Goal-Based Automation of Peer-to-Peer Electricity Trading

Jordan Murkin, Ruzanna Chitchyan and David Ferguson

Abstract As the uptake of microgeneration increases, the centralised model of electricity generation will be significantly altered. One of the new models that is being investigated is the notion of a peer-to-peer electricity market in which prosumers can market their electricity exports to any other household. This paper investigates this model and proposes a new algorithm for automating the sale and purchase of electricity in this market aiming to optimise the market while providing increased control to householders.

Keywords P2P electricity trading · Trading automation · Goal optimisation

1 Introduction

Microgeneration is the generation of small amounts of electricity by individual households (e.g., via solar panels) and non-energy specialist small businesses (e.g., farmers via wind turbines). Uptake of the microgeneration in the UK is steadily increasing (DEC 2016) and with it new market models are necessary to make full use of this newly available generation capability. One of the new models that has gained a lot of research and practical interest with increased *distributed generation* (i.e., generation by other than centralised power plants) is the concept of *peer-to-peer (P2P) electricity trading* (Power Ledger 2016; Power Matcher 2017; Brooklyn

J. Murkin (✉) · R. Chitchyan
Department of Informatics, University of Leicester, Leicester, UK
e-mail: jpm45@le.ac.uk

R. Chitchyan
e-mail: rc256@le.ac.uk

J. Murkin · D. Ferguson
EDF Energy, London, UK
e-mail: david.ferguson@edfenergy.com

Microgrid 2017) in which excess microgeneration can be sold by the producer directly to another consumer instead of to the grid.

A P2P market model would bring a number of benefits over the current system; prosumers would be able to market their excess generation, possibly providing an extra income stream. While consumers would see a far greater choice over their energy source, with the ability to pick the exact source of their supply. However, should the participants be required to manually manage sales, the required time commitments and need to acquire trade-related domain knowledge would likely outweigh any benefit for the majority of users. Consequently, it is necessary to automate the trading process, and with automation also comes potential optimisation.

This paper will explore the methods of automation with the objective to provide increased choice to consumers and incentivise the purchase of locally generated electricity to increase the value proposition of microgeneration. The intent being to reduce transmission losses caused by long distance energy distribution and increase the use of renewable technologies, both aiming to increase energy efficiency in the energy industry. Thus the contribution this paper provides is to propose:

- An electricity market structure that enables the automation of P2P trading;
- A new algorithm to automate P2P electricity trading;
- Evaluation of performance of this algorithm using simulations.

This paper investigates the different ways in which electricity trading over a P2P market could be automated and the different optimisations it makes possible. It will start with a background of the current research that has taken place in this area, the ways that a peer-to-peer market may appear and ends by proposing a new matching algorithm that could be used to optimise this future market.

2   Background

Previous research has explored different methods of household trading inside the electricity market. The authors of (Yaagoubi and Mouftah 2015a, b) used a modified regret matching as a method to match buyers and sellers individually, allowing for price preferences and taking into account transmission losses in the network. However, in (Yaagoubi and Mouftah 2015a) each buyer can only receive energy from one seller and price is the only consideration. Another game theoretic approach was used in (Wang et al. 2014) which focused primarily on the increased utilisation of energy storage using an auction based approach to determine price. A microgrid focused trading mechanism (Luo et al. 2014) has also been suggested that relies on sharing energy schedules between participants so as to match producers and consumers at the correct time to ensure electricity is utilised fully. This research expands on these with the addition of multiple preferences for each participant in the network, the ability for buyers to purchase from multiple sellers and vice versa and proposes a new market matching approach designed to find an optimal solution while reducing overheads.

While automation is the focus, it is important to first detail how the market would be structured and function in a P2P trading environment. The structure of the market will also inform the choice of trading algorithm type.

The current electricity market divides each day into 30 min periods which electricity can be traded for. This research kept the same approach in which trading is separated into clearly defined periods allowing for regular trading to continue as well as providing a structure to the automation. It is currently undecided whether these fixed trading periods will remain at 30 min or whether a shorter time period will be used. To decide this, a number of simulations will be run with different trading periods to identify if there is a benefit to shortening these periods.

2.1 Structure

To begin, it is necessary to define the structure of these trading periods and the trading procedure. In the current electricity market, trade contracts are determined in advance between parties and controlled by those involved. To remove the complexity of the electricity market and make it accessible to the average consumer, contracts will instead be organised by the market itself. This means that from a consumer's point of view the only interaction with the system is the setup of their account and the configuration of their preferences.

With the market establishing contracts between participants, there are two main approaches that can be taken. The first and more traditional is to automate the trades of each participant individually. The second is to have a single entity that performs a matching of all buyers and sellers together, negotiating all contracts at the same time. The first serves to benefit individuals to the maximum potential, but may not find the optimal solution for the network as a whole. The second allows for matching to be decided on a market level, reducing overheads. This method though, is generally more open to abuse by having a single entity controlling all contracts being created. As this research is built upon distributed ledger technology (DLT), as described in (Murkin et al. 2016), which supports trust by design, this risk of abuse can be avoided. For this reason, the second option was chosen. This left a number of choices available, summarised in Table 1.

The choice of which option to use would also depend on which algorithm would be used to match buyers and sellers together. There are many types of matching algorithm. In these (in general), sellers would announce what they have to sell and buyers would announce what they want to purchase. A system would then match buyers and sellers based on price and/or preferences, using such mechanisms as regret matching, sharing electricity schedules, auctions or rank order lists.

While option 4 supports a truly open market, where a product can be listed for any price and any buyer can instantly purchase that product, it does not apply well to electricity mainly due to network latencies playing such a large role in sales.

Table 1 Market design options

Option	Seller	Buyer	Market	Notes
1.	List sales	List buy orders	Match buy orders to sales	Overheads for matching all buyers against all sellers
2.	List sales	Directly assign buy orders to sales	Determine which buyer gets which sale	Reduced overheads for matching buyers against specific sellers
3.	List sales	Bid on sales	Bids are frozen at the end of trading and the highest bids win	High bidding overheads under large-scale trading
4.	List sales	Purchase sales	First come first served basis	Network latencies will have large effect on sales

An auction mechanism as described in option 3 would likely allow significant control to users of the system, but was removed due to the overhead that would be created if used on a large scale.

The final options, 1 & 2, are quite similar, but again due to possible increased overhead for minimal improvement, option 1 was removed and the project settled on a variation of rank order list matching to fit option 2.

This type of matching typically sees two groups, each member of each group has a list defining their preference over the parties in the other group. These two lists are then used to determine how the participants from each group are matched. This will be detailed further in the algorithm section.

With the trading process defined, it is then possible to structure the trading periods by splitting them into defined sections. These can be seen in (Fig. 1).

The purpose of these sections is described below:

Sell: In this period participants will place their sell orders.

Buy: In this period buyers would list interest against the sales created in the Sell period.
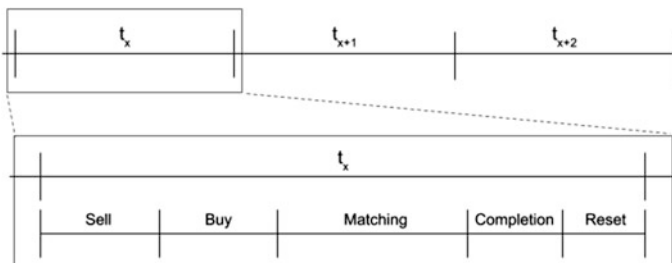
Fig. 1 Proposed structure for P2P trading periods

Matching: This period will deal with the matching aspect of the system, wherein buyers would be matched to specific sales.

Completion: This section has been added due to the DLT nature of this platform. Time is given to ensure that transactions are completed successfully.

Reset: This section again has been added due to the DLT nature of the platform. This section will reset all necessary fields and arrays ready for the next time period.

3 Algorithm Design

With a market structure in place and an algorithm picked for the matching procedure the next stage was to implement it. This section details the design decisions that took place while implementing this algorithm and the considerations that need to be made when automating trades for a P2P electricity market.

3.1 Considerations

An integral element of this system is the distance charge. It was added to the system to serve two purposes. Firstly, to pay for the infrastructure necessary for the electricity network to function including any other obligations that need to come through the sale of electricity. Secondly, to differentiate each sale from one another. Whereas in a traditional commodity market, from a consumer's point of view each unit is thought of to be identical to enable the market to function correctly. The distance charge serves the opposite purpose, intending to differentiate each sale. The reason for this is that while electricity is identical from all sources, transmission loss and other associated costs are largely dependant on the distance it travels.

3.2 Preferences

A primary objective of this research is to provide a mechanism to increase choice over electricity source. The current system already allows choice over price, but there is limited choice that accurately reflects real-time wholesale prices. Generation type is already partially supported with suppliers selling products that are backed by renewable sources. Distance is not supported at all, as in the current system electricity is seen as a commodity in which each unit is identical from a consumer's point of view.

To incorporate these into the algorithm, each participant is given a method to specify a number of preferences that allowed them to control the way in which they were matched by the market. The preferences included for sellers are (i) Minimum price preference and (ii) Distance preference. Preferences for buyers are

(i) Maximum price preference, (ii) Distance preference, and (iii) Energy type preference.

These preferences are used as a means to create scores between the market participants creating the rank order lists for each.

3.3 Scoring

The scoring module of the algorithm dictates how matches are made and so was designed as an independent module allowing it to be replaced whenever needed allowing easy testing of different calculation methods. Below, the calculations used for the current implementation are shown.

In this section a number of letters will be used throughout to improve readability of equations. Firstly, B and S will be used to represent buyer and seller respectively and these letters may be appended to other letters, such as Pb referring to buyer's price preference and Ps referring to sellers price preference.

Other letters:

P Price preference
D Distance preference
E Energy generation type
C_d Distance charge
EP() a function returning the energy type preference given an energy type E
d distance

The scoring system used needs to incorporate the preferences that were previously mentioned. For each of these preferences a different method was used, this section will go through each. To begin it is first necessary to calculate the distance between the participants to allow for the score to be calculated. This was implemented by adding a latitude and longitude to each participant and then using a great-circle calculation to determine the distance between the two points.

3.3.1 Price Preference

The price preference is used as a method to create an initial score between the participants. By using the price preference of both the buyer and seller and including the distance charge we can create a basic score that also reflects the maximum price that the buyer is willing to pay for the given sale.

$$score = P_b - (d \times C_d)$$

This current score only takes into account buyer preference. To incorporate seller preference as well, a check is added to ensure that the maximum price that a buyer

is willing to pay is above the minimum price they are looking for. This means discarding all scores in which $score < P_s$.

3.3.2 Energy Type Preference

Allowing householders to specify the energy type they prefer provides a significant benefit over the current system in which electricity sources are indistinguishable. This preference allows individual preference to be given for each microgeneration energy type: solar, micro CHP, wind, hydro, and anaerobic digestion. Each of these individual preferences is normalised to 1, thus the energy type preference is included by multiplying the score as follows $score = score \times EP(E_s)$.

3.3.3 Distance Preference

To incorporate the buyer's distance preference a simple check is added to determine if the seller is outside of the buyer's specified distance. If the seller is farther than this distance, a multiplier is applied to the score so as to reduce the score more as the distance increases. The reason for this is to avoid creating an exact range within which a buyer searches for sales, instead allowing sellers outside to still be considered.

$$score = score, d <\, = D_b$$

$$score = score \times (D_b \div d), d > D_b$$

The same mechanism can then be applied to incorporate the seller's distance preference.

$$score = score, d <\, = D_s$$

$$score = score \times (D_s \div d), d > D_s$$

3.3.4 Final Score Calculation

After each section described above had been incorporated this gave a final method for calculating the score between two participants as:

$$score = (P_b - (d \times C_d)) \times EP(E_s), d <\, = D_b, d <\, = D_s$$
$$score = (P_b - (d \times C_d)) \times ((D_s \div d) + EP(E_s)) \div 2, d <\, = D_b, d > D_s$$
$$score = (P_b - (d \times C_d)) \times ((D_b \div d) + EP(E_s)) \div 2, d > D_b, d <\, = D_s$$
$$score = (P_b - (d \times C_d)) \times ((D_b \div d) + (D_s \div d) + EP(E_s)) \div 3, d > D_b, d > D_s$$

3.3.5 Implementation

As noted, our algorithm is adapted from a traditional rank order list matching. Here:

1. Seller exports are listed as sales
2. Buyers are assigned to all sales for which their preferences suit and a score is assigned between the two. This score acts as a method to create rank order lists for both the buyer and the seller. For buyers, this is the list of sales they are assigned with the respective scores for these sales. For sales this is the list of interested buyers and their respective scores.
3. The highest scoring buyer for each sale is assigned as an initial match to this sale.
4. The highest score of all initial matches for each buyer is chosen to realise a sale.
5. The algorithm then repeats from step 3 until there are no buyers assigned to a sale, no more sales or there is no remaining consumption to satisfy from buyers.
6. This algorithm performs an individual based optimisation but can easily be modified to incorporate other types of optimisation (e.g. a score maximising algorithm to allow for market level optimisation).

4   Preliminary Evaluation

To assess the performance of this algorithm a JavaScript simulation was implemented. As we aim to assess the performance of the algorithm (and not to accurately simulate an electricity market) a number of generators were created to automate the creation of the accounts and simulate their generation and consumption. Simulation configuration: The accounts were randomly generated for each simulation to ensure that the algorithm functioned correctly regardless of the accounts it used. They were assigned a random location inside an arbitrary area in the UK, with latitude between 50.956,870 and 52.438,562 and longitude between −2.386,779 and 0.292,914. A percentage of these accounts were assigned to be prosumers. Price preferences were randomised; for prosumers, minimum sell price set between 46p/kWh and for consumers maximum purchase price set between 14–16 p/kWh. Distance preference was set for all accounts between 5–10 km. Energy type preferences for consumers were randomised between 0.5–1 and then normalised to 1 between all preferences with the highest preference being set to 1. Lastly generation type for prosumers was randomised to reflect figures from UK microgeneration installations (DEC: Monthly Central Feed-in Tariff register statistics—Statistical data sets—GOV.UK 2016).

Simulations were run for 16 different scenarios, where the percentage of prosumers in the market ranged from 5 to 20% in 5% increments, and the number of participants ranged from 500 to 2000 in increments of 500. In all scenarios the prosumers generated between 5–10 kWh of electricity and consumers consumed between 1–6 kWh and the distance charge was set to 0.2 p/km/kWh. Each scenario was run 50 times and then averaged and the results can be seen below.
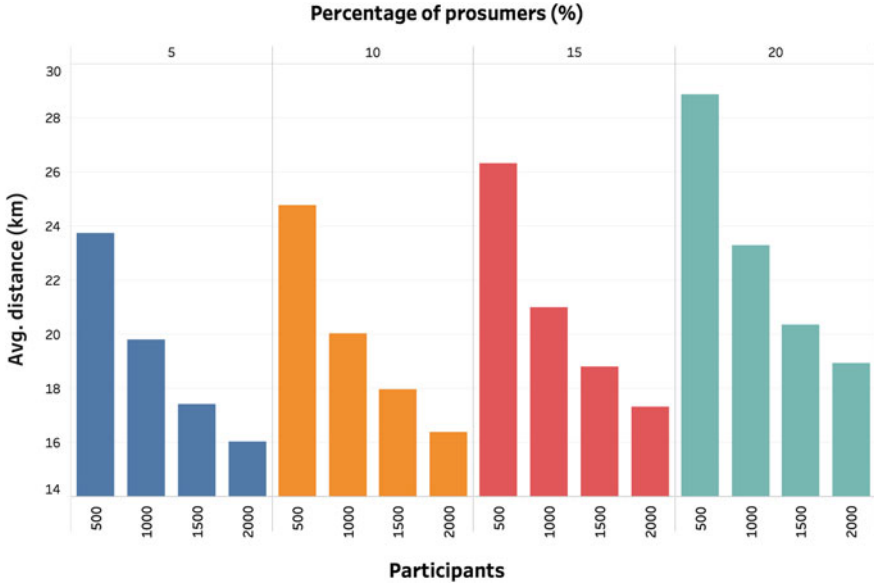
Fig. 2 The average distance between sales for each scenario

Simulation results: Here we assess the algorithm and its performance.

Figure 2 shows how the distance between sales changes through the different scenarios. It can be seen that an increase in the number of participants in the market decreases the average distance between sales. This is likely due to an increase in the density of participants and shows that the algorithm is performing as expected, incentivising the use of local generators.

A point to be noted is that as the percentage of generators increased, the distance between sales also increased. This was a more surprising result, but could be explained by a decrease in the density of buyers causing an increase in the average distance between buyers and sellers.

Another important characteristic of the algorithm can be seen in Fig. 3, which shows how the price per kWh is affected through each scenario. This figure shows that the algorithm reflects the standard economic model of supply and demand. As the number of participants (demand) increases, competition increases and this increases the average price paid. At the same time it can be seen that with an increased number of prosumers (supply) a decrease in the average price occurs.

Figures 4, 5 and 6 assess the performance of the algorithm. It can be seen that as the number of participants increases, the number of sales, the number of rounds and the time taken all increase linearly. The time taken though increases with a significant gradient and thus could pose a problem with a large number of participants. As the current simulation is running synchronously it may be possible to reduce through parallelisation and caching mechanisms to reduce calculations per round if it was to be scaled to millions.
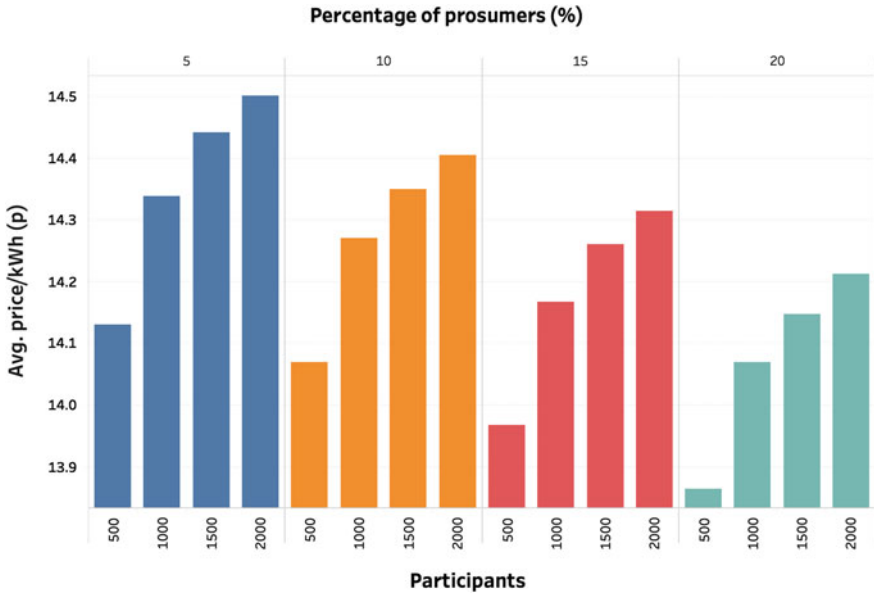
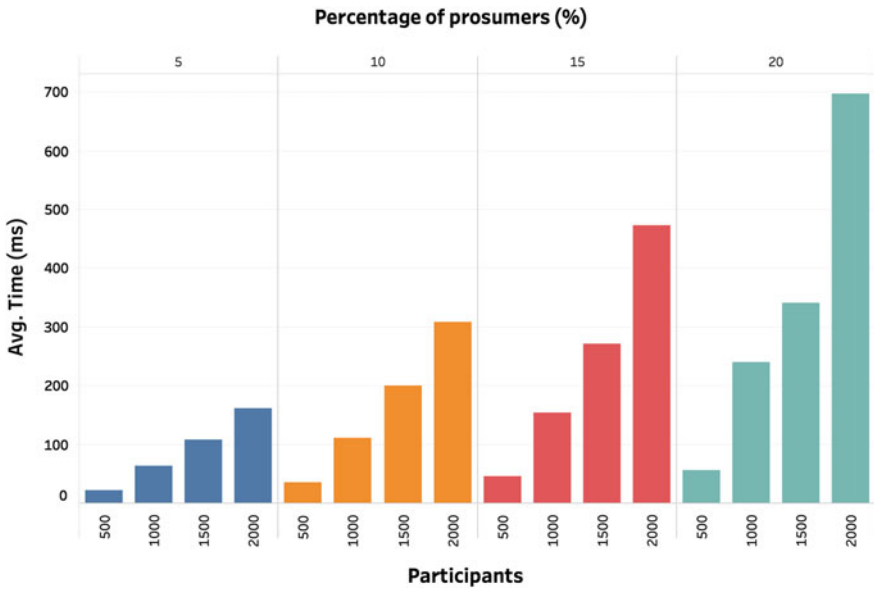**Fig. 3** The average price paid per kWh in each scenario

Fig. 4 Average time taken to converge for each scenario
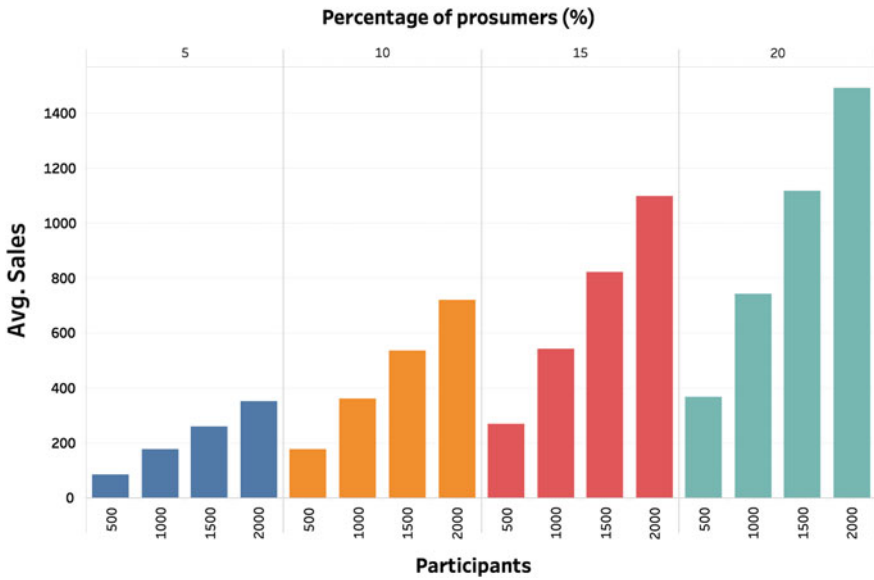
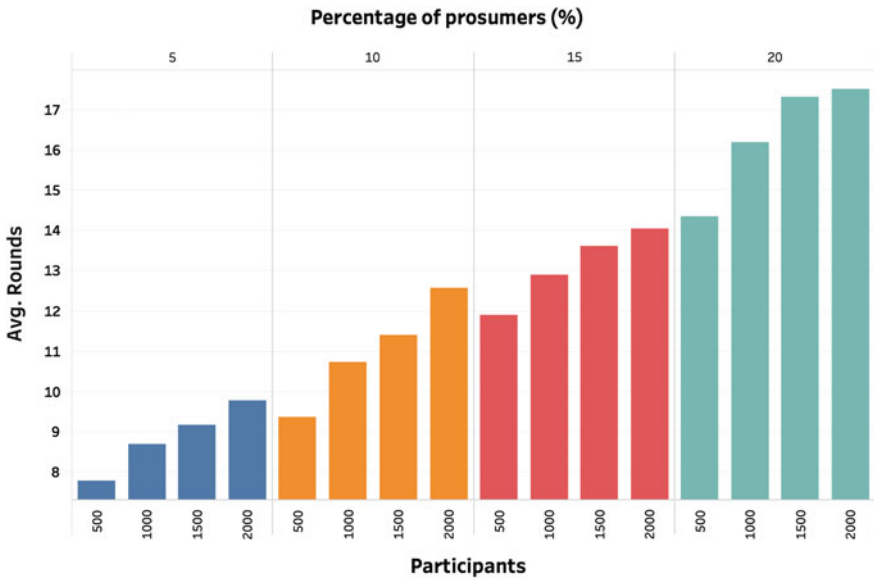**Fig. 5** The number of sales that took place in each scenario

Fig. 6 Average number of rounds taken to converge for each scenario
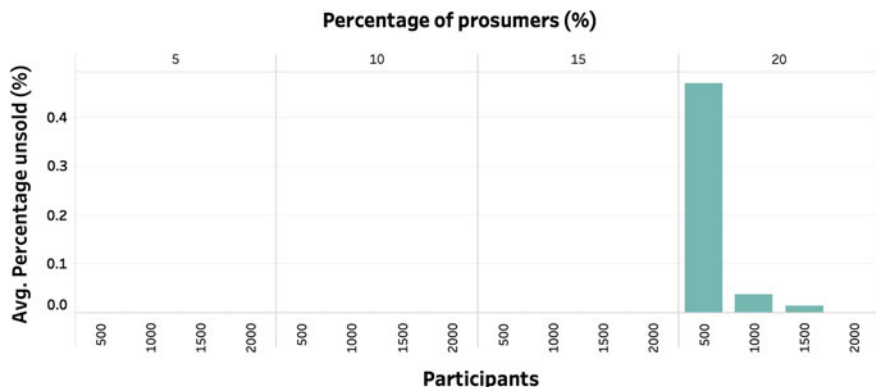
Fig. 7 Average percentage of unsold generation for each scenario

The final important piece of information is the amount of unsold generation throughout the simulations. With 5%, 10% and 15% of participants as prosumers there was no unsold generation, but as the share of prosumers increased to 20% a small amount of generation was unsold Fig. 7. Yet, as the number of participants increased further, this unsold generation again dropped to zero. This is likely due to the large area used in simulation, leading to large distance between participants, with some generation unsold due to high distance charge effect.

These evaluations allow us to conclude that the algorithm is functioning as originally intended, incentivising the purchase of locally produced electricity, minimising the electricity that is left unsold and encouraging competition. We have also seen that while at a low number of participants (thousands) the algorithm will converge in acceptable times, large scale use (millions) would likely require some alterations.

5   Conclusion and Future Work

This paper has presented a new approach to buyer and seller matching in a P2P electricity market. Through simulations we have shown that the algorithm performs acceptably, reducing unsold electricity to almost zero, incentivising the purchase of locally produced electricity and accounting for individual preferences, but may require alterations for large-scale implementation.

As this research continues: real locational data will be added to the simulation to help locate areas that would benefit most from this type of trading mechanism and remove the unrealistic spacing between households; and energy data will be added to conclude the quantitative benefits to both buyers and sellers that can be achieved using this algorithm.

The next stage of this research will investigate alternative scoring mechanisms and algorithms to look beyond purely individual optimisation to allow for group based optimisation and optimisation of the market itself.

References

Brooklyn Microgrid. http://brooklynmicrogrid.com/ Accessed 28 Apr 2017

DECC: Monthly Central Feed-in Tariff register statistics—Statistical data sets—GOV.UK (2016)

Flexiblepower Alliance Network: The PowerMatcher Suite. http://flexiblepower.github.io/. Accessed 28 Apr 2017

Luo Y, Itaya S, Nakamura S, Davis P (2014) Autonomous cooperative energy trading between Prosumers for Microgrid systems, 693–696

Murkin J, Chitchyan R, Byrne A (2016) Enabling peer-to-peer electricity trading. In: 4th International Conference on ICT for Sustainability

Power Ledger-A New Decentralised Energy Marketplace—Where Power meets Blockchain (2016)

Wang Y et al (2014) A game-theoretic approach to energy trading in the smart grid. IEEE Trans Smart Grid 5(3):1439–1450

Yaagoubi N, Mouftah HT (2015a) Energy trading in the smart grid: a game theoretic approach. In: IEEE SEGE, pp 1–6

Yaagoubi N, Mouftah HT (2015b) A distributed game theoretic approach to energy trading in the smart grid. In: IEEE EPEC, pp 203–208