

Trading Agents for the Smart Electricity Grid

Perukrishnen Vytelingum, Sarvapali D. Ramchurn,
Thomas D. Voice, Alex Rogers, and Nicholas R. Jennings
Intelligence, Agents, Multimedia Group, School of Electronics and Computer Science,
University of Southampton, UK
{pv, sdr, tdv, acr, nrj}@ecs.soton.ac.uk

ABSTRACT

The vision of the Smart Grid includes the creation of intelligent electricity supply networks to allow efficient use of energy resources, reduce carbon emissions and are robust to failures. One of the key assumptions underlying this vision is that it will be possible to manage the *trading* of electricity between homes and micro-grids while coping with the inherent real-time dynamism in electricity demand and supply. The management of these trades needs to take into account the fact that most, if not all, of the actors in the system are self-interested and transmission line capacities are constrained. Against this background, we develop and evaluate a novel market-based mechanism and novel trading strategies for the Smart Grid. Our mechanism is based on the Continuous Double Auction (CDA) and automatically manages the congestion within the system by pricing the flow of electricity. We also introduce mechanisms to ensure the system can cope with unforeseen demand or increased supply capacity in real time. Finally, we develop new strategies that we show achieve high market efficiency (typically over 90%).

Categories and Subject Descriptors

J.4 [Social and Behavioural Sciences]: [Economics]

General Terms

Economics

Keywords

Market Mechanism, Continuous Double Auction, Smart Grid

1. INTRODUCTION

The creation of the “Smart Grid” is widely recognised as one of the most important challenges faced by developed countries this century [8]. The vision includes, but is not limited to, the creation of intelligent electricity supply networks that use energy resources efficiently, reduce carbon emissions and increase robustness to failures [8, 9, 15]. In the UK specifically, the *2008 Climate Change Act* mandates a 32% reduction of carbon emissions by 2020 and 80% by 2050. To this end, in recent years, new technologies such as smart meters and micro-grids (where electricity is generated and

used within a local network which may or may not be part of the larger grid) have been developed. These technologies, coupled with energy storage technology and embedded green energy generators (e.g., biomass, wind power, and solar), will support a complete decentralisation of the supply and management of electricity that will be more efficient and reduce reliance on fossil fuels.

One of the key assumptions underlying this vision of the Smart Grid is that it will be possible to manage the *trading* of electricity between homes and micro-grids, while coping with the inherent real-time dynamism in electricity demand and supply. Moreover, the management of these trades needs to take into account the fact that the actors in the system are self-interested. This means, they may misrepresent their preferences (e.g., amount of electricity required, the capacity they can supply and prices they would accept) in order to maximise their profit. Within current electricity markets, such attempts to “game” the system are often mitigated by strict regulation and audits. However, this approach would be impractical if every household is a potential supplier and/or consumer. In such systems it is also critical that trades take into account the limited capacity of *transmission* lines. When these lines are overloaded, they may break down and cause, in the best case, damage to the system and a minor blackout or, in the worst case, a massive blackout.¹ These significant challenges call for the need to build decentralised autonomous systems that are self-organising and achieve high levels of efficiency. In this context, the multi-agent systems paradigm has been advocated as one of the main approaches to building such systems. In particular, market-based approaches have been particularly successful at achieving high levels of efficiency when having to deal with large decentralised systems composed of self-interested agents [5]. These market mechanisms are particularly useful when agents prefer to keep their preferences private (as opposed to revealing them to a trusted centre that optimises the state of the entire system) and only have a local view of the environment. While building large scale systems in this way has been particularly successful (e.g., stock markets, Internet auctions), so far, applications to the electricity grid have been limited.

Against this background, in this paper we develop and evaluate a novel market-based mechanism and novel agent-based trading strategies for the Smart Grid. Our mechanism naturally manages the self-interested actions of the participants, while guaranteeing a high level of surplus and ensuring that transmission lines are never overloaded. In more detail, we consider that each node in the electricity network can contain buyers or sellers (e.g., individuals, whole neighbourhoods, or generators) that aim to buy or sell electricity on a day-ahead basis (most markets are run a day ahead

Cite as: Trading Agents for the Smart Electricity Grid, Vytelingum *et al.*, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. 897-904
Copyright © 2010, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

¹The largest power outage in history, in the Northeast US in 2003, resulted in the total loss of electrical power to 55M people and was attributed to the failure of a single overhead power line that sagged due to excessive thermal heating and touched nearby vegetation.

as large generators have physical limits on how quickly they can change their supply rate). Our approach differs from mechanisms that currently exist in the wholesale electricity market in that we do not assume that buyers and sellers, which we term agents, will truthfully reveal their reserve prices and consumption/generation pattern. Specifically, our mechanism is based on the Continuous Double Auction (CDA) which allows agents to make offers continuously in the market and improve upon these until a transaction is possible (i.e. a match between a buyer and a seller is found). While the CDA has been shown to be very efficient (in surplus maximisation) for the trading of goods or services (e.g., in the Nasdaq and NYSE stock market), little is known as to what its performance would be in electricity markets which are significantly different from traditional applications. First, in contrast to typical goods that go directly from seller to buyer, electricity flows along the paths of least resistance to any node in the network. Hence, when transmission lines are congested, it is important that more profitable transactions are prioritised. Given this, we implement a congestion pricing scheme for electricity flow in transmission lines. Second, if agents use more electricity than they bought one day-ahead (since the flow of electricity cannot be controlled), generators in the system have to cope with the real-time demand in order to guarantee that the system stays balanced (i.e. supply meets demand). To this end, we design novel *balancing mechanisms* that ensure buyers pay a fair price for their unexpected demand and generators are also fairly compensated for accommodating such sudden increases in demand. In so doing, the market guarantees the best deal for all agents even if they fail to predict accurately their demands and supply. Moreover, we show that agents cannot game the balancing mechanism, as doing otherwise, they are guaranteed to make a loss.

In more detail, this paper advances the state of the art as follows:

1. We provide a new electricity market mechanism for self-interested agents. Our mechanism manages congestion within the system by pricing the flow of electricity computed by DC flow approximation.
2. We introduce a novel a balancing scheme that ensures the system can cope with unforeseen demand or increased supply capacity. Using our scheme, unmatched offers in the day-ahead market are used to respond to changes in supply and demand at the most competitive prices. More importantly, such prices are generated in *real-time* and ensure agents bid truthfully in the market.
3. We provide a method for evolving transmission line congestion prices over time to increase the efficiency of our proposed electricity market. This method allows our market mechanism to achieve close to optimal allocations.
4. We provide new measures, which we term Dynamic Locational Marginal Prices (DLMPs) that indicate levels of endemic congestion in areas of the network and can be used as a guide for future infrastructure improvements.

When taken together, this is the first attempt to create a scalable and efficient electricity market that ensures security of supply (i.e., is resilient to failures and dynamic demand and supply).

The rest of this paper is organised as follows. Section 2 contains a review of the related work. In Section 3 we describe mathematical models that motivate our proposed electricity market. The market mechanism is explained in detail in Section 4 and Section 5 describes our trading agents. Section 6 evaluates our market mechanism and agents empirically and Section 7 concludes.

2. RELATED WORK

Research in Smart Grid technologies has leapt forward in recent years (see [8, 9] for overviews). Furthermore, agent based auction simulations have been used to model existing electricity networks [1, 4] and the macro-economics literature has examined their design [17, 11]. However, there has been little research into designing market rules for the Smart Grid where capacity constraints and automated network management (i.e., ensuring secure transmission) become serious issues due to the complexity of managing the increased number of market participants [6].

In more detail, the leading attempt to create an intelligent agent based market system under capacity constraints is the AMES Wholesale Power Market Test Bed [14]. Their work is based on the concept of Locational Marginal Pricing whereby generators in the system are paid according to location in the network and given the transmission line capacity constraints. AMES models an electricity market and computes Locational Marginal Prices (LMPs) for allocations computed a day ahead of the actual consumption for 48 half-hour settlement periods. In their model, bids are linear price sensitive demand and supply preference curves, combined with fixed, price insensitive demands. Their model closely approximates recently introduced electricity markets in the US (New Hampshire) and New Zealand. However, as they acknowledge, the impact of the agents (particularly generators) misreporting their preference curves can significantly reduce the efficiency of the system. In this paper we do not make any such assumptions, however we do use the AMES method to compute the optimal allocation against which we compute the efficiency of our experimental results [18].

Our work is also related to most research on Continuous Double Auctions with automated bidding agents. Previous studies on such systems have found them to be highly efficient — often averaging at only a few percent away from optimality. This is true even for “zero-intelligence” (ZI) agents that only bid randomly above their limit price [10], as well as for agents that have much more intelligent strategies [3]. In this case, the two bidding strategies we use are ZI [10] which is the known baseline (since they have the simplest behaviour) and AA [19], the best performing algorithm (shown to reach efficiencies of 99.9% in static environments) in the literature respectively. Finally, our work also relates to other congestion control schemes, in particular Dual algorithms for computer networks [12]. However, these algorithms are not directly applicable as they assume simple node responses, whereas our proposed mechanism runs a more complex CDA. Moreover, flow in electricity networks is governed by physical laws as opposed to network routing mechanisms.

3. BACKGROUND

We consider a system where there are a set of agents that can be both buyers $b \in B$ and sellers $s \in S$. The electricity network is a graph composed of nodes $n^1, n^2, \dots \in N$ and transmission lines $t \in T$ and is noted as $G = (N, T)$. A given node n is a point at which agents reside and can either generate or consume electricity that induce flows on the lines connected to it. A transmission line is a pair of nodes $t = (n, n')$ where $n, n' \in N$. Now, each buyer or seller is located on one of the nodes (e.g., individual houses, or a generator) and there may be multiple buyers and sellers at each node (e.g., neighbourhoods, multiple wind turbines). We denote as n_b and n_s the nodes where a buyer b resides and seller s resides respectively. Alternatively, we specify as $B_n = \{b|n_b \in N\}$ as the set of buyers at node n and $S_n = \{s|n_s \in N\}$ as the set of sellers at node n . We also, denote the transmission lines connected to a given node n as $T_n = \{t|t = (n, n') \in T\}$.

3.1 Properties of Buyers and Sellers

Each buyer has a *fixed* demand for electricity $q_b^{fixed} \in \mathbb{R}^+$ (i.e., for which it is ready to pay any price). Each buyer also has a marginal cost function that dictates how much it is willing to pay for a given quantity $q_b \in [0, q_b^{max}]$ beyond their fixed demand, where $q_b^{max} + q_b^{fixed}$ is the maximum amount of electricity it needs. Similarly, each seller's cost function says how much the seller is willing to sell a quantity $q_s \in [0, q_s^{max}]$ for where q_s^{max} is the maximum amount of electricity it can generate. We consider typical marginal cost functions (see [13]) for buyers and sellers as follows:

- Buyer b has cost function $p_b = c_b - d_b q_b$, where constants $c_b, d_b \in \mathbb{R}^+$ and d_b is usually very small compared to c_b . In particular, d_b represents b 's price sensitivity to increasing the quantity it is buying.
- Seller s has cost function $p_s = x_s + y_s q_s$, where constants $x_s, y_s \in \mathbb{R}^+$. In particular, y_s represents the seller's increasing costs with increasing production.

When buyers and sellers at different nodes trade electricity, they generate power flows in the transmission lines throughout the network where the amount of flow in each line is determined according to the properties of the line and the voltage generated at each node in the network. We discuss these properties in more detail next.

3.2 Properties of Transmission Lines

Electricity generators in the network will typically generate an alternating current (AC) [14]. However, for tractability, we compute the AC power flow using the DC flow approximation, as is common in the study of power networks [18]. In the DC flow approximation, each transmission line carries power according to the properties of the line and voltage angles (created by an alternating voltage) applied at its ends.² In more detail, a transmission line has a reactance $r_t \in \mathbb{R}^+$ which dictates how much power will flow through it given the angle difference at its end points. If $\delta_n, \delta_{n'} \in \mathbb{R}^+$ are voltage angles (in radians) at the different ends of a transmission line $t = (n, n')$, then the power flow in the line is given by $q_t = (\delta_n - \delta_{n'})/r_t$.

Now, in traditional electricity networks, transmission lines are owned by heavily regulated operators that charge a fixed fee to the users to *connect* to the network. In this paper, we look at the novel approach of having transmission line owners charge users according to the amount of electricity they transport for them. This is preferable to a flat connection fee as it prices low profit transactions out of the market, thus, prioritising more profitable transactions when resources are limited and increasing the overall market efficiency. However, it has never been possible because electricity networks typically date back more than 50 years and are not equipped with appropriate information networking capabilities. However, in the future the Smart Grid will be network-enabled which, in turn, will allow transmission lines to detect the effect of each node's input/output in the system and hence charge each user for its usage of the line. The system presented in Section 4.2 gives an idea of how these flows might be computed. We assume that all lines are owned or maintained by a *network manager* that applies some form of congestion pricing to the network [12]. The network manager endows each transmission line with price function $p_t = w_t + z_t \alpha_t |q_t|^{\alpha_t}$, where $w_t, z_t, (\alpha_t - 1) \in \mathbb{R}^+$ are constants defined per line and $q_t \in [-q_t^{max}, q_t^{max}]$ where q_t^{max} represents the maximum that the line can carry. Lines capacity is limited due to the physical properties of copper wires in them, they heat up with increasing flow and

²This relies on the assumptions that resistances in transmission lines are small and voltage magnitudes stay close to some fixed value. These assumptions are not unreasonable, with typical net power loss due to electrical resistance over the UK transmission network at less than 3% [11].

may sag excessively if overheated. Hence, we assume transmission lines charge $p_t = \infty$ for $|q_t| \geq q_t^{max}$. Note, since electricity can flow in any direction along a line, it is beneficial for transmission lines to *pay* agents that create counter flows (i.e., demand and supply at its endpoints that reverse the flow of electricity to some degree) across them since this reduces the total flow.

Having described the properties of all the actors in the system, we next describe how, traditionally, the efficient allocation is computed in a day-ahead electricity market and how the market copes with unexpected demand or supply intra-day. In so doing we establish the benchmark for the mechanisms we develop in this paper.

3.3 Allocations and the Balancing Mechanism

In this section we detail two main aspects of traditional electricity markets. First, we elaborate on how they typically compute the efficient allocation given the reported day-ahead cost functions and quantities of all actors in the system. Second, we discuss the balancing mechanism used to charge agents when they do not conform to their stated day-ahead consumption and generation profiles.

3.3.1 Computing the Efficient Allocation

To compute the maximally efficient allocation, we must assume that all agents have reported their cost functions and buyers have also revealed their fixed demand. This extends the model provided in the AMES testbed in order to compute the optimal allocation of electricity in the network based on the cost functions specified by all the actors in the system. In particular, we add the marginal cost functions of the transmission lines as another penalty to the objective function in the AMES model. The model consists of a convex optimisation problem which describes the goal of efficiency maximisation subject to modelled physical constraints on the system. More formally, we maximise³

$$\sum_{b \in B} q_b (c_b - d_b q_b) - \sum_{s \in S} q_s (x_s + y_s q_s) - \sum_{t \in T} q_t (w_t + z_t q_t^{\alpha_t}), \quad (1)$$

subject to $|q_t| \leq q_t^{max}$ for all $t \in T$ and

$$\sum_{b \in B_n} q_b + q_b^{fixed} - \sum_{s \in S_n} q_s + \sum_{t \in T_n} q_t = 0 \forall n \in N, \quad (2)$$

Here, the first constraint restricts the quantity that can flow in a line and constraint (2) ensures that the total amount of power entering a node is the same as the amount leaving. Since flows are calculated according to voltage angles differences, as a reference point we also specify that the voltage angle at node n^1 is zero, $\delta_{n^1} = 0$.

If we set $w_t = z_t = 0$ for all $t \in T$, then the optimal value of this is the maximum efficiency possible for given buyer and seller agents' preferences. All measurements of efficiency we use are given in reference to this maximum. When congestion prices are non-zero, the market efficiency of the optimal allocation may be obtained by taking the optimal value of (1) and adding $\sum_{t \in T} q_t (w_t + z_t q_t^{\alpha_t})$. This optimal allocation efficiency acts as the benchmark against which will we evaluate our market mechanism.

As we stated above, existing studies of the power market that computes optimal allocations rely on the actors in the system reporting their cost functions truthfully a day ahead. While this may be plausible for the network manager whose main goal is to ensure security of supply, it is unrealistic for the individual buyers and sellers who are simply interested in maximising their profits by either charging a higher price (for sellers) or requesting a lower price (for buyers). Moreover, since the market works on a day-ahead basis, agents may misreport due to the difficulty of accurately predicting

³We used IBM ILOG CPLEX to implement and solve the optimisation problem.

their preferences ahead of time. For example, many sources of renewable generation (such as wind and solar power) are inherently variable. We consider how such issues are currently managed.

3.3.2 The Balancing Mechanism

So far, we have considered what happens in an electricity market a day ahead of actual consumption, in which agents must submit their required quantities for every half-hourly period of the day. This means they need to accurately predict their consumption pattern or generation capacity one day-ahead in order to minimise their costs (i.e., by trading exactly what they need to). If their predictions are wrong, the agents consume or produce more or less than what the allocation allowed them to. For example, sudden cold weather will induce buyers to use up more than they expected (and bought in the day-ahead market). Usually, if demand is not as expected, suppliers have to cope by generating more than they were expecting to in order to keep the system stable.⁴ Hence, generators have to be compensated for their extra production and buyers have to pay for their extra demand.

The pricing of this extra demand is usually catered for by what is termed a *balancing mechanism*. This mechanism usually assumes the existence of a *pool rate* (e.g. as in the UK) which dictates the cheapest cost of generating an extra unit of power at the point when it was needed (see [11] for more details). This rate is used to charge extra demand and reward extra supply. However, there are several issues with existing balancing mechanisms. First, they are usually run independently of the day-ahead allocation process and ignore the bids that were submitted in the mechanism a day-ahead. Hence, the pool rate computed in a balancing mechanism is not the best price that agents could receive or pay, and this may decrease the efficiency of the system. Second, the pool rate can be easily gamed by the agents in the system by simply understating their generation capacity or overstating their demanded levels and therefore they get compensated for over generation and under consumption intra-day. Third, and not least, the balancing mechanism is run as an *off-line* process that involves significant auditing effort that will simply not scale to the number of transactions envisioned in the Smart Grid.

Against this background, in the next section we introduce an electricity market mechanism to remedy the above issues. Our mechanism does not require agents to be truthful or willing to reveal their preferences to a centre. The mechanism also ensures that the flows in the system are secure (i.e., no lines are overloaded) and that the balancing mechanism can scale up to large numbers of agents in real-time.

4. ELECTRICITY MARKET DESIGN

Our market mechanism is composed of three main parts:

1. **The Trading Mechanism** — this dictates the rules of interactions of the agents trading in the system.
2. **The Security Mechanism** – this computes the flows generated in the system by each trade and informs the market mechanism of the transmission line charges for every trade in the system.
3. **The Online Balancing Mechanism** — this uses information generated within the market mechanism to settle prices for extra demand and supply intra-day in real-time.

⁴If demand is higher than supply, generators will start to slow down and the frequency of the voltage will drop to low levels which might be harmful to generators and devices connected to the grid. The only way to bring the frequency back up is to increase generation in the system or cut power to the loads.

The combination of these three mechanisms define the complete context within which we develop new *trading agents* that can automate the trading procedure and use advanced trading strategies (in Section 5) to maximise the profit of each individual buyer or seller they represent. Moreover, as we show later in Section 6 these trading strategies also generate a very high level of system-wide efficiency while guaranteeing secure supply and they do not rely on any regulation to ensure the system efficiency is not affected by agents misreporting their preferences (i.e., this is left to the effect of open market competition). Moreover, our online balancing mechanism ensures that the agents cannot game the prices given to them in real-time.

We detail each component of our mechanism in the following subsections and thoroughly empirically evaluate it in section 6 to show that it can generate high levels of efficiency when the participating trading agents have both advanced and simple strategies.

4.1 The Trading Mechanism

Here, we describe the protocol of our trading mechanism, i.e. the rules that define the exchange process between buyers and sellers in the market to manage their offers to buy (bids) and sell (asks). Our approach is based on the use of the CDA which we detail next.

4.1.1 The Continuous Double Auction

The CDA allows multiple buyers and sellers to compete in a market for homogeneous goods, as opposed to Single-Sided Auctions such as the English Auction (with a single seller) or the Dutch Auction (with a single buyer). Specifically, the CDA lasts a fixed period of time, known as the *trading period* (at the end of which the market closes and no more offers are accepted) and is termed as continuous because it allows traders to submit offers at any time during a trading period and the market clears continuously (i.e., whenever a new transaction is possible between an accepted bid and ask). Traders are allowed to submit two types of offers, the elastic *limit order* (to buy or sell with a price constraint) and the inelastic *market order* (to buy or sell an exact quantity at any price or nothing at all). Note that market orders are usually cleared immediately (given no price constraints) if there are enough unmatched offers in the orderbook (where all bids are recorded). Information about the market state is made public to all market participants through the *bid and ask orderbooks* where the accepted bids and asks are listed, respectively. The current market price (i.e. the latest transaction price) is also published in the traditional CDA format.⁵

We design our electricity market as a variant of the CDA. Specifically, we provide a trading mechanism⁶ to allow the allocation of electricity between multiple consumers and producers over a transmission network, subject to the critical constraint of a secure flow (i.e., the flow through each transmission line, as a result of power allocation between consumers and generators, is within its capacity). To this end, in the following subsections we define the three key parts of our trading mechanism, namely; (i) the quote-accepting policy which dictates what offers get accepted or rejected in the market, (ii) the market clearing procedure which matches the offers to buy and sell electricity and, (iii) the information revelation policy which dictates what information gets revealed to the agents to incentivise competition in the system.

4.1.2 The Quote-Accepting Policy

This policy defines how we decide which offers to accept or reject in the market. To elaborate on this policy, we first define the types

⁵There exist many variants of the CDA, based on different market protocols — see [7] for more details.

⁶As with the traditional CDA, buyers and sellers are allowed to submit bids and asks at any time during a trading period.

of offers that can be submitted in the market and make-up of the bid and ask orderbooks.

1. **Market-order Bid:** Typically, buyers have some amount of fixed demand (e.g., to heat their house in the winter, or light up streets). To meet this, buyer b can submit an inelastic market-order bid, $bid_b^M = bid(b, q_b, n_b)$ to buy *exactly* q_b units of power at *any* price at any time during a trading period.
2. **Limit-order Bid:** Buyers can also have price sensitive demand (e.g., turning on their washing machine, or charging their home storage device). Thus, buyer b , located at node n_b can submit an elastic limit-order bid, $bid_b^L = bid(b, q_b, p_b, n_b)$ to buy up to q_b units of power at a maximum unit price of p_b at any time during a trading period.
3. **Limit-order Ask:** Seller s located at node n_s can submit an elastic ask, $ask_s^L = ask(s, q_s, p_s, n_s)$ to sell up to q_s units of power at a minimum unit price of p_s at any time during a trading period. In general, sellers could also have market-orders (as buyers above) but in the electricity domain, generators do not sell fixed amounts of electricity at any price.
4. **Bid orderbook:** Unmatched bids are queued in a bid orderbook $orderbook_{bid}$, ordered by first decreasing bid prices such that the higher and, thus, more desirable offers (from a seller's perspective) are at the top of the orderbook and, second, by earliest arrival times.
5. **Ask orderbook:** Unmatched asks are queued in an ask orderbook $orderbook_{ask}$, ordered by increasing ask prices such that the lower and, thus, more desirable offers (from a buyer's perspective) are at the top of the orderbook and, second, by earliest arrival times.

Electricity Market							
Market Price: 36.50, Volume=501.49							
Buy Orders				Sell Orders			
Buyer	q_b	p_b	n_b	Seller	q_s	p_s	n_s
b_8	1.9	36.9	n_8	s_7	5.350	36.44	n_7
b_9	0.849	36.62	n_9	s_4	2.350	36.76	n_4
b_2	7.2	12.98	n_2	s_6	9.0	37.44	n_6
				s_1	8.05	38.04	n_1
				s_9	7.1	38.33	n_9

Table 1: Bid and Ask orderbooks.

Against this background, we adopt the *NYSE quote-accepting policy* which dictates that only bids and asks that improve on themselves are accepted in the market. That is, a buyer cannot decrease the bid price of any of its accepted bids and, conversely, a seller cannot increase the ask price of any of its accepted asks. This effectively speeds up the allocation process. Note that because market orders are not constrained by price, they are usually immediately accepted and matched, unless there are insufficient quantities to match such an order. In such a case, they are usually placed at the top of the orderbooks, with unmatched market orders sorted by earliest arrival times.

4.1.3 The Market Clearing Procedure

The market clears continuously. Now, by definition, orderbooks contain the unmatched offers in the market. Whenever the state of the orderbooks changes (i.e., with a new offer or an offer being improved on), the market attempts to clear by finding the best $match(bid_b, ask_s, q_{secure})$ (i.e. a buyer b and a seller s are willing to transact q_{secure} units of power) in the orderbooks (based

on algorithm 4.1.3) by iterating down the orderbooks and trying to match each bid with each ask until the bid price is less than the ask price and no match is possible. When the best match is found, the market clears the matched bid and ask and a transaction occurs between the successful buyer and seller. Before doing so however, the mechanism needs to check that the flow is *secure* and hence obtain the cost of transmitting q_{secure} in the network. How this is achieved is discussed in the security mechanism in section 4.2. For now, we will assume that for a trade between agents b and s , the security mechanism returns $cost_{q_{secure}}^{b,s}$ as the transmission line cost for the transaction (if q_{secure} overloads any line the $cost_{q_{secure}}^{b,s} = \infty$).

Algorithm 1 The Clearing Algorithm

Require: $orderbook_{bid}, orderbook_{ask}$

- 1: **for all** $bid_b \in orderbook_{bid}$ **do**
 - 2: **for all** $ask_s : orderbook_{ask}$ **do**
 - 3: $match_{i,j} \leftarrow match(bid_b, ask_s, q_{secure})$
 - 4: **end for**
 - 5: $ask_s^* \leftarrow \arg \min_{ask_s \in orderbook_{ask}} (p_s + cost_{q_{secure}}^{b,s})$
 - 6: **subject to:** $q_{secure} > 0$
 - 7: **clear** bid_b and ask_s^* for q_{secure} units at a price defined by our market pricing.
 - 8: **end for**
-

Thus, the market prices a transaction between a matched bid bid_b and a matched ask ask_s using the traditional κ -pricing method. Specifically, a buyer pays $bid_b - \kappa (bid_b - ask_s - cost_{q_{secure}}^{b,s})$, where $\kappa \in (0, 1)$, a seller is paid $ask_s + (1 - \kappa) (bid_b - ask_s - cost_{q_{secure}}^{b,s})$, where $\kappa \in (0, 1)$ and $cost_{q_{secure}}^{b,s}$ is distributed by the security mechanism among the transmission lines as payment proportional to their additional flow (see section 4.2 for more details). Note that because a transaction can result in counterflow in a transmission line, it can happen that $cost_{q_{secure}}^{b,s}$ is negative. While at the end of the trading period, the cumulative payments to transmission lines will equal their charging rates, during the trading period, trades that result in counterflow (as explained in section 3.2) can be more profitable for traders as transmission costs will be less or even negative (i.e. being paid to use the transmission network). Thus, the mechanism effectively incentivises traders to decongest the network.

4.1.4 The Information Revelation Policy

As discussed earlier, the orderbooks are usually either completely or partially public (i.e., a certain depth of the orderbook is visible). In our mechanism, to favour a decentralised approach, the orderbooks, and all transaction prices, are made completely public. Now, the outstanding bid and ask (i.e., the best and highest bid and the best and lowest ask) are usually key information that can be extracted from the market. The outstanding bid represents the price a seller should accept to transact and, conversely, the outstanding ask represents the price a buyer should offer to transact. In our electricity market, these prices no longer hold as matches between bids and asks are subject to secure flow over the network (i.e., include transmission line costs). This prevents buyers and sellers from making informed decisions about what to bid next in the market. Hence, it is important to give indications to buyers and sellers about the types of transactions they might make with other agents at the same or other nodes in the network based on transmission constraints. To this end, we extend the notion of LMPs (see section 2) to provide an indication of the prices at each node.

We define two terms to replace the outstanding bids and asks, namely a Dynamic Locational Marginal Price (DLMP) for buyers,

$DLMP_{n_b}^B$ at node n_b and $DLMP_{n_s}^S$, which are the minimum and the maximum price a buyer b and seller s at nodes n_b and n_s would need to accept to trade respectively. These DLMPs include transmission costs. The DLMPs are computed as follows:

1. $DLMP_n^B$ is calculated as the lowest cost of buying +1 unit of power from the unmatched sellers and the cost of transmission of an additional unit (from the security mechanism — see section 4.2). As with our matching in our market clearing, we first find the cheapest *secure* matches, by iterating through the unmatched asks in the ask orderbook, that can sell +1 unit of power. Our $DLMP_n^B$ is the sum of the cheapest matches (that can supply exactly +1 unit of power) at each node n . However, if $q_{secure} \leq 1$, we set the $DLMP_n^B$ as infinite as the buyer could never acquire at least +1 unit however high it bids.
2. $DLMP_n^S$ is calculated as the highest price of selling +1 unit of power to the unmatched buyers at node n and the transmission cost for the additional unit (from the security mechanism — see section 4.2). As with our matching in our market clearing, we first find the highest *secure* matches, by iterating through the unmatched bids in the bid orderbook, that can buy +1 unit of power. Our $DLMP_n^S$ is the sum of these highest matches (that can buy exactly +1 unit of power). However, if $q_{secure} \leq 1$, we set the $DLMP_n^S$ as 0 as the seller could never sell at least +1 unit even at price 0.

The market *individually* publishes the $DLMP_n^B$ to all buyers and $DLMP_n^S$ to all sellers within node $n \forall n \in \mathcal{N}$. In the next section, we show how to compute the cost of transmission using congestion pricing principles and DC flow approximation of the flows in the network.

4.2 The Security Mechanism

Here we show how to determine whether a match between a bid and an ask can occur and what transmission cost to charge to the transaction. Thus, given $bid(b, q_b, p_b, n_b)$ and $ask(s, q_s, p_s, n_s)$, we need to determine a quantity q_{secure} in the range $q_{secure} \in [0, \min(q_b, q_s)]$ (since the seller can never buy or sell more than they offer or demand) that the network can handle between them. In the next section we first compute q_{secure} based on capacity constraints and given transmission line charges. Then we show how to create such charges and how they help to maximise efficiency in the system.

4.2.1 Making Secure Transactions

Given the current state of the network where some trades have happened, for every line $t \in T$ we compute the quantity q_t that flows in each transmission line using the DC power flow model. Now, the DC power flow equations state that for each line $t \in T$, the flow is given by $q_t = (\delta_n - \delta_{n'})/r_t$ (see section 3.2) where δ_n and $\delta_{n'}$ are the voltage angles at its end nodes n, n' . Given values for the amount of net load or generation at each node $n \in N$, if the net power transfer is zero, or if the value for one node is left as an unknown, then the DC power flow equations have a unique solution. Furthermore, as it is linear, the map which takes the vector of net power flow at each node to the set of net power flows through each transmission line may be expressed as a matrix. To calculate the net power flow at a node $n \in N$ it is only required to sum $q_b - q_s$ for all $b \in B_n$ and $s \in S_n$. Thus, to calculate the effect of a trade of amount q between buyer b and seller s , it is only necessary to multiply the given matrix and a vector with q in the entry for the node where b is located and $-q$ in the entry for the node where s is located, and zeros elsewhere. Since the given matrix may be pre-calculated from knowledge of the network topology and transmission line properties, this calculation can be very fast.

Having obtained q_t for every t in the network, we can compute the net effect of the transaction of $q_b = \min(q_i, q_j)$ and $q_s = -\min(q_b, q_s)$ at nodes n_i and n_j respectively. To this end we compute the effect of $q_b = 1$ and $q_s = -1$ on every line and let θ_t be the resulting flow of power in t . Then we can compute, for all lines: $q^* = \arg \max_{q'} \left(\int_0^{q'} \theta_t dq_t \right)$ subject to: $q_t^* \leq q_t^{max}$ for all $t \in T$. If the q_t^* exists (i.e., the constraint is satisfied for all lines in the network), $q_{secure} = q^*$. Then, given the new flows q'_t for every t , generated in the system using the same procedure as above, we can compute the overall marginal cost for one unit of power in the transaction between b and s as:

$$\begin{aligned} cost_{q_{secure}}^{b,s} &= cost^{b,s}(\min(q_b, q_s)) \\ &= \frac{\sum_{t \in T} (q'_t - q_t) \cdot cost_t(q_t)}{\min(q_b, q_s)} \end{aligned}$$

where $cost_t(q)$ is the integral of $p_t = w_t + z_t(\alpha_t - 1)|q_t|^{\alpha_t}$ with respect to q . If lines will be overloaded by any q , then the bid and ask cannot be matched as the cost of the transaction is infinite.

4.2.2 Transmission Line Pricing

When there is no congestion pricing, that is $w_t = z_t = 0$ for all $t \in T$, then the optimal allocation gives the best possible market efficiency given the buyer and seller agents' preferences. However, in that case the CDA is unlikely to perform well, because many different traders are competing for limited resources. Congestion pricing, as detailed above, can improve efficiency by pricing lower profit transactions out of the market, thus prioritising more profitable transactions. From convex optimisation theory [2] we know that if w_t is chosen correctly for all $t \in T$, then the optimal allocation is the not only the same whether or not capacity constraints are followed, but it also coincides with the optimal allocation when there is no congestion pricing. Since CDAs are typically very efficient for unconstrained markets, we would expect our market mechanism to perform very well under such circumstances.

In order to find optimal congestion pricing, we adapt the transmission line fees over multiple iterations of the market. Each line independently adjusts its pricing in a decentralised way with the intention of charging more if it is congested and less if it is uncongested. Specifically, we have a constant $\gamma = 0.95$ such that at the end of a market iteration, for each $t \in T$, if $|q_t| > \gamma q_t^{max}$ then w_t is increased by 5% and if $|q_t| < \gamma q_t^{max}$ then w_t is decreased by 5%. The motivation for this approach is that, if the CDA can get close to optimal allocation at each iteration, this process is similar to a steepest descent method for solving the dual problem to market efficiency maximisation [2].

Having defined how transactions would be managed in the day-ahead market, we now examine the more dynamic intra-day transactions and specify a novel online balancing mechanism.

4.3 The Online Balancing Mechanism

At the end of the day-ahead trading period, no more transactions are possible, with the orderbooks showing the unmatched bids and asks from buyers and sellers still willing to trade during the real-time transmission. Because demand and supply cannot be accurately predicted, demand and supply need to be *balanced* in real-time for stability (see Subsection 3.3).

During the balancing stage, the market is responsible for *balancing* the unpredicted demand and supply with respect to the market-allocated ones. It does so in the following way:

1. $+\Delta demand_i$ is additional power required by a buyer b_i . Because it is already being drawn from the network, it needs to be settled as an inelastic bid, i.e. a market-order bid submitted in the bid orderbook.

2. $-\Delta demand_i$ is extra power demanded by a buyer b_i . To cover the *short position*,⁷ the buyer needs to increase the market demand with a market-order to sell $\Delta demand_i$ regardless of price.
3. $-\Delta supply_j$ is extra power that seller s_j could not supply. To cover the *short position*, the buyer needs to increase the market supply by submitting a market-order bid to buy $\Delta supply_j$.

These market-order bids and asks are placed at the top of the bid and ask orderbooks respectively (because these traders *have* to buy and sell at any price). The market then clears these orders with matches only based on secure quantities (as in our clearing algorithm) and not price. $+\Delta demand_i$ and $-\Delta supply_j$ are then priced at $DLMP^B$ while $-\Delta demand_i$ is priced at $DLMP^S$ that represent the best prices (in real-time) in the market.

Now, $DLMP^S \leq o_{bid} < o_{ask} \leq DLMP^B$ (where o_{bid} and o_{ask} are outstanding bids and ask respectively). At the end of the trading period, because the intra-marginal (i.e. the cheaper) sellers have been allocated, the extra-marginal (i.e. the more expensive) ones are left such that $DLMP^B$ is now very high. Similarly, because the intra-marginal (i.e. the richer) buyers have been allocated, leaving the poorer ones unallocated, $DLMP^S$ is very low.

Errors in demand and supply quantities result in the buyer having to buy $+\Delta demand_i$ at much higher prices and covering for $-\Delta demand_i$ at very low prices, making losses in both cases. Similarly, the seller needs to cover $-\Delta supply_j$ at higher prices than he was paid for supplying $\Delta supply_j$ and, thus, making a loss. This clearly shows it is in the best interests of the agents to truthfully reveal and accurately predict their demand or generation capacity a day ahead. We omit a more formal proof due to lack of space.

5. THE TRADING AGENTS

Given the market protocols, trading agents have to strategise in the market based on their private preferences. Now, because their preferences are defined as continuous functions (see Section 3), we discretise the traders' preferences as piecewise constant incremental cost curves [13] to form a set of endowments (pairs of quantity q and limit price ℓ). Given their endowments, we will describe, in the next subsection, how a buyer and a seller strategise in the market by forming a bid based on its limit price ℓ_i^B and an ask based on its limit price of ℓ_j^S respectively using trading strategies based on the AA and ZI strategies.

5.1 The ZI Strategy

We first consider the Zero-Intelligence strategy. As developed by Gode and Sunder [10], the ZI strategy randomises over the whole space of offers allowed in the system, with a bid price p_i drawn from a uniform distribution between 0 and the limit price ℓ_i^B and the ask price from a uniform distribution between the limit price ℓ_j^S and some arbitrarily high price that would always be accepted in the market. Because it makes random and uninformed decisions (by ignoring all market information), ZI is rationally the baseline strategy, providing a lower bound on the efficiency of the system.

5.2 The AA-EM Strategy

We next consider the more efficient AA strategy [19]. Due to properties of our market mechanism a single market equilibrium price [16] at which the market can optimally clear no longer exists. In particular, different nodes have different $DLMPs$ and the best buyers may not trade with the best sellers because of the power flow constraints. As a consequence of this, the AA strategy, which

⁷Covering a position means buying or selling power such that the market demand equals supply. E.g. a buyer who overbids for $q + \Delta q$ and draws only q needs to sell Δq .

is based on the micro-economic theory of transaction prices convergence to an equilibrium, is not appropriate for our system. To this end, we develop a novel variant of the AA (called AA-EM) that is tailored to Electricity Markets.

Specifically, we make a number of fundamental changes.⁸ Because the outstanding bid and ask are no longer valid (see Subsection 4.1.4), the AA-EM trader instead considers $DLMPs$ as a target of what it requires to transact. Furthermore, because there is no equilibrium, transaction prices are no longer a good indication for the 'equilibrium price' (which represents the expected transaction price in a standard CDA). Instead, transaction prices within the same node (and, hence, subject to the same power flow constraints) are a better measure for equilibrium price. Thus, we use a weighted average of transaction prices (with significantly more weight given to transactions within the same node) to calculate that equilibrium price in our variant of the AA. Finally, because a buyer might have the outstanding bid which is higher than the outstanding ask (that typically result in a transaction in a standard CDA) and still not transact because of power flow constraints, the AA-EM has an additional bidding rule. In more detail, if a buyer or seller has the outstanding bid or ask (i.e. it is at the top of the orderbook), it sets its target as the $DLMP^B$ and $DLMP^S$ respectively.

6. EMPIRICAL EVALUATION

In this section, we empirically evaluate our trading mechanism using the ZI and AA-EM strategies for a number of configurations for the electricity network. Specifically, we draw the parameters a, b, c, d for buyers and sellers from uniform distribution and create a random network (with a square lattice) with 16 nodes, with a buyer and a seller at each node.

In our first set of experiments, we consider the performance of our mechanism with different capacity constraints in the system, by varying the average capacity constraint of each transmission line. Our results in Figure 1 show that the structure of our mechanism is fairly efficiency, with a *lower bound* efficiency ranging from 88% to 96% of the optimal (when considering the baseline Zero-Intelligence Strategy). Now, when we employ the more intelligent AA-EM strategy, our system efficiency (the total profits of all traders) increases, varying then from 92% to 99% of the optimal. As expected of the system, the efficiency drops as capacity decreases, reaching a point, with a capacity tending towards 0, where only buyers and sellers within the same node can trade with each other. Our mechanism indeed degrades well in such a situation, with our mechanism ensuring system security at all times. Furthermore, we empirically demonstrate that the efficiency of the optimal allocation is reduced if it is calculated from misreported agent preferences (as predicted in Section 1). In particular Figure 1 shows how the optimal allocation is worse than AA (and ZI in less congested networks) when malicious traders start *shading* their bids by up to 10%.

Next, we evaluate our mechanism within different standard graph topologies that exist in an electricity network (e.g., fully connected, a ring, a sparse and a small-world network). As expected, from Figure 2, we can see that AA-EM outperforms ZI in these different topologies, with efficiencies of above 90% of the optimal. One interesting observation is that the structure of mechanism tends to be less efficient in ring and sparse topologies. This is to be expected, as the lack of interconnectivity inhibits non-local trading and thus reduces competition within the market. In such cases, the gap between ZI and AA-EM performance is greatest and the benefit of *market intelligence* though better strategies is accentuated.

⁸Because of space constraints, we omit the specific rules of the

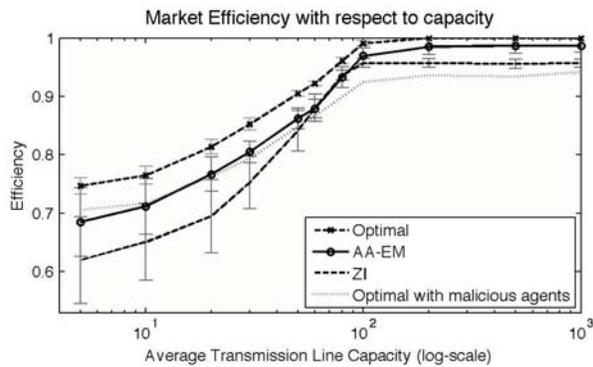


Figure 1: System Efficiency against average capacity of transmission lines. The optimal is calculated according to Equation 1. Malicious traders shade their bids or asks by up to 10%.

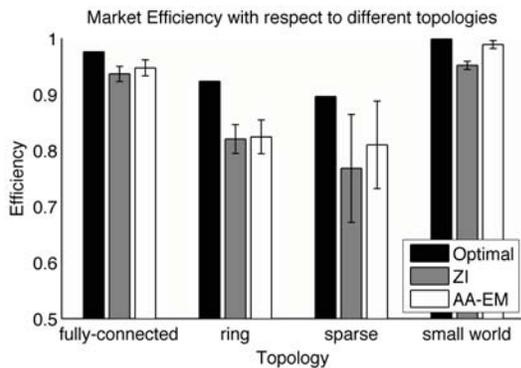


Figure 2: System Efficiency given different topologies (namely a fully connected, a ring, a sparse and a small-world network).

Finally, in our last set of experiments, we validate the learning mechanism of the transmission lines. From Figure 3, the system efficiency, on average, increases slowly over the trading days. As expected from the manner transmission lines adapt the charges (see Subsection 3.2), the congested lines will gradually increase the charges while the non-congested ones will decrease their charges. Because of these new charges, the richer buyers and cheaper sellers are more likely to trade and generate more surplus, which then increases the efficiency of the system.

7. CONCLUSIONS

We have proposed a novel electricity market which operates in a non-cooperative setting and strictly avoids overloading transmission lines. Furthermore, we have provided a novel online balancing mechanism that is tightly coupled to our market mechanism that ensures agents cannot game the day-ahead market in order to make profits in the intra-day transactions. We have demonstrated the high level of efficiency this system achieves in a variety of simulated environments and provided novel trading strategies that can generate up to 99% efficiency in the market and establish a lower bound of 86% when simple behaviours are used in the system.

Our market mechanism contains novel components which advance the state of the art in electricity market design. We have also developed effective automated trading strategies, which could be

AA-EM though they can easily be recreated from its description.

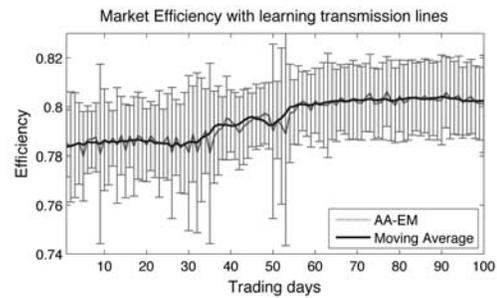


Figure 3: System Efficiency with transmission lines learning.

adopted by users of the Smart Grid to maximise their benefits. Furthermore, we have explored ways for future network operators to manage congestion, protect against transmission line overloads and plan for further infrastructure improvements.

Further open research questions that remain are how the mechanism's stability and performance are affected by random effects and changes in the market, how to further expand the scalability and functionality of the system and whether transmission losses or measurement error can have a significant negative impact on efficiency. In the future, we also intend to develop a real-time simulation of the electricity market and empirically evaluate losses of traders trying to game the system or with poor prediction of their demand and supply.

8. REFERENCES

- [1] J. Bower and D. W. Bunn. A model-based comparison of pool and bilateral market mechanisms for electricity trading. *The Energy Journal*, 21(3):1–29, 2000.
- [2] S. P. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [3] N. T. Chan, B. D. LeBaron, T. Poggio, and A. W. Lo. Agent-based models of financial markets: A comparison with experimental markets. In MIT Sloan Working Paper No. 4195-01, 2001.
- [4] P. Chand, G. Grozev, P. da Silva, and M. Thatcher. Modelling australia's national electricity market using NEMSIM. In *Proceedings of SimTect*, Melbourne, 2008.
- [5] S. H. Clearwater. *Market-based Control: A Paradigm for Distributed Resource Allocation*. World Scientific, 1996.
- [6] R. A. F. Currie, G. W. Ault, R. W. Fordyce, D. F. MacLeman, M. Smith, and J. R. McDonald. Actively managing wind farm power output. *IEEE Transactions on Power Systems*, 23(3):1523–1524, 2008.
- [7] D. D. Davis and C. A. Holt. *Experimental Economics*. Princeton University Press, Princeton, NJ, 1993.
- [8] U. S. Department Of Energy. Grid 2030: A National Vision For Electricity's Second 100 Years, 2003.
- [9] R. Galvin and K. Yeager. *Perfect Power: How the MicroGrid Revolution Will Unleash Cleaner, Greener, More Abundant Energy*. McGraw-Hill Professional, 2008.
- [10] D. K. Gode and S. Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119–137, 1993.
- [11] C. Harris. *Electricity Markets: Pricing, Structures, and Economics*. Wiley, 2005.
- [12] F. Kelly. Charging and rate control for elastic traffic. *European transactions on Telecommunications*, 8(1):33–37, 1997.
- [13] D. S. Kirschen and G. Strbac. *Fundamentals of Power System Economics*. John Wiley and Sons, 2004.
- [14] H. Li and L. Tesfatsion. Development of open source software for power market research: The ames test bed. *Journal of Energy Markets*, 2(2):111–128, 2009.
- [15] D. MacKay. *Sustainable energy: without the hot air*. UIT, Cambridge, 2009.
- [16] A. Mas-Colell, W. Whinston, and J. Green. *Microeconomic Theory*. Oxford University Press, 1995.
- [17] D. Newbery. Electricity liberalisation in Britain: the quest for a satisfactory wholesale market design. *Energy Journal*, pages 43–70, 2005.
- [18] J. Sun and L. Tesfatsion. Dc optimal power flow formulation and solution using quadprog. In ISU Economics Working Paper No. 06014, 2007.
- [19] P. Vytelingum, D. Cliff, and N. R. Jennings. Strategic behaviour in continuous double auctions. *Artificial Intelligence Journal*, 172 (14):1700–1729, 2008.