

Automated Trading for Smart Grids: Can It Work?

Barry Laffoy, Saraansh Dave, Mahesh Sooriyabandara

► **To cite this version:**

Barry Laffoy, Saraansh Dave, Mahesh Sooriyabandara. Automated Trading for Smart Grids: Can It Work?. David Hutchison; Takeo Kanade; Madhu Sudan; Demetri Terzopoulos; Doug Tygar; Moshe Y. Vardi; Gerhard Weikum; Wilfried Elmenreich; Falko Dressler; Vittorio Loreto; Josef Kittler; Jon M. Kleinberg; Friedemann Mattern; John C. Mitchell; Moni Naor; Oscar Nierstrasz; C. Pandu Rangan; Bernhard Steffen. 7th International Workshop on Self-Organizing Systems (IWSOS), May 2013, Palma de Mallorca, Spain. Springer, Lecture Notes in Computer Science, LNCS-8221, pp.1-13, 2014, Self-Organizing Systems. <10.1007/978-3-642-54140-7_1>. <hal-01291487>

HAL Id: hal-01291487

<https://hal.inria.fr/hal-01291487>

Submitted on 21 Mar 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Automated Trading For Smart Grids: Can It Work?

Barry Laffoy^{1,2}, Saraansh Dave^{1,3}, and Mahesh Sooriyabandara¹

¹ Toshiba Research Europe Ltd, Telecommunications Research Laboratory

² Department of Computer Science, University of Bristol

³ Industrial Doctorate Centre in Systems, University of Bristol
(saraansh.dave), (mahesh.sooriyabandara)toshiba-trel.com

Abstract. This paper applies basic economic principles which have been developed in financial markets to a future smart grid scenario. Our method allows for autonomous bidding for electricity units to create an emerging market price for electricity. We start with replicating the popular Zero-Intelligence-Plus algorithm and setting it in a electricity supplier-consumer scenario. We identify significant weaknesses of applying this in an electricity market especially when intermittent sources of energy are present or when the supplier to consumer ratio is very small. A new algorithm (ZIP-260) is proposed which includes a measure of fairness based on minimising the deviation across all un-matched demand for a given period. This approach means that no consumer in the system is constantly experiencing an electricity supply deficit. We show and explain how market conditions can lead to collective bargaining of consumers and monopolistic behaviour of suppliers and conclude with observations on automated trading for smart grids.

1 Introduction

Financial markets have been used in recent history as a robust method for matching buyers and sellers. Since the late 1990's computers have been increasingly used as intelligent agents to trade on behalf of human operators at a much higher frequency. High frequency trading has been bolted onto pre-existing system regulation and this may enhance the effects where billions of dollars have been lost in a span of just 15 minutes [1] [2]. The root cause of some of these 'dips' such as the flash crash of May 2006 have not been unilaterally agreed on but have been classed as a systemic failure [1]. However, financial markets can be said to be a very efficient method for solving resource allocation problems in decentralised systems with multiple actors.

Smart grids carry the promise of real-time energy management, diagnostics and fault correction with the aim of integrating renewable energy sources and decarbonised transport solutions [3]. The smart grid is characterised as integrating information and communication technologies with the electricity network to create a responsive and adaptable network. The primary challenge with intermittent energy sources and real-time balancing is trying to match demand patterns to those of supply. In a future where decarbonisation of energy is paramount, the ability to respond to demand fluctuations will become more difficult. Demand response [4] as a method of addressing this balancing issue can be described as reflecting the value of electricity at a point in time given the state of the system.

This view resembles the resource allocation problems found in real-time trading markets and we propose to investigate the application of algorithmic trading as

an efficient mechanism to manage electricity demand for households with variable electrical supply. We apply the Zero-Intelligence-Plus (ZIP) algorithm [5], a well known automatic trading algorithm, to this problem scenario. Initially the algorithm is directly mapped onto a scenario including electricity suppliers and domestic consumers. After evaluating its performance, changes are made to allow for a more fair consumption distribution. Even after promoting fairness in our system we identify inefficient market behaviour and monopolistic behaviour when there is a supply surplus. This observation of different behaviour between deficit and surplus states of the market lead us to implementing a 2-state switching algorithm for the suppliers which accounts for this. We identify some key limitations of applying the ZIP algorithm and provide regulatory and algorithmic requirements that are needed for real-time pricing of electricity when trying to balance supply and demand.

This paper is organized as follows; §2 covers relevant literature, §3 outlines the system model, §4 details the simulation scenarios used and results, §5 presents our analysis and §6 concludes on the work and outlines areas for further research.

2 Background Literature

Demand response has been classified as incentive based or priced based programs (IBP or PBP respectively). IBP includes more direct aspects such as direct load control or emergency demand response while PBP utilises more market oriented solutions such as real time pricing or time of use pricing. Time of use pricing has been shown to create a bias of consumer behaviour [6] which causes ‘rebound’ peaks to be seen during the off-peak tariff period. Clearly this method is limited in its approach and can have undesirable consequences when the solution is scaled up. Game-theoretic approaches show how the market can converge to a Nash equilibrium while improving the aggregate demand curve [7] but are somewhat limited as they usually assume a central coordinating agent (e.g. the utility) which makes it difficult to implement on a large scale. Real time pricing is more favourable towards the principle of using price as a signal of the value of electricity at that specific moment in time but there has been little work in the application and understanding of real time trading in such a system. The closest work in the smart grid domain applies transmission constraints and novel balancing mechanisms to trading agents which perform bidding and acceptance of electricity units [8].

More directly relevant work falls into the domain of agent-based computational economics [9] and continuous double auctions (CDA). CDA’s are a market system where buy and sell prices are quoted continuously and asynchronously for commodities and is widely used around the world. Previous work on CDA’s applied to automatic trading show that very simple behaviour of agents have high levels of allocative efficiency [10] and can be improved to also provide a high level of profit dispersion (an indicator of market equality) [5]. The Zero Intelligence (ZI) traders’ [10] private information only included the limit price for that trader, but by giving the agents an internal, mutable expected profit margin and memory of the last shout on the market, simple machine-learning techniques could give the traders more sophisticated behaviour [5]. The resulting Zero-Intelligence Plus (ZIP) agents would observe each shout on the market and remember if the last one was a bid or an offer, for how much, and whether or not it resulted in a trade. Based on this, the agent would update its internal profit margin and shout price.

This research also falls into the domain of current electricity markets which operate differently depending on country specific contexts. In the U.K. contracts are

agreed between those demanding electricity and generators in half-hourly periods. When there is a difference between the contracted amount and generator supply the system operator (National Grid) is responsible for real-time balancing through bids and offers [11]. This basic model may have to be reviewed to allow for smart grid implementation such that the effect of wholesale prices can be observed by the consumer and thus influence demand patterns. Currently there is a buffer between real-time consumption and the price of electricity. The role of the system operator will still apply as the real-time management must be controlled to avoid grid failures.

Directly applying the ZIP algorithm needs careful thought and experimentation as there are a few fundamental differences between electricity and financial markets. Financial markets have liquidity which assumes that a seller is able to sell their stock on the market if they wish to do so at a given price, additionally, a seller is able to withhold stock in anticipation of better prices in the future. In an electrical market, it is not possible to withhold a unit for a better price as electricity cannot be stored. Financial markets also have a large number of buyers and sellers which is in stark contrast to electricity generators and consumers where there is a much bigger difference. These differences make the application of ZIP to electricity markets non-trivial but also enhances the view that a real-time solution should be used in order to allow the system to respond at a similar rate, especially as more renewable energy is integrated into the grid.

3 System Model

This section outlines the scenario we use to test our algorithms and measure the behaviour of the system. Our simulation consists of 100 households and two electricity suppliers; fossil fuel (Supplier F) and wind power (Supplier W) suppliers. Supplier F is modelled as having constant power output reflecting a non-renewable energy source (e.g. coal or nuclear) while Supplier W has a variable power output modelled as following a Weibull distribution [12]. By including suppliers with different output profiles, we will be able to investigate whether the real-time (half-hourly) changes in supply affect the market price and whether the mechanism is efficient in allocating resources. The households are able to bid for units of electricity through the exchange (which could be organised by the utility) where Suppliers F and W are able to sell units of electricity. Households bid for a total number of units every half hour depending on their individual consumption profile whilst the suppliers are able to sell units every half hour depending on their plants output. Empirical smart meter data measured at half-hourly intervals in an Irish study [13] is used as the demand profile for households in the simulation. This also represents the future scenario of a decarbonised electricity supply where centrally generated base load capacity will support intermittent and decentralised sources. Both the heterogeneity in demand profiles and changes in aggregate supply and demand will create an environment with different market states which will test the algorithms' performance. A genetic algorithm was used to search the parameter space in order to find an optimal solution (we note that due to the nature of genetic algorithms this may well be a local optimum) where each genome expresses a phenotype in the form of a vector $[\mu, \beta, \lambda, \rho, \epsilon]$. The fitness function can be chosen according to the aim of the algorithm; a popular choice is Smith's α [14] which is a measure for how well the market converges to an equilibrium price. This measure did not promote fairness over the agents, especially in a situation where there is an overall deficit of supply. As a result we decide to use a measure which minimises the demand deficit variance across all the households (9). In the

the remainder of this section, §3.1 introduces the original ZIP algorithm and §3.2 explains the limitations of directly applying this to electricity markets and the significant changes we made to the algorithm and §3.3 covers price and fitness functions.

3.1 The ZIP Algorithm

This sub-section is a recap and overview of the ZIP algorithm as first proposed by Cliff [5]. ZIP agents have an internal, mutable expected profit margin (μ_i) which can be used with the current shout price (λ_i) to calculate the current shout price ($p_{i,t}$) as shown in (1). Based on this, it might alter its internal profit margin. The rate at which an agent's profit margin converges to the observed shout price is determined by a simple version of the Widrow-Hoff delta rule for machine learning. This gives an update rule between the times of t and $t+1$ as shown in (2), where $\Delta_{i,t}$ is defined in (3).

$$p_{i,t} = \lambda_i \times (1 + \mu_{i,t}) \quad (1)$$

$$\mu_{i,t+1} = (p_{i,t} + \Gamma_{i,t}) / \lambda_i - 1 \quad (2)$$

$$\Delta_{i,t} = \beta_i \times (\tau_{i,t} - p_{i,t}) \quad (3)$$

$$\tau_{i,t} = q_{i,t} \times (1 + \rho) + \epsilon \quad (4)$$

$$\Gamma_{i,t+1} = \lambda_{i,t} \times \Gamma_{i,t} + (1 - \lambda_{i,t}) \times \Delta_{i,t} \quad (5)$$

Here, β_i is the agent's learning rate, which must be added to the private information of each agent. $\tau_{i,t}$ is the target price, towards which the shout price is being moved and $\Gamma_{i,t}$ is the learning momentum. There is a distinction between the target price and the most recent market shout price, $q_{i,t}$ as otherwise the agents will no longer test the market if $p_{i,t} \approx q_{i,t}$, since this would lead to a value of $\Delta_{i,t}$ close to zero and thus $\mu_{i,t+1} \approx \mu_{i,t}$. The ZIP agents calculate the target value using a randomly generated relative multiplier (ρ) of the last shout price, plus a randomly generated absolute perturbation (ϵ) shown in (4). Finally, the agents were given a learning momentum coefficient $\Gamma_{i,t}$, so that when updating prices, they could take into account the history of price updates. This dampens the impact of a series of price updates in alternate directions, and reinforces the effect of a series all in the same direction.

Thus the complete behaviour of a ZIP agent is defined by the initial profit assumption μ and the values of β, λ, ρ , and ϵ . One variant of the ZIP algorithm (ZIP60) [15] allowed agents to distinguish up to six different scenarios under which prices can be changed and create a rule-set for each. The ZIP algorithm also allows buyers and sellers to start with different assumptions hence embracing heterogeneity that is present in most trading environments.

3.2 ZIP-260

Unlike ZIP, we do not impose any scenario under which prices cannot be updated. To increase flexibility the agent considers five binary factors relating to the previous shout at every time step:

1. Am I still active?
2. Am I buying or selling?
3. Was the last shout accepted or not?
4. Was the last shout a bid or an offer?
5. Would I have traded at the price of the last shout?

This leads to thirty-two distinct scenarios where each scenario is associated with a four element vector $[\beta, \lambda, \rho, \epsilon]$ which defines the price update rule as shown in (6).

$$\Delta_t = [p_t - (q_t \times (1 + \rho) + \epsilon)] \times (1 - \lambda) \times \beta + \lambda \times \Delta_{t-1} \quad (6)$$

Each element of the vector is set by a uniform random distribution between a maximum and minimum provided at the outset. This means that each simulation has $2^5 \times 4 \times 2 = 256$ parameters associated with price updates. An additional four parameters are associated with the initial profit assumption which leads to this algorithm being named ‘ZIP-260’, in keeping with the naming convention of the ZIP variants.

One major difference between using ZIP in finance and in electricity markets is the notion of trading at a loss. In finance, it is possible to withhold selling a unit to avoid trading at a loss if required. In electricity systems, once a unit of electricity is produced it must be distributed and consumed. Clearly, a supplier could store the energy but currently this has limited large-scale availability. As such, we assume the supplier has no storage capabilities and so must dispose of the electricity generated in some way once it is produced. We also assume that the supplier would sell the electricity at a loss instead of incurring the costs associated with disposing of generated electricity. ZIP shows this behaviour because there is a strict budget constraint on the agents in the form of the limit price, λ_i , which sets an absolute value on the price they are willing to meet. This represents the maximum value that a consumer will buy at and a minimum price that suppliers will sell at. Removing this constraint would return the ZIP algorithm to behaviour characterised by the ZI agents [10]. To resolve this issue we allow agents to enter a loss accepting state in certain situations. After updating the price as described earlier (6), a second simple algorithm using a Markov chain controls the state of the agent. If a trade was achieved at time t then all the agents will revert to profit-seeking behaviour. However, if a shout at time $t - 1$ was unmet then there is a small probability (P_L) that an agent will return to a loss accepting state. The internal profit assumption remains unchanged in this state but they accept any price that is quoted on the market. Figure 1 shows the distribution of unmatched supply and demand with and without the loss accepting mechanism. Positive values of units before simulation indicate a supplier while a negative value of units indicates a consumer. Figures 1(c) and 1(d) shows how well the resources are allocated. In these simulations, there is an overall deficit in supply.

3.3 Price and Fitness Functions

Two general pricing schemes are used throughout the paper; constant and varying demand price. The price for the i th agent’s j th unit is given by $\lambda_{i,j}$. The constant price function is given by (7) whilst the varying price is given by (8) where u_i is a utility function parameter. The suppliers price has been set to be a constant value of 1.5, which, while being simplistic is justified by considering only one trading day where the cost of fossil fuel powered energy and wind powered energy is assumed to be stable.

$$\lambda_{i,j} = 2.5 \quad (7)$$

$$\lambda_{i,j} = 1.20 + \frac{10}{9+j} + u_i(0.05) \quad (8)$$

Throughout the paper two measures of fitness are used, Smith’s α [14] and we introduce a measure of unit variance (σ^2) which we define in (9) where N is




Fig. 1: Distribution of units to agents before and after the simulation runs for a single day, with and without a loss accepting algorithm. Positive units indicate an excess of supply whilst negative units indicate an excess of demand.

the number of consumers and x is a vector representing the units that an agent demands and \bar{x} is the population mean.

$$\sigma^2 = \frac{1}{N} \sum_i^N (x_i - \bar{x})^2 \quad (9)$$

4 Simulation Scenarios and Results

We start by defining the supply patterns for Supplier F and Supplier W. In the case of Supplier F we simplified the plant to supply a constant and highly predictable power output, mirroring a fossil fuel plant. On the other hand, Supplier W has been modelled using a Weibull distribution with a probability distribution function as shown in (10). Where c , the scale parameter, and k , the shape parameter, take values $5.428ms^{-1}$ and 1.4128 [12] respectively. The power generated, $P(v)$, by a turbine is not a direct mapping from this distribution. A turbine has a cut-in speed (v_{cut-in}), below which it will not generate any power; a rated speed (v_{rated}), at which it produces maximum power; and a cut-out speed ($v_{cut-out}$), above which it will not produce power (for safety reasons). Between the cut-in and rated wind-speeds, it is assumed here that the power generated is a linear interpolation between zero and the rated power output as shown in (11). For the simulation analysis Supplier W was assumed to own a small wind farm consisting of five

turbines. Each turbine is modelled as behaving according to (10) and (11) with no interactions between turbines. MICON turbine specifications [16] as shown in Table 1 were taken to simulate power output for Supplier W sampled half-hourly over a 24 hour period. This gives the wind generation a more unpredictable production pattern as would be the case when integrating renewable energy to the grid.

$$f(v, c, k) = \frac{k}{c} \left(\frac{v}{c}\right)^{k-1} e^{-\left(\frac{v}{c}\right)^k} \quad v > 0; c > 0; k > 0 \quad (10)$$

$$P(v) = \begin{cases} 0 & 0 \leq v < v_{cut-in} \\ P_{rated} \times \frac{v - v_{cut-in}}{v_{rated} - v_{cut-in}} & v_{cut-in} \leq v < v_{rated} \\ P_{rated} & v_{rated} \leq v < v_{cut-out} \\ 0 & v \geq v_{cut-out} \end{cases} \quad (11)$$

Table 1: MICON TURBINE PARAMETERS


Parameter	Value
v_{cut-in}	$4ms^{-1}$
v_{rated}	$14ms^{-1}$
$v_{cut-out}$	$25ms^{-1}$
P_{rated}	$200kW$

4.1 ZIP Performance


We implement the ZIP60 algorithm using a fixed price mechanism with Supplier F as the sole source of electricity. Figure 2(a) shows the demand and supply positions before and after trading for the entire day whilst Figure 2(b) shows the corresponding average price that trades were executed at. This indicates there are two key issues with the simulation, firstly, the fixed price function is very simplistic and allows monopolistic tendencies. Secondly, it indicates a problem with the α measure itself in very uneven market conditions; as long as there is a consistent excess of demand over supply the equilibrium price will favour the supplier. While not shown here, a simulation for only Supplier W showed the same behaviour.

4.2 ZIP-260

We first implement the ZIP-260 algorithm with a fixed price function and without the loss-accepting behaviour. Figures 3(a) and 3(b) show that ZIP-260 behaves similarly to ZIP60 when optimised for α which means that by changing the fitness function we have not adversely affected the market mechanism. Figure 4 shows ZIP60's behaviour when optimised for α and σ^2 under a variable pricing function which achieves worse α performance than the ZIP-260 algorithm shown in Figure 5. Having shown that the ZIP-260 algorithm is in fact an improvement on the ZIP60 variant, and that the fitness function σ^2 is preferred to α as it encourages all available units to be traded, we implement the ZIP-260 algorithm as outlined in §3.2 and introduce σ^2 as the fitness parameter as well as a variable pricing function. We examine a system where both suppliers are operating to investigate whether fossil fuels can complement wind in this algorithmic market. Figure 6 shows the supply and demand positions before and after trade has been completed. In this




(a) Fossil fuel market.




(b) Equilibrium prices-fossil fuel.

Fig. 2: Units traded and average trading price over one day.




(a) ZIP 60 behaviour.




(b) ZIP-260 behaviour.

Fig. 3: The trading activity for hour 12 of the fossil fuel market for both ZIP60 and ZIP-260 algorithms.

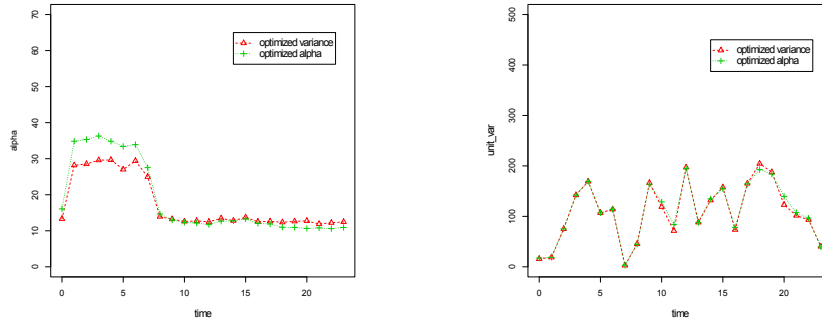


(a) α measure for fossil fuel market.



(b) σ^2 measure for fossil fuel market.

Fig. 4: Performance of the ZIP60 fossil fuel market for the two fitness measures under variable pricing.



(a) α measure for fossil fuel market. (b) σ^2 measure for fossil fuel market.

Fig. 5: Performance of the ZIP-260 fossil fuel market for the two fitness measures under variable pricing.

scenario we deliberately ensure that there is an excess of supply in the market in order to test the algorithm’s ability to distribute all units of supply. Figure 7 shows the trading activity and price equilibria for different trading periods of the system when the probability of accepting a loss (P_L) is set at 0.1. The consumers are constantly shouting zero priced bids which the loss-accepting suppliers are matching. This shows that the consumers are using their collective bargaining power to set the price very low because they have learned that the suppliers will accept it. We then reduce the probability of accepting a loss ($P_L = 0.001$) and observe the trades and equilibria at the same periods, shown in Figure 8. This now shows a more equitable market with the equilibria following the price curve of the consumers.






Fig. 6: Supply and demand positions of the joint market shown before and after trading has occurred.

5 Analysis


We show that the direct application of automated trading in the traditional sense of profit maximizing is actually not appropriate for electricity markets as it leads to




(a) Trades at hour 5.



(b) Equilibria at hour 5.




(c) Trades at hour 19.




(d) Equilibria at hour 19.

Fig. 7: Joint supplier scenario where $P_L = 0.1$


monopolistic behaviour by the suppliers. The imbalance of consumers and suppliers seriously affects the performance of the algorithm as the suppliers display monopolistic tendencies. To overcome this issue, we introduced a notion of unit variance as a fitness function which optimises for an equal distribution of outstanding demand units across all the consumers which allows for all available supply units to be allocated, as shown by Figure 1. We compare the performance of the new ZIP-260 algorithm with ZIP60 and see that both measures of performance (α and σ^2), are greatly improved which can be observed by comparing Figures 4 and 5. Unfortunately, when this is implemented in a joint supplier scenario with the probability to accept a loss set at 0.1 ($P_L = 0.1$), we see the consumers learning to abuse this loss making behaviour by consistently shouting zero priced bids. This indicates that the collective bargaining power can have a large effect such that the suppliers are not being compensated at all for the electricity they produce. This can be explained by considering the cumulative effect of the learning algorithm used in the price update for the agents. An increased acceptance of lower prices on the market causes consumers to update their prices towards that lower price. This is compounded by the frequency with which the suppliers accept loss making prices, until a price of zero is reached as shown by Figure 7. This effect can be alleviated by reducing P_L to 0.001, as shown in Figure 8 where the bid prices now follow a similar pattern to the consumers' variable




(a) Trades at hour 5.



(b) Equilibria at hour 5.



(c) Trades at hour 19.



(d) Equilibria at hour 19.

Fig. 8: Joint supplier scenario where $P_L = 0.001$.

price curves. This highlights a high sensitivity to P_L which may indicate that it is an unsuitable solution or this value should be allowed to evolve over the trading period.

6 Conclusion

Using ZIP like algorithms directly leads to non-functioning and undesirable market effects as the ratio of consumers to suppliers is uneven. To allow for this imbalance we have introduced a notion of loss-acceptance in the suppliers as well as aiming to minimise the unit variance of un-matched demand across all consumers. It was hoped that these additions would allow market prices to emerge based on the level of supply in the market which varies due to the wind powered supplier in the joint supplier scenario. However, we found that the market now becomes sensitive to the probability of loss acceptance by the suppliers where a value of 0.001 was found to promote fairer trading prices by not promoting collective bargaining power by consumers. Based on the extensive simulations and scenarios we have investigated we have to conclude that neither ZIP or ZIP-260 is, at this stage, a suitable candidate for algorithmic trading of electricity in real time. However, as we have noted, there are signs that a fair market mechanism can be created but this requires further research and testing on a wide variety of market situations.

References

1. Cliff, D.: The Flash Crash of May 6 2010: WTF? Technical report, Department of Computer Science, University of Bristol, Bristol (2010)
2. Nuti, G., Mirghaemi, M., Treleven, P., Yingsaeree, C.: Algorithmic Trading. *Computer* **44**(November) (2011) 61–69
3. Amin, S.M., Wollenberg, B.F.: Toward a Smart Grid. *IEEE Power & Energy Magazine* **3**(5) (2005) 34–41
4. Albadi, M.H., El-Saadany, E.F.: Demand Response in Electricity Markets: An Overview. 2007 IEEE Power Engineering Society General Meeting (June 2007) 1–5
5. Cliff, D.: Minimal-Intelligence Agents for Bargaining Behaviors in Market-Based Environments. Technical report, Hewlett-Packard Laboratories, Bristol (1997)
6. Ramchurn, S.D., Vytelingum, P., Rogers, A., Jennings, N.: Agent-Based Control for Decentralised Demand Side Management in the Smart Grid. In: International Conference on Autonomous Agents and Multiagent Systems AAMAS, Taiwan (2011) 2–6
7. Mohsenian-Rad, A.H., Wong, V.W.S., Jatskevich, J., Schober, R., Leon-garcia, A.: Autonomous Demand-Side Management Based on Game-Theoretic Energy Consumption Scheduling for the Future Smart Grid. *IEEE Transactions on Smart Grid* **1**(3) (2010) 320–331
8. Vytelingum, P., Ramchurn, S.D., Voice, T.D., Rogers, A., Jennings, N.R.: Trading Agents for the Smart Electricity Grid. In van der Hoek, Kaminka, Lesperance, Luck, Sen, eds.: 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010), Toronto (2010) 897–904
9. Tesfatsion, L., Judd, K.L., eds.: *Handbook of Computational Economics*. North-Holland
10. Gode, D.K., Sunder, S.: Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality. *The Journal of Political Economy* **101**(1) (1993) 119–137
11. Elexon: The Electricity Trading Arrangements: A Beginner’s Guide. Technical Report July, Elexon (2009)
12. Qian, K., Zhou, C., Li, Z., Yuan, Y.: Benefits of energy storage in power systems with high level of intermittent generation. In: 20th International Conference on Electricity Distribution. Number 0358, Prague (2009) 8–11
13. : Electricity Smart Metering Customer Behaviour Trials (CBT) Findings Report. Technical report, The Commission for Energy Regulation (2011)
14. Smith, V.L.: An Experimental Study of Competitive Market Behavior. *Journal of Political Economy* **70**(2) (1962) 111–137
15. Cliff, D.: Zip60: an enhanced variant of the zip trading algorithm. In: E-Commerce Technology, 2006. The 8th IEEE International Conference on and Enterprise Computing, E-Commerce, and E-Services, The 3rd IEEE International Conference on. (June 2006)
16. Yeh, T.h., Wang, L.: A Study on Generator Capacity for Wind Turbines Under Various Tower Heights and Rated Wind Speeds Using Weibull Distribution. *IEEE Transactions on Energy Conversion* **23**(2) (2008) 592–602