Improving Zero-Intelligence Plus for Call Markets

Jannick Hemelhof[1], Mihail Mihaylov[1,2], and Ann Nowé[1]

[1] Vrije Universiteit Brussel, Pleinlaan 2, 1050 Elsene, Belgium
[2] i.LECO, Kleinhoefstraat 6, 2440 Geel, Belgium

Abstract. Double auctions have been widely employed and studied throughout history. Two particular variants are most commonly employed: The Call Market (CALL), also known as the Periodic Double Auction, and the Continuous Double Auction (CDA). While numerous automated trading strategies exist for the Continuous Double Auction, there is a lack of high performing strategies for CALL. The former auction variant is becoming increasingly popular in the context of energy-related auctions in Smart Grids. Therefore, there is a need for efficient trading strategies. This paper explores whether a well-performing trading strategy designed for CDA, namely Zero-Intelligence Plus (ZIP) can be used in CALL. We first study the performance of the ZIP trader in CALL without any modifications. We then design several strategies and demonstrate that we can significantly improve the performance of ZIP in CALL while retaining the market's high efficiency. As a result, our modified ZIP trader can be employed by autonomous agents, e.g. for trading energy in a CALL in the Smart Grid domain.[3]

Keywords: Zero-Intelligence Plus · Call Markets · Continuous Double Auctions · Automated Trading Agents.

1 Introduction

An auction type that has seen a rise in popularity in recent years is the Call Market (CALL), also referred to as a Periodic Double Auction. The CALL mechanism collects incoming orders (also called *shouts*) in an order book and attempts to clear as many of them as possible at specific time intervals. Upon reaching such an interval the auction is *called*, meaning a call price c_p is calculated and bids with price $\geq c_p$ and asks with price $\leq c_p$ are all matched. This contrasts with the Continuous Double Auction (CDA) mechanism which attempts to continuously clear incoming orders. CALL has been utilized in modern financial exchanges [10] and in electricity markets [14, 20]. Automated traders are expected to become ever more important in Smart Grids applications, as they allow individual households to better manage their energy consumption and injection. One cannot expect these households to manage the purchase and selling of electricity on their own. This responsibility will be given to an autonomous intelligent trader, fully modelling the preferences of the household, which will participate in the electricity auction in name of the household. While CDA are quite efficient for high liquidity assets (e.g. currency) CALL are very well suited for low liquidity assets such as e.g. energy [10]. As the latter is therefore envisioned to be employed in Smart Grids [11, 19], having a wide array of adequately performing trading strategies is vital. **However, the availability of highly-performing and well-described trading strategies for this popular auction type is severely lacking.** Most automated traders have been designed for CDA and they have not been tested in a CALL setting.

In this paper we therefore present the following contributions:

- Zero-Intelligence Plus (ZIP) is a highly-performing trading strategy [6] designed for CDA settings. We extensively investigate how it performs in a CALL setting.
- We design several strategies aiming to improve the performance of ZIP in CALL. We demonstrate that our best performing strategies greatly increase the original ZIP performance in a CALL setting.

2 Call Market

Similarly to the traditional CDA such as the one described by Vytelingum et al. [18] traders in a CALL submit shouts which are collected and matched. However, contrary to the CDA, these are not matched immediately but rather at a specified point time. Upon reaching such a predetermined point the auction is *called*, meaning a *call price* c_p is calculated and bids with price $b \geq c_p$ and asks with price $a \leq c_p$ are matched. There is no standard definition for the c_p calculation and various methods have been employed by authors:

Satterthwaite and Williams [16] All submitted shouts are sorted in increasing order $s_1 \leq s_2 \leq ... \leq s_{2m}$ with $c_p = s_{m+1}$ and m representing the number of participating buyers (the authors assume there are as many sellers as buyers).

Arifovic and Ledyard [1] All bids are ordered $b_1 \geq b_2 \geq ... \geq b_N$ while asks are ordered $a_1 \leq a_2 \leq ... \leq a_N$. The authors then define k as the highest number such that $b_k \geq a_k$ after which they define $c_p = (Z + z)/2$ with

$$Z = min\{b_k, a_{k+1}\}$$

$$z = max\{a_k, b_{k+1}\}$$

Another way of calculating c_p would be to utilize the shout price which clears the largest number of units. If there would be a tie the average of possible call prices would then be utilized.

3 Related work

When looking at available traders for CALL settings we encounter a barren landscape. The majority of high-performing trading strategies such as ZIP [6], Adaptive Attitude (AA) [13] and Adaptive Aggressiveness (AAgr) [18] were all designed for CDA settings and would thus need modifications to function properly in CALL. Traders specifically designed for the latter are very hard to find, with brokers developed for the PowerTAC competition [11] coming closest. While the broker agents can be considered as CALL traders, they are specifically tailored towards the PowerTAC competition setting, employing various aspects made available by the simulation. AgentUDE utilizes energy consumption patterns in order to set shout prices while TacTex incorporates imbalance fees (coming from the balancing market) for its pricing strategy. Due to the complexity of the competition, and the incorporation of various data sources for pricing strategies in a wholesale market, the broker parts responsible for participating in the wholesale market cannot easily be generalized towards CALL.

The work of Chowdhury et al. [5] is a first attempt at decoupling the wholesale market aspect from the rest of the competition, with the authors describing some strategies for short-term energy markets (in this context a CALL) used to balance demand on the power grid. While being more generic than other wholesale market strategies the design is still centred around this particular structure rather than a more generic CALL design.

ZIP has also been adapted and tested in the context of first- and second-price sealed-bid auctions [3, 2]. However, those are out of the scope of this article.

4 Zero-Intelligence Plus

The Zero-Intelligence Plus (ZIP) trader as described by Cliff and Bruten in their 1997 paper [6] was introduced in the context of a CDA as a response (and improvement) over the Zero-Intelligence (ZI) trader proposed by Gode and Sunder [8]. Unlike the ZI trader, ZIP *"employs an elementary form of machine learning"* [12] and achieves a performance significantly closer to available human data when compared to ZI traders.

Each ZIP agent keeps track of a profit margin *"which determines the difference between its limit price and the shout price to be submitted"* [22]. For sellers the limit price represents the minimum amount the agent needs to receive for a unit in order to consider a trade. Buyers have a similar limit price where it represents the maximum amount the agent is willing to pay for a unit. If the agent had a successful transaction in the previous round it then increases its profit margin. Failed transactions decrease the profit margin, respectively. Our definition of a round is identical to that of [18]:

> *"A **trading round** is the period during which bids and asks are submitted until there is a match and a transaction occurs." [18]*

Most importantly, this profit margin update is performed after *each* successful (or failed) transaction – whether own transaction or that of another trader. Thus, all ZIP agents update their respective profit margin after each transaction.

For the sake of completeness, we include here the complete set of rules for the ZIP strategy. For **ZIP sellers** the following rules are specified [6]:

- If the last shout was accepted at a price q
  1. Any seller who asked a price lower or equal to q ($p_i \leq q$) raises its profit margin.
  2. If the last shout was a bid any seller who asked a price higher or equal to q ($p_i \geq q$) lowers its profit margin.
- Else
  1. If the last shout was an offer any seller who asked a price higher or equal to q ($p_i \geq q$) lowers its profit margin.

For **ZIP buyers** the following rules are specified [6]:

- If the last shout was accepted at a price q
  1. Any buyer who bid a price higher or equal to q ($p_i \geq q$) raises its profit margin.
  2. If the last shout was an offer any buyer who asked a price lower than or equal to q ($p_i \leq q$) lowers its profit margin.
- Else
  1. If the last shout was a bid any buyer who asked a price lower than or equal to q ($p_i \leq q$) lowers its profit margin.

The profit margin is updated using the Widrow-Hoff with momentum learning rule as introduced by [21]:

$$\Delta_i(t) = \beta_i * (\tau_i(t) - p_i(t)) \tag{1}$$

Here β_i is the employed learning rate, p_i the price at which the trader submitted its shout and τ_i the so-called *target price* which is calculated based on the shout that was most recently submitted to the auction. After a trader observes any submitted shout at a particular time t it updates its profit margin μ_i according to equation 2 (copied from [4]).

$$\mu_i(t+1) = \frac{p_i(t) + \Gamma_i(t+1)}{l_i - 1} \tag{2}$$

with l_i representing the trader's limit price and $\Gamma_i(t+1)$ being calculated through equation 3 (copied from [4]).

$$\Gamma_i(t+1) = \gamma_i(t) + (1 - \gamma_i(t)) * \Delta_i(t) \tag{3}$$

where γ_i represents the so-called *momentum coefficient*. Finally, when the trader has to submit a shout at time t it will calculate a shout price through equation 4.

$$p_i(t) = l_i * (1 + \mu_i(t)) \tag{4}$$

5 Experimental setup

The ZIP strategy will be tested in three different auctions:

Non-persistent CDA An implementation which follows the auction design described in the original ZIP paper [6] and the more recent work of [18]. Characterized by the fact that shouts do not persist between rounds we will refer to this auction type as *npCDA*.

Persistent CDA Based upon the auction described by [17] we try to provide a more realistic (when compared to real-life auctions such as the New York Stock Exchange) auction setting where shouts are not removed upon submission of an improved bid/ask but rather persist until they are matched. We will refer to this auction type as *pCDA*.

Call Market An example implementation, henceforth referred to as *CALL*, which follows the design of the CALL described by [1].

To compare the performance of the ZIP strategy in the above three auctions we use the following two objectives, which are commonly employed:

Market Efficiency How much of the theoretically available surplus was acquired by traders participating in the auction [7, 17, 9]. Ideally this would be 100%, implying that all traders are trading at equilibrium prices (i.e. the price at which demand equals supply). It is defined as [9]:

$$\frac{\sum_i p_a^{(i)}}{\sum_i p_e^{(i)}}$$

where $p_a^{(i)}$ represents the actual surplus generated by a trader i over the entire trading period while $p_e^{(i)}$ represents the theoretical surplus generated by a trader i if it had traded its goods at equilibrium price.

Market Surplus How much surplus (i.e. profit) was generated over the course of an auction [15, 17, 9, 13], where we focus on the Average Market Surplus (AMS):

$$\frac{\sum_i p_a^{(i)}}{i}$$

This allows us to reason about the amount of surplus we could expect from a typical ZIP trader when being deployed in a particular auction setting.

6 Experiments

6.1 Employed parameters

We implemented the ZIP traders as specified by [6] employing the same parameter configurations:

β Each trader generates a random value from a uniform distribution $U(0.1, 0.5)$.
γ Randomly generated for each trader from a uniform distribution $U(0, 0.1)$

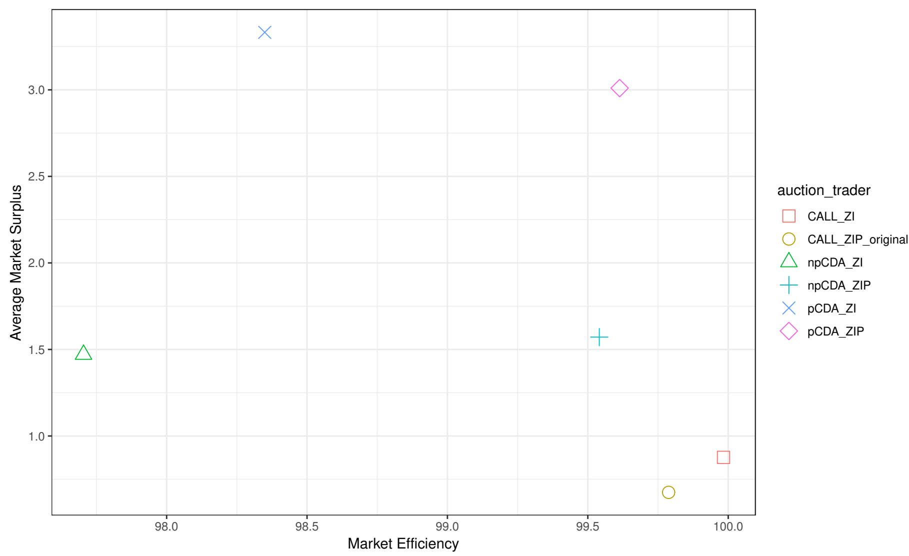

Fig. 1. Comparing the performance of various auction/trader combinations. Results are median values obtained after 1000 runs with each run ending after 1000 transactions were submitted.

μ Randomly generated for each trader from a uniform distribution $U(0.05, 0.35)$
c Perturbations were also randomly generated for each trader from a uniform distribution $U(0, 0.05)$

All participating traders have a limit price, in this scenario again randomly generated from a uniform distribution $U(0.75, 3.5)$. All auction types have 200 participating traders, 100 buyers and 100 sellers, with shouts concerning one single unit. Results are always averaged over 1000 runs.

6.2 Original ZIP performance in all auctions

For our first experiment we implemented the ZIP traders as specified in Section 6.1. Performance was measured as the AMS and observed market efficiency as obtained after 1000 transactions submitted. ZI traders serve as a "baseline" since these do not employ any intelligence. Results can be seen in Figure 1. ZIP traders outperform their ZI counterpart in npCDA which matches results published in other research [6]. Looking at pCDA we see that, although being more efficient, pCDA_ZIP obtains a smaller AMS than pCDA_ZI which can be attributed to the fact we employ the original ZIP implementation designed for a npCDA (a modified version adapting it for pCDA was proposed by [17]). This was done to ensure we compare the *same* trader, thereby providing a fair playing field. At the same time we emphasise the need for adaptation when utilizing ZIP in a new auction setting. The latter behaviour can also be observed in our CALL: original ZIP actually performs worse than ZI on both performance metrics.

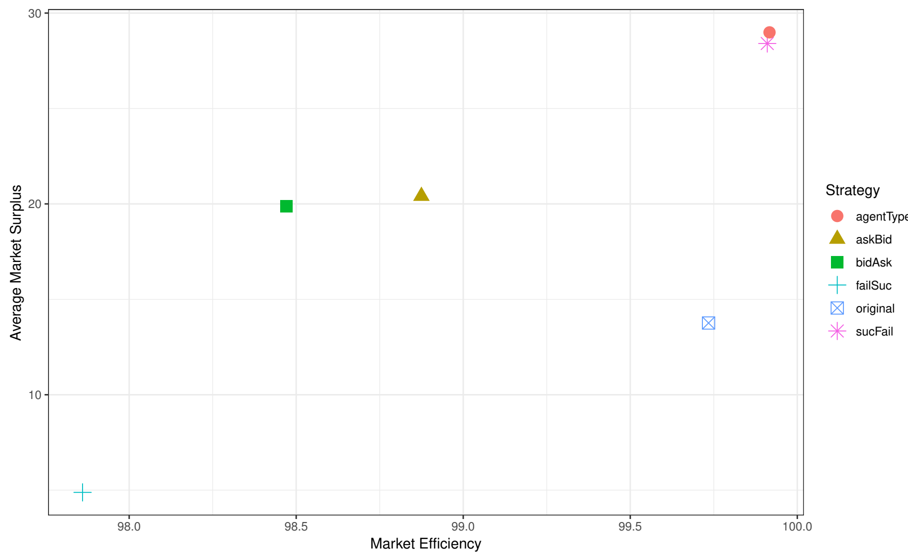

Fig. 2. Comparing the performance of our ZIP strategies in CALL. Results are median values obtained after 1000 runs with each run ending after 100 market calls.

The low performance of the ZIP strategy in CALL can be attributed to how shouts are received by traders. In a CDA, participating traders receive information of exactly one shout at a time, i.e. whether the submitted shout was matched or not. In CALL setting, participating traders receive information of *all* submitted shouts right after the market was called. One supposes that the order in which these shouts are presented to ZIP agents after the CALL is called influences how their profit margins (and hence new shouts) are updated. For this particular experiment, after calling the market all shout information was provided to traders in the same random order in which they were submitted to the auction.

6.3 Improving CALL-ZIP performance

As mentioned before, ZIP traders in a CDA setting update their behaviour one shout at a time, while their CALL counterparts receive a *list* of all successfully matched shouts and failed shouts after the market is called. A relevant question is then: How, and in what order, should this list be provided to participating traders, in order to maximize Market Efficiency and Average Market Surplus? We designed several strategies specifying the sequence in which shouts are presented to ZIP traders after the market is called:

original Order in which shouts were submitted (i.e. random).
bidAsk All bids are provided first followed by the asks.
askBid All asks are provided first followed by the bids.

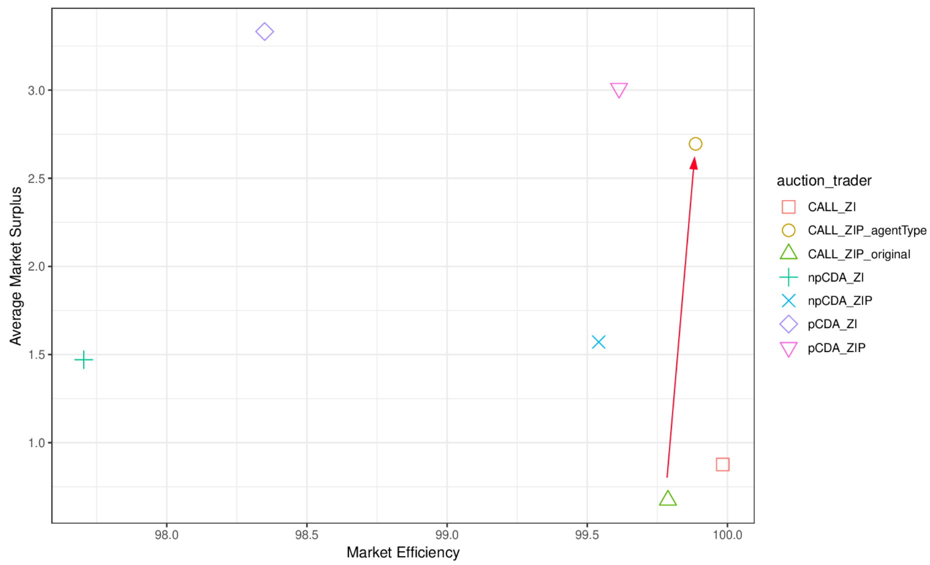

Fig. 3. Comparison of our best CALL_ZIP strategy with default ZIP implementations.

sucFail Successful shouts are provided first followed by the failed ones.
failSuc Failed shouts are provided first followed by the successful ones.
agentType Variant of sucFail where ZIP buyers only receive bids while ZIP sellers only receive asks.

For our second experiment we compared the performance of these six strategies in terms of AMS and Market Efficiency as obtained after 100 market calls, with results visible in Figure 2. The performance of *sucFail* and *agentType* is interesting to see, with both strategies significantly outperforming the others. While sucFail and agentType achieve comparable performance, we will focus on the latter as it reduces the amount of computational work for each trader.
In our third and final experiment we can then compare the performance of our best performing strategy, i.e. agentType, with the original ZIP implementations for both CDA settings. Results can be seen in Figure 3: The increase in performance (marked by the arrow) is evident, achieving levels comparable with the best performing CDA implementation (pCDA_ZIP).

7 Conclusions

We demonstrated that the original ZIP trading strategy performs poorly with respect to average market surplus, when employed in a CALL. We developed several strategies with the aim to increase the performance of ZIP traders to a level comparable with those observed in both CDA settings. Two of our strategies, namely *sucFail* and *agentType* greatly improve the average market surplus of ZIP in CALL and result in a slightly higher CALL efficiency. The resulting performance of our modified ZIP agent in CALL becomes comparable to that observed in both the npCDA and pCDA settings. Autonomous traders can thus use our modified ZIP agent when trading energy in a CALL in the Smart Grid domain.

Future work will focus on further improving the best strategies, e.g. by testing additional orders of bids/asks in the agentType strategy. We will also repeat the above experiments on real data from electricity markets in the context of micro-grids.

References

1. Arifovic, J., Ledyard, J.: Call market book information and efficiency. Journal of Economic Dynamics and Control **31**(6), 1971–2000 (Jun 2007)
2. Bagnall, A., Toft, I.: Autonomous Adaptive Agents for Single Seller Sealed Bid Auctions. Autonomous Agents and Multi-Agent Systems **12**(3), 259–292 (May 2006)
3. Bagnall, A.J., Toft, I.E.: Zero intelligence plus and Gjerstad-Dickhaut agents for sealed bid auctions. In: Workshop on Trading Agent Design and Analysis, part of international conference on autonomous agents and multiagent systems (AAMAS-2004). pp. 59–64 (2004)
4. Cai, K., Gerding, E., McBurney, P., Niu, J., Parsons, S., Phelps, S.: Overview of CAT: A Market Design Competition Version. Technical Report ULCS-09-005, Department of Computer Science, University of Liverpool, Liverpool, UK (Jan 2009)
5. Chowdhury, M.M.P., Kiekintveld, C., Tran, S., Yeoh, W.: Bidding in Periodic Double Auctions Using Heuristics and Dynamic Monte Carlo Tree Search. In: Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence. pp. 166–172. International Joint Conferences on Artificial Intelligence Organization, Stockholm, Sweden (Jul 2018)
6. Cliff, D., Bruten, J.: Minimal-Intelligence Agents for Bargaining Behaviors in Market-Based Environments. Hewlett-Packard Labs Technical Reports (Aug 1997)
7. Gjerstad, S., Dickhaut, J.: Price Formation in Double Auctions. Games and Economic Behavior **22**(1), 1–29 (Jan 1998)
8. Gode, D.K., Sunder, S.: Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality. Journal of Political Economy **101**(1), 119–137 (Feb 1993)
9. He, M., Leung, H.f., Jennings, N.R.: A fuzzy-logic based bidding strategy for autonomous agents in continuous double auctions. IEEE Transactions on Knowledge and Data Engineering **15**(6), 1345–1363 (Nov 2003)
10. Kehr, C.H., Krahnen, J.P., Theissen, E.: The Anatomy of a Call Market. Journal of Financial Intermediation **10**(3-4), 249–270 (Jul 2001)
11. Ketter, W., Collins, J., Reddy, P.: Power TAC: A competitive economic simulation of the smart grid. Energy Economics **39**, 262–270 (Sep 2013)
12. Ma, H.: Adaptive Agents for Sequential Vickrey Auctions with Homogeneous Goods. In: 2008 International Conference on Computational Intelligence for Modelling Control & Automation. pp. 482–487. IEEE, Vienna, Austria (2008)

13. Ma, H., Leung, H.F.: An adaptive attitude bidding strategy for agents in continuous double auctions. Electronic Commerce Research and Applications **6**(4), 383–398 (Dec 2007). https://doi.org/10.1016/j.elerap.2006.12.003, http://www.sciencedirect.com/science/article/pii/S1567422306001207
14. Nicolaisen, J., Petrov, V., Tesfatsion, L.: Market power and efficiency in a computational electricity market with discriminatory double-auction pricing. IEEE Transactions on Evolutionary Computation **5**(5), 504–523 (Oct 2001)
15. Park, S., Durfee, E.H., Birmingham, W.P.: An adaptive agent bidding strategy based on stochastic modeling. In: Proceedings of the third annual conference on Autonomous Agents. pp. 147–153. ACM (1999)
16. Satterthwaite, M.A., Williams, S.R.: The Rate of Convergence to Efficiency in the Buyer's Bid Double Auction as the Market Becomes Large. The Review of Economic Studies **56**(4), 477 (Oct 1989)
17. Tesauro, G., Das, R.: High-performance Bidding Agents for the Continuous Double Auction. In: Proceedings of the 3rd ACM Conference on Electronic Commerce. pp. 206–209. EC '01, ACM, New York, NY, USA (2001)
18. Vytelingum, P., Cliff, D., Jennings, N.R.: Strategic bidding in continuous double auctions. Artificial Intelligence **172**(14), 1700–1729 (Sep 2008)
19. Wang, Y., Saad, W., Han, Z., Poor, H.V., Baar, T.: A game-theoretic approach to energy trading in the smart grid. IEEE Transactions on Smart Grid **5**(3), 1439–1450 (May 2014)
20. Weidlich, A., Veit, D.: A critical survey of agent-based wholesale electricity market models. Energy Economics **30**(4), 1728–1759 (Jul 2008)
21. Widrow, B., Hoff, M.E.: Adaptive Switching Circuits. In: 1960 IRE WESCON Convention Record, Part 4, pp. 96–104. Institute of Radio Engineers, Cambridge, MA, USA (1960)
22. Yue, C., Mabu, S., Hirasawa, K.: A bidding strategy using genetic network programming with adjusting parameters for large-scale continuous double auction. In: 2012 IEEE International Conference on Systems, Man, and Cybernetics (SMC). pp. 1896–1902. IEEE, Seoul, Korea (South) (Oct 2012)