to the maximum utility u_i^{max} that can be gained from trading with all partners combined. This represents the market power of the agent. The agent also defines a reservation value rv_i, which is the minimum utility the agent will accept. The bid in consequent rounds is determined as a linear function of the round number r that decays to the minimum utility when the deadline d is reached. Formally, this function is defined as

$$o_i(r) = rv_i + (u_i^{max} - rv_i) \cdot (1 - \frac{r}{dl}) \qquad (17)$$

The agent will accept a bid, if a received bid has a value higher or equal to its last bid.

It is important to observe this decentralised P2P trading process is much more computationally intensive than the approach managed by a central market maker. Essentially, at each time step, each agent sends and accepts/rejects offers from k peers or negotiation partners (for example, in our simulations $k = 5$), after which the accepted offer with the highest GT in the market is cleared, a coalition forms and the process repeats. In the centralised clearing process, also one coalition is formed per time step, but the market maker selects and clears the deal with the highest GT, leading to much less offers being exchanged. Of course, we mostly look at market efficiency, but decentralised negotiation is computationally more expensive, especially if implemented in a decentralised system, such as a blockchain.

6. Experimental comparisons using real world datasets

To compare the effectiveness of peer-to-peer trading in a realistic community, we consider two scenarios, using data from real prosumers of two energy trials in the United Kingdom. The first scenario uses anonymised data from the Thames Valley Vision trial, where we compare the methods in an average-size community of 200 households [32]. The second scenario makes use of a much larger set of prosumers, using anonymised data from the Low Carbon London trial, providing data from several thousand households. Thus we can check whether the results hold up in larger settings as well as the effect that different compositions of a smaller community may have on the effectiveness of trading. In both scenarios, we model the energy assets of a prosumer as follows

- Each prosumer has access to a lithium-ion battery. The price of the battery was fixed to £150/kWh and its lifetime to 20 years. The battery life cycle data was obtained from [33]. For each prosumer, the battery capacity was computed as the battery capacity that minimises the prosumer's bill when used for self-consumption only.

- Each prosumer has access to a share of a wind turbine, where the share was computed by minimising the yearly bill given no access to any storage. The price of the wind turbine was set to £1072 /kW and its lifetime to 20 years. The wind speed to power curve has been interpolated using data for a typical community wind turbine (i.e. Enercon E-33).

Furthermore, both scenarios consider a flat import tariff of 16 p/kWh (i.e. 16 pence or 0.16 British pounds per kWh) and an export tariff of 0 p/kWh. We consider the simulation of these energy communities for a 1 year period with 30-min granularity, resulting in 365 ∗ 48 = 17520 time steps. Finally, we note that during the simulations conducted with real load profiles (from the two communities described above) and the European low voltage test feeder network provided by IEEE [34], we found no grid constraints were violated. As a result, the trades or contracts with the highest gains were always accepted by the grid operator.

6.1. Scenario 1: Thames Valley Vision

The first scenario considers a simulation of 200 households. The demand data has been retrieved from the Thames Valley Vision trial [32].

Fig. 4(a) shows the Gains from Trade achieved by the community as more and more contracts (and thus P2P trading coalitions) are established. This figure displays the convergence of the GT towards the theoretical limit of maximal GT possible in that community (this limit is reached when every single prosumer in the community participates in trading). As expected, this graph shows that it eventually does converge to this theoretical limit as the number of P2P energy contracts increases.

However, we find that only a small number of P2P energy contracts are actually required, for both the centralised matching and clearing models and the decentralised negotiation model. Both models achieve close to the maximum possible Gains from Trade after about 1000 established contracts, while with a community of 200 prosumers there would be a theoretical total of $\frac{200^2}{2} - 200 = 19900$ different pairings. This might suggest that peer-to-peer models achieve maximal efficiency relatively quickly and that only a small number of bilateral contracts need to be established. Between centralised matching and clearing, the difference between convergence is insignificant. This supports that the converging behaviour is generally applicable to different peer-to-peer trading models.

Fig. 4(b) shows the convergence of the Gains from Trade in terms of prosumer participation. Notably, also here we do need all prosumers to participate in peer-to-peer trading to achieve maximum efficiency. A noticeable portion of the gains can be achieved with less participation,
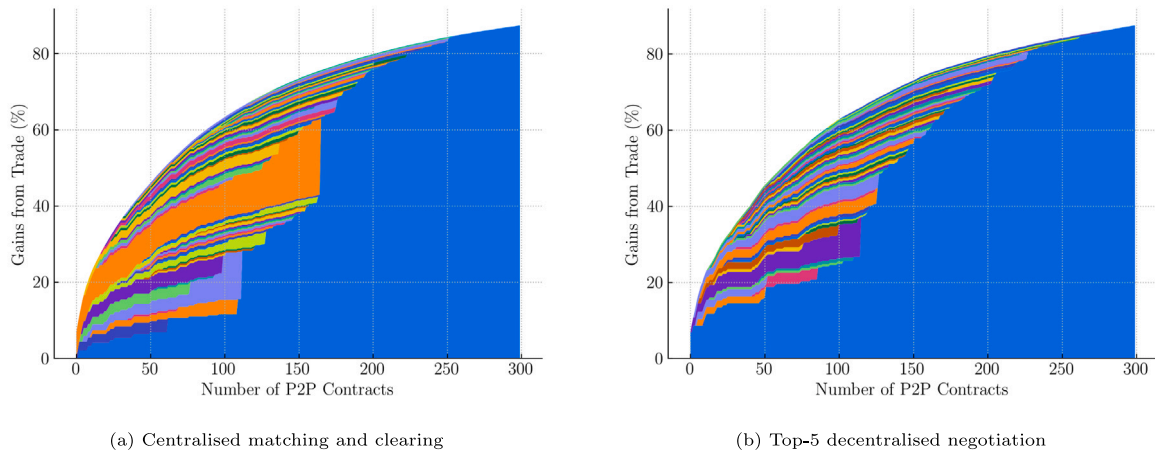
(a) Centralised matching and clearing   (b) Top-5 decentralised negotiation

Fig. 5. Gains per energy trading coalition as a function of the number of P2P contracts between prosumers. Different colours represent coalitions formed by P2P contracts.
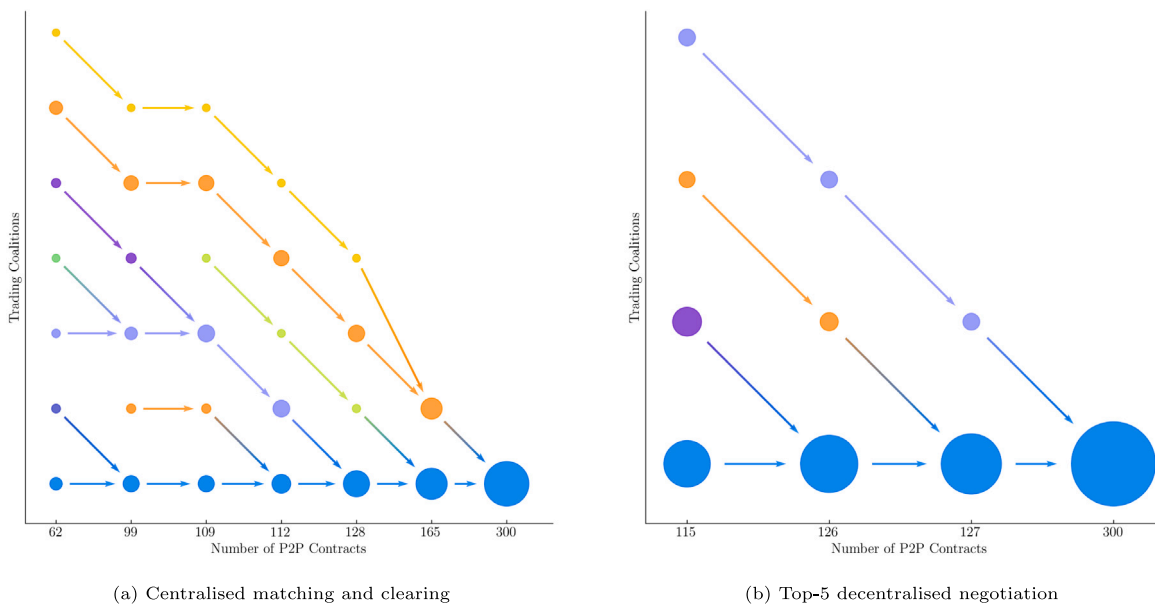

(a) Centralised matching and clearing   (b) Top-5 decentralised negotiation

Fig. 6. Merging of different energy trading coalitions based on the number of P2P contracts established. The area of the disks represents the total gains from trade the trading coalition achieves. While a new coalition merge happens in each step, only time steps where coalitions that contribute at least 3% of the Gains from Trade are merged are shown in the illustration, and only the coalitions that contribute at least 2% to the maximal GT are shown.

as 60% of the gains are reached with 30% of prosumers' participation. Furthermore, the marginal contribution provided by increased prosumer participation diminishes quickly for both models, which can mostly be explained by the fact that the set up proposed in this work first selects the trades with the highest gains for the community, leaving trades with lower gains for later rounds. Both models have similar behaviour as they both use optimal schedules computed by the joint profile optimiser.

To investigate further why the marginal contribution of a contract diminishes so quickly and why both models show similar behaviours, we take a closer look at the dynamics of the peer-to-peer trading mechanism, by looking at the trading coalitions that form between prosumers in each negotiation/matching round. A trading coalition can be formed when prosumers have agreed on a contract with each other either directly, but it is also formed when a prosumer is indirectly connected by another prosumer, e.g. prosumer i has a contract with prosumer j and prosumer i has a contract with prosumer k, then the trading coalition would be $\{i, j, k\}$, and the profile optimiser will jointly optimise the assets of the 3 prosumers.

Fig. 5 shows the Gains of Trade for each trading coalition, where the x-axis represents the number of contracts accepted and the y-axis

represents the total GT achieved. Different colours represent a different disjoint trading coalition. Merging of coalitions can be seen in the figure when one coalition of a certain colour that has accepted the contract is taken over by the colour of the coalition that has offered the contract. Initially, we see that there are many small coalitions, but eventually, every prosumer is included within the same trading coalition, i.e. forming the grand coalition. Here we can see clear differences between the dynamics experienced, by the centralised matching and clearing approach as opposed to the decentralised negotiation approach. The centralised matching approach creates a larger trading coalition before merging them together into the grand coalition, while for decentralised negotiation trading coalitions stay relatively small.

In Fig. 6, we show the merging process of the trading coalitions for both centralised matching and decentralised negotiation. The x-axis represents the time steps in which merging happened in terms of how many P2P contracts were established. The area of the disk represents the size of the coalition, and the colours correspond to the different coalitions in Fig. 5. Only coalitions with significant size (contributing to at least 2% of Gains from Trade) and their merges are presented in the figure. Here we can see that the centralised matching approach
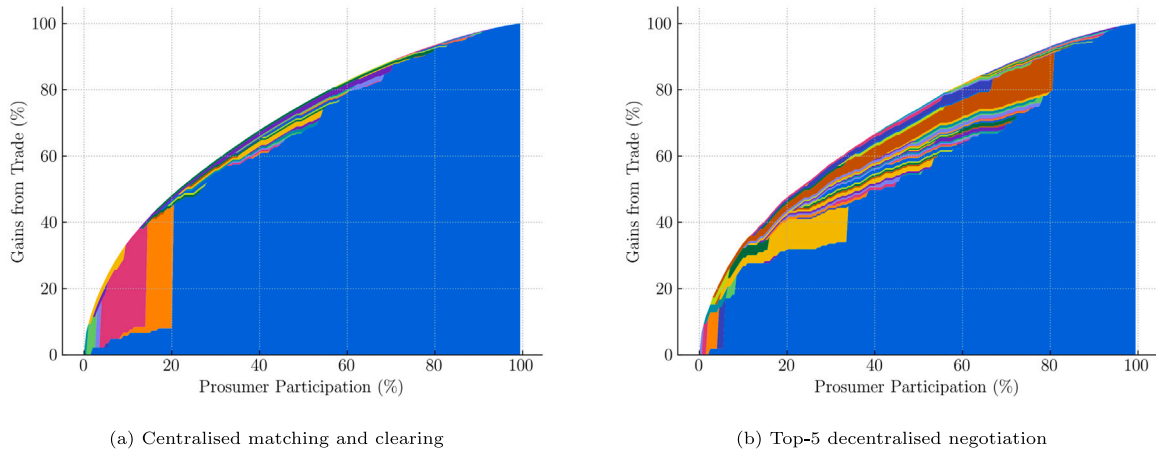
(a) Centralised matching and clearing

(b) Top-5 decentralised negotiation

Fig. 7. Gains per trading coalition as a function of prosumers that participate in trading. Different colours represent coalitions formed by P2P contracts.
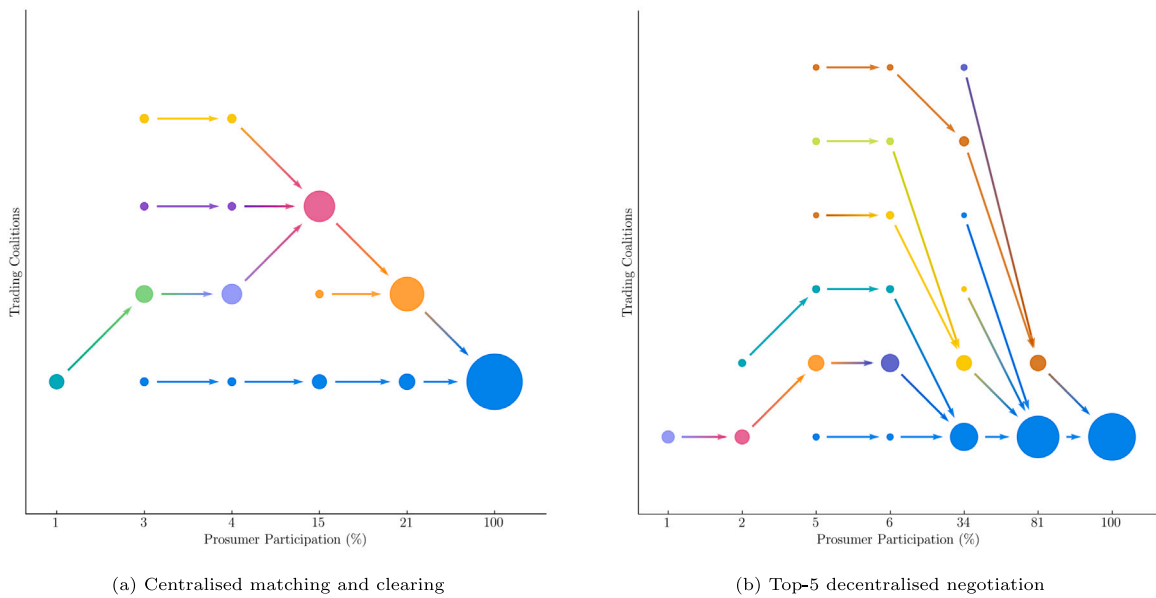

(a) Centralised matching and clearing

(b) Top-5 decentralised negotiation

Fig. 8. Merging of different energy trading coalitions based on the percentage of prosumers that participate in trade. The size (area) of the disks represents the total Gains from Trade that specific trading coalition achieves. Only timesteps where coalitions that contribute at least 5% of the GT are included in the visualisation, and only the coalitions that contribute at least 1% to the gains from trades are shown.

creates larger and more trading coalitions. This seems to suggest that centralised matching prefers contracts between prosumers that are already part of a trading coalition, before merging them into the larger trading coalition. Furthermore, when the trading coalitions are merging into the largest coalition, the majority of the potential Gains from Trade has already been achieved. This might suggest that the gains primarily come from matching the prosumers that have very different demand profiles since those contracts are established first. After those gains have been realised, the trading coalitions have compounded enough minor differences in demand profiles for them to lead to a significant gain. For decentralised negotiation, while this effect can also be observed, since lots of small trading coalitions form, they tend to not grow larger. This is likely due to these small trading coalitions being quite diverse as opposed to the larger coalitions, giving them more market power and therefore laying claim to a larger fraction of the benefits. Larger coalitions have less market power and are therefore more likely to accept a contract proposed by a small coalition, leading to small trading coalitions merging into the largest.

Fig. 7 shows the Gains from Trade of trading coalitions related to prosumer participation. The *x*-axis represents the percentage of prosumer participation in the P2P model which continues to increase

as more contracts are established, and the *y*-axis represents the total GT achieved, as percentage of the maximal. The colours represent trading coalitions, similarly to Fig. 5. Here we observe that for the centralised matching and clearing there are fewer trading coalitions, while for decentralised negotiation we observe a lot more trading coalitions.

In Fig. 8, we show the trading coalitions of significant size (contributing to at least 5% of the maximal GT) and their merging processes, where the size of the disk represents coalition size and the colour corresponds to the coalitions from Fig. 7. Here we can also see for centralised matching that there are a small number of significantly sized trading coalitions and that they tend to merge into the largest coalition. Yet, the merging process for decentralised negotiation involves a lot more groups. These effects seem quite different from those observed when looking at the dynamics when looking at the accepted contracts, However, they can be explained by the fact that the most diverse consumers first form a trading coalition, leading to prosumers that join at a later stage either joining the diverse consumers in the larger coalition, or having to form themselves a smaller trading coalition. We see that some of these smaller coalitions do form, however, the majority do join this larger coalition since no other single consumer is diverse enough. For decentralised negotiation, this plays a smaller role, since
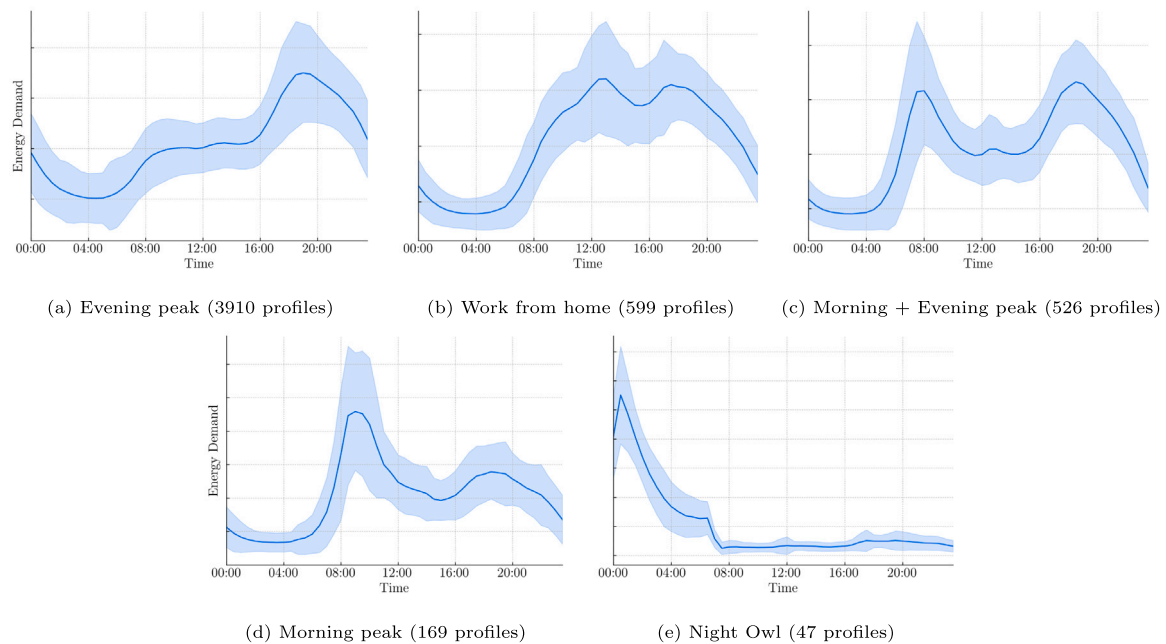
(a) Evening peak (3910 profiles) (b) Work from home (599 profiles) (c) Morning + Evening peak (526 profiles)

(d) Morning peak (169 profiles) (e) Night Owl (47 profiles)

Fig. 9. Average daily demands for different consumer profile clusters. Error bars represent the standard deviation of the demand profiles within the cluster.

these diverse consumers are more likely to not join together initially due to them having a higher market power and being more likely to form smaller trading coalitions.

6.2. Scenario 2: Low Carbon London

In the second scenario, we consider a larger dataset with over 5000 households. The demand data has been retrieved from the Low Carbon London trial [35]. This dataset provides a larger and more representative set of households for the London region. However, since realistically sized communities behind an LV substation often are smaller sized, e.g. around 50–200 [6], instead of considering the whole dataset at once, we create experimental scenarios with realistic diversity and sizes closely fitting the distributions in the large data set.

To ensure that these scenarios reflect the whole community correctly, we apply stratified sampling of consumer profiles from the larger pool. A widely used method for identifying these strata in energy demand modelling is clustering on energy demand profiles [36,37]. In this study, we use K-means clustering on daily averaged consumption in the winter months of the prosumers. The full methodology for performing the clustering is described in Appendix B. From our clustering analysis, we identify and select 5 distinct demand profiles from the resulting clusters. These are shown in Fig. 9.

Analogous to earlier works, we find that the majority of profiles follow an evening peak pattern, where a noticeable peak in consumption can be observed during the early evening, often as a result of prosumers coming home from work. Fig. 9(b) shows another much occurring pattern, where energy demands stay consistent throughout the day while still being low during the night. Since this is likely an effect of prosumers working from home and therefore using energy throughout the day, we call this group 'work from home'. The final large group consists of prosumers that consume large quantities of energy during the morning and evening, likely pertaining to users that have a similar lifestyle to the 'Evening peak' group, but also use a lot of energy in the mornings. The remaining two clusters do not have this peak consumption during the evening but at different moments during the day. The 'morning peak' group consumes the majority of energy during the early mornings, while the 'night owl' group consumes the majority of their energy overnight.
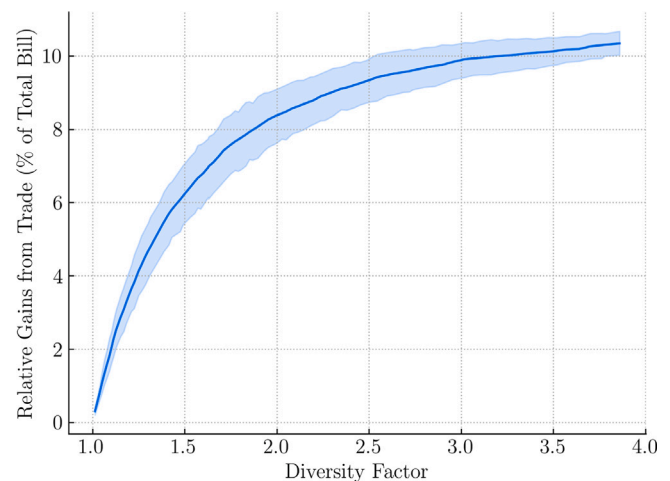

Fig. 10. Influence of diversity factor of the energy community on total Gains from Trade.

Experiments for convergence of the Gains from Trade and the trading dynamics show results very similar to those found for the Thames Valley Vision data set, which is not surprising, given that both trials study consumers from roughly the same area of the UK (southern England/London area).

However, we also explore scenarios where change the composition of the community (in terms of demand profiles or patterns), and here we find slightly different results. To quantify how the composition changes, we evaluate the efficiency with regard to the diversity of the community. The diversity of a community can be expressed using the diversity factor.

Fig. 10 shows the Gains from Trade as a percentage of the total bill compared to the diversity factor. Multiple (i.e. 15) different communities have been generated for several (i.e. 100) diversity values. We find that the gains from trade increase with diversity. Diverse communities contain prosumers that have many different demand profiles, making it more likely that another prosumer can cover residual demands.
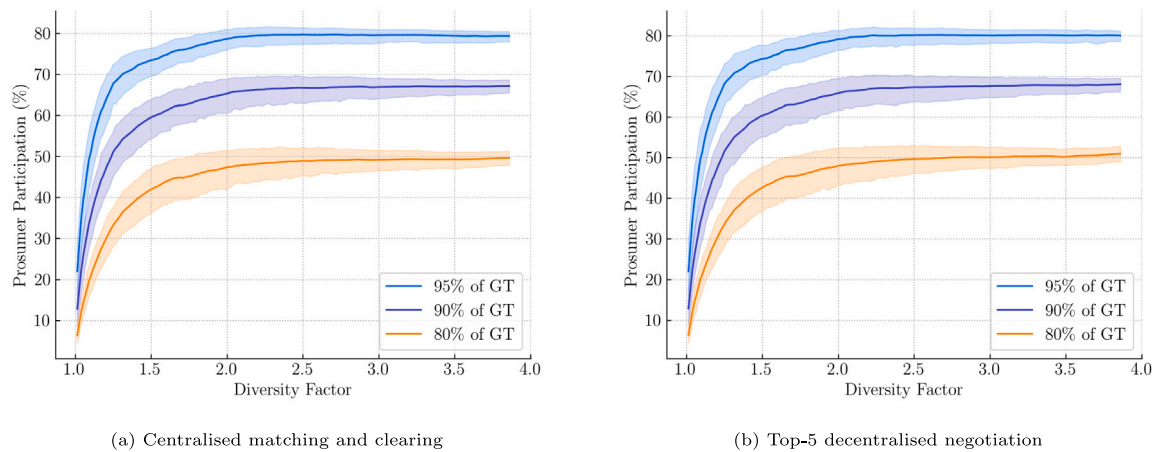
(a) Centralised matching and clearing   (b) Top-5 decentralised negotiation

Fig. 11. Influence of diversity factor of the energy community on the prosumer participation required for achieving the majority of the gains from trade.

However, this effect is less noticeable the more diverse a community is.

Fig. 11 shows the prosumer participation required for different fractions of the Gains from Trade. Interestingly, these thresholds grow quickly for non-diverse communities. The difference effect of increasing diversity for an already decently diverse community seems to have negligible effects.

6.3. Discussion

Results from both the Thames Valley Vision and Low Carbon London datasets suggest that only a small number of contracts need to be established to realise the bulk of the potential Gains from Trade from establishing a P2P or community energy trading scheme. Since generation profiles between prosumers are very similar due to the locality of the community, only requiring a small number of contracts can be attributed due to the majority of the benefits being achieved by prosumers that have significantly different demand profiles from the other agents. After these prosumers establish contracts, they likely still have some excess generation and residual demand, however, these values are closer to the mean and, therefore, less of a factor for future contracts. Similarly, we can see that for the majority of the Gains from Trade, not all prosumers need to participate, with those prosumers that exhibit common consumption behaviour having less marginal contribution. This points to the benefits of considering the exact demand profiles in promoting participation of prosumers in P2P trading schemes.

The convergence of the maximal Gains from Trade happens at a fast initial rate and then levels off, for both forms of market organisation (centralised vs. peer-to-peer models). However, trade dynamics do reveal that there is a difference in how both models achieve similar gains. Both of these models prioritise different contracts in the beginning, when most contracts are likely to be formed, with the peer-to-peer model resulting initially in a larger number of small coalitions.

These results are further confirmed by the scenarios generated using the large-scale Low Carbon London data set. Furthermore, the Low Carbon London data set provides some insights on how this effect generalises to other compositions of communities. Experiments show that while the Gains from Trade increase based on the diversity of the community, prosumer participation required for realising the majority of the benefit stays fairly consistent, with around 50% of the prosumers being required for 80% of the maximal potential GT. However, a lack of diversity within a community suggests that the majority of the gains can be achieved by a very small number of contracts. This makes sense, as only those that differ significantly from the majority benefit significantly from participating in peer-to-peer trading.

In addition to these quantitative results, it is important to discuss the replicability of these energy communities in real implementations. In [38], the acceptability of such schemes of energy communities was assessed through a survey filled by 617 citizens from Finland, and showed that only 20% of the respondents were not interested in participating in an energy community, whereas the incentive of citizens interested where mostly the potential economic gains through buying and selling electricity, but also the reduction of their environmental impact, and the reduction of their dependence to big energy companies. However, it was shown that energy communities should not require any change in end-users' consumption nor any technical challenge. This feedback is similar to the one that was collected within the ReFLEX (Responsive FLEXibility) project in the Orkney Islands, one of the largest Smart Grid demonstrator in the UK [39], that highlighted the fact that social acceptability is a critical driver for energy communities adoption, and that energy trades should be transparent for the end-users. This can be achieved by enabling a post-delivery local energy market clearing (centralised or decentralised) that would result in validated trades that will be deducted from the energy supplier bills of the prosumers and consumers. This enables a seamless integration of energy communities within the current energy system and only requires an energy community manager to operate the local energy market and send the resulting trades to the DSO or the energy suppliers, as is the case in several countries in Europe [40]. However, further work on social acceptability and end-users' preferences for a centralised or fully decentralised approach should be conducted through surveys and focus groups.

7. Conclusions & future work

Peer-to-peer trading has a lot of potential in local energy markets. However, there is no consensus on the best organisation of a local energy market. Centralised market clearing models (such as double auction-type models) seem to be more efficient for extracting the total social benefit, while decentralised peer-to-peer models provide more flexibility and individual choice. In this paper, we examine the efficiency of peer-to-peer markets vs. centralised markets, in terms of contracts established and market participation using large-scale data from several trials in the UK. We study this by using 2 models grounded in earlier research, each respectively representing centralised and decentralised markets. Furthermore, to ensure that the results are representative, we develop a new framework that uses a joint profile optimiser to allow for simulating large-scale peer-to-peer processes.

Results suggest that irrespective of the model used peer-to-peer markets can achieve maximum benefits using a very small number of contracts, i.e. less than 10% of the possible P2P contracts are required. Furthermore, a majority of the benefits can be extracted even when

only a fraction of the community participates in trade, e.g. only about 50% needs to trade for about 80% of the benefits. This effect is more noticeable given a less diverse community, requiring only a small fraction of the community to trade. However, less diverse communities also have less potential for trading, as the overall benefits increase when diversity of prosumer demand profiles increases.

Results in this work are based on two large-scale public datasets from the UK. However, future work may consider replicating these results for other regions around the world, as consumption habits and generation profiles are likely to differ by location. Furthermore, we include only prosumers that have access to a wind turbine, but the effect of solar panels, heat pumps and other forms of renewable generation could be considered in future work. Further investigation could be made for a more complex community, consisting of individuals that exclusively consume energy and prosumers that use different kinds of renewable energy sources such as photovoltaic.

As a result of our experiments, we find that the marginal Gains from Trade from established contracts is high in the beginning (for the most promising prosumers in the community), but then tapers of drastically as more prosumers are added. While costs of establishing contracts and P2P coalition are not explicitly modelled here (e.g. administrative, IT or convenience costs), it may be that the costs associated with establishing such a contract outweighing the benefits it can provide, for many prosumers. In future work, these costs could be explicitly modelled and used to derive an equilibrium point, where establishing more contracts would be counterproductive. Furthermore, this point may be used to determine a practical limit for a fraction of benefits that can be extracted from peer-to-peer energy trading.

Another promising extension – especially important for practical application – is extending the community energy model to consider physical network constraints, such as voltage and power limits (an aspect considered only in a limited number of prior works, such as [41–43]. Physical constraints may make specific energy contracts more or less valuable depending on the logistics and network used costs of transporting the energy, and thus influence the participation in the peer-to-peer trading scheme for prosumers at specific times.

Finally, another technology that is increasingly used for peer–peer and transactive energy systems are blockchain-enabled smart contracts [44–47]. Implementing our energy exchange and coalition formation protocols on a blockchain (e.g. through a smart contract) could also be a relevant avenue for further work.

CRediT authorship contribution statement

Ying Zhang: Writing – original draft, Visualization, Validation, Software, Resources, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Valentin Robu:** Writing – review & editing, Writing – original draft, Validation, Supervision, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Conceptualization. **Sho Cremers:** Writing – review & editing, Visualization, Validation, Software, Investigation. **Sonam Norbu:** Validation, Software, Resources, Methodology, Investigation, Data curation. **Benoit Couraud:** Writing – review & editing, Resources, Methodology, Investigation, Formal analysis. **Merlinda Andoni:** Writing – review & editing, Methodology, Investigation, Formal analysis. **David Flynn:** Writing – review & editing, Supervision, Resources, Investigation. **H. Vincent Poor:** Conceptualization, Project administration, Supervision, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Appendix A. Battery usage

Additional symbols

$SoC_\%(t)$	SoC level of the battery at time t [%]
$SoC_\%^{max}$	Maximum SoC of the battery [%]
$SoC_\%^{min}$	Minimum SoC of the battery [%]
$p^{bat,max}$	Maximum (dis)charging power of the battery [%kW]
Δt	Duration of time period t [hour]
η^c	Charging efficiency of battery
η^d	Discharging efficiency of battery
DoD	Depth of discharge of battery [%]
$DF^{regular}$	Depreciation factor by regular cycles
$DF^{irregular}$	Depreciation factor by irregular cycles
L	Set of irregular cycles
l	for irregular cycles
$SoC_\%^{l,Start/End}$	Starting/ending SoC of cycle l [%]
$N_{cycles}^{DoD,max}$	Maximum number of cycles allowed at specific DoD, provided in manufacturer specification
$n_{cycles}^{DoD,regular}$	Number of regular cycles with the specific DoD

Battery control algorithm

In this appendix, we provide the details of the heuristic-based battery control algorithm. Overall, our battery control algorithm follows the method first described in [22] (and also used in [25,29], and interested readers can also consult that paper for the full details.

The charging of the battery is constrained by the maximum state of charge (SoC) level of the battery, $SoC_\%^{max}$, and the maximum (dis)charging power of the battery, $p^{bat,max}$. Note that the SoC level in this section indicates the battery's level of charge relative to its capacity in percentage, whereas in the main paper SoC level refers to how much of the battery is charged in kW. Therefore, the SoC level in percentage is denoted as $SoC_\%$ in the remaining of section for clarity. The constraints of the battery for any time t while charging are defined as follows.

$$SoC_\%(t) \leq SoC_\%^{max} \tag{A.1}$$

$$\left| p^{bat}(t) \right| \leq p^{bat,\ max} \tag{A.2}$$

Similar constraints apply while discharging the battery. The battery may only be discharged while the SoC level has not reached the

minimum SoC level, $SoC_\%^{min}$. Furthermore, the discharging power may not exceed the maximum discharging power. These constraints are expressed as the following.

$$SoC_\%(t) \geq SoC_\%^{\min} \tag{A.3}$$

$$\left| p^{\text{bat}}(t) \right| \leq p^{\text{bat, max}} \tag{A.4}$$

Given the battery constraints above, the heuristic algorithm operates the battery according to certain rules. When the locally generated power exceeds the power demand, this excess power charges the battery with the charging efficiency η^c, while the SoC level has not reached its maximum capacity. When the excess power cannot be used to charge the battery due to the capacity limit or the charging power of the battery, the remaining power is sold to the grid (note that his algorithm does not consider uncertainty in future demand and supply, unlike e.g. [48]). The updated power and SoC level of the battery and exported energy to the grid $e^s(t)$ at time t with excess power are calculated as the following.

$$p^{\text{bat}}(t) = -\min(\min((g(t) - d(t)), p^{\text{bat, max}}), \frac{SoC_\%^{\max} - SoC_\%(t-1)}{\eta^c \Delta t}) \tag{A.5}$$

$$SoC_\%(t) = SoC_\%(t-1) - \eta^c p^{\text{bat}}(t)\Delta t \tag{A.6}$$

$$e^s(t) = (g(t) - d(t) + p^{\text{bat}}(t))\Delta t \tag{A.7}$$

Where Δt is the length of one time step in hours. In this study, 1/2 hour time steps were used. The exported energy is sold to the grid with export tariff $\tau^s(t)$, generating a profit for the prosumer(s).

When the generated power falls short of the demand, on the hand, the discharged power from the battery with discharging efficiency η^d is first used to cover the deficit. If the discharged battery is not enough to meet the demand, additional power $e^b(t)$ is imported from the grid with the import tariff $\tau^b(t)$. During the power deficit, the battery state and imported energy are determined as follows.

$$p^{\text{bat}}(t) = \min(\min((d(t) - g(t)), p^{\text{bat, max}}), \frac{\eta^d}{\Delta t}(SoC_\%(t-1) - SoC_\%^{\min})) \tag{A.8}$$

$$SoC_\%(t) = SoC_\%(t-1) - \frac{p^{\text{bat}}(t)}{\eta^d}\Delta t \tag{A.9}$$

$$e^b(t) = (d(t) - g(t) - p^{\text{bat}}(t))\Delta t \tag{A.10}$$

Battery degradation

In this section, we describe the computation of the battery's depreciation factor (DF) in Eq. (4), as adapted from the work of Norbu et al. [22]. Frequent charging and discharging of the battery, especially when the depth of discharge is deep, can accelerate degradation significantly. Manufacturers of batteries often specify a battery cycle life on a provided datasheet. In this study, we use the lithium-ion battery cycle life data specified by Xu et al. [33] and Roman et al. [49]. This cycle life specifies the expected number of charge/discharge cycles per depth of discharge (DoD) that can undergo before the performance drops below operable levels. A shortened lifetime of the battery can have an impact on the depreciation cost, hence the inclusion of the depreciation factor provides a more realistic estimate of the community cost.

The total depreciation factor of the battery can be found by computing the depreciation factors of regular and irregular cycles. A regular cycle starts its discharging phase from an SoC level of 100%, and charged back to 100% SoC level. The depreciation factor of regular cycles is then computed as follows:

$$DF^{\text{regular}} = \sum_{DoD=0\%}^{100\%} \frac{n_{cycles}^{\text{DoD,regular}}}{N_{cycles}^{\text{DoD,max}}} \tag{A.11}$$
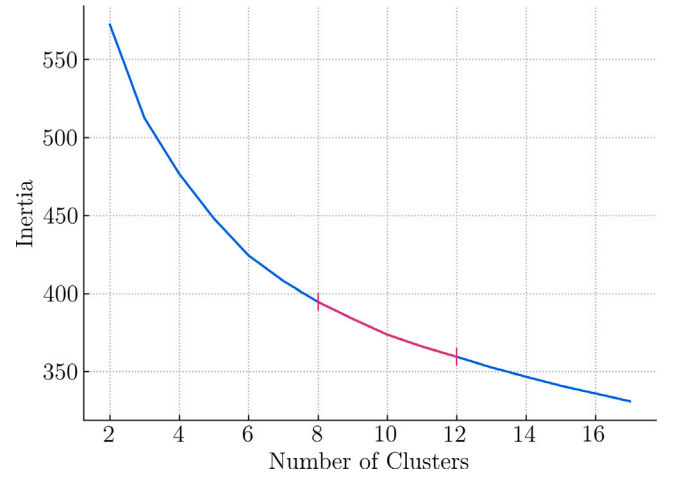

Fig. B.12. Elbow plot of K-means clustering on daily average demand profiles.

where $n_{cycles}^{\text{DoD,regular}}$ is the number of regular cycles (starting from 100% SoC) with the DoD value, and $N_{cycles}^{\text{DoD,max}}$ is the lifetime of the battery provided by the battery manufacturer in terms of the number of regular cycles the battery can go through with the DoD value.

However, not all cycles can start from 100% SoC level. Cycles that have a starting point other than 100% SoC are called irregular cycles. Regular and irregular cycles can have the same DoD, e.g., a regular cycle that starts discharging at 100% SoC level until it reaches 60% then charged back to 100% and an irregular cycle with 80% starting SoC discharging until 40% SoC level and charged back to 80% SoC, both cycles have the same DoD of 40%. Yet, regular and irregular cycles have different impacts on the depreciation factor, and hence they are computed separately. Furthermore, an irregular cycle can either be a full or half cycle. A full cycle consists of both discharging and charging phases, whereas a half cycle is only one of the charging/discharging phases. The rain-flow cycle counting algorithm [22,31] is used to classify and count the cycles as regular or irregular, as well as full or half.

For all the irregular cycles L during the evaluation period, we compute the deprecation factor as follows:

$$DF^{\text{irregular}} = \sum_{l \in L} n_l \times \left| \frac{1}{N_{cycles}^{\text{DoD}^{eq}(SoC_\%^{l,Start}),max}} - \frac{1}{N_{cycles}^{\text{DoD}^{eq}(SoC_\%^{l,End}),max}} \right| \tag{A.12}$$

where $SoC_\%^{l,Start}$ and $SoC_\%^{l,End}$ are the starting and ending state of charges, respectively, for the irregular cycle l, and n_l represents the type of the cycle (1 for full cycle, and $\frac{1}{2}$ for half cycle). Then, $N_{cycles}^{\text{DoD}^{eq}(SoC_\%^{l,Start}),max}$ corresponds to the maximum number of cycles the battery can perform for $\text{DoD}^{eq}(SoC_\%^{l,Start})$, i.e. a depth of discharge equivalent to a cycle starting at 100% $SoC_\%$ and ending at $SoC_\%^{l,Start}$. This is computed as the following:

$$\text{DoD}^{eq}(SoC_\%^{l,Start}) = 100 - \left(\frac{SoC_\%^{l,Start}}{SoC_\%^{\max}} \times 100 \right) \tag{A.13}$$

We compute $N_{cycles}^{\text{DoD}^{eq}(SoC_\%^{l,End}),max}$ using a similar notion.

Finally, Given the depreciation factors of regular and irregular cycles, the total depreciation factor of the battery is the sum of the two components:

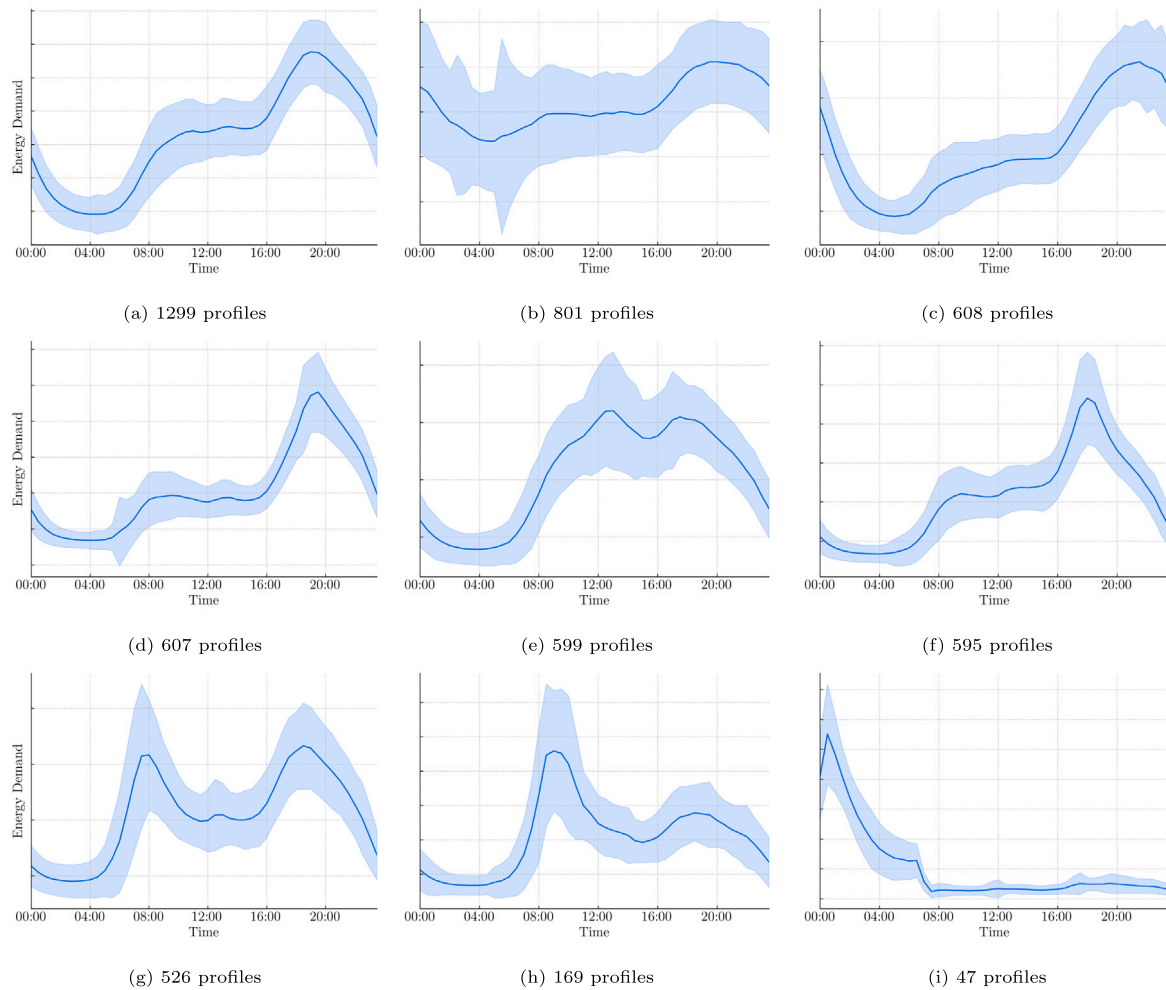$$DF = DF^{\text{regular}} + DF^{\text{irregular}} \tag{A.14}$$

Fig. B.13. Average daily demands for different consumer profile clusters. Error bars represent the standard deviation of the demand profiles within the cluster.

Appendix B. Clustering & consumer profiles

In this section, the method and steps used for clustering consumer profiles used in Section 6 are detailed. The data to be clustered originally consists of 48 half hours for 365 days, creating a demand profile with $365 \cdot 48 = 17520$ points and is generated from the Low Carbon London trial [35]. However, due to the high dimensionality of the data, clustering directly on the profiles suffers from the curse of dimensionality [50]. Instead, we reduce the number of points to consider by reducing the demand profile to a recognisable consumption pattern. Categorising consumption patterns is often done using daily average data from the winter months by utility companies and other studies [51–53]. We apply a similar approach by considering the demands on the prosumers from December to March. Furthermore, demand data from the weekends (Friday to Sunday) and holidays are filtered out, as prosumers may display different consumption patterns. Then, the remaining data from December to March is averaged over the days. Each prosumer is then represented by a 48-dimensional vector representing the average 48 half hours of their typical daily energy demand. As the interest mainly lies in grouping the daily consumption pattern of the prosumers, the demand vectors are normalised using the L1-norm. These vectors are then clustered using K-means clustering. To determine the suitable number of clusters, the elbow method and the silhouette method [54] were used.

First, the elbow method identifies a range of reasonable cluster numbers. Fig. B.12 shows the result of the method. In this plot, the x-axis presents the number of clusters, and the y-axis is the average

inertia of the clusters representing the sum of the squared distance to the centroid of the belonging cluster.

Based on this analysis, we first identify a range of k values, i.e. 8–12, in which the inflexion point could be situated. This is further refined using the silhouette method, from which we find that the best fit value to be $k = 9$ demand profile classes. These are shown in Fig. B.13.

For our P2P market simulation, we aimed to further restrict the number of clusters to a smaller number of more distinctive clusters, such as to be able to illustrate the effects of demand diversity. Looking in more depth at the mean demands of the 9 identified clusters in Fig. B.13, it can be seen that multiple clusters show the evening peak consumption pattern (i.e. profiles (a), (c) and (d) are rather similar), although there are some variations in which time the evening peak is located or the amplitude of the peak.

To reduce computation in the market simulation, all of these evening peak clusters are combined to create one large 'evening peak' cluster. This finally results in the 5 clusters shown in Fig. 9 in Section 6, which are used for the market simulation experiments.

Analogous to earlier works, we find that the majority of these profiles follow an evening peak pattern, where a noticeable peak in consumption can be observed during the early evening, often as a result of prosumers coming home from work. Fig. 9(b) shows another frequent pattern, where energy demands stay consistent throughout the day while still being low during the night. Since this is likely an effect of prosumers working from home and therefore using energy throughout the day, we call this group 'work from home'. The final large group consists of prosumers that consume large quantities of

energy during the morning and evening, likely pertaining to users that have a similar lifestyle to the 'Evening peak' group, but also use a lot of energy in the mornings. The remaining two clusters do not have this peak consumption during the evening but at different moments during the day. The 'morning peak' group consumes the majority of energy during the early mornings, while the 'night owl' group consumes the majority of their energy overnight.

References

[1] Tushar W, Saha TK, Yuen C, Morstyn T, McCulloch MD, Poor HV, et al. A motivational game-theoretic approach for peer-to-peer energy trading in the smart grid. Appl Energy 2019;243:10–20. http://dx.doi.org/10.1016/j.apenergy.2019.03.111.

[2] Tushar W, Nizami S, Azim MI, Yuen C, Smith DB, Saha T, et al. Peer-to-peer energy sharing: A comprehensive review. Found Trends Electr Energy Syst 2023;6(1):1–82. http://dx.doi.org/10.1561/3100000031.

[3] Nizami S, Tushar W, Hossain M, Yuen C, Saha T, Poor HV. Transactive energy for low voltage residential networks: A review. Appl Energy 2022;323:119556. http://dx.doi.org/10.1016/j.apenergy.2022.119556.

[4] Capper T, Gorbatcheva A, Mustafa MA, Bahloul M, Schwidtal JM, Chitchyan R, et al. Peer-to-peer, community self-consumption, and transactive energy: A systematic literature review of local energy market models. Renew Sustain Energy Rev 2022;162:112403. http://dx.doi.org/10.1016/j.rser.2022.112403.

[5] Schwidtal J, Piccini P, Troncia M, Chitchyan R, Montakhabi M, Francis C, et al. Emerging business models in local energy markets: A systematic review of peer-to-peer, community self-consumption, and transactive energy models. Renew Sustain Energy Rev 2023;179:113273. http://dx.doi.org/10.1016/j.rser.2023.113273.

[6] Lucas A, Prettico G, Mengolini A, Fulli G, Gangale F. Distribution system operators observatory : From European electricity distribution systems to reference network. Publications Office; 2016, http://dx.doi.org/10.2790/471701.

[7] Sousa T, Soares T, Pinson P, Moret F, Baroche T, Sorin E. Peer-to-peer and community-based markets: A comprehensive review. Renew Sustain Energy Rev 2019;104:367–78. http://dx.doi.org/10.1016/j.rser.2019.01.036.

[8] Lee W, Xiang L, Schober R, Wong VWS. Direct electricity trading in smart grid: A coalitional game analysis. IEEE J Sel Areas Commun 2014;32(7):1398–411. http://dx.doi.org/10.1109/JSAC.2014.2332112.

[9] Sioshansi FP. Evolution of global electricity markets: New paradigms, new challenges, new approaches. Academic Press; 2013, http://dx.doi.org/10.1016/C2012-0-00444-9.

[10] Wang Z, Yu X, Mu Y, Jia H. A distributed peer-to-peer energy transaction method for diversified prosumers in Urban Community Microgrid System. Appl Energy 2020;260:114327. http://dx.doi.org/10.1016/j.apenergy.2019.114327.

[11] Pinto T, Faia R, Ghazvini MAF, Soares J, Corchado JM, Vale Z. Decision support for small players negotiations under a transactive energy framework. IEEE Trans Power Syst 2019;34(5):4015–23. http://dx.doi.org/10.1109/TPWRS.2018.2861325.

[12] Saxena K, Abhyankar AR. Agent based bilateral transactive market for emerging distribution system considering imbalances. Sustain Energy Grids Netw 2019;18:100203. http://dx.doi.org/10.1016/j.segan.2019.100203.

[13] Guo Z, Pinson P, Wu Q, Chen S, Yang Q, Yang Z. An asynchronous online negotiation mechanism for real-time peer-to-peer electricity markets. IEEE Trans Power Syst 2022;37(3):1868–80. http://dx.doi.org/10.1109/TPWRS.2021.3111869.

[14] Imran K, Zhang J, Pal A, Khattak A, Ullah K, Baig SM. Bilateral negotiations for electricity market by adaptive agent-tracking strategy. Electr Power Syst Res 2020;186:106390. http://dx.doi.org/10.1016/j.epsr.2020.106390.

[15] Etukudor C, Couraud B, Robu V, Früh W-G, Flynn D, Okereke C. Automated negotiation for peer-to-peer electricity trading in local energy markets. Energies 2020;13(4):920. http://dx.doi.org/10.3390/en13040920.

[16] Chakraborty S, Baarslag T, Kaisers M. Automated peer-to-peer negotiation for energy contract settlements in residential cooperatives. Appl Energy 2020;259:114173. http://dx.doi.org/10.1016/j.apenergy.2019.114173.

[17] Dang J, Lu Y, Zhao P, Jafari M. Wholesale power trading through concurrent multiple-issue negotiation. Trans Inst Meas Control 2011;33(3–4):386–405. http://dx.doi.org/10.1177/0142331208100101.

[18] Kalbantner J, Markantonakis K, Hurley-Smith D, Akram RN, Semal B. P2PEdge: A decentralised, scalable P2P architecture for energy trading in real-time. Energies 2021;14(3):606. http://dx.doi.org/10.3390/en14030606.

[19] Khorasany M, Paudel A, Razzaghi R, Siano P. A new method for peer matching and negotiation of prosumers in peer-to-peer energy markets. IEEE Trans Smart Grid 2021;12(3):2472–83. http://dx.doi.org/10.1109/TSG.2020.3048397.

[20] Huang H, Nie S, Lin J, Wang Y, Dong J. Optimization of peer-to-peer power trading in a microgrid with distributed PV and battery energy storage systems. Sustainability 2020;12(3):923. http://dx.doi.org/10.3390/su12030923.

[21] Long C, Zhou Y, Wu J. A game theoretic approach for peer to peer energy trading. Energy Procedia 2019;159:454–9. http://dx.doi.org/10.1016/j.egypro.2018.12.075.

[22] Norbu S, Couraud B, Robu V, Andoni M, Flynn D. Modelling the redistribution of benefits from joint investments in community energy projects. Appl Energy 2021;287:116575. http://dx.doi.org/10.1016/j.apenergy.2021.116575.

[23] Robu V, Chalkiadakis G, Kota R, Rogers A, Jennings NR. Rewarding cooperative virtual power plant formation using scoring rules. Energy (Elsevier) 2016;117:19–28. http://dx.doi.org/10.1016/j.energy.2016.10.077.

[24] Kota R, Chalkiadakis G, Robu V, Rogers A, Jennings NR. Cooperatives for demand side management. In: Proceedings of the 20th European conference on artificial intelligence. IOS Press; 2012, p. 969–74. http://dx.doi.org/10.3233/978-1-61499-098-7-969.

[25] Cremers S, Robu V, Zhang P, Andoni M, Norbu S, Flynn D. Efficient methods for approximating the Shapley value for asset sharing in energy communities. Appl Energy 2023;331:120328. http://dx.doi.org/10.1016/j.apenergy.2022.120328.

[26] Robu V, Vinyals M, Rogers A, Jennings NR. Efficient buyer groups with prediction-of-use electricity tariffs. IEEE Trans Smart Grid 2018;9(5):4468–79. http://dx.doi.org/10.1109/TSG.2017.2660580.

[27] Pumphrey K, Walker SL, Andoni M, Robu V. Green hope or red herring? Examining consumer perceptions of peer-to-peer energy trading in the United Kingdom. Energy Res Soc Sci 2020;68:101603. http://dx.doi.org/10.1016/j.erss.2020.101603.

[28] Moret F, Baroche T, Sorin E, Pinson P. Negotiation algorithms for peer-to-peer electricity markets: Computational properties. In: 2018 power systems computation conference. 2018, p. 1–7. http://dx.doi.org/10.23919/PSCC.2018.8442914.

[29] Couraud B, Norbu S, Andoni M, Robu V, Gharavi H, Flynn D. Optimal residential battery scheduling with asset lifespan consideration. In: 2020 IEEE PES innovative smart grid technologies Europe. 2020, p. 630–4. http://dx.doi.org/10.1109/ISGT-Europe47291.2020.9248889.

[30] Ofgem. FIT scheme closure. 2020, https://www.ofgem.gov.uk/publications/faq-fit-scheme-closure. [Accessed 26 October 2023].

[31] Ke X, Lu N, Jin C. Control and size energy storage systems for managing energy imbalance of variable generation resources. IEEE Trans Sustain Energy 2015;6(1):70–8. http://dx.doi.org/10.1109/TSTE.2014.2355829.

[32] UKERC Energy Data Centre. New thames valley vision - End point monitors. 2017, https://ukerc.rl.ac.uk/DC/cgi-bin/edc_search.pl?WantComp=147. [Accessed 26 October 2023].

[33] Xu B, Oudalov A, Ulbig A, Andersson G, Kirschen DS. Modeling of lithium-ion battery degradation for cell life assessment. IEEE Trans Smart Grid 2018;9(2):1131–40. http://dx.doi.org/10.1109/TSG.2016.2578950.

[34] European Low Voltage Test Feeder. Reseources | IEEE PES Test Feeders. 2023, https://site.ieee.org/pes-testfeeders/resources. [Accessed 26 October 2023].

[35] UK Power Networks. SmartMeter energy consumption data in London households. 2014, https://data.london.gov.uk/dataset/smartmeter-energy-use-data-in-london-households. [Accessed 26 October 2023].

[36] Jeong HC, Jang M, Kim T, Joo S-K. Clustering of load profiles of residential customers using extreme points and demographic characteristics. Electronics 2021;10(3):290. http://dx.doi.org/10.3390/electronics10030290.

[37] Kwac J, Flora J, Rajagopal R. Household energy consumption segmentation using hourly data. IEEE Trans Smart Grid 2014;5(1):420–30. http://dx.doi.org/10.1109/TSG.2013.2278477.

[38] Tuomela S, Hänninen T, Ruokamo E, Iivari N, Kopsakangas-Savolainen M, Svento R. Energy community preferences of solar prosumers and electricity consumers in the digital energy ecosystem. In: Trading in local energy markets and energy communities: concepts, structures and technologies. Springer International Publishing; 2023, p. 113–35. http://dx.doi.org/10.1007/978-3-031-21402-8_4.

[39] Couraud B, Andoni M, Robu V, Norbu S, Chen S, Flynn D. Responsive FLEXibility: A smart local energy system. Renew Sustain Energy Rev 2023;182:113343. http://dx.doi.org/10.1016/j.rser.2023.113343.

[40] Frieden D, Tuerk A, Antunes AR, Athanasios V, Chronis A-G, d'Herbemont S, et al. Are we on the right track? Collective self-consumption and energy communities in the European union. Sustainability 2021;13(22):12494. http://dx.doi.org/10.3390/su132212494.

[41] Guerrero J, Chapman AC, Verbič G. Decentralized P2P energy trading under network constraints in a low-voltage network. IEEE Trans Smart Grid 2019;10(5):5163–73. http://dx.doi.org/10.1109/TSG.2018.2878445.

[42] Gebbran D, Mhanna S, Ma Y, Chapman AC, Verbič G. Fair coordination of distributed energy resources with volt-var control and PV curtailment. Appl Energy 2021;286:116546. http://dx.doi.org/10.1016/j.apenergy.2021.116546.

[43] Norbu S, Couraud B, Robu V, Andoni M, Flynn D. Modeling economic sharing of joint assets in community energy projects under LV network constraints. IEEE Access 2021;9:112019–42. http://dx.doi.org/10.1109/ACCESS.2021.3103480.

[44] Kirli D, Couraud B, Robu V, Salgado-Bravo M, Norbu S, Andoni M, et al. Smart contracts in energy systems: A systematic review of fundamental approaches and implementations. Renew Sustain Energy Rev 2022;158:112013. http://dx.doi.org/10.1016/j.rser.2021.112013.

[45] Andoni M, Robu V, Flynn D, Abram S, Geach D, Jenkins D, et al. Blockchain technology in the energy sector: A systematic review of challenges and opportunities. Renew Sustain Energy Rev 2019;100:143–74. http://dx.doi.org/10.1016/j.rser.2018.10.014.

[46] Couraud B, Robu V, Flynn D, Andoni M, Norbu S, Quinard H. Real-time control of distributed batteries with blockchain-enabled market export commitments. IEEE Trans Sustain Energy 2022;13(1):579–91. http://dx.doi.org/10.1109/TSTE.2021.3121444.

[47] Andoni M, Robu V, Flynn D. Blockchains: Crypto-control your own energy supply. Nature 2017;548(7666):158. http://dx.doi.org/10.1038/548158b.

[48] Ströhle P, Gerding EH, de Weerdt MM, Stein S, Robu V. Online mechanism design for scheduling non-preemptive jobs under uncertain supply and demand. In: Proceedings of the 2014 international conference on autonomous agents and multi-agent systems. International Foundation for Autonomous Agents and Multiagent Systems; 2014, p. 437–44, https://dl.acm.org/doi/10.5555/2615731.2615803.

[49] Roman D, Saxena S, Robu V, Pecht M, Flynn D. Machine learning pipeline for battery state-of-health estimation. Nat Mach Intell 2021;3(5):447–56. http://dx.doi.org/10.1038/s42256-021-00312-3.

[50] Bellman RE. Dynamic programming. Princeton, NJ: Princeton University Press; 2010, http://dx.doi.org/10.1515/9781400835386.

[51] Li K, Ma Z, Robinson D, Ma J. Identification of typical building daily electricity usage profiles using Gaussian mixture model-based clustering and hierarchical clustering. Appl Energy 2018;231:331–42. http://dx.doi.org/10.1016/j.apenergy.2018.09.050.

[52] Guo Z, Zhou K, Zhang X, Yang S, Shao Z. Data mining based framework for exploring household electricity consumption patterns: A case study in China context. J Clean Prod 2018;195:773–85. http://dx.doi.org/10.1016/j.jclepro.2018.05.254.

[53] Antonopoulos I, Robu V, Couraud B, Flynn D. Data-driven modelling of energy demand response behaviour based on a large-scale residential trial. Energy AI 2021;4:100071. http://dx.doi.org/10.1016/j.egyai.2021.100071.

[54] Yuan C, Yang H. Research on K-value selection method of K-means clustering algorithm. J 2019;2(2):226–35. http://dx.doi.org/10.3390/j2020016.