Software Agents in Multiple Adaptive Double Auction Markets

Lampros C. Stavrogiannis and Pericles A. Mitkas
Department of Electrical and Computer Engineering
Aristotle University of Thessaloniki
GR541 24, Thessaloniki, Greece
lstavr@issel.ee.auth.gr, mitkas@eng.auth.gr

***Abstract*—The prevalence of on-line auctions has stimulated the interest of both the economists and computer scientists in an effort to understand and improve their rules. One famous form of these economic mechanisms is the double auction and, more specifically, the continuous double auction which is commonly used in today's major stock exchanges worldwide. Software agents constitute a promising tool for the study and implementation of novel double auction mechanisms. TAC Market Design tournament makes an attempt to study the competition among such dynamically adjusted market institutions trying to attract potential traders while maximizing their profit. In this paper we describe the tournament and discuss about the importance of the global competitive equilibrium in its economy. Moreover, we provide a simple, yet effective estimation technique for the latter that our entrant, Mertacor, has utilized during the games of 2008.**

***Keywords*-Double Auction, Market Based Control, Mechanism Design, Trading Agent Competition**

I. Introduction

The invasion of the Internet in our daily life has given rise to new, innovative applications of e-commerce. Common examples include electronic marketplaces and online trading systems that most of the major stock exchanges use to trade futures, options, equities, and their derivatives.

The *double auction* (DA) is an auction where multiple buyers and sellers are able to make committed offers to buy and sell units of a commodity and then accept similar offers. Besides its prevalence in financial markets, many variants of the DA have been successfully applied as a solution to a multitude of resource allocation problems [1], [2]. The importance of DAs lies in the fact that they exhibit a high *allocative efficiency* (ratio of traders' actual profit to their theoretical maximum profit) with the implementation of very simple rules.

Mathematicians, economists and computer scientists have long used game theory to analyze simple forms of this mechanism [3]–[5] but their findings have been criticized for being of scant relevance to practical scenarios due to their strict assumptions, like traders' full rationality. Moreover, the dynamics of the continuous DA presents an important obstacle in any pure theoretical approach. This led to the adoption of simulation techniques where human subjects trade to verify the effectiveness of the mechanism [6]. The need to reduce experimental cost as well as the increasing requirements for a fully controlled environment inroduced the use of multi-agent systems in this kind of experiments, a field known as *Agent-based Computational Economics* [7].

The majority of the research on DAs focuses on the bidding strategies of the traders. However, scientists have recently turned their attention to the rules and the protocols of this mechanism. Most of the relevant literature deals with isolated markets which operate free of charge. Nevertheless, in today's global economy each country's market institutions compete with each other as well as with the remainder stock exchanges worldwide. Having recognized this, scientists from the universities of Liverpool and Southampton, and Brooklyn College introduced TAC Market Design (or CAT) tournament[1] in 2007, in a joint effort to study the impact of this globalization on trading.

The rest of this paper is organized as follows: Section 2 provides a description of the CAT tournament. In section 3 we introduce the notion of the global competitive equilibrium. We also present a method for its estimation and shortly describe how our entrant has utilized this technique for the games of 2008. Section 4 presents the results of our experiments. Finally, a brief summary concludes the paper in Section 5.

II. TAC Market Design Tournament 2008

The CAT game consists of two distinct entities: *trading agents* (or *traders*) and *specialists*. Each trader may be either a buyer or a seller willing to exchange goods, whereas each specialist represents the auctioneer of a DA market where these traders will trade. Trading agents are provided by the organizers and specialists are designed by the competition entrants. There is also a server responsible for the communication among the clients and the timing of the games. The platform of the tournament is JCAT[2], a client-server implementation of the Java Auction Simulator API (JASA[3]), providing additional support for the operation of multiple markets [8].

Traders are equipped with a *market selection strategy* and a *trading strategy*. The former specifies the specialist to register with for their trades and is typically based on their

[1] http://www.marketbasedcontrol.com
[2] http://jcat.sourceforge.net
[3] http://jasa.sourceforge.net

profit from the market. The trading strategy determines their bidding behavior in the market and follows one of the four extensively studied strategies in the DA literature:

1) *ZI-C* [9]: These trading agents exhibit zero rationality, selecting their offers randomly from a uniform distribution, but are not allowed to trade at a loss.
2) *ZIP* [10]: ZIP traders try to remain competitive in a market by adjusting their profit margin according to current market conditions.
3) *RE* [11]: RE trading agents make an attempt to mimic human behavior on trading, using recent profit as a reward in a learning algorithm.
4) *GD* [12]: These traders consider the history of executed transactions and submitted offers and form a belief function based on which they select their preferred offers.

Every trader is endowed with a set of goods to trade and a *private value* (the maximum amount willing to purchase or the minimum accepted sale's price for buyers and sellers respectively) for each of them. Both strategies and endowment constitute personal information which is not revealed to the competitors during the game.

Each game of CAT comprises several virtual trading days, each of which is further divided in trading rounds of fixed duration. At the beginning of each day, specialists announce their fees and traders must decide upon which market to select for the rest of this day. Traders' offers are *single-unit* (every offer expresses their desire to trade one unit of the good) and *persistent* (once accepted, offers remain active until they result in transactions or the end of the day is reached). Traders' private values are drawn from an unknown distribution at the start of the game and remain constant for the rest of it.

The daily evaluation of the entrants consists of three parts: (i) the *market-share*, which is the percentage of the total traders' population registered in the market, (ii) the *profit-share*, which is the ratio of the daily profit a specialist obtains to the profit of all specialists, and (iii) the *transaction success rate* (TSR), which is the percentage of the offers accepted that result in transactions. The daily score of each specialist is the mean value of the above metrics. Assessment commences and terminates in randomly selected trading days and total score is the sum of the scores across these days [13].

III. The Global Competitive Equilibrium

A. Background

According to the theory of microeconomics, the aggregate demand and supply curves of a market are expected to meet at a point (pair of price and quantity) which is called the *Competitive Equilibrium* (CE) of the market. If all the transactions are cleared at the price of the CE then the market's allocative efficiency is maximized.

The *global CE* is the CE of the equivalent single global market where all of the buyers and sellers would trade had it not been their splitting due to the existence of multiple specialists. In an efficient global allocation only *globally intra-marginal traders* (buyers and sellers with private values above and below the price of the global CE respectively) are allowed to transact.

However, the diffusion of the traders in the various markets presents the opportunity for the *globally extra-marginal traders* (buyers and sellers with private values below and above the price of the global CE respectively) to transact either for the reason that they might be intra-marginal ones for the market registered with or because of the inability of the specific auctioneer's rules to prevent extra-marginal trades, thus leading to a drop in both global and market's allocative efficiency. It is therefore to the entrant's interest to identify this equilibrium and coordinate its transaction prices with it.

B. Estimation

Most of the CE estimation techniques to our knowledge try to predict the price of this point using the history of the transaction prices observed in the market [14], [15], and so are highly dependent on the pricing strategies of the opponent specialists for the case of the global CE. In this paper we follow a different approach.

More specifically, our entrant continually keeps track of the highest bids (buy offers) and lowest asks (sell offers) submitted in its market. The prices of these offers constitute the closest observable estimation of their private values. However, our specialist has only access to a restricted number of trading agents, the ones registered with its own market. To overcome this limitation our contestant exploits the possibility of subscription to the opponent markets, gaining access to the offers submitted in its competitors (four different opponents are observed each day due to the volume of the offers), and thus accelerating the estimation process. When a sufficient number of traders have been explored (80% of the total trader population for the games of 2008), Mertacor uses the 4-heap algorithm [16] to estimate the global CE price in time $O(L \log L)$, where L is the number of the offers recorded. Our estimation method is based on the observation that when there is a sufficiently large population of traders, the real global CE (p^*, q^*) and the reported global CE (p'^*, q'^*) almost coincide, as illustrated in Figure 1.

C. Mertacor Specialist

We now discuss how Mertacor, which was placed 5^{th} in the finals of CAT 2008, utilizes the information gained from the estimation technique mentioned above.

The first responsibility of a specialist is to select the offers that will possibly lead to transactions. This is the task of the *accepting policy*. Our entrant records the number of daily transacted items M_i per trading agent i over the last days.

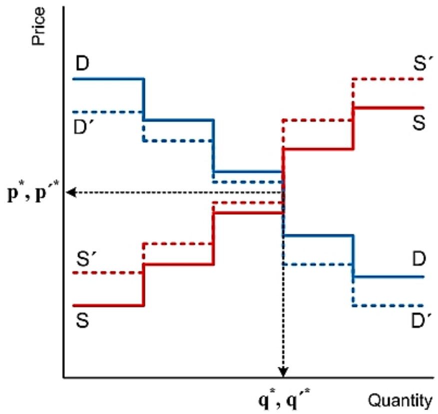

Figure 1. An example of a global DA market. DD and SS are the true global aggregate demand and supply curves and D'D', S'S' the corresponding reported curves.

This number provides us with an estimation of the traders' *entitlement*, i.e. the number of the goods that they are willing to trade. If n is the number of our daily traders, for the first $N = max_i(M_i), i \in \{1, ..., n\}$, trading rounds Mertacor implements a *global equilibrium beating accepting policy*, according to which the prices of the qualified offers must be better (higher for buyers and lower for sellers) than the global CE price, allowing only globally intra-marginal trades to take place.

Accepted offers are maintained in a so called order book. An entrant must then select the pairs of offers that will lead to transactions according to its *matching policy*. In the case of Mertacor this policy implements the 4-heap algorithm discussed above, such that the highest bids placed are matched with the lowest qualified asks.

The *clearing policy* and *pricing policy* determine the time and price of the transactions respectively. For the first N rounds all of the transactions are cleared at the end of each round and their price is set at the global CE price, so that a trader's profit is proportional to the difference of its offer's price and that of the global CE. For the remaining trading rounds, Mertacor acts like a modified continuous DA market, so a transaction is executed as soon as there is an available qualified pair of bids and asks. Our modification lies in the pricing strategy, which uses a variation of the *side-biased pricing policy*, introduced by IAMwildCAT in 2007 [15], such that globally intra-marginal traders are given a higher priority. According to this policy, if b_j is the price of the j^{th} bid accepted and a_l is the price of the corresponding l^{th} ask price, then the transaction price p will be:

$$p = kb_j + (1 - k)a_l \qquad (1)$$

The parameter k is equal to the ratio of the mean number of buyers to trading agents registered with our specialist over the last five days and is limited in the interval $[0.3, 0.7]$. When a trade between a globally intra-marginal and a globally extra-marginal trader takes place, k obtains one of the two extreme values in this interval so as to provide more profit to the former agent.

Finally, the *charging policy* selects the type and amount of fees a specialist should charge trading agents for the services provided in its market. Mertacor implements a *limited score-based charging policy*. According to this policy, our entrant estimates the scores of its opponents and tries to beat them in time intervals proportional to their score differences. There are four different kinds of fees in the tournament (a *registration fee* for the registration with the specialist, an *information fee* to obtain additional information from opponent markets, a *shout fee* for each offer submitted, a *transaction fee* for each transaction conducted, and a *profit fee*, which is a percentage of the trader's profit from a transaction). Nevertheless, we have decided to charge only a small profit fee, so that only profitable traders pay for our entrant's services, and we limit its value in the interval $[0.2, 0.3]$ to avoid a possible loss of significant portion of market-share.

IV. EXPERIMENTS

We conducted a number of experiments to assess the effectiveness of our technique over each of the trading strategies in CAT. More specifically, we have used 6 different specialists obtained from the TAC agent repository[4] and 240 trading agents (120 buyers and 120 sellers). All traders follow the same trading strategy (ZI-C, ZIP, RE, GD) and an ϵ-greedy market selection strategy (ϵ = 0.1, α = 1) [17]. Moreover, traders' private values were drawn from the same distribution, which was either a uniform $U(50, 150)$ or a normal $N(100, 278)$ distribution, as is commonly assumed in the experimental economics literature. The duration of the games was 500 trading days and each day consisted of 10 trading rounds. Each experiment was repeated 10 times. Our results are not statistically significant, considering the length of the games (four hours approximately) which presents an important limitation for such an outcome.

Figures 2a - 2d illustrate the probability density functions (PDFs) of the real and estimated traders' private values for each type of bidding strategies and each distribution, which were obtained using the Parzen window method [18]. As can be seen, our technique accurately acquires the real PDFs for both the normal and the uniform distribution of private values in the first three cases (ZIP, ZI-C, RE). We observe a small divergence for the case of the GD trading strategy, which is justified by the fact that these agents strategically wait to submit their offers, thus resulting in a lack of submitted offers from a noticeable number of trading agents (approximately 17% of the total trader population for our experiments), so there is no possible way to estimate their preferences for the goods traded.

[4]http://www.sics.se/tac/showagents.php

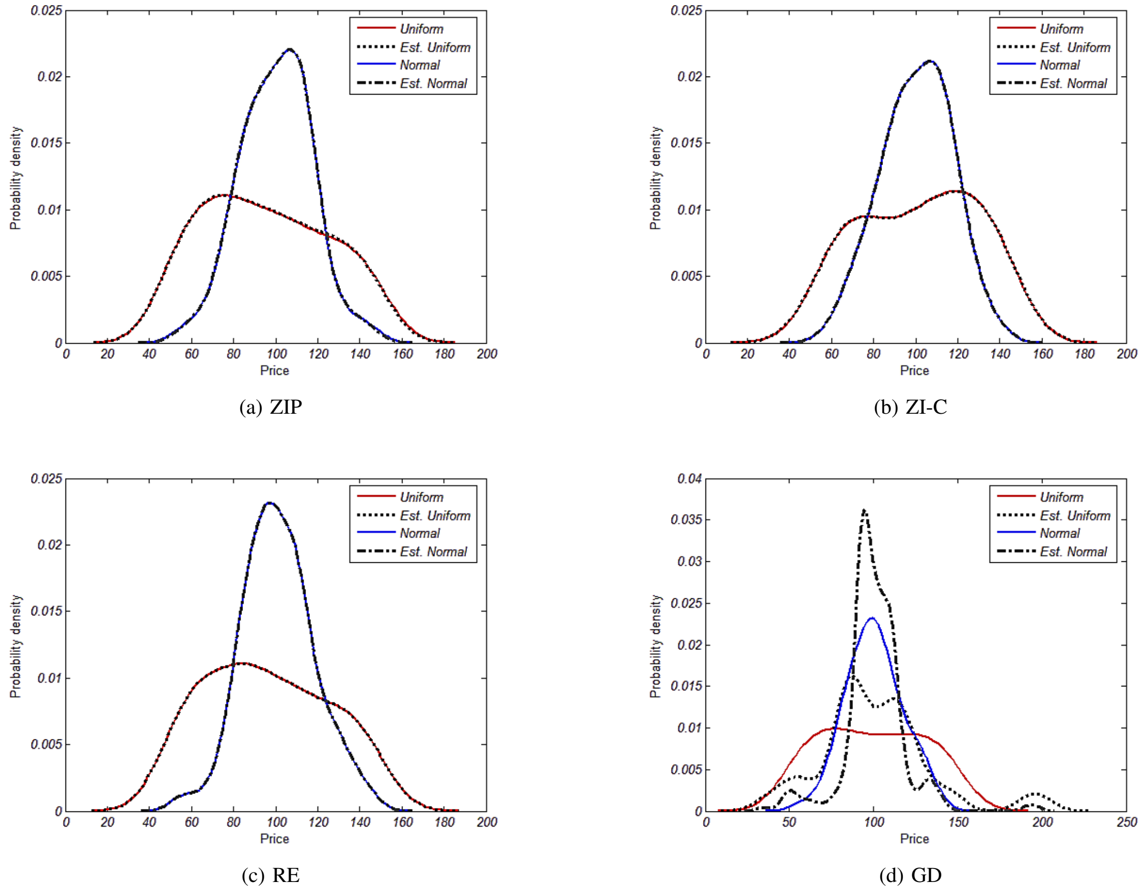

(a) ZIP (b) ZI-C

(c) RE (d) GD

Figure 2. Real and estimated private value distributions for each trading strategy.

Table I shows the values of the absolute percentage error (APE) for our estimation over the four trading strategies and the two private value distributions discussed above. These values verify once again the effectiveness of our method as well as its robustness. The results are slightly better for the normal distribution with an expected maximum APE less than 0.8% in every case. Moreover, it seems that our technique works best with the ZIP traders, which yield the smallest mean APE. Finally, despite the minor problems with the GD traders, the mean APE for these trading agents is less than 0.18% for both distributions.

Table I
ABSOLUTE PERCENTAGE ERROR OF THE ESTIMATED GLOBAL COMPETITIVE EQUILIBRIUM PER TRADING STRATEGY AND DISTRIBUTION OF PRIVATE VALUES.

		Absolute Percentage Error (%)			
		ZIP	ZI-C	RE	GD
$U(50,150)$	Mean	0.017	0.055	0.116	0.163
	Min	0.002	0.007	0	0.128
	Max	1.835	1.238	2.196	0.352
$N(100,278)$	Mean	0.004	0.044	0.057	0.177
	Min	0	0.006	0	0.043
	Max	0.214	0.529	0.766	0.483

V. Conclusions And Future Work

In this paper we have shortly described TAC Market Design tournament as well as our agent's policies. We have also introduced a successful way of estimating the global competitive equilibrium, which constitutes the most valuable component of our specialist's strategy, and have experimentally confirmed its efficacy.

The importance of this point is twofold. From a market designer's perspective, the successful approximation of this equilibrium may help him meet his design objectives. From the trading agent's view, this estimation might provide the opportunity to obtain novel bidding strategies for multiple double auction markets. Moreover, the global competitive equilibrium could be utilized from an *arbitrageur* (trader that exploits the price difference of the same good exchanged in multiple markets, buying it low and then selling it high) to identify the most profitable stock exchanges for its trades irrespective of the markets' pricing policies implemented.

While we have demonstrated that our technique is sound for homogeneous populations of traders, we must further investigate its effectiveness when buyers and sellers follow different trading strategies. It is also useful to see the influence of the specialists' strategy mix selected on our results. Finally, we plan to conduct experiments where traders' private values are affiliated, as this is true in many practical scenarios.

REFERENCES

[1] R. K. Dash, P. Vytelingum, A. Rogers, E. David, and N. R. Jennings, "Market-based task allocation mechanisms for limited capacity suppliers," *IEEE Transactions on Systems, Man, and Cybernetics - Part A*, vol. 37, no. 3, pp. 391–405, 2007.

[2] J. Gomoluch and M. Schroeder, "Market-based resource allocation for grid computing: A model and simulation," in *Proceedings of the First International Workshop on Middleware for Grid Computing*, 2003, pp. 211–218.

[3] K. Chatterjee and W. Samuelson, "Bargaining under incomplete information," *Operations Research*, vol. 31, no. 5, pp. 835–851, 1983.

[4] J. H. Kagel and W. Vogt, "Buyer's bid double auctions: Preliminary experimental results," in *The Double Auction Market: Institutions, Theories and Evidence*, D. Friedman and J. Rust, Eds. Cambridge, MA, USA: Perseus Publishing, 1993, ch. 10, pp. 285–305.

[5] M. A. Satterthwaite and S. R. Williams, "The bayesian theory of the k-double auction," in *The Double Auction Market: Institutions, Theories and Evidence*, D. Friedman and J. Rust, Eds. Cambridge, MA, USA: Perseus Publishing, 1993, ch. 4, pp. 99–123.

[6] V. L. Smith, "An experimental study of competitive market behaviour," *Journal of Political Economy*, vol. 70, no. 2, pp. 111–137, 1962.

[7] L. Tesfatsion, "Agent-based computational economics: Growing economies from the bottom up," *Artificial Life*, vol. 8, no. 1, pp. 55–82, 2002.

[8] J. Niu, K. Cai, S. Parsons, E. Gerding, P. McBurney, T. Moyaux, S. Phelps, and D. Shield, "JCAT: A platform for the TAC market design competition," in *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems*. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, 2008, pp. 1649–1650.

[9] 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*, vol. 101, no. 1, pp. 119–137, February 1993.

[10] D. Cliff and J. Bruten, "Minimal-intelligence agents for bargaining behaviors in market-based environments," Hewlett-Packard Labs, Tech. Rep. HPL-97-91, 1997.

[11] A. E. Roth and I. Erev, "Learning in extensive form games: Experimental data and simple dynamic model in the intermediate term." *Games and Economic Behavior*, vol. 8, pp. 164–212, January 1995.

[12] S. Gjerstad and J. Dickhaut, "Price formation in double auctions," *Games and Economic Behavior*, vol. 22, no. 1, pp. 1–29, January 1998.

[13] E. Gerding, P. McBurney, J. Niu, S. Parsons, and S. Phelps, "Overview of CAT: A market design competition," Dept. of Computer Science, Univ. of Liverpool, Tech. Rep. ULCS-07-006, 2007.

[14] J. Niu, K. Cai, S. Parsons, and E. Sklar, "Reducing price fluctuation in continuous double auctions through pricing policy and shout improvement," in *Proceedings of the 5th International Joint Conference on Autonomous Agents and Multiagent Systems*. New York, NY, USA: ACM, 2006, pp. 1143–1150.

[15] P. Vytelingum, I. Vetsikas, B. Shi, and N. Jennings, "IAMwildCAT: The winning strategy for the TAC Market Design competition," in *Proceedings of the 18th European Conference on Artificial Intelligence*, 2008, pp. 428–432.

[16] P. R. Wurman, W. E. Walsh, and M. P. Wellman, "Flexible double auctions for electronic commerce: Theory and implementation," *Decision Support Systems*, vol. 24, no. 1, pp. 17–27, 1998.

[17] R. Sutton and A. Barto, *Reinforcement Learning: An Introduction*. Cambridge, MA, USA: MIT Press, 1998.

[18] A. W. Bowman and A. Azzalini, *Applied smoothing techniques for data analysis: The kernel approach with S-Plus illustrations*, ser. Oxford statistical science series. Walton Street, Oxford OX2 6DP, UK: Oxford University Press, 1997, vol. 18.