A Long-Term Profit Seeking Strategy for Continuous Double Auctions in a Trading Agent Competition

Dionisis Kehagias[1,2], Panos Toulis[1], Pericles Mitkas[1,2]

[1] Aristotle University of Thessaloniki,
Department of Electrical and Computer Engineering,
54 124, Thessaloniki, Greece
[2] Centre for Research and Technology Hellas,
Informatics and Telematics Institute,
57 001, Thermi, Greece
diok@iti.gr, ptoulis@olympus.ee.auth.gr, mitkas@eng.auth.gr

Abstract. This paper presents a new bidding strategy for continuous double auctions (CDA) designed for Mertacor, a successful trading agent, which won the first price in the "travel game" of Trading Agent Competition (TAC) for 2005. TAC provides a realistic benchmarking environment in which various travel commodities are offered in simultaneous online auctions. Among these, entertainment tickets are traded in CDA. The latter, represent the most dynamic part of the TAC game, in which agents are both sellers and buyers. In a CDA many uncertainty factors are introduced, because prices are constantly changing during the game and price fluctuations are hard to be predicted. In order to deal with these factors of uncertainty we have designed a strategy based on achieving a pre-defined long-term profit. This preserves the bidding attitude of our agent and shows flexibility in changes of the environment. We finally present and discuss the results of TAC-05, as well as an analysis of agents performance in the entertainment auctions.

1 Introduction

The advent of Internet and accompanying networking infrastructures has significantly contributed to the development of electronic commerce today. As more computational and networking resources become available to the users, electronic transactions move to more sophisticated ways of process automation. Autonomous agents that participate in online trading environments with the goal to increase revenue for humans, represent such an advanced paradigm of process automation. Designing effective bidding strategies for agents that participate in uncertain competitive auction environments, is typically based on the optimization of an objective function of the profit or savings they materialize. As uncertainty increases in the environment, the development of appropriate heuristics becomes a compulsory task for trading agent design. The measurement of agents' performance is commonly provided by a benchmark platform, which encounters the clients' utility accomplished, as well as the expenditure costs. The international Trading Agent Competition (TAC) [10] provides one of the most popular and realistic benchmark environments where a number of autonomous trading agents compete to each other in order to assembly travel packages on behalf of a number of clients. Goods are procured in multiple online simultaneous and interrelated auctions of different types. The "travel" (or "classic") game scenario of TAC involves three types of such auctions; a) continuous one-sided that sell flight tickets, b) ascending multi-unit auctions for booking hotel rooms, and c) continuous double auctions (CDA) for entertainment tickets.

In this paper we present *Mertacor*, an agent that ended up first in the finals of TAC for 2005 (TAC-05). In particular, we describe the bidding strategy developed for entertainment auctions, providing at the same time a generic bidding framework for the CDA environment. The latter is the most common variation of double auctions [2] as it is applied in many real life cases, the most typical of which is the stock market. In simple double auctions sellers advertise their offered services or items at prices called asks, while buyers respond, according to their preferences over the available auctioned resources, by posting their desired buying prices called bids. A CDA supports transactions between buyers and sellers that may occur continuously over a specific trading period. In this auction setup, buyers and sellers are allowed to continuously update or even withdraw their bids or asks at any time throughout the trading period [5]. We especially focus on the CDA bidding component of Mertacor, because this is the most generic and non game-specific part of its compound bidding strategy. Thus, this paper's main contribution is the introduction of a new bidding strategy for CDA, realized in the form of entertainment auctions in the context of TAC.

The paper is structured as follows. First, in Section 2 we give an overview of the trading simulation environment on which the developed methodologies are applied. Section 3 reviews related work in trading agent design. Next, Section 4, describes the details of the bidding strategy for CDA deployed by our agent. In Section 5 we present an analysis of the agent performance results taken from TAC-05 environment. Finally, Section 6 concludes the paper.

2 The Trading Agent Competition

The TAC (`http://www.sics.se/tac`) provides a competitive trading environment in which each participating agent operates with the goal of assembling travel packages on behalf of eight clients. Each package refers to a 5-day period travel and consists of a round-trip flight, a hotel reservation and tickets for three different entertainment events. The clients have separate preferences over the arrival and departure dates, the type of hotel and entertainment events they wish to visit, which are randomly assigned to each client at the beginning of the game. The objective of each agent is to maximize the total satisfaction of its clients. In the TAC simulation environment, all three kinds of commodities (flights, hotels and entertainment tickets) are sold in simultaneous online interrelated auctions of three different types, running over a game, which lasts for 9 minutes. These are described in the following.

Flight auctions: There is only one airline company that sells tickets in single seller continuous one-sided auctions, which close at the end of the game. Each auction sells tickets for a particular day and direction, whereas an unlimited number of seats are available. Prices in flight auctions are updated according to a random walk process.

Hotel auctions: There are two hotels in which clients can stay between the arrival and departure dates. The one is more preferable than the other, thus it is expected to be more expensive. Hotel rooms are traded in standard ascending multi-unit 16^{th} price English auctions, which close at randomly determined times in the last 8 minutes of the game. In each auction 16 rooms are offered for each combination of hotel and night.

Entertainment auctions: Entertainment tickets are traded in continuous double auctions, which are held between the participants during the game. Each agent holds a randomly chosen number of tickets from the beginning of the game and can be either a buyer or a seller. Entertainment ticket auctions clear continuously. On clearing, bids match immediately. A bid that does not completely match remains standing in the auction.

The score that the agent receives at the end of each game is calculated as the utility minus the expenditure costs. The utility function in its general form is:

$$Utility = 1000 - travelPenalty + hotelBonus + FunBonus \qquad (1)$$

Apart from tackling with the utility optimization problem, TAC participants need to also deal with many uncertainty factors introduced by the different nature of each auction types and the interrelations that hold between them. For example, the agents need to acquire flight tickets, hotel rooms and entertainment tickets so that are all consistent with the preferred arrival and departure dates. Moreover, the agents may advance their performance and bidding accuracy when they deploy price prediction mechanisms in their decision-making process. Indeed, price prediction in TAC has been thoroughly used for efficient decision-making various agent-development teams have developed many forecasting methods [10].

Regarding the entertainment ticket auctions, one way to deal with uncertainty is to deduce how much an agent valuates a particular item. This involves an appropriate representation of the profit, which is expected to be obtained from every transaction. Moreover, agents should also take compound decisions, including how much and when to bid and to preserve a consistent bidding behavior in all auctions.

3 Related Work

Since the beginning of the competition in 2000, the TAC problem attracted many participants from different countries and organizations. ATTac-2000 agent [8] made the first key contribution to this challenging area. The intricacies of the game were clarified and attacked in a systematic way. ATTac-2000 was the first agent who won the TAC. The notion of *marginal utility* was then recognized to play an important role in the TAC game framework. In fact it has been proven that bidding marginal values in sequential auctions with deterministic prices is an optimal strategy and a fairly satisfactory one in the TAC environment [3].

Since the first TAC, teams concentrated their efforts mainly on developing effective ways of price prediction. Due to space limitations we only refer here to those approaches that most influenced our work. A novel price prediction method was designed for agent SouthamptonTAC [4], which was a competent in TAC-01 and TAC-02, by the application of fuzzy techniques. Regarding the entertainment auctions, the same agent deploys a strategy, according to which it submits offers (bids or asks) based on the number of items returned by an allocation optimization procedure, driven by Linear Programming (LP) [6]. In order to determine the price to be offered, the agent calculates the bid value using the equation $bid = V - \phi(t)$, where V is the valuation of the item and $\phi(t) > 0$, is a time-dependent descending function. In the case that the agent acts as a seller, the ask price is determined by the equation: $ask = V + \phi(t)$. The $\phi(t)$ function represents the profit that the agent receives as the auction progresses. Being a descending function, $\phi(t)$ leads to a profit, which decreases as the auction reaches its end. SouthamptonTAC team participated in TAC-05 with agent Dolphin. Agent Walverine [1] on the other hand provides an analytical approach relied on the principles of a competitive economy for the hotel and flight auction, while it handles the CDAs as completely unpredictable by the use of heuristics.

Another approach was proposed by agent whitebear [9], which simply, but also interestingly enough, used average prices for the prediction of the hotel auctions closing prices. The main strategy of agent whitbear for CDAs is to buy/sell the entertainment tickets it needs/does not need at a price equal to the current bid/ask plus an increment step. This behavior is modified at the early stages of the game, where the agent intends to buy tickets at low prices, even if it does not really need them. The agent adopts this tactic in order to increase its flexibility in the market. Moreover, when whitebear deduces that its competitor will increase its profit, it retains transactions. This bidding attitude relaxes at the later stages of the auction. Making extensive experimentation on mixing of different ‘boundary’ strategies and keeping the design as simple as possible, whitebear proved to be the most robust agent in the short TAC classic game history.

LearnAgents [7] also uses an LP model for allocating the acquired items to all the clients the agent serves. LP calculates the optimal allocation of the acquired goods given the buying price of each item. Although this approach works for the flight and hotel auctions, its application in the domain of the entertainment auctions is non-trivial. LearnAgents tries to make predictions about how the transaction prices will evolve during the game. For this reason it preserves a very active bidding behavior.

4 A Strategy for Entertainment Auctions

Entertainment auctions are continuous double sided auctions, where agents receive new price quotes every 30 seconds and bids are processed continuously. Agents involved in CDAs are both buyers and sellers. Although the TAC entertainment auctions adhere to the continuous double auction protocol applied in the stock market, they formulate a more simplified auction environment for two reasons. First, the number of eight participants is significantly lower than the ones met in a typical stock market. Second, the agents remain adherent to their initial plans about acquiring the desired tickets. This is not always the case in a real stock market, where traders may deploy totally unpredicted bidding behaviors. However, the general model of the bidding mechanism presented in this paper applies on any CDA environment with the appropriate parameterization. In particular, the structure of the decision algorithm deployed by Mertacor can be applied as is in the generic CDA case. The only aspect that needs to be taken into account is to change the values of the various variables that determine the decisions to be taken. These are clarified in the description of the Mertacor's bidding algorithm in the remaining section.

As it was previously mentioned, the main problem that a bidding strategy design copes with is the efficient estimation of the unique private valuation that an agent assigns to a particular entertainment ticket at a given moment. Each item in the CDA environment is intended to be acquired by only one client that the agent represents. Each client has a preference over each ticket, expressed by a real number, randomly drawn in the interval [0, 200]. This value represents the bonus that the agent receives by buying one ticket of the desired type for its client. Although assigning a value to each ticket is a TAC-specific manner for evaluating the available items in the auction, it also reflects the client preferences in the real stock market, which is a specific instance of CDA. Indeed, in the stock market a client desires to buy stocks instead of entertainment tickets at a particular price range. The bonus value imposes a constraint to agents according to which, buying a ticket at a price higher than the bonus value is not a preferable action. Although such a statement is intuitively correct, is not a sufficient criterion for the evaluation of a potential transaction. If, for instance, the bonus value for a particular ticket is 120, then buying it at 100 seems undoubtedly a profitable choice, which results in a profit of 120-100=20. However, if the corresponding seller values it for 40, the aforementioned transaction results in a profit of 100-40=60 > 20 for the competitor, hence to a loss for the buyer. Thus, understanding how much the other competitors valuate the auctioned items is a critical strategic element that may significantly improve an agent's performance.

Mertacor's bidding mechanism in entertainment auctions is based on a simple and consistent algorithm that aims to achieve a long-term profit at the end of the auction based on the following hypothesis. Assume that a seller in a CDA wants to sell a resource at a price that will lead to a satisfactory profit M. A successful seller should be flexible, i.e., to accept transactions at different prices. Therefore, there must be some tolerance ranges around the value of M where a transaction is still accepted. Let us also assume that the seller, after finding the best buy offer, accepts a transaction whose profit is R. Then, if R is close to M the buyer decides to sell the ticket, otherwise (if R is far from M) it prefers to send a new sell offer to the potential buyers.

The goal of our seller agent is not the achievement of a profit in every single transaction, because this would result in the completion of very few transactions, but the collection of a positive mean profit from all the transactions it completes.

It is very critical for the agent to posses an effective manner to calculate profit, based on the information available in the CDA environment. When the agent is prepared to complete a transaction it has to firstly evaluate the auctioned item. This involves the calculation of the ticket valuation V, based on the client preferences. E.g., if a client desires to acquire a particular ticket, this will be highly valuated. Otherwise it may have a low or no value to buy. The next step for the agent is to deduce the potential profit that will make from a specific transaction. Apart from the calculated valuation of the item that the agent desires to buy, it also knows the current bid or ask price, if it is a buyer or a seller respectively. An intuition about the notion of profit that stems form the generic CDA environment is that a buyer agent makes a high profit if it buys at a very low price compared to how much it truly evaluates the desired item. Similarly, a seller agent makes a high profit if it sells much higher than it believes the true price of the item is. The question here is "how much it sells?" and the answer depends on the specific CDA environment and the allowed price ranges in this environment. Thus, in general, if a seller asked for a price of P_b for a specific auctioned item, a buyer would make a high profit if its valuation V about the item was $V >> P_b$, e.g. $V = 2P_b$, and then the profit would be $V – P_b = P_b$. In order to achieve at least such a profit, the agent should assign to its *expected profit* a value of $V – 2P_b$. Thus, the agent submits a buying bid, only if its expected profit is positive.

The notion of profit has been defined in Mertacor in the following manner.

a) If the agent needs to buy a ticket whose value has been estimated equal to V and the current bid price is P_b , then *profit* is defined as: $profit = V - 2 \cdot P_b$.
b) If the agent wants to sell the ticket it holds, the value of the ticket is V, and the current ask price is P_a, then *profit* is given by the equation $profit = P_a - 2 \cdot V$

Defining profit in this way is a substantial step for the evaluation of both the buying and selling transactions, because it guarantees that profit gained by the competitor it will be relatively close to the one gained by Mertacor. The above definition of profit is valid in the TAC-specific CDA. If we would like to apply our bidding algorithm on the generic CDA case, this definition should be adopted accordingly. Mertacor's selling strategy for the entertainment auctions is implemented by the procedure *MertacorSellStrategy* illustrated in Fig. 1. We describe its functionality in what follows.

The procedure iterates over all tickets that the agents possesses. The *target* variable represents the long-term average profit that the agent aims at achieving. This is set to a pre-specified value for each of the entertainment tickets. Typical values for *target* lie on [5, 10]. The variable *mean* is set equal to the current mean value of profit gained over the previously completed transactions. This is calculated by the function *getMeanValue()* in row 3. Variable M, which is given by the following equation:

$$M = A \cdot target + B \cdot mean \tag{2}$$

is used to determine the range of the profit sought. After experiments, the weights A and B were chosen to be A=0.7 and B=0.3. These values show that our agent seeks for a profit highly influenced (70%) by the value of *target*. In order to keep our agent adherent to this goal, we impose (row 5) the variable M to only take values in the range [$a_1 \cdot target$, $a_2 \cdot target$]. Suitable values for the parameters a_2 and a_1 in the TAC environment are $a_1 = 1/2$, $a_2 = 3/2$. For this reason, we use function *relocateM*() (row 6), which returns values in the range [min {*target*, *mean*}, max {*target*, *mean*}]. In order to calculate current ticket's valuation V, our agent calls function *calcVal()*, in row 8, which makes use of a LP model. Since we are interested in always selling higher than a determined reserve price V_o, if $V<V_o$ we set $V=V_o$ in row 9. This prevents Mertacor from selling at very low prices. The choice of upper bound V_o for Mertacor was V_o = \$40. In the next step Mertacor calculates the profit it would make in a potential transaction by using the equation

$$profit = P_a - 2V \tag{3}$$

This preserves our agent's goal, which is to keep $E\{profit\}=target$, where the $E\{.\}$ operator denotes the mean value.

```
PROCEDURE MertacorSellStrategy
 1:  FOR each entertainment ticket in possession
 2:      Assign a pre-specified value to target;
 3:      mean ← getMeanValue();
 4:      M ← A*target + B*mean;
 5:      IF M ≤ (1/2)*target OR M ≥ (3/2)*target THEN
 6:         M ← relocateM();
 7:      END IF
 8:      V ← calcVal();
 9:      IF V < V_o THEN V ← V_o; END IF
10       profit ← P_a – 2*V;
11       M_t ← w(t)*M;
12       R_a ← M;
13       R_b ← 3*M/2;
14       R_c ← rand(M/2, M);
15       IF profit ≥ M_t-R_a THEN sellTicket();
16       ELSE IF profit ≥ M_t-R_b THEN ask (M_t-R_a+2*V);
17       ELSE ask (M_t - R_c +2*V);
18       END IF
19   END FOR
```

Fig. 1. The selling strategy deployed by Metacor in the entertainment auctions.

As it was previously stated, M is the profit that Mertacor assumes to be satisfactory. Any CDA market offers different opportunities for a bidder to make profit with respect to different times in the game. Thus, M is actually a time-dependant variable. For this purpose we define in row 11 a new variable $M_t=w(t)M$ that introduces a time-depended low bound for the desired profit. The $w(t)$ function is graphically represented in Fig. 2. From this graph we can quantitatively monitor the time in the game at which Mertacor seeks for the highest profit. The form of the $w(t)$ function is determined based on the specific requirements of the auction environment. For instance, in TAC the duration of all auctions is 9 minutes, while the hotel auctions close every minute on the minute. Thus, in the middle of the game half of the hotel auctions are closed. At this point uncertainty about how feasible is to assemble valid travel packages reaches its maximum value. After the fifth minute, the game approaches its end, since all hotel auctions for at least one particular day will be closed. This is the reason why we have chosen the form of $w(t)$ depicted in Fig. 2. The peak that appears in the middle of the game (4.5 minutes) represents the maximum profit sought at that time. The demand for profit decreases and then remains constant at a relatively low value as the game reaches its end. The specific $w(t)$ is based on the intuition that big uncertainty results in big profit.

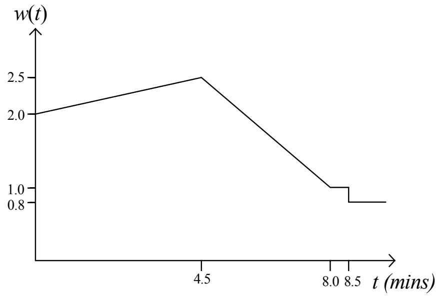

Fig. 2. The $w(t)$ function plotted over time (in minutes).

Next, we define three variables, namely R_a, R_b and R_c, in rows 12, 13 and 14 respectively. The variable R_a determines the case where R is near M**.** This case indicates that there is an explicit interest from buyers and the transaction is still satisfactory. The variable R_b is used to handle a second case, where R is neither near nor far from the desired profit M**.** It represents the case where the buyer is interested in the procured resource but the transaction is not yet satisfactory. The variable R_c is defined to cover the case where there is no real interest from the buyers' side to obtain the auctioned resource. In this case our agent sends an offer that results in a bigger profit than in the first case. Function *rand()* in row 14 returns a random number uniformly distributed in the range [M/2, M].

In order to decide when to bid, our agent uses the three decision rules shown in rows 15-17. According to these rules, Mertacor sells the ticket it currently possesses if the profit to be achieved is M_t-R_a or above. This is done by function *sellTicket()*, in row 15. Otherwise, it asks for those prices that will result in a profit equal to or bigger than either M_t-R_a or M_t-R_c. Our agent submits its ask price (derived from Eq. (3)) for both cases, by invoking the *ask()* function in rows 16 and 17, respectively.

Since $R_c \leq R_a$ the expected profit will be at least equal to ($w(t)$-1)M**.** From Fig. 2 we can see that $w(t) \geq 1$ for most of the time. During the last 30 seconds of the game $w(t)$ is fixed to value 0.8 and the profit becomes -0.2M, which is negative. Thus, if M has a big value, a transaction occurred during the last 30 seconds of the game, will lead to a big loss. In addition to this, transactions occurring at times close to the middle of the game are not so likely to happen, because the quantity ($w(t)$-1)M is relatively big. This will lead to a decrease of the overall average profit. In addition, from Eq. (3) it is derived that M will also decrease. On the other hand, if M is small, transactions that lead to a positive profit are more likely to happen, while negative profit transactions will result in a small loss. The aforementioned mechanism is highly adaptive to changes in the market environment, since when a big profit is assumed the number of transactions is reduced and vice versa. The buying bidding strategy of Mertacor is completely symmetrical to the selling strategy.

The decision mechanism described above ensures that each increase/decrease of the *M* variable results in certain conditions in the environment that strive *M* to the opposite direction. These conditions become more active as the *M* variable is monotonic. Preserving such a dynamic equilibrium for *M* makes Mertacor acting in an autonomous manner, adopting a realistic bidding behavior.

5 Benchmarking Results

In the 6th TAC, the agents were evaluated according to the average score they gained at each round. The TAC servers calculate the score of each agent when a game finishes and this is equal to the utility minus expenditure costs. The utility is given by equation (1). For each round the average score determines the performance metric for all agents. TAC-05 consisted of 4 rounds. In the first (qualifying) round of TAC all eleven agents (for the complete list of participants, please see: `http//www.sics.se/tac`) participated in 600 games, running for almost two weeks. Our agent ranked fourth, gaining a score of 3918.45, which was 240.33 below the top score achieved by agent whitebear05. In this round Mertacor employed a greedy bidding strategy to deal with the entertainment auctions. Next, in the seeding round 680 games were played using the same server configuration. At the beginning of this round we introduced the bidding algorithm presented in this paper, assigning the relatively high value of 12 to the *target* parameter. In this round Mertacor improved its overall performance. It finished third with the score of 4033.32, managing to reduce its distance from the top score agent whitebear05 to 135.29. In the semi-final round ten participants competed in 56 games. Mertacor only fixed some bugs, compared to its version in the previous round and it finally retained its performance by gaining a score of 4023.88.

Eight out of the ten competitors who participated in the semi-finals were invited to the final round. 40 games in the final round were concurrently played in two TAC servers, namely tac1 and tac2, resulting in 80 games. Compared to the semi-final round, Mertacor improved its entertainment bidding strategy, because it fine-tuned its parameters. In particular, we decided to lower the value of the *target* parameter to 5. This proved to be a critical intervention that significantly increased Mertacor's performance and ranked it as the top-scoring agent of TAC-05. Fig. 3 shows the average scores of the four top-scoring agents in the various rounds of TAC.

Apart from the overall benchmarking results provided by TAC, we have also conducted an additional analysis of the agents' performance regarding only the entertainment auctions. In this respect, we measured the profits that the agents made in their transactions, in all games of the final round. The results of this analysis are shown in Fig. 4. For each agent we measured the average profit collected in buying, selling and

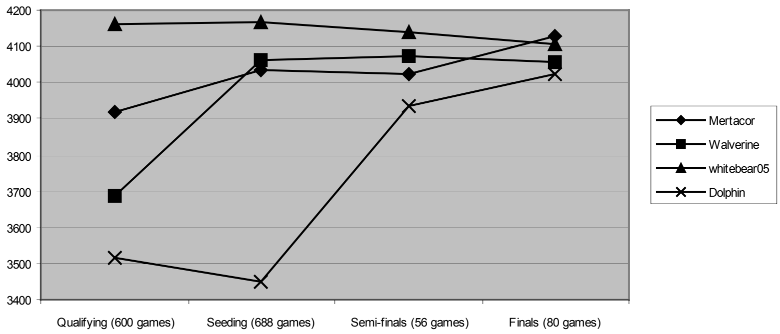

Fig. 3. Performance of the four top scoring agents during the competition rounds. The number of games played in each round is denoted.

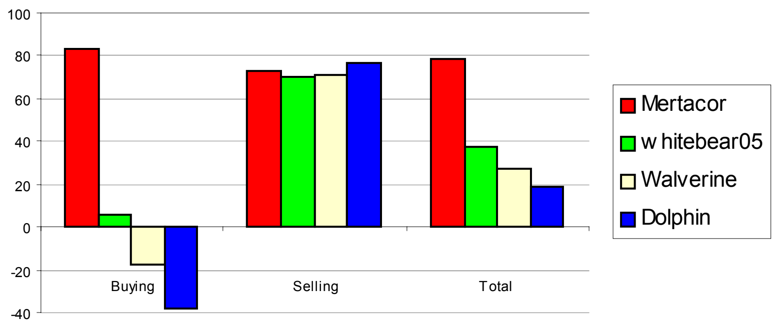

Fig. 4. Average Profit collected by the four top-scoring agents in the final round when they participated in buying, selling and in both types of transactions. The results correspond to data from TAC servers tac1 and tac2.

both types of transactions. Mertacor's performance with respect to buying transactions is extremely high compared to the other competitors'. This shows that the proposed algorithm is particularly efficient in buying transactions. Mertacor performs almost as well as its competitors in selling transactions. In the latter case the best performing agent is Dolphin. Noticeably, Mertacor is the only agent that manages to receive a positive profit when participating in buying transactions. This aspect of Mertacor's bidding algorithm resulted in an overall better performance of Mertacor in entertainment auctions.

6 Conclusions

This paper presented a bidding strategy for CDA, which was designed for agent Mertacor, the first finalist in the TAC-05. We analyzed our bidding/selling strategy for CDA and provided the details of its internal mechanisms. Our strategy was designed with the goal to achieve a long-term profit specified by a target value. This strategy proved to be robust and easily adaptable to market fluctuations. Our agent, which exploited a combination of strategies, accomplished an outperforming score in the TAC. The element that boosted Mertacor's performance was the CDA bidding component, since a fine-tuning of their parameters resulted in a better performance. The bidding algorithm presented in this paper proved to be valid for the TAC-specific CDA setup. Slight modifications of the proposed strategy allow it to be applied on the generic CDA environment. The strong point of the strategy is that its main decision mechanism is applicable on any CDA setup. In order to build an efficient strategy tailored to the specific needs of any particular CDA we need to define the correct values for the *target*, as well as the R_a, R_b, R_c parameters.

References

1. Cheng, S.-F., Leung, E., Lochner, K.M., O'Malley, K., Reeves D.M., Schvartzman L.J., and Wellman, M.P: Walverine: A walrasian trading agent. In Proceedings of the Second international Joint Conference on Autonomous Agents and Multiagent Systems, ACM Press, New York, NY, (2003) 465-472
2. Friedman, D. and Rust, J.: The Double Auction Market: Institutions, Theories and Evidence, Addison-Wesley (1992).
3. Greenwald, A., and Boyan, J.: Bidding Under Uncertainty: Theory and Experiments. In Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence (2004), 209-216
4. He, M., and Jennings, N.R.: Designing a Successful Trading Agent: A Fuzzy Set Approach. IEEE Transactions on Fuzzy Systems, 12, 3 (2004) 389-410
5. Kagel, J.H. and Vogt, W.: Buyer's bid double auctions; Preliminary Experimental Results, in the Double Auction Market: Institutions, Theories, and Evidence, Friedman D. and Rust J. (ed.) Addison-Wesley, (1991) 285-305.
6. Kolman, B., and Beck, R.: Elementary Linear Programmig with Applications. Academic Press Inc. London, (1995)
7. Sardinha, J.A.R.P., Milidiú, R.L., Paranhos P. M., Cunha P. M., Lucena C.J.P.: An Agent Based Architecture for Highly Competitive Electronic Markets. In Proceedings of the 18th International FLAIRS Conference, Clearwater Beach, Florida, USA, May 16-18 (2005)
8. Stone, P., Schapire, R., Littman, M., Csirik, J. and McAllester, D.: Decision-theoretic bidding based on learned density models in simultaneous, interacting auctions, Journal of Artificial Intelligence Research, 19 (2003) 209-242
9. Vetsikas, I.A. and Selman, B.A.: Principled study of the design tradeoffs for autonomous trading agents. In Proceedings of the Second international Joint Conference on Autonomous Agents and Multiagent Systems. ACM Press, New York, NY (2003) 473-480
10. Wellman, M.P., Wurman, P.R., Oapos M.K., Bangera, R. Lin, S.-D., Reeves, D., Walsh, W.E.: Designing the market game for a trading agent competition. IEEE Internet Computing, 5, 2 (2001) 43-51