Downloaded from orbit.dtu.dk on: Nov 14, 2019

DTU Library

Online matching and preferences in future electricity markets

Esch, Helge Stefan; Moret, Fabio; Pinson, Pierre; Marin Radoszynski, Andrea

Published in:
Proceedings of the Nineteenth Yale Workshop on Adaptive and Learning Systems

Publication date:
2019

Document Version
Peer reviewed version

Link back to DTU Orbit

Citation (APA):
Esch, H. S., Moret, F., Pinson, P., & Marin Radoszynski, A. (2019). Online matching and preferences in future electricity markets. In *Proceedings of the Nineteenth Yale Workshop on Adaptive and Learning Systems*

General rights
Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Online Matching and Preferences in Future Electricity Markets

Helge S. Esch, Fabio Moret, Pierre Pinson, Andrea M. Radoszynski

***Abstract*—Electricity markets are to be rethought in view of the context of deployment of distributed energy resources, new enabling technologies and evolving business models. Future market mechanisms should have no barrier to entry, while being scalable and giving the possibility to accommodate asynchronicity. Consequently, we propose here to use online matching algorithms, relying on various types of continuous double auctions. They allow agents to trade electricity forward contracts while expressing preferences and being continuously matched as new orders come. Such markets can accommodate agents and trades of any size and characteristics. We eventually concentrate on naive greedy and pro-rata matching algorithms. A discrete double-auction is used as a benchmark. The double auctions are generalized to account for preferences. A case-study application allows us to discuss the computational properties and optimality of the various approaches. An upper bound on the sub-optimality of online matching algorithms, compared to an offline double auction, is also provided.**

***Index Terms*—Electricity markets, Peer-to-peer trading, Order (online) matching, Greedy algorithms,**

I. Introduction

The increasing deployment of renewable energy sources, combined with other distributed energy resources like storage devices (e.g., electric vehicles and household batteries) challenges the status quo in electricity markets. While the operation of power systems has gradually considered this evolution towards a more distributed and dynamic system, electricity markets have not yet adapted to this new context. However, there has been a strong push, both in academia and industry, towards more consumer-centric structures for those electricity markets, using community-based and peer-to-peer concepts [1]–[3]. This may certainly have been supported by the prospects that many see with the use of blockchain as an enabler of new business models in the electric energy sector [4].

Future electricity markets are to be thought as markets that have no barrier to entry for small players like prosumers and that operate with direct trading between sellers and buyers, in contrast with pool approaches. For instance, peer-to-peer electricity markets rely on large sets of bilateral contracts, to be negotiated among all agents involved in electricity exchange e.g. generators, consumers, prosumers and storages. This possibly extends to having system operators in the loop in order to enforce operational constraints [5] and to account for sensitivities of the next transactions to be settled on network operation [6]. Future electricity markets may then rely on two families of negotiation mechanisms, with various levels of decentralization. On the one hand, distributed and consensus-based optimization approaches generalize the clearing principles that exist today in electricity pools [7], [8]. On the other hand, matching methods consider networks of contracts to be simultaneously negotiated in a multi-bilateral fashion [9]. Whatever the approach, those electricity markets allow for new business models and for expressing preferences (related to e.g. distance and electricity type). However, as for any type of peer-to-peer systems, scalability and asynchronicity are some of the key issues to be considered. This is while the optimal resource allocation should also be analysed from a fairness based perspective. Our main objective is to propose an approach to future electricity markets that allows for asynchronicity, with good scaling properties, while still allowing to express preferences. In addition, operational constraints and network aspects are overlooked here, though extensions towards such considerations may build on e.g. [5], [6].

While distributed and consensus-based optimization approaches are inherently relevant for decentralized electricity markets, they require substantial efforts to accommodate scalability and asynchronicity [10]. In contrast, our proposal involves a central matching platform: it is not as decentralized but still relies on bilateral trades among agents. We concentrate on matching approaches and a class of online matching algorithms for forward contracts based on continuous double auctions, from naive greedy (somewhat equivalent to zero-intelligence trading) to pro-rata based approaches. These approaches inherently meet the requirements for scalability and accommodating asynchronicity. The matching engine allows matching buy and sell orders on the fly, also allowing to express preferences, and with various degrees of optimality. The price to pay, however, is some form of sub-optimality compared to optimization approaches. One of our contributions is to give an upper bound on the optimality loss of online matching algorithms compared to their discrete double auction counterparts.

After describing our matching market framework, in terms of the agents, contracts, timeline, etc. in Section II, all aspects of the online matching approaches are introduced and discussed in Section III. Those include the naive greedy and pro-rata based approaches. The optimization approach used as a benchmark is also presented. Subsequently, particular emphasis is placed in Section IV on the question of preferences and how they can be handled within our online matching framework. The case study of Section V allows us to illustrate the workings of the various approaches, while allowing

H. S. Esch, F. Moret, P. Pinson and A. M. Radoszynski are with the Technical University of Denmark, Department of Electrical Engineering, Centre for Electric Power and Energy, Kgs. Lyngby, Denmark (emails: {fmoret,ppin}@elektro.dtu.dk). The work is partly supported by the Danish ForskEL and EUDP programmes through the Energy Collective project (grant no. 2016-1-12530). Pierre Pinson would like to thank the Isaac Newton Institute for Mathematical Sciences (Cambridge, UK) for support and hospitality during the programme Mathematics of Energy Systems (EPSRC grant number EP/R014604/1) and his related Simons Fellowship.

to underline the advantages and drawbacks of the various approaches. Eventually, Section VI closes the paper with a number of conclusions and perspectives for future work.

II. Matching Market Framework

A. Timeline

The overall market framework resembles that of intra-day electricity markets in Europe (see [11]). A major difference though is that, while in European electricity markets participation is restricted to fairly large agents only (hence the term *wholesale* market), access to the market is here granted to all potential agents of the electric power system, without minimum trade volume. Interacting through this market are a set of agents, which may all be buyers and sellers at any time. To discard arbitrage opportunities and potential strategic behaviour, it is assumed that an agent does not both buy and sell within a given trading session.

The market is then organized as a continuous double-auction: this implies that buyers and sellers may come at any time prior to the period of delivery with their offers, to then be matched as they come. This contrasts with the approach in e.g. forward wholesale electricity markets, where all supply and demand offers are matched at once, a fairly long time before delivery. The timeline for our matching market is illustrated in Figure 1. The physical exchange of electricity is discretized into delivery time periods t_i^D. The granularity of the market corresponds to the duration between each delivery time period (as for the example 5-minute resolution considered in our case study). Exchanges, also referred to as trades, for that delivery period t_i^D can be contracted forward in time throughout a trading session running from an opening time t_i^o and to a closure time t_i^c, $t_i^c \leq t_i^D$. We write $T_i = [t_i^o, t_i^c]$ the trading session and its duration.

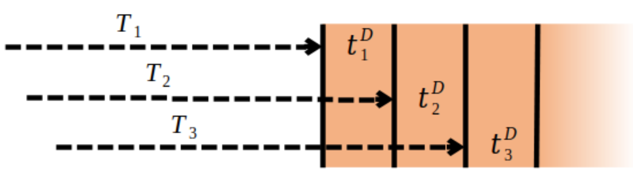

Fig. 1. Market trading time and period of delivery.

In the following since we only concentrate on a given trading session T_i for delivery period t_i^D and then omit i indices, to lighten notations.

B. Orders and Limit Order Book

The matching is organized through a centralized matching engine which market players can communicate with to submit and cancel orders.

Definition 1 (order)**.** *An order $\omega(q_\omega, p_\omega, t_\omega^s)$ is fully characterized by a (standing) quantity $q_\omega \in \mathbb{R}$, a limit price $p_\omega \in \mathbb{R}^+$ and a submission time $t_\omega^s \in T_i$.*

The limit price p_ω is to be understood as a minimum price to receive if on the supply side, and a maximum price to pay if on the demand side. The positiveness of p_ω is assumed for simplicity, though this may readily be relaxed. In order to be considered in the matching process, an order must be submitted within the trading session, hence $t_\omega^s \in T_i$.

Buying orders β, referred to as *bids*, are indicated by a negative quantity, $q_\omega < 0$ while selling offers α, referred to as *asks*, imply a positive quantity, $q_\omega > 0$. Following the definition of *limit orders* in [12], each bid is guaranteed not to be matched at a price higher than its limit, while each ask is guaranteed not to be matched at a price lower than its limit.

Definition 2 (Limit Order Book)**.** *At a given time t, the Limit Order Book (often abbreviated LOB) $\mathcal{L}_s(t)$ consists of the union of the set $\mathcal{A}_s(t)$ of standing asks and the set $\mathcal{B}_s(t)$ of standing bids, i.e.,*

$$\mathcal{L}_s(t) = \mathcal{A}_s(t) \cup \mathcal{B}_s(t), \quad \forall t \tag{1}$$

At any given time t in the limit order book $\mathcal{L}_s(t)$, the difference between the highest standing bid $b^*(t)$ and the lowest standing ask $a^*(t)$ defines the market (bid-ask) spread $s(t) = a^*(t) - b^*(t)$.

C. Matching principles

The aim of a matching engine is to pair orders $\omega(q_\omega, p_\omega, t_\omega^s)$ based on their previously introduced characteristics.

Definition 3 (matching)**.** *A matching (also referred to as trade) $\tau(p^\tau, \Omega^\tau, \boldsymbol{q}_\omega^\tau)$ is defined by a set Ω^τ of all orders ω involved, a vector $\boldsymbol{q}_\omega^\tau$ of quantities expressing the amount with which each order $\omega \in \Omega^\tau$ participates in τ, as well as a price p^τ.*

Ω^τ must at least hold one ask and one bid, but can possibly be made up of multiple orders on both sides. Each $\omega \in \Omega^\tau$ receives (asks) or pays (bids) the same marginal price p^τ per involved unit e.g. kWh or MWh.

A matching is to respect limit orders, i.e., in terms of limit prices comprising a minimum price for bids and a maximum price for asks. In addition, in line with the basic principle of matching supply and demand, any matching is to be balanced.

Definition 4. *A matching $\tau(p^\tau, \Omega^\tau, \boldsymbol{q}_\omega^\tau)$ is said to be balanced when the quantities $\mathbf{q}_\omega^\tau$ involved are such that*

$$\mathbf{1}^\top \mathbf{q}_\omega^\tau = \mathbf{0} \tag{2}$$

with $\mathbf{1}$ a vector of ones with same length as $\boldsymbol{q}_\omega^\tau$.

Consequently, the power system dispatch resulting from such matching will be consistent with the necessary supply-demand equilibrium in power system at any time. The limit orders considered may be partially matched: in that case, the quantity q_ω^τ by which an order participates in a trade can be any value in between zero and the order standing quantity q_ω, $q_\omega^\tau \in [0, q_\omega]$. Consequently, an order standing quantity q_ω must be updated after ω is involved in a trade. It is to be seen as a time-dependant variable which is to evolve throughout the trading session, even though time indices will not be used for that quantity in the following. Necessarily, the standing quantity q_ω can never change its sign (a bid cannot become an ask, as well as the opposite) and its absolute value can only decrease over time.

III. Online Matching through Continuous Double Auctions

Matching may be performed following alternative principles. Here the approaches considered all rely on double auctions. They are described for the case of a single delivery period of interest t_i^D, and thus over a trading session T_i. While our aim is to propose online matching algorithms based on continuous double auctions eventually, a discrete double auction is first introduced as a benchmark. Two types of continuous double auctions are then presented: *(i)* a naive greedy approach, and *(ii)* a pro-rata based approach.

A. Discrete double auction as a benchmark

In the discrete case, the trading session is split into regular time intervals of duration Δt. Over a given time interval, all incoming orders are collected. The limit order book is consequently composed of both incoming orders and those pending orders that were not matched at previous time interval (if any). The matching τ is obtained at once at the end of the time interval by solving a market clearing optimization problem. The matching is only performed if the bid-ask-spread is negative.

The market clearing problem is formulated as a linear optimization problem, with the objective to maximize social welfare i.e. combining the surplus of both buyers and sellers. This reads as

$$\max_{\mathbf{q}_\omega^\tau} \quad \sum_{\omega \in \mathcal{L}} -q_\omega^\tau \, p_\omega \tag{3a}$$

$$\text{s.t.} \quad 0 \leq q_\omega^\tau \leq q_\omega, \quad \forall \omega \in \mathcal{A}_s \tag{3b}$$

$$q_\omega \leq q_\omega^\tau \leq 0, \quad \forall \omega \in \mathcal{B}_s \tag{3c}$$

$$\sum_{\omega \in \mathcal{L}} q_\omega^\tau = 0, \qquad [p^\tau] \tag{3d}$$

The quantities q_ω^τ are constrained between 0 and the standing order q_ω for the asks, as in (3b), and between the standing order q_ω and 0 for the bids, as in (3c). The balanced matching is imposed through (3d). Solving (3) yields the set of quantities $\mathbf{q}_\omega^\tau$ of the matching, with the set Ω^τ of matched orders including all orders ω such that $q_\omega^\tau \neq 0$. In view of that linear program, the equilibrium price p^τ (dual variable of the balance constraint 3d) defines the price for the trade: all sellers (resp. buyers) are to pay (resp. receive) p^τ per unit of energy exchanged.

B. Continuous double auction: The naive greedy case

In a continuous double auction, instead of considering time intervals, orders are matched in a continuous manner. This means that, at any time t during the trading session, if a new order comes and the bid-ask-spread is negative, a matching τ will be performed.

The naive greedy algorithm is in essence similar to zero-intelligence trading. When a new order ω comes, it is readily matched with another order ω' standing in a subset $\mathcal{L}_\omega$ of the limit order book $\mathcal{L}$. $\mathcal{L}_\omega$ includes all offers that are

- on the opposite side of ω (a bid for an ask, and inversely), and
- whose price limit $p_{\omega'}$ is feasible to be matched with ω without a loss.

For an extensive description and discussion, see [13]. The algorithm is referred to as naive since the standing offer ω' used for matching is randomly chosen among all standing offers in $\mathcal{L}_\omega$. In terms of matching outcome, one has $\Omega^\tau = \{\omega, \omega'\}$.

The new order ω is considered the price taker and accepts the limit price of the standing order ω' that was pending in $\mathcal{L}$. ω' is referred to as the price maker. The matching price is then $p^\tau = p_{\omega'}$. Finally, as only two orders are involved, the resulting trade quantities are

$$q^\tau = \min\{|q_\omega|, |q_{\omega'}|\} \tag{4a}$$

$$q_\omega^\tau = q^\tau \operatorname{sign}(q_\omega) \tag{4b}$$

$$q_{\omega'}^\tau = q^\tau \operatorname{sign}(q_{\omega'}) \tag{4c}$$

The standing quantities for ω and ω' are updated after the matching τ is executed. As long as orders have standing quantities different from zero, they remain in the limit order book $\mathcal{L}$ and may be matched at a later point in time. The naive greedy algorithm executes further matches until either incoming quantity reach 0, or until there are no more feasible matches available. In that latter case, it joins remaining pending orders in $\mathcal{L}$.

C. Continuous double auction: The pro-rata case

As for the naive greedy algorithm, at any time t over the trading session, the pro-rata algorithm aims at matching an incoming order immediately upon its arrival. For that, when a new order ω comes, the same subset $\mathcal{L}_\omega$ of the limit order book $\mathcal{L}$ is used. The main difference, however, is that standing orders in $\mathcal{L}_\omega$ are prioritised based on their limit prices. That is, all bids at $b^*(t)$ or asks at $a^*(t)$ are given priority over orders at lower or higher prices, respectively. Consequently, any incoming order is guaranteed to be matched at the best available price. We write Ω' be the set of all price-maker orders, which has a single element in case there is only one standing order at the best price level, and with higher cardinality if there are multiple orders standing at the same best price level.

The incoming order ω is the price taker and accepts the price of the price maker(s) Ω'. Therefore, one has

$$p^\tau = \begin{cases} a^*(t), & \text{if } \Omega' \subset \mathcal{A} \\ b^*(t), & \text{if } \Omega' \subset \mathcal{B} \end{cases}, \tag{5}$$

and $\Omega^\tau = \{\omega, \Omega'\}$.

The cumulative traded quantity is defined as the maximum quantity that can be exchanged. The individual trade quantities on the maker side are then split proportionally [14]. This yields

$$q^\tau = \min\left\{|q_\omega|, \left|\sum_{\omega' \in \Omega'} q_{\omega'}\right|\right\} \tag{6a}$$

$$q_\omega^\tau = q^\tau \operatorname{sign}(q_\omega) \tag{6b}$$

$$q_{\omega'}^\tau = q^\tau \frac{q_{\omega'}}{\left|\sum_{\omega' \in \Omega'} q_{\omega'}\right|}, \quad \forall \omega' \in \Omega' \tag{6c}$$

The standing quantities for ω and $\omega' \in \Omega'$ are updated after the matching τ is executed. As long as orders have standing quantities different from zero, they remain in the limit order book $\mathcal{L}$ and may be matched at a later point in time. The pro-rata algorithm executes further matches until either incoming quantity reach 0, or until there are no more feasible matches available. In that latter case, it joins remaining pending orders in $\mathcal{L}$.

D. Upper bound on sub-optimality gap of online matching approaches

Matching orders in an online (causal) manner, hence using only the information available at the time of arrival of a new order, implies an obvious loss of optimality compared to an offline market clearing, where all the information on orders is considered. We refer to as offline market clearing the case of the discrete double auction with only one time interval to collect offers covering the whole trading session. It is intuitively expected that the naive greedy matching algorithm should be more sub-optimal than the pro-rata matching algorithm, since orders are matched randomly. To compute an upper bound for this sub-optimality gap, we first simplify the orders to be of unit quantity. This does not lose generality, as each order can be split in multiple orders at the same price level but of unit quantity. We then define a best and a worst possible sequence of arrivals for the social welfare to be optimized.

The social welfare is maximised, as in the offline market clearing, when the highest bids are matched to the lowest asks until the lowest bid-ask spread is negative. This can be readily understood by considering than in a offline market clearing, the equilibrium is found after ranking the bids and asks in decreasing and increasing price order, respectively.

Definition 5 (best arrival sequence)**.** *Given a set of orders $\omega \in \mathcal{L} = \mathcal{A} \cup \mathcal{B}$ with $\mathcal{A}$ and $\mathcal{B}$ the sets of all asks and bids over the trading session, respectively, let $\mathcal{T}_B$ be the sequence of arrivals such that*

$$\mathcal{T}_B = \{t_\omega | \max(t^s_{a\in\mathcal{A}}) < \min(t^s_{b\in\mathcal{B}}), t^s_{b'} < t^s_{b''} \text{ if } p_{b'} > p_{b''}\} \quad (7)$$

Then $\mathcal{T}_B$ is one of the possible sequence of arrivals that yields the best (optimal) social welfare.

Definition 5 implies that all asks arrive first and then all the bids arrive in decreasing order of price. In this way, we make sure that the bids with higher prices are matched to the asks with lower prices until the incoming bids have lower prices than the standing asks. The orders matched will be the same as the ones selected from an offline algorithm, even if with a different pricing algorithm, yielding the optimal social welfare. With a similar approach, we identify the sequence of arrivals that yields the worst possible social welfare, when all the asks are standing and the bids arrive with increasing price order.

Definition 6 (worst arrival sequence)**.** *Given a set of orders $\omega \in \mathcal{L} = \mathcal{A} \cup \mathcal{B}$ with $\mathcal{A}$ and $\mathcal{B}$ the sets of all asks and bids over the trading session, respectively, let $\mathcal{T}_W$ be the sequence of arrivals such that*

$$\mathcal{T}_W = \{t_\omega | \max(t^s_{a\in\mathcal{A}}) < \min(t^s_{b\in\mathcal{B}}), t^s_{b'} < t^s_{b''} \text{ if } p_{b'} < p_{b''}\} \quad (8)$$

Then $\mathcal{T}_W$ is one of the possible sequence of arrivals that yields the worst social welfare.

Definition 6 relies on a sequence of arrivals such that the lowest asks are matched to the lowest bid (conditionally on $p_b \geq p_a$) minimizing the social welfare generated. The upper bound of the sub-optimality gap can then be linked to the ratio between worst and best social welfare.

Proposition 1. *The sub-optimality of matching orders in an online manner compared to an offline one is upper bounded by*

$$\epsilon = 1 - \frac{\int_0^1 \left[\Phi_b^{-1}(x) - \Phi_a^{-1}(x)\right]^+ dx}{\int_0^1 \left[\Phi_b^{-1}(1-x) - \Phi_a^{-1}(x)\right]^+ dx} \quad (9)$$

where $\Phi^{-1}_{a/b}$ is the inverse cumulative distribution function of the ask and bid price, respectively.

A proof of the above proposition is given in the Appendix. One has to notice that Proposition 1 applies for continuous cumulative distribution functions. In case of a discrete approximation, whenever $\Phi_b^{-1}(x) \leq \Phi_a^{-1}(x)$ there exist possible matches that generate a social welfare that are not considered in (9). Additionally, the expression of the upper bound, in particular for the pro-rata matching algorithm, is conditional on the probability of the sequence $\mathcal{T}_W$ (or any other sequences that yield to the worst social welfare) to happen. Therefore, especially in case of large number of orders, the upper bound in expectation can be substantially increased.

Remark 1. *The upper bound Ξ of the suboptimality gap in expectation for the pro rata algorithm is*

$$\Xi = 1 - \pi \frac{\int_0^1 \left[\Phi_b^{-1}(x) - \Phi_a^{-1}(x)\right]^+ dx}{\int_0^1 \left[\Phi_b^{-1}(1-x) - \Phi_a^{-1}(x)\right]^+ dx} \quad (10)$$

with π the probability of the sequence $\mathcal{T}_W$ (or any other sequences yielding the same matches) to happen.

As for the naive greedy algorithm, the matches yielding the worst social welfare could happen also with other sequences of arrivals, since the match is selected randomly among all the feasible standing offers, increasing the value of the expected upper bound in (10).

IV. Accounting for Preferences

An appealing feature of consumer-centric electricity markets is to allow expressing preferences that enter the matching process and contribute to the price formation [8], [15]. This is motivated by the heterogeneous views of electricity consumers on the importance of certain attributes, e.g. type of electricity generation source and localization [16]. We therefore generalize here the double auction approaches introduced previously to also account for preferences. This is done by considering the possibility for an order to own certain attributes and to require attributes through the matching process.

A. Attributes and sub-markets

Preferences may be accommodated by organizing sub-markets that bids and asks can be granted access to, if they fit the required attributes. An attribute is defined as a “characteristic” of electricity, that may be used as a basis for differentiation. Let us denote by $\Lambda = \{\lambda_j\}$ this set of attributes, with $j = 1, \dots, n_\lambda$, where each attribute is a binary variable, $\lambda_j \in \{0, 1\}$. $\Lambda^{\mathcal{A}_s}$ is a subset of attributes that can only be owned by asks and may be required by bids (e.g. green producer). $\Lambda^{\mathcal{B}_s}$ is a subset of attributes that can only be owned by bids and may be required by asks, $\Lambda = \Lambda^{\mathcal{A}_s} \cup \Lambda^{\mathcal{B}_s}$. In case ownership of an attribute is relevant for both bids and asks, it must be defined in both subsets. For example, an attribute $smbus$ that indicates ownership of a "small business" certificate must be defined once as $smbus.ask \in \Lambda^{\mathcal{A}_s}$ and in addition as $smbus.bid \in \Lambda^{\mathcal{B}_s}$.

Based on this idea, for an order ω, let $\Lambda_\omega^{\text{own}}$ be a set of owned attributes and $\Lambda_\omega^{\text{req}}$ a set of required attributes from a match partner. For every ask and bid,

$$\Lambda_\omega^{\text{own}} \begin{cases} \subseteq \Lambda^{\mathcal{A}_s}, & \omega \in \mathcal{A}_s \\ \subseteq \Lambda^{\mathcal{B}_s}, & \omega \in \mathcal{B}_s \end{cases}, \quad \Lambda_\omega^{\text{req}} \begin{cases} \subseteq \Lambda^{\mathcal{B}_s}, & \omega \in \mathcal{A}_s \\ \subseteq \Lambda^{\mathcal{A}_s}, & \omega \in \mathcal{B}_s \end{cases} \tag{11}$$

Taking any possible combination of the binary elements of Λ yields a set I of 2^{n_λ} sub-markets. Each sub-market $i \in I$ is defined through exactly one tuple of binary attribute indicators λ_j. A binary indicator $\gamma_{\omega i}$ is finally introduced to indicate whether the order ω may access sub-market i.

Remark 2. *A match τ happens in exactly one sub-market i. In order for two orders ω and ω' to be matched, both ought to have access to that sub-market i, hence $\gamma_{\omega i} = \gamma_{\omega' i} = 1$.*

B. Priorities

Since each sub-market is identified through exactly one attribute combination, the matching will ensure that preferences are respected. Attributes and eventually sub-markets may be prioritized based on a willingness to pay more (bids) or a willingness to be paid less (asks), if a match occurs in given sub-markets.

Definition 7 (priority-augmented order)**.** *A priority-augmented order $\omega(q_\omega, t_\omega^s, \Lambda_\omega^{own}, \gamma_\omega, \mathbf{p}_\omega)$ is an order additionally characterized by a set of owned attributes Λ_ω^{own}, a vector $\boldsymbol{\gamma}_\omega$ of sub-market access indicators and a vector $\mathbf{p}_\omega$ of prices for all sub-markets.*

The vector $\boldsymbol{\gamma}_\omega$ constrains participation to certain sub-markets and hereby require attribute ownership from a match partner. Every sub-market an order participates in ($\gamma_{\omega i} = 1$), must be given a specific price $p_{\omega i}$. For an ask, the lower $p_{\omega i}$, the higher the priority for sub-market i. For a bid, the higher $p_{\omega i}$, the higher the priority. Different sub-markets may be assigned with the same price in which case they are prioritized equally. In the following, when referring to "attribute requirement" only the highest priority is considered, while if referring to "attribute priorities" all are considered.

C. Algorithms with priorities

Both the naive greedy and the pro-rata matching algorithms now intend to match an incoming order in a descending manner from highest to lowest prioritized sub-markets. However, this structured prioritization only applies when an order arrives at the matching engine and acts as a market taker. Once an order becomes a market maker in the LOB, it is always subject to incoming market takers. Nonetheless, the prioritization is indirect since the order is offering better limit prices for higher priorities.

The original market clearing optimization problem (3) is updated with sub-market individual prices for each order. There is now an individual trade τ_i clearing at p^{τ_i} in each sub-market. p^{τ_i} is given by the dual variable of the balance equations in (12b). The formulation becomes

$$\max_{\mathbf{q}_\omega^\tau} \quad \sum_{\omega \in \mathcal{L}} \sum_{i \in \mathcal{I}} -q_\omega^{\tau_i}\, p_{\omega i} \tag{12a}$$

$$\text{s.t.} \quad \sum_{\omega \in \mathcal{L}} q_\omega^{\tau_i} = 0, \quad [p^{\tau_i}], \quad \forall i \in I \tag{12b}$$

$$\sum_{i \in \mathcal{I}} q_\omega^{\tau_i} \le q_\omega, \quad \forall \omega \in \mathcal{A}_s \tag{12c}$$

$$q_\omega \le \sum_{i \in \mathcal{I}} q_\omega^{\tau_i}, \quad \forall \omega \in \mathcal{B}_s \tag{12d}$$

$$0 \le q_\omega^{\tau_i} \le q_\omega \gamma_{\omega i}, \quad \forall \omega \in \mathcal{A}_s, i \in I \tag{12e}$$

$$q_\omega \gamma_{\omega i} \le q_\omega^{\tau_i} \le 0, \quad \forall \omega \in \mathcal{B}_s, i \in I \tag{12f}$$

It is important to note that while naive greedy and pro-rata matching iteratively intend to match from highest to lowest priority, the optimization problem maximizes overall social welfare. Hence, in the optimization it could happen that an order is matched at a lower priority in case it generates a higher social welfare.

V. Application and Case-study

Emphasis is placed here on an application example for a single trading session, in order to thoroughly analyse the workings of the various double auctions proposed. Since simulating different type of offer sequences, our analysis permits to look into the variability of the outcomes. After introducing the test case, results are first given by taking a system-level point of view, then followed by the agent-level one.

A. Test case characteristics

Let us consider a matching market with a resolution of 5 minutes, and with trading sessions of 5 minutes right before the time of delivery. A single trading session in simulated and analysed here. We define a set Λ with 5 attributes, $\Lambda = \{smbus.bid,\ comkey.bid,\ smbus.ask,\ comkey.ask,\ green.ask\}$. Here, $smbus$ stands for an order coming from a certified small business, $green$ for electricity generated by renewable energy sources and $comkey$ is a key to identify a member of a specific community e.g. residing in Helsingør in Denmark. This set of attributes translates to having 32 sub-markets. An example of a submarket would be the tuple $[01011]$, which can be accessed by any bid that at least owns the community access key and any ask that comes from a green electricity generator inside the community.

Based on this setup, a sequence of 5.000 orders for the trading session are generated. Quantities are such that

$$q_\omega \sim 1000\, \beta(1, 3), \tag{13}$$

while 50% of all orders are bids, others are asks. 20% of all orders own the $smbus$ attribute (which then becomes $smbus.bid$ or $smbus.ask$ depending on the type of offer), 20% of all orders own the $comkey$ attribute ($comkey.ask$ and $comkey.bid$ respectively), 50% of all asks own the $green.ask$ attribute.

In parallel, a limit price p_ω^{inf} for the lowest priority (no attribute requirement) for each order is generated such that

$$p_\omega^{\text{inf}} \sim \begin{cases} 50\ \beta(2,10) + 43, & \omega \in \mathcal{A} \\ 50\ \beta(10,2) + 7, & \omega \in \mathcal{B} \end{cases}, \tag{14}$$

while then, for the highest priority of an order, a willingness to pay (if ask) or willingness to accept (if bid) is added to the above limit price, i.e.,

$$p_\omega^{\text{sup}} \sim p_\omega^{\text{inf}} + \begin{cases} -40\ \beta(1,10), & \omega \in \mathcal{A} \\ 40\ \beta(1,10), & \omega \in \mathcal{B} \end{cases}, \tag{15}$$

In terms of distribution of priorities, 30% of all orders require the $smbus$ attribute for their highest priority, 70% of all bids require the $green.ask$ attribute for their highest priority, 90% of all orders that own the $comkey$ attribute also require it for their highest priority. For all orders, an evenly distributed random number of priorities is generated between the highest and lowest ones.

Three types of cases are consequently considered: a first case without preferences with p_ω^{sup} as a limit price (so all orders trade in the same sub-market without constraints), a second case with priorities and a third case in which the highest priority is a requirement (referred to as ”attribute requirement”). Each simulation is run 100 times with the same data set, but with a different and random time of arrival for each order. Results for the discrete double auction are shown for three clearing intervals Δt (3s, 30s and 300s). Since the trading session is of 5 minutes, the optimization with a 300s clearing interval has full information of the entire order set. It is hence expected to be the benchmark approach that will yield the most optimal results.

B. System-level results

Let us first look at overall social welfare for the various cases considered. Those social welfare results are gathered in Figure 2. The results are normalized by the social welfare of the discrete double auction with full information on the entire order set.

For the discrete double auction case, the effect of shortening the clearing intervals is to lower social welfare, since working with subsets of offers over the trading session, and obtaining matches that are less optimal than if considering the whole order set at once. The decrease in social welfare is also more important for the case of using attributes, with in terms of requirement or with priorities. Generally, using attributes here also lowers social welfare since, as for shorter clearing intervals, leading to clearing markets with subset of offers. Both, the pro-rata and the naive greedy algorithm yield a lower social welfare in the priority run compared to the attribute requirement. This could be due to the fact that lower welfare generating matches are taking liquidity from the order book. However, this may be impacted by the way the priorities and prices in the test-cases were set-up.

The quantities traded through the alternative matching approaches are shown in Figure 3, where quantities are normalized by the overall volume of offers in that trading session. There, the online matching algorithms are dependent of the arrival of orders and their instant matching. This makes that, generally, higher volumes are traded with those online matching algorithms, though yielding trades that will not give the highest social welfare over the trading session. Then the more optimal the matching is, the lower the quantities traded.

C. Agent-level results

In view of the definition of offers, the agent-level analysis concentrates on the priority satisfaction, and prices actually paid (or received) by those agents. Those prices are weighted by quantities when offers get split by the matching algorithms.

The frequencies of satisfied priorities are illustrated in Figure 4. Here, 0 indicates the highest priority. Whatever the matching approach and clearing interval for the discrete case, a large share of agents and offers get their first priority, and if not their second one. Online matching algorithms are slightly worse than the discrete one for the agents to obtain their highest priorities. Those results may quantitatively be highly case-dependent, though we expect this effect to be qualitatively similar.

Finally, the distributions of prices per agent are gathered in Figure 5. At a first glance, the spread of prices is much higher for online matching algorithms, and especially the naive greedy one. Without attribute and a clearing interval of 5 minutes, a unique price is obtained. This is not the case when having smaller clearing intervals, and/or if considering attributes (required or preferred). We expect, however, that the price in expectation (i.e., here over the 100 replicates) to be the same for the various matching algorithms, given the attribute setup. This price may not be the same in expectation with the different ways of considering, or not, attributes.

VI. Conclusions

Online matching inherently possess interesting properties for application in consumer-centric electricity markets, since having good scaling properties, while accommodating asynchronicity. Naive greedy and pro-rata matching algorithms can also be modified to account for preferences, based on required or preferred attributes. Despite introducing some suboptimality, those have the advantage of simplicity of implementation and scalability over discrete double auctions based on periodic market clearing.

It is of utmost interest to further investigate those online matching approaches, to allow for more complicated offers e.g. multi-period with blocks and profiles, to accommodate some of the technical constraints of agents in electricity markets. In addition, the implementation of online matching algorithms should be re-thought in a decentralized fashion to potentially suit the requirements of peer-to-peer electricity markets. Finally, as for other electricity market types, the way those markets interact with system operators is to be studied,

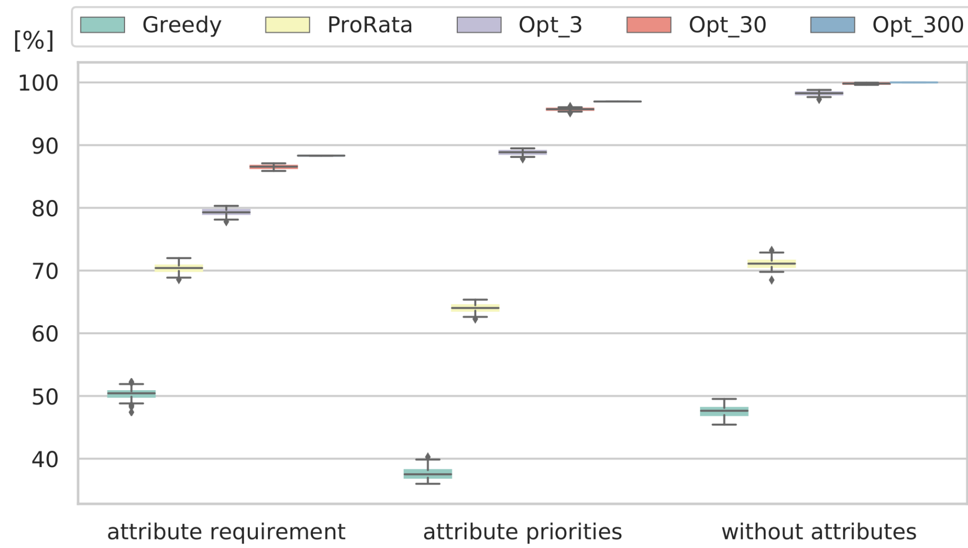

Fig. 2. Normalized social welfare for the various matching approaches, potential consideration of attributes, a well as clearing intervals for the discrete double auction (3s, 30s, and 300s).

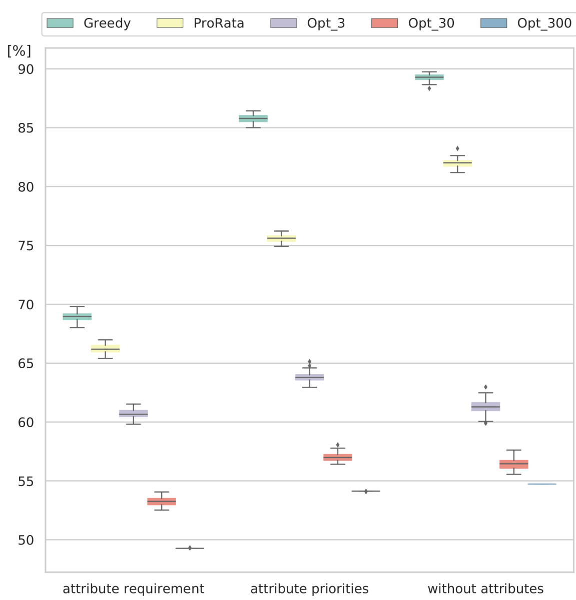

Fig. 3. Normalized quantities traded through the various matching approaches, potential consideration of attributes, a well as clearing intervals for the discrete double auction (3s, 30s, and 300s).

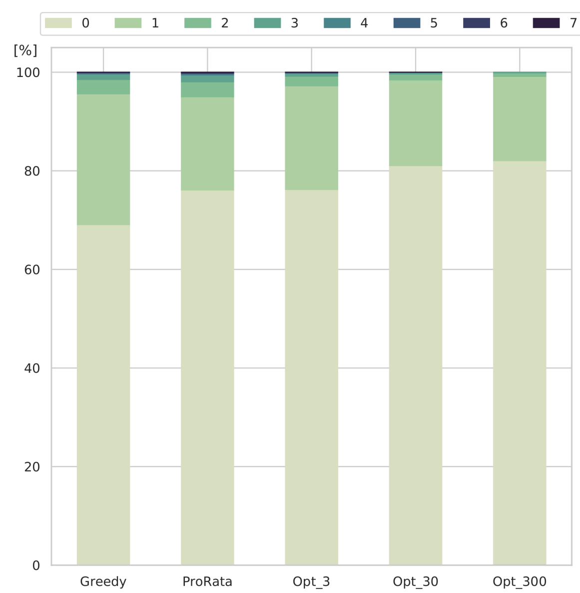

Fig. 4. Frequencies of priorities being satisfied for the various matching algorithms, a well as clearing intervals for the discrete double auction (3s, 30s, and 300s).

so as to ensure that the dispatch obtained on the power grid remains operationally feasible.

References

[1] Y. Parag and B. Sovacool, "Electricity market design for the prosumer era," *Nature Energy*, vol. 1, pp. 1–6, 2016.

[2] T. Morstyn, N. Farrell, S. J. Darby, and M. Mcculloch, "Using peer-to-peer energy-trading platforms to incentivize prosumers to form federated power plants," *Nature Energy*, vol. 3, 02 2018.

[3] T. Sousa, T. Soares, P. Pinson, F. Moret, B. Thomas, and E. Sorin, "Peer-to-peer and community-based markets : A comprehensive review," *Renewable & Sustainable Energy Reviews*, vol. 104, pp. 367–378, 2019.

[4] M. Andonia, V. Robu, D. Flynn, S. Abram, D. Geach, D. Jenkins, P. McCallum, and A. Peacock, "Blockchain technology in the energy sector: A systematic review of challenges and opportunities," *Renewable & Sustainable Energy Reviews*, vol. 100, pp. 367–378, 2019.

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

[6] J. Guerrero, A. Chapman, and G. Verbic, "Decentralized p2p energy trading under network constraints in a low-voltage network," *IEEE Transactions on Smart Grid*, 2019, available online.

[7] F. Moret and P. Pinson, "Energy collectives: A community and fairness based approach to future electricity markets," *IEEE Transactions on Power Systems*, 2019, available online.

[8] E. Sorin, L. Bobo, and P. Pinson, "Consensus-based approach to peer-to-peer electricity markets with product differentiation," *IEEE Transactions on Power Systems*, vol. 34, no. 2, pp. 994–1004, 2019.

[9] T. Morstyn, A. Teytelboym, and M. D. Mcculloch, "Bilateral contract

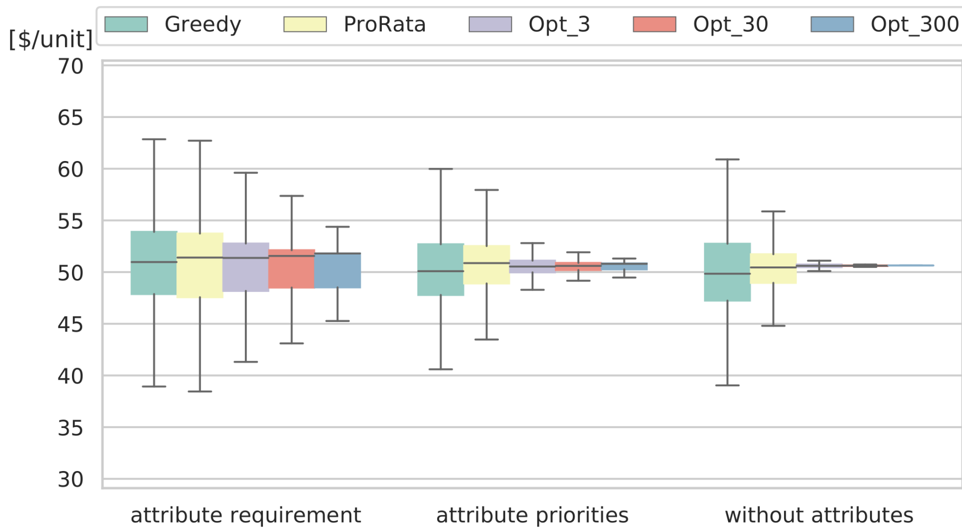

Fig. 5. Prices per agent for the various matching approaches, potential consideration of attributes, a well as clearing intervals for the discrete double auction (3s, 30s, and 300s).

networks for peer-to-peer energy trading," *IEEE Transactions on Smart Grid*, vol. 10, no. 2, pp. 2026–2035, March 2019.

[10] F. Moret, T. Baroche, E. Sorin, and P. Pinson, "Negotiation algorithms for peer-to-peer electricity markets: Computational properties," in *Proceedings of 2018 Power Systems Computation Conference (PSCC)*, 2018.

[11] K. Neuhoff, N. Ritter, A. Salah-abou-El-Enien, and P. Vassilopoulos, "Intraday markets for power: Discretizing the continuous trading?" *DIW discussion Paper*, 2016.

[12] M. D. Gould, M. A. Porter, S. Williams, M. McDonald, D. J. Fenn, and S. D. Howison, "Limit order books," *Quantitative Finance*, vol. 13, no. 11, pp. 1709–1742, 2013.

[13] A. Blum, T. Sandholm, and M. Zinkevich, "Online algorithms for market clearing," *Journal of the ACM*, vol. 53, no. 5, pp. 845–879, 2006.

[14] H. Martin and S. Otterson, "German intraday electricity market analysis and modeling based on the limit order book," in *2018 15th International Conference on the European Energy Market (EEM)*, June 2018, pp. 1–6.

[15] T. Morstyn and M. McCulloch, "Multi-class energy management for peer-to-peer energy trading driven by prosumer preferences," *IEEE Transactions on Power Systems*, pp. 1–1, 2018.

[16] Y. Yang, H. S. Solgaard, and W. Haider, "Value seeking, price sensitive, or green? analyzing preference heterogeneity among residential energy consumers in denmark," *Energy Research & Social Science*, vol. 6, pp. 15 – 28, 2015.

Appendix A
Proof of Proposition 1

Let $\phi_a(p)$ and $\phi_b(p)$ be the probability density function (pdf) of asks and bids respectively. Moving from a discrete to a continuous domain, as the asks are all unitary, one can build the supply curve (as in a common market clearing) as the inverse of the cumulative distribution function (cdf) of the ask prices $\Phi_a^{-1}(x)$, with $x \in [0,1]$. In fact, the supply curve is created by ordering the asks with increasing prices and cumulatively summing the quantities. Following the same reasoning, the demand curve can be linked to the cdf of the bid prices $\Phi_b(p)$. In the case of best social welfare, the bids arrives in decreasing order, therefore the demand curve of the correspondent market clearing is proportional to $\Psi_b^{-1}(x)$, with $\Psi_b(p) = 1 - \Phi_b(p)$. In case of the worst social welfare, the bids arrives in increasing order, therefore the demand curve of the correspondent market clearing is proportional to $\Phi_b^{-1}(x)$. See Figure 6 for a graphical representation of both cases.

We can therefore express the best social welfare (SW_B) and the worst social welfare (SW_W) as

$$SW_B = M \int_0^1 \left[\Psi_b^{-1}(x) - \Phi_a^{-1}(x)\right]^+ dx \qquad (16)$$

$$SW_W = M \int_0^1 \left[\Phi_b^{-1}(x) - \Phi_a^{-1}(x)\right]^+ dx \qquad (17)$$

with M the maximum cumulative quantity used to scaling the demand and supply function to a unitary domain and $[\cdot]^+$ the positive part function. The upper bound of the suboptimality (ϵ) gap becomes

$$\epsilon = 1 - \frac{SW_W}{SW_B} = 1 - \frac{\int_0^1 \left[\Phi_b^{-1}(x) - \Phi_a^{-1}(x)\right]^+ dx}{\int_0^1 \left[\Psi_b^{-1}(x) - \Phi_a^{-1}(x)\right]^+ dx} \qquad (18)$$

concluding the proof by using $\Psi_b^{-1}(x) = \Phi_b^{-1}(1-x)$.

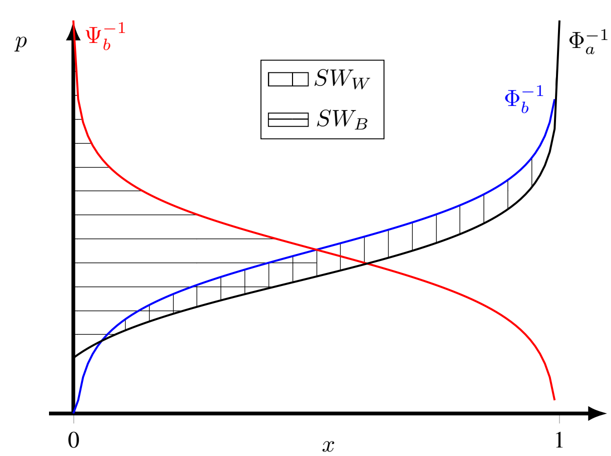

Fig. 6. Inverse cumulative distribution function of the ask prices $\Phi_a^{-1}(x)$, bid prices $\Phi_b^{-1}(x)$ and bid prices yielding best social welfare $\Psi_b^{-1}(x)$.