AUTONOMOUS AGENTS FOR PARTICIPATING IN MULTIPLE ONLINE AUCTIONS

Patricia Anthony, Wendy Hall, Viet Dung Dang, Nicholas R. Jennings
Department of Electronics and Computer Science
University of Southampton
Highfield, Southampton SO17 1BJ, UK
email: {pa99r, wh, vdd00r, nrj}@ecs.soton.ac.uk

Abstract

The increasing number of online auctions poses a big challenge to e-consumers, especially to those who are actively looking for good deals. In this paper, we present the design of an autonomous agent that can alleviate some of these problems by participating across multiple online auctions (in particular, English, Dutch, and Vickrey auctions). The agent makes decisions on behalf of the consumer and endeavours to guarantee the delivery of the item according to the user's preferences. Our agent monitors and collects information from the ongoing auctions and determines which auction it wishes to participate in. The decision on how much to bid in the selected auction is made based on a series of tactics and strategies. The proposed bidding algorithm has been implemented in a simulated marketplace environment and its performance has been evaluated empirically.

1 Introduction

Over the last few years, the number of online auction houses has increased tremendously. To date there are more than 500 auction houses that conduct business online[1]. In 1998, the total revenue for both business-to-consumer and consumer-to-consumer auctions was $USD1.4 billion and it is estimated to reach $USD19.0 billion in 2003[2]. Some examples of popular online auction houses include eBay, Amazon, Yahoo!Auction, Priceline, UBid, and FirstAuction. The types of auction that are conducted vary from site to site, but the most popular ones are English, Dutch, first-price sealed bid and second-price sealed bid (Vickrey). In an English auction, the auctioneer begins with the lowest acceptable price and bidders are free to raise their bids successively until there are no more offers to raise the bid. The winning bidder is the one with the highest bid [McAfee and McMillan, 1987]. The Dutch auction is the converse of the English one; the auctioneer calls for an initial high price, which is then lowered progressively until there is an offer from a bidder to claim the item. In the first-priced sealed bid, each bidder submits his offer for the item independently without any knowledge of the other bids. The highest bidder gets the item and he pays a price equal to his bid amount. Finally, a Vickrey auction is similar to a first-price sealed bid auction, but the item is awarded to the highest bidder at a price equal to the second highest bid [Vickrey, 1961].

Due to the proliferation of these online auctions, consumers are faced with the problem of monitoring multiple auction houses, picking which auction to participate in, and making the right bid to ensure that they get the item under conditions that are consistent with their preferences. These processes of monitoring, selecting and making bids are time consuming. The task becomes even more challenging when the individual auctions have different start and end times. Moreover, auctions can last for a few days or even weeks. To assist consumers in this task, some online auctions provide bidding robots that act on their behalf. The robots send updates to the consumers from time to time. However, this service only operates within the hosting auction site. In this case, the consumer only needs to supply the maximum bid value to the bidding robot, and it will then automatically increment the bid value progressively while the auction is ongoing. It stops bidding when the maximum price is reached.[3] While freeing the consumer to a certain extent, this type of facility limits the choice of auctions that a consumer can participate in. If the consumer wishes to purchase the item, he has to wait until the auction is concluded before he bids in another auction to

[1] http://bidfind.com
[2] Forrester Research 1999
[3] http://www.ebay.com

avoid getting multiple items. In addition, the auction site host could use information about the consumer to cheat him by manipulating the auction. This is possible through shill bidding where bids are deliberately placed to drive up the price of an item [McAfee and McMillan, 1992]. In this case, the consumer ends up paying a price equal (or very close) to his maximum bid.

In view of these complexities, consumers tend to focus on a single auction of their choice. Unfortunately, winning in that auction does not necessarily mean that they get the best deal. They may have received a better price in another auction. The losers may also have had better luck in another auction as well. From the seller's perspective, the lock in of bidders is also a disadvantage since those agents that value the good highly may not be amongst this set. This, in turn, restricts the expected revenue from the auction.

To address these shortcomings, a service that is starting to emerge is that of an auction search engine. These allow the consumer to monitor multiple concurrent auctions, but they leave the actual bidding decision to the consumer[4]. While this certainly increases the consumer's knowledge of the global marketplace, it does not solve the problem of reducing the amount of time that has to be spent on-line. Moreover, deciding what amount to bid for an item requires an intelligent decision where the consumer needs to come up with a strategy to work out the bid value. In many cases, the outcome of this decision making is that the consumer is trapped with the *winner's curse* phenomenon where he pays more than the actual value of the item [Klemperer, 1999]. From this analysis, it can be seen that time and decision making are the two most important factors in online auction environments. Given this, we believe it is desirable if both processes can be automated leaving the consumer free to do other tasks.

Such automation can be achieved by building a software agent that acts on behalf of a consumer and is empowered with trading capabilities such as the ability to search online auctions, negotiate with sellers and make purchases autonomously. A buyer agent that participates in multiple online auctions needs to possess the ability to decide in which auction it should participate and then to determine the appropriate bid value for the desired item. In the face of multiple auctions, this decision is not straightforward. It is dependent on many factors including the type of auction that the agent is participating in, the number of ongoing auctions, the amount of time allocated for it to deliver the item, the behaviour of the agent itself and other participants, and the consumer's valuation of the item. There are standard models of auction behaviour that an agent can follow when participating in an auction [Klemperer, 1999], but these models are applicable to single shot auctions only. Thus, for our purposes, these models are not appropriate since we are dealing with multiple simultaneous auctions and with multiple auction protocols. In short, the standard results from auction theory are not applicable.

Against this background, we detail and evaluate the decision-making procedures of an autonomous agent that can participate in multiple online auctions to assist consumers in purchasing their required item. The agent makes decisions on behalf of the consumer and endeavours to guarantee the delivery of the item. The agent must ensure that it never bids above the reservation price (the maximum amount that the consumer is willing to pay) and it tries to get the item in a manner that is consistent with the consumer's preferences (e.g. at the earliest time, at the lowest price, or with the maximum chance of succeeding). To make decisions, the agent generates a set of tactics and strategies that it can use, based on the existing environment. The agent participates in three different types of auction protocol: English, Dutch and Vickrey. The first-price sealed bid auction is not considered here because of its similarities to the Dutch auction [McAfee and McMillan, 1987]. The main contribution of this work is that we develop a novel algorithm that enables an agent to bid across multiple simultaneous auctions, each of which may be employing a different protocol. In addition, we detail the implementation of this algorithm and we present a preliminary evaluation of its operational performance characteristics.

The remainder of the paper is structured in the following manner. In the next section, we describe our electronic marketplace scenario. Section 3 describes the bidding algorithm. Our initial experiments and results are presented in Section 4. Section 5 discusses related work and Section 6 describes our plans for future work in this area.

2 The Electronic Marketplace Simulation

The simulated electronic marketplace consists of a number of auctions that run concurrently (see Figure 1)[5]. There are three types of auctions running in the environment: English, Dutch and Vickrey. The English and Vickrey auctions have a finite start time and duration generated randomly from a standard probability distribution, the Dutch auction has a start time but no pre-determined end time. The start time

[4] http://www.omnibot.com/

[5] The marketplace is a virtual simulation in that it is supposed to represent all the auctions that are selling the desired item anywhere on the Internet. It is obviously a simplification, since by grouping them in this way we can focus on the agent's bidding strategy (our main aim) and we do not have to worry about the problem of finding all the auctions that sell the desired item, semantic interoperability due to heterogeneous information models and ontologies, latency due to network traffic variability, or interfacing our software to proprietary commercial auction houses.

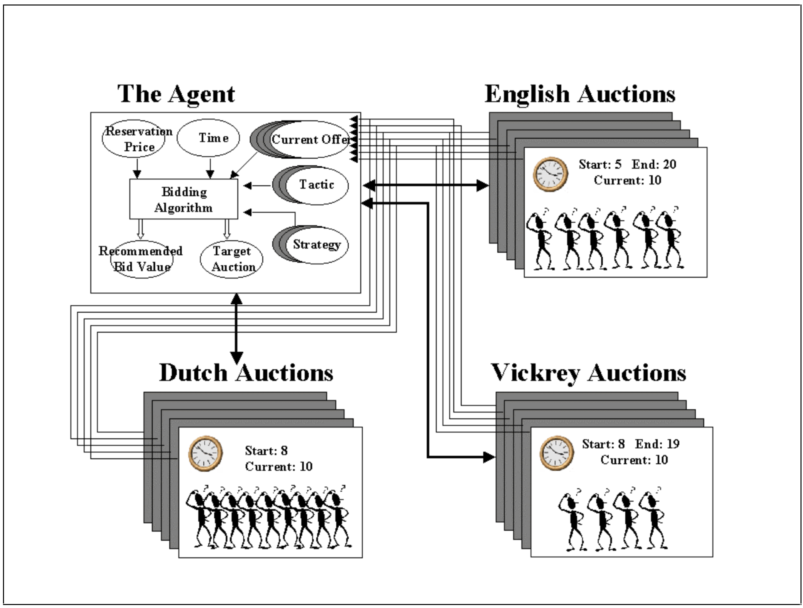

Figure 1 - The Marketplace Simulator

and the end time varies from one auction to another. At the start of each auction (irrespective of the type), a group of random bidders are generated to simulate other auction participants. These participants operate in a single auction and have the intention of buying the target item and possessing certain behaviour. They maintain information about the item they wish to purchase, their private valuation of the item (reservation price), the starting bid value and their bid increment. These values are generated randomly from a standard probability distribution. Their bidding behaviour is determined based on the type of auction that they are participating in. In an English auction, they start bidding at starting bid value; when making a counter offer, they add their bid increment to the current offer (provided the total is less than the reservation price), and they stop bidding when they acquire the item or when their reservation price is reached. In a Dutch auction, they wait until the offer value is equal to their reservation price before making an offer. Finally, in a Vickrey auction, they bid at their reservation price. These strategies are based on the dominant strategies of the respective one-shot single auctions [Sandholm, 1999].

The auction starts with a predefined starting value; a small value for an English auction and a high value for a Dutch auction. There is obviously no start value for a Vickrey auction. Offers and counter offers are accepted from bidders who are picked randomly from the group of bidders in that particular auction. These processes are repeated until the reservation price is reached or until the end time for that auction is reached. The winner in an English auction is the bidder with the highest bid value at the end of the auction. In a Dutch auction, when no offer is received from the bidders, the value is reduced (based on a fixed decrement value) and the whole process is repeated again. The item is sold when a bidder agrees to buy the item at the offer price. If there is more than one bidder who is interested at the same price, the item will be sold to the bidder who offered to buy the item first. There may be cases where there is no offer from the bidders at all throughout the auction. In this situation, the auction terminates when the decremented offer reaches the reservation price. Bidders in a Vickrey auction submit their bid values before the end of the auction. Bids are opened at the end of the auction and the winner is the one who offered the highest price. If there is a tie, the winner is the bidder who submits the earliest bid.

The marketplace is flexible and can be configured to take up any number of auctions and any value of discrete time. We assume that all the auctions in the marketplace are auctioning the item that the consumers are interested in. Our bidder agent is allowed to bid in any of the auctions at any time when the marketplace is active. The objective of the bidder agent is to participate across the multiple auctions, bid in the auctions and deliver the item to its consumer in a manner that is consistent with their preferences. The bidder agent is given a deadline by when it needs to obtain the item. The bidder agent utilises the available information to make its bidding decision; this includes the consumer's reservation price, the time it has left to acquire the item, the current offer of each individual auction, and its set of tactics and strategies. The reservation price is derived from the item's closing price distribution, observed from past auctions. The tactics and strategies are the main constituents that drive the agent's behaviour in making the bidding decision (these are described later in the paper). The output of the bidding decision is the auction the agent should bid in and the recommended bid value that it should bid in that auction. If the agent does not purchase the item by its deadline, it returns to the consumer for further instructions.

3 Designing the Agent's Bidding Strategy

To ensure that the bidder agent operates effectively in the marketplace, it needs to possess a strategy to ensure that it obtains the item within the given time in a manner consistent with the consumer's preferences. Here the bidding strategy for the agent is modelled on the idea of negotiation decision functions as proposed by Faratin et al. [Faratin et al., 1998]. The original model defined a range of strategies that an agent can employ to generate initial offers and counter offers in a two party negotiation. This model identifies the key constituents that drive an agent's negotiation behaviour and defined a single tactic to deal with each of them. The agent's overall behaviour is then the amalgamation of these different facets, weighted by their relative importance to the user. Mapping this to an auction environment, the bidder agent needs to decide the new bid value based on the current offer price. The current offer can be treated as the offer and the new bid value can be treated as the counter offer. Negotiation is over when the auction terminates or when the bidder's reservation price is reached (bidder walks out of the negotiation process).

Firstly, we will present our notation. Let t be the current universal time across all auctions, where $t \in \mathcal{T}$, and $\mathcal{T}$ is a set of finite time intervals. Let t_{max} be the time by when the agent must obtain the good (i.e. $0 \le t \le t_{max}$), and let A be the list of all the auctions that will be active before time t_{max}. At any time t, there is a set of active auctions $L(t)$ where $L(t) \subset A$. Let $E(t)$, $D(t)$, and $V(t)$ be the set of active English, Dutch and Vickrey auctions, respectively, where $E(t) \cap D(t) = \phi$, $D(t) \cap V(t) = \phi$, $E(t) \cap V(t) = \phi$, and $E(t) \cup D(t) \cup V(t) = L(t)$. Each auction $i \in A$, will have its own start time, σ_i, and its own end time η_i where $i \in E(t) \cup V(t)$. Let λ be the agent's bid increment value, and p_r be its reservation price for the target item. Let $Item\_NA$ be a boolean flag to indicate whether the target item has already been purchased by the agent. We assume that the value of p_r is based on current reliable market price observed from past auctions, and that the marketplace is offering the item that our agent is interested in.

With these definitions in place, the algorithm for the bidding agent is summarised in Figure 2. Since each auction has a different start and end time, the bidder agent needs to build an active auction list to keep track of all the auctions that are currently active in the marketplace. We define an active auction as one that has started but not reached its end time. The agent identifies all the active auctions and gathers relevant information about them. It then calculates the maximum bid it is willing to make at *the current time* using the agent's strategy (described later in the paper). This current maximum bid, by definition, will always be less than or equal to p_r. Based on the value of the current maximum bid, the agent selects the potential auctions in which it can bid and calculates what it should bid at this time in each such auction. The auction and corresponding bid with the highest expected utility is selected from the potential auctions as the target auction. Finally, the agent bids in the target auction.

while $(t \le t_{max})$ *and* $(Item\_NA = true)$
    Build active auction list
    Calculate current maximum bid using the agent's strategy
    Select potential auctions to bid in, from active auction list
    Select target auction as one that maximises agent's expected utility
    Bid in target auction using current maximum bid as reservation price at this time
endwhile

Figure 2 - Algorithm for the Bidding Agent

3.1 Calculating the Current Maximum Bid

At any given time t, the agent needs to determine its current maximum bid. This bid is defined as the maximum value the agent is willing to offer at the current moment in time. There are several factors that the agent needs to take into consideration when calculating this value. One is the remaining time that it has left to acquire the item. Thus, a key determinant of what price to offer depends on how much time it has left to bid. For example, the agent may decide to make a low bid value when it has a lot of time left, and as the remaining time decreases, the agent bids closer to its reservation price. Another consideration is the number of remaining auctions that the agent can bid in. Here the agent's behaviour is similar to the time constraint, whereby it may choose to maximise its chances of success by bidding close to the reservation price when the number of auctions is small. The level of desperateness to obtain the item is another consideration that the agent may need to take into account. If the agent is desperate to get the item, it bids aggressively as soon as it starts to ensure that it maximises its chances of getting the item. This level of aggressiveness is influenced by the agent's desperateness for the item. The opposite of this behaviour is that of an agent who is looking for a bargain. If an agent wishes to get the item at a bargain price, it starts bidding at a very low price and eventually bids its reservation price when it has very little time left. The set of considerations of the remaining time left, the remaining auctions left, the desire to get a bargain and the level of desperateness are here referred to as the *bidding constraints*. The agent uses some combination of these constraints in order to determine its maximum bid value at the current moment in time. Our model is open in that if there was another aspect that the consumer wanted the agent to consider, then this could easily be added in as a new bidding constraint. Exactly which constraints are used in a given situation is determined by the consumers and their preferences. Thus, a consumer who wants to ensure it receives the item as quickly as possible would place the greatest weight on the time until deadline and the desperateness tactics, whereas a consumer who is looking to minimise the price he pays would value the bargain tactic most highly.

More formally, let C be the set of considerations that the agent takes into account when formulating a bid and j represent the individual bidding constraints, where $j \in \{1,...|C|\}$. Let Δt denote the remaining time left for the agent to bid (i.e. $t_{max} - t$), and Δa denote the number of auctions left in the marketplace. Let μ denote the agent's desire for a bargain, where $\mu \in [0,1]$ (where 1 is actively looking for bargain and 0 is not actively looking for a bargain), and ε denote the agent's level of desperateness for the item, where $\varepsilon \in [0,1]$ (where 1 is very desperate and 0 is less desperate). For each of constraint $j \in C$, there is a corresponding function $f_j(t)$, which suggests the value to bid based on that constraint. These individual constraints are then combined using a function F to produce the agent's overall position. Examples for F include weighted average, max, or min.

At a given time, the agent may consider any of the bidding constraints individually or it may combine them depending on the situation (what the agent sees as being important at that point in time). In this work, if the agent combines multiple bidding constraints, it allocates weights to denote their relative importance. Thus, let $w_j(t)$ be the weight on constraint j at time t, where $\forall j \in C$, $0 \le w_j(t) \le 1$, and $\sum_{j \in C} w_j(t) = 1$.

The current maximum bid value for the agent at time t, is then calculated as:

$$M(t) = \sum_{j \in C} w_j(t) f_j(t).$$

The agent uses a set of polynomial functions (drawn from Faratin et al.'s negotiation functions) to calculate the bid value based on a single bidding constraint. Here this set of functions is referred to as the tactics. In the current implementation, the four tactics are remaining time, remaining auctions, desire for bargain and desperateness. The definition of each of these is given below.

3.1.1 The Remaining Time Tactic

This tactic determines the recommended bid value based on the amount of time remaining for the agent. Assume that the agent is bidding at time $0 \le t \le t_{max}$. The agent bids closer to p_r as t approaches t_{max}, and it eventually bids at p_r when $t = t_{max}$. To calculate the bid value at time t, the following expression is used:

$$f_{rt} = \alpha_{rt}(t) p_r$$

where $\alpha_{rt}(t)$ is a polynomial function of the form:

$$\alpha_{rt}(t) = k_{rt} + (1 - k_{rt})\left(\frac{t}{t_{max}}\right)^{1/\beta}$$

k_{rt} is a constant that when multiplied by the size of the interval determines the value of the starting bid of the agent in any auction. By varying the value of $\alpha_{rt}(t)$, a wide range of time dependent functions can be defined from those that start bidding near p_r quickly, to those that only bid near p_r right at the end, to all possibilities in between. The only condition is that $0 \le \alpha_{rt}(t) \le 1$, $\alpha_{rt}(0) = k_{rt}$, $\alpha_{rt}(t_{max}) = 1$, and $0 \le k_{rt} \le 1$.

Using the polynomial function defined earlier, different shapes of curves can be obtained by varying the values of β. This represents an infinite number of possible tactics, one for each value of β. In this tactic, β is drawn from $\Re^+$, where $1/200 \le \beta \le 1000$. When $\beta < 1$, the tactic maintains a low bid value until the deadline is almost reached, where this tactic concedes by suggesting the reservation price as the recommended bid value. The other extreme is when $\beta > 1$; the tactic starts with a bid value close to the reservation price and quickly reaches the reservation value long before the deadline is reached.

3.1.2 The Remaining Auctions Tactic

This tactic is broadly similar to the remaining time tactic; the agent bids closer to p_r as the number of remaining auctions approaches 0 (since it is running out of opportunities to purchase the desired good).

f_{ra} has the same form as f_{rt} and α_{ra} is defined as follows:

$$\alpha_{ra} = k_{ra} + (1 - k_{ra})(\frac{c(t)}{|A|})^{1/\beta}$$

Most of these terms are similar to α_{rt}, the only difference being that $c(t)$ is the list of all the auctions that have been completed between time 0 and time t. β is again drawn from $\Re^+$, where $1/200 \le \beta \le 1000$.

3.1.3 The Desire for Bargain Tactic

This tactic is employed when the agent is motivated to try and obtain a bargain. The agent keeps the λ to a minimum as it progresses from $t = 0$ to t_{max}, but eventually bids its reservation price when t_{max} is reached. To determine the bid value for this tactic, the agent considers the minimum bid value for the target item across all the auctions in the marketplace. At a given time t, newly started English auctions have low current bid values and Dutch auctions have very high current bid values. On the other hand, when auctions are terminating, English auctions typically have high current bid values and Dutch auctions have low current bid values. Vickrey auctions do not have information on the bid values since bids are sealed and they are only opened at the end time. To deal with these points, the minimum bid value is calculated by taking into consideration the current bid value and the proportion of time left in the auction. These values are summed and averaged with respect to the number of active auctions at that time.

Let $v_i(t)$ be the current highest bid value in an auction i at time t, where $i \in L(t)$, and $\omega(t)$ be the minimum bid value for the agent at time t where:

$$\omega(t) = \frac{1}{|L(t)|}\Big(\sum_{1 \le i \le |L(t)|} \frac{t - \sigma_i}{\eta_i - \sigma_i} v_i(t) \Big)$$

The bid value is then calculated using the expression:

$$f_{ba} = \omega(t) + \alpha_{ba}(t)(p_r - \omega(t))$$

where $\alpha_{ba}(t)$ is defined as:

$$\alpha_{ba}(t) = k_{ba} + (1 - k_{ba})(t / t_{max})^{1/\beta}$$

Assume that $\alpha_{ba}(t)$ is similar to the polynomial function discussed in the first two tactics, but this time, $0.1 \le k_{ba} \le 0.3$, the minimum value of β equals 1/200 and the maximum value of β equals 0.5. These values reflect the fact that an agent that is looking for a bargain should never bid with $\beta > 1$ because this would inflate the agent's bid well before the deadline. In contrast, an agent that is looking for a bargain (with $\beta < 1$) maintains a low bid value until the deadline is almost reached, where it will then suggest p_r as the recommended bid value. By conceding with a recommended bid value of p_r, the agent tries to ensure that it still successfully acquires the item even if it did not succeed in getting a bargain.

3.1.4 The Desperateness Tactic

This tactic is employed when the agent is desperate to get the item. The agent bids close to p_r at $t = 0$, and eventually bids at p_r when t_{max} is reached. In this tactic, the agent utilises the minimum bid value and the polynomial function described in subsection 3.1.3 but with a slight variation to the value of β, where $1.67 \le \beta \le 1000$ and $0.7 \le k_{de} \le 0.9$. The values picked for k_{de} are high since a desperate agent starts bidding at a value that is near to p_r. With these minor variations, f_{de} is the same as f_{ba} and α_{de} is the same as α_{ba}.

3.2 Selecting Potential Auctions and the Target Auction

The agent selects potential auctions if and only if it is not holding the highest bid in an English auction, or it has not placed a bid in a Dutch or a Vickrey auction. This is to ensure that the agent does not acquire more than one of the target item. The agent selects the potential auctions by considering values for the current maximum bid for each active auction. In the English auctions, this is carried out by taking those auctions that are close to their end time, in which the current bid value when added to the bid increment is less than or equal to the current maximum bid. The agent's new bid value is the current bid plus the bid increment. Only English auctions that are close to their end time are picked to maximise the agent's chances of winning. If the agent currently holds a bid in an English auction that still has a long time to complete, it will not be able to participate in other auctions until it loses out to another bidder or until the auction terminates. There are several potential outcomes when an auction terminates; the agent loses out to another bidder; the agent's bid value may be less that the reservation price (in which case there will be no winner); or the agent wins. If either of the first two situations occur, the agent loses out in that it wasted a lot of time in one auction, thus reducing its chances of participating in other auctions. The potential Dutch auctions, in which the agent may bid, are those with current bids that are less than the current maximum bid. Here, the agent's new offer is the current bid for that particular Dutch auction. The potential Vickrey auctions in which the agent may bid are those that are terminating at the current time and the agent's bid value is its current maximum bid value. This selection of timing is based on the same reasoning as that of the English auction.

If there is only one auction in the potential auction list, that auction is picked as the target auction. If there are multiple auctions, the agent calculates the expected utility for each of these potential auctions. By definition, the expected utility is the product of the probability of the agent winning in that auction at the given bid value and the value of the agent's utility function. The auction with the highest expected utility for the agent's bid value is picked as the target auction. Here, the probability of winning is highly dependent on the type of auction and the agent's bid value. In English and Vickrey auctions, the closer the bid value is to the reservation price the higher the probability of the agent winning (i.e. the probability of the agent winning in the auction with a bid value p_r is close to 1). Here we assume that the agent's reservation price is selected based on the current reliable market price, observed from past auctions and that if the agent bids at this reservation price it is likely that it will win the auction. Thus, the probability function is, in fact, the probability of winning given that the agent's reservation price is competitive to the prevailing market condition. In a Dutch auction, this is slightly different. The probability of winning in the auction with bid value of the reservation price or lower is very close to 1 due to the decreasing nature of the bid values. If the bid value is at p_r, it is almost certain that the agent will win the auction (unless there are clashes). This is similar when the value is lower than p_r. In more detail, let $P_i(v)$ be the probability of winning in auction i if the agent bids with the value v. The expected utility for an auction i with a bid value v is calculated as:

$$u_i(v) = P_i(v)U_i(v), \text{ where } U_i(v) = 1 - (\frac{v}{p_r})^{1/\beta}$$

The β used here is the same as the one used in generating the polynomials for the tactics. The utility function for each potential auction is calculated by dividing the payoff amount with p_r. The utility value is higher when the payoff ($p_r - v$) is high (value greater than 0), and it is lower when the payoff is low (value close to or equals to 0).

There is a possibility that our agent may make a counteroffer in an English auction as soon as other participants outbid it. This occurs when the value of the expected utility is in favour of the previous target auction. When this happen, our agent has the advantage, in that it implicitly forces the other agents to move towards their reservation price.

4 Experimental Evaluation

To evaluate the performance of our agent using the bidding algorithm described above, we undertook an empirical evaluation. Here we defined three control models as a basis for comparison. These models simulate three plausible modes of behaviour in online auctions. In the first model (control C1), the agent randomly joins one auction and stays there until its reservation price has been reached or until the auction's end time is reached. The agent does not move to any other auction. This is a similar strategy to most current bidding agents. Our second control agent deploys a greedy strategy (control C2). This agent picks the auction that has the closest end time where the current bid value for the item does not exceed the reservation price. If there is more than one possibility, it makes a random choice. The agent stays there until the reservation price has been reached or until the auction is over. If it has not purchased the item, it then moves to another auction using the same selection technique and repeats the process until its allocated time is over. In the last model (control C3), an agent picks an auction randomly from all the active auctions and bids in that auction. It stays there until its reservation price is reached or until the auction is over. If the agent is not successful, it randomly picks another auction and repeats the process until the allocated time is over or it successfully purchases the item.

Our experiments consist of 100 runs for the proposed and control models. These experiments were run in an environment with $t_{max} = 20$, between 3 and 10 English, Dutch and Vickrey auctions running concurrently, and for each auction, there are between 2 and 10 participants. We deployed different strategies for the bidding agent by varying the values of the weights for the tactics. In this particular experiment, the polynomial functions for the tactics are:

- $\alpha_{rt}(t)$ had $k_{rt} = 0.6$ and $\beta = 4$ (giving a high start bid value which quickly reaches p_r)
- $\alpha_{ra}(t)$ had $k_{ra} = 0.6$ and $\beta = 4$ (giving similar behaviour as above)
- $\alpha_{ba}(t)$ had $k_{ba} = 0.3$ and $\beta = 0.3$, (giving a low level desire for bargain behaviour)
- $\alpha_{de}(t)$ had $k_{de} = 0.7$ and $\beta = 5$ (giving a low level of desperateness behaviour).

We defined six different strategies for our bidding agent to use. The first four (*RT*, *RA*, *BA*, *DE)* are based on a single tactic strategy. The last two strategies use combinations of the weights for all the four tactics: the COM1 strategy uses equal weighting among the four tactics, and COM2 uses variable weights ($w_{rt} = 0.4$, $w_{ra} = 0.2$,

$w_{ba} = 0.1$, and $w_{de} = 0.3$). These values were chosen to model a behaviour of an agent that values time as the most important facet, and is desperate to get the item, but is not interested in a bargain or the remaining auction tactics. The control models are labelled as C1, C2, and C3. All the agents use the same reservation price and the same time for t_{max}.

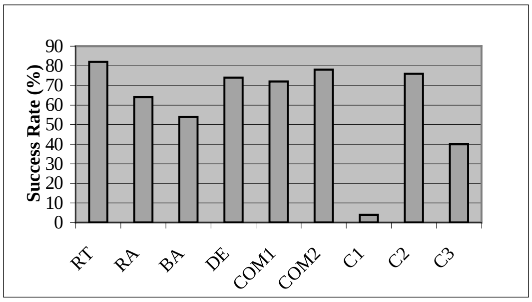

Figure 3 - Success Rate Comparisons

Figure 3 shows the performance of all the agents in terms of their success rate. The success rate is defined as the number of times, as a percentage, the agent is successful in obtaining the item. Strategy RT and COM2 performed best, followed by C2, DE, COM1, and RA. The performance of the agents that used a single tactic can be associated to the value of k and β. Agent RT achieved a success rate higher than 80% because it starts with a bid value close to the reservation price and quickly goes to its reservation price. Agent RA, did not perform as well as agent RT (even though they have a similar pricing function) because multiple auctions may terminate at each discrete time t. At each time step, the agent works out the number of auctions that are terminating at that current time, and uses this information to work out the value of $c(t)$. Unfortunately, the agent has no way of knowing whether there are any Dutch auctions that are terminating at that point in time (Dutch auctions do not have end times), and this adversely affects the performance of agent RA. The performance for the agent that deploys BA is low due to its bargaining nature. The agent maintains a low bid until the deadline has almost expired and finally concedes with a bid value of *RP*. Agent DE's performance is very close to that of RT, since the values for its k_{DE} and β are similar. The performance of agent COM1 indicates that with equal weighting, the success rate is close to the average success rate of those four single tactics combined (as would be expected). COM2's performance is slightly better than COM1 since heavier weights were placed on RT and DE, where both RT and DE performed well as a single tactic. The performance of C2 is good, since it takes every opportunity to bid in a terminating auction without considering other issues like payoff and time left to bid. C1 and C3 perform poorly due to their simplified behaviour of picking any auction randomly.

Figure 4 relates closely to the success rate performance of all the agents in the experiment. It shows the agents' final bid values (the bid values at which they acquired the item) as percentages of the average closing prices for all the auctions in the marketplace that closed within the period that the agents were active. It can be clearly seen that our six agents performed better than the three control models, since their final bid values are lower than any of the final bid values of the control models. The final bid value is considered as a very important factor in our analysis since it determines the closing price of a particular auction and the resulting payoff that the agent expects to get upon acquiring the item. The agents' behaviours in the selection of bid values are described in more detail in the next paragraph. This result leads us to conclude that our agents successfully acquire the item at a price lower than or equal to the reservation price, which is subsequently lower than the average closing price in all the auctions that closed within the period that they were active.

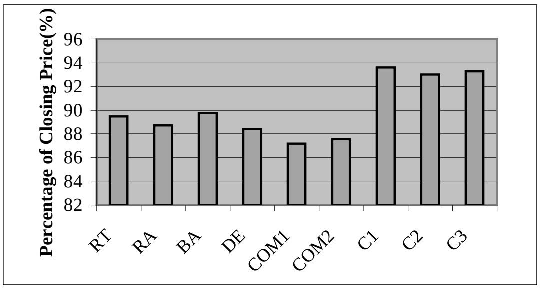

Figure 4 - Percentage of Closing Price Comparisons

Figure 5 shows the average payoff for each agent participating in the experiment. Payoff is calculated by subtracting the agent's bid value (when it successfully acquired the item) from the reservation price. The average payoff is the summation of all the payoffs divided by the number of times the agent successfully acquires the item. As can be seen, all our agents performed better than the control models except for DE. DE's low average payoff can be attributed to the fact that the agent is desperate to get the item. It starts bidding at a bid value close to its reservation price and quickly reaches this reservation price. Its goal is to try and get the item as early as possible. The high average payoff by our other agents clearly shows that payoff is an important criteria that the agent needs to take into consideration when bidding in any auction. Agent RT, RA, and COM2 achieved the highest average payoff, followed by BA and COM1. The performance of these agents (RT, RA, COM2, and COM1) can be related to the average time taken by the agents to acquire the item (see Figure 6). As an example, the time spent by agent RT to obtain the item is short, which means the recommended bid value at the time of acquisition is much lower than the reservation price. The closer the agent's acquisition time to t_{max} the closer its bid value to the reservation price. This is the same case with agent RA, COM2, and COM1. BA's performance is acceptable (even though it is the one that actively looks for a bargain, hence its payoff should be higher) since it continually tries to looks for a bargain and finally concedes when the auction is close to its end time. At this time, the agent bids at its reservation price resulting in a lower payoff. The performance of the three control models is poor because of their disregard for the payoff issue. Their main goal is to get the item at a price lower than or equal to the reservation price.

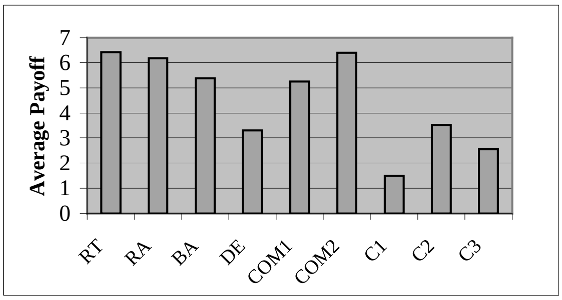

Figure 5 - Average Payoff Comparisons

Finally, Figure 6 shows the average time taken by the agents to acquire the item. Our agents spend between $t = 10$ and $t = 18$ to get the item, whereas the control models spend between $t = 8$ and $t = 16$. This indicates that the control models try to take any opportunity to bid in a particular auction as soon as possible. Agent BA spent the longest average time to acquire the item since it tries to look for possible bargains in the auctions. In the case of DE, it is more interested in obtaining the item as soon as possible, thus less time is required. This is a similar situation for agent RT, where it views time as the most important consideration. RA, COM1 and COM2 have similar average times since time is not a major consideration in their reasoning.

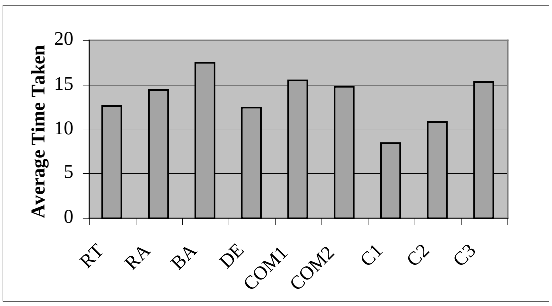

Figure 6 - Average Time Taken Comparisons

When taken together, the preliminary results indicate that allowing agent to participate in multiple auctions is an effective strategy for meeting the consumer's objectives. Our agent consistently outperformed C1 in terms of success rate, average payoff, and average time taken to acquire the item.

5 Related Work

There have been several attempts to design sophisticated and efficient bidding strategies for agents participating in online auctions. The Recursive Modeling Method (RMM) uses a decision theoretic paradigm of rationality, where an agent makes decisions based on what they think the other agents are likely to do, and what the other agents think about them and so on [Gmytrasiewicz and Durfee, 1995]. The downside to this approach is that not all the information in the recursive model may be relevant to the agent; there may be cases where this information does not influence the agent's decision-making at all. It is also possible that little or no information is available for the agent to use in the types of environment considered in this work. Faratin et al's model is broadly similar to the one defined in this paper [Faratin et al., 1998]. However, there are also several important differences between one-to-one negotiations and multiple auctions. Chief amongst these, are the type of tactics that are considered relevant and the aspect of the domain that need to be reflected in these tactics.

The possibility based approach to designing bidding strategies views the strategies as a decision made under uncertainty [Garcia et al., 1998; Gimenez-Funes et al., 1998]. They rely on the possibilistic-based decision theory to model a buyer agent's behaviour. This work was implemented in the FM96.5 (a Java based electronic auction house) that focuses on Dutch auctions only. This approach differs from ours in two ways. Firstly, it operates on one auction protocol only. Secondly, it generates a possibility distribution based on previous similar market situations. The single auction approach used by this model is not applicable in the environment that we are working in. Moreover, our approach uses a probability based on current market value rather than relying on the similarities of previous market histories. Park et al. proposed an adaptive bidding strategy called the p-strategy that is based on stochastic modelling and uses reinforcement learning to make the strategy adaptive [Park et al., 1999]. This approach was implemented using the University of Michigan Digital Library (UMDL) auction based on a continuous double action protocol and the strategy was designed for a seller agent. The model is designed for implementation in a continuous double auction protocol, hence it is inappropriate to the work described here.

Preist proposed an algorithm design for agents that participate in multiple simultaneous English auctions [Preist et al., 2000a; Preist et al., 2000b]. The algorithm proposes a co-ordination algorithm to be used in the environment where all the auctions terminate simultaneously, and a learning method to tackle auctions that are terminating at different times. This work is designed specifically for purchasing multiple items in multiple English auctions, and it is not applicable in an environment where multiple auctions with multiple protocols are used. BiddingBot is a multi-agent system that supports users in attending, monitoring and bidding in multiple auctions through a process called co-operative bidding [Ito et al., 2000]. This approach demonstrates how agents can cooperate and work together to do the bidding process in multiple auctions. It consists of one leader agent and several bidder agents, where the leader agents acts as the coordinator and facilitator of the whole bidding process. Bidding is done by exchanging messages between the user, the leader agent, and the bidder agents. However, the main problem with this approach is that the agents do not actually make the bidding decision. This decision is left to the user. Thus, the agents do not have full autonomy and the decision-making process is slow since the agent needs to interact with the user from time to time.

The first trading agent competition (TAC-2000)[6], provided a platform for agent designers to allow their autonomous agents to compete with one another in online simultaneous auctions for complimentary and substitutable goods. The key feature of TAC is that it required autonomous bidding agents to buy and sell multiple interacting goods in auctions of different types [Greenwald and Stone, 2001]. Each participant agent is a simulated travel agent, with the goal of assembling a number of travel packages for its 8 clients. Each client is characterised by a random set of preferences for the possible arrival and departure dates, hotel rooms and entertainment tickets. The objective of a TAC agent is to maximise the total satisfaction of its customers (i.e., the sum of the customer utilities). The competition attracted 22 participants, in which 12 qualified for the semi-finals and finals. One of the interesting outcomes from this competition was the fact that the designs of these agents were motivated by a wide variety of research interests including machine learning, artificial life, experimental economics and real-time systems. ATTac–2000, which emerged as the winner, used a principled bidding strategy that included several elements of adaptivity. This feature allows the agent to cope with a variety of possible scenarios during the competition. Another top finisher, RoxyBot, was based on heuristic search techniques. Aster, the third placed agent, used a cost estimation framework that can respond to strategic behaviour of competing agents. Finally UmbcTAC, placed more importance on the network load (in that it adapted to network performance more frequently than competing agents and received more frequent updates). All the TAC agents are involved with two basic activities; bidding and allocation. The agents need to determine the most profitable allocation before they decide on what goods to bid. If they fail to obtain the good, they need to recalculate the optimal allocation. This scenario is somewhat different from ours in that we concentrate on the bidding strategies to obtain one particular item rather than worrying about the complementary goods that need to be bundled with the desired item. Moreover in TAC, the type of auctions that are conducted are continuous one-sided auctions, standard English ascending multi-unit auctions and continuous double auctions as opposed to our environment that runs simultaneous standard English auctions, Dutch auctions and Vickrey auctions.

[6] http://tac.eecs.umich.edu

6 Conclusions and Future Work

This paper presented a bidding algorithm that can be used for an agent to participate in multiple online auctions with multiple protocols. We initially described the environment in which the bidding algorithm is implemented. The bidding algorithm itself is based on multiple tactics, that each deal with a single facet of the agent's reasoning. These tactics are then combined in order to give the agent's overall view at a given moment in time. Our preliminary experimental results demonstrate the potential of our approach.

Our main line of work is to further explore the development of strategies for our bidding agent. Since the environment in which the multiple auctions are running is highly dynamic, we intend to extend this work by using an evolutionary approach based on genetic algorithms (GAs). GAs will be used to determine the relative success of these strategies and how these strategies can evolve over time to better fit their environment. The performance of the agent is very much influenced by the strategy employed which, in turn, relates to the values of k and β *in* the given tactics and the weights for each tactic when combined. Different strategies may perform well in some environments but may perform poorly in another. The numbers of strategies that can be employed are endless and the search space is huge. To address this issue, we intend to use GAs to search for the most successful strategies in predefined environments in an off-line fashion. This knowledge can then be codified into an agent's online reasoning behaviour so that it can determine the best strategy to employ in the prevailing circumstances.

References

[Faratin et al., 1998] P. Faratin, C. Sierra, and N. R. Jennings. Negotiation decision functions for autonomous agents. *International Journal of Robotics and Autonomous Systems 24 (3-4) pp. 159-182*, 1998.

[Garcia et al., 1998] P. Garcia, E. Gimenez, L. Godo, and J. A. Rodriguez-Aguilar. Possibilistic–based design of bidding strategies in electronic auction. In *Proceedings* *of the Thirteen European on Artificial Intelligence (ECAI-98),* Brighton, UK, 1998, pp. 575-579.

[Gimenez-Funes et al., 1998] E. Gimenez-Funes, L. Godo, J. A. Rodriguez-Aguilar and P. Garcia-Calves. Designing bidding strategies for trading agents in electronic auctions. In *Proceedings of the Third International Conference on Multi-Agent Systems (ICMAS' 98),* Paris, France, 1998, pp. 136-143.

[Gmytrasiewicz and Durfee, 1995] P. J. Gmytrasiewicz and E. H. Durfee. A rigorous, operational formalization of recursive modeling. *In Proceedings of the First International Conference on Multi-Agent Systems (ICMAS),* Menlo Park, California, 1995, pp. 125-132.

[Greenwald and Stone, 2001] A. Greenwald and P. Stone. Autonomous Bidding Agents in the Trading Agent Competition. *IEEE Internet Computing, March-April 2001*, pp. 52-60.

[Ito et al., 2000] T. Ito, N. Fukuta, T. Shintani, and K. Sycara. BiddingBot: A multiagent support system for cooperative bidding in multiple auctions. *In Proceedings of the Fourth International Conference on Multi-Agent Systems (ICMAS'2000),* Boston, USA, 2000 pp. 399-400.

[Klemperer, 1999] P. Klemperer. Auction Theory: A Guide to the Literature. *Journal of Economic Surveys Vol .13 (3),* 1999.

[McAfee and McMillan, 1987] R.P. McAfee and J. McMillan. Auctions and bidding. *Journal of Economic Perspectives*, pp. 699-738, June 1987.

[McAfee and McMillan, 1992] R.P. McAfee and J. McMillan. Bidding rings. *American Economic Review, Vol. 82 No. 3*., pp. 579-599, June 1992.

[Milgrom and Weber, 1982] P. R. Milgrom and R. J. Weber. Theory of Auctions and Competitive Bidding. *Econometrica,* pp. 1089-1122, September 1982.

[Park et al., 1999] S. Park, E. H. Durfee, and W. P. Birmingham. An adaptive agent bidding strategy based on stochastic modeling. *In Proceedings of the Third International Conference on Autonomous Agents (Agents '99),* Seattle, Washington, 1999, pp 147-153.

[Preist et al., 2000a] C. Preist, A. Byde, and C. Bartolini. Economic Dynamics of Agents in Multiple Auctions. International Conference on Autonomous Agents, Montreal Canada, 2001.

[Preist et al., 2000b] C. Preist, C. Bartolini, I. Phillips. Algorithm design for agents which participate in multiple simultaneous auctions. *In Agent Mediated Electronic Commerce III (LNAI).* Springer-Verlag: Berlin, German, pp. 139-154, 2001.

[Sandholm, 1999] T. Sandholm. Distributed rational decision making. In the textbook *Multiagent Systems: A Modern Introduction to Distributed Artificial Intelligence, G. Weiß, ed.,* MIT Press, pp. 201-258, 1999.

[Vickrey, 1961] W. Vickrey. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance,* 18:8-37, 1961.