Autonomous Trading Agent Design in the Presence of Tradeoffs

Ioannis A. Vetsikas
Computer Science Dept., Cornell University
Ithaca, NY 14853, USA
vetsikas@cs.cornell.edu

Bart Selman
Computer Science Dept., Cornell University
Ithaca, NY 14853, USA
selman@cs.cornell.edu

ABSTRACT

In previous work we have introduced a principled methodology for systematically exploring the space of bidding strategies when agents participate in a significant number of simultaneous auctions, and thus finding an analytical solution is not possible. We decompose the problem into sub-problems and then use rigorous experimentation to determine the best partial strategies. In this paper we clarify and extend our methodology. We discuss our agent design for TAC 2003 and furthermore the changes to our agent as a result of the rule changes in TAC 2004. We also present a "full" set of experiments for determining an overall "optimal" strategy in the 2003 and 2004 Trading Agent Competition (TAC). Our agent was created by using the results of this methodology and has consistently been the top-scoring agent in several rounds of the TAC competition.

Keywords

agent-mediated electronic commerce, bidding agents, determining bidding strategies, electronic marketplaces, simultaneous auctions

1. INTRODUCTION

Auctions are becoming an increasingly popular method for transacting business either between individuals over the Internet (e.g. eBay) or even between businesses and their suppliers. While optimal strategies for bidding in an auction for a single item are known, in practice agents (or humans) are mostly interested in multiple items. They wish to bid in several auctions in parallel for multiple substitutable or complementary goods. Goods are called complementary (resp. substitutable) if the value of acquiring both together is higher (resp. lower) than the sum of their individual values. The fact that the value of each good to an agent depends on other goods is what makes this problem particularly hard. There had been relatively few studies about agents bidding in multiple simultaneous auctions and they mostly involve bidding for substitutable goods (e.g. [1], [3], [2]). However the introduction of the Trading Agent Competition (see [9]) has given momentum to the research in this field by providing researchers with a challenging benchmark domain which incorporates several elements found in real marketplaces.

In [5] we presented a principled methodology for systematically exploring the space of bidding strategies when it is not possible to find an equilibrium solution, because the agent participates in a significant number of simultaneous auctions. Whereas most other teams decided that decomposing the bidding problem completely was probably not to their best interest, because of the obvious dependencies among the bids and concentrated to a significant degree on learning a good predictor for the auction closing prices (e.g. ATTac used a boosting-based method [4]), we decided to explore "in the opposite direction"; we decompose the problem into subproblems and then use rigorous experimentation to determine the best partial strategies. In this paper we clarify the procedure that determines what experiments to run in order to explore the strategy space, we present and justify the changes we incorporated into our agent as a result of the 2004 TAC rule changes, and we present a "complete" set of experiments for determining an overall "best" strategy in the Trading Agent Competition (TAC) both under the 2003 and 2004 rules. These are the primary contributions of the paper. Furthermore we extend the methodology by adding strategies that are based on equilibria that we can compute or approximate and in particular we briefly examine the Bayes-Nash equilibria that exist in the case that the agent must decide on the price of its bid in an auction that can close at one or more predetermined times; thus we generate one more strategy for our methodology inspired by this analysis. In addition we continue to incorporate other ideas into our agent (e.g. examining the effect of better price predictions) and we present some results that indicate that it is possible to reduce the number of games required by our methodology by computing the effects that the random parameters have on the agent scores.

This paper is organized as follows. In section 2 we present the problem of trading in simultaneous auctions and our methodology. In section 3 we present the tradeoffs present in the TAC 2003 game and the strategies available to the agent and our analysis of the equilibria for a particular tradeoff. In section 4 we present the updates that we implemented in our agent as a result of the TAC rule changes in 2004. In section 5 we present all the experiments we did in order to explore the strategy space both for the 2003 and 2004 version of the TAC game and how these results were evaluated.

2. TRADING GOODS IN SIMULTANEOUS AUCTIONS

2.1 Problem setting

The general problem setting that we deal with involves several autonomous agents, which wish to trade commodities in order to acquire the goods that they need. There is a predefined time window during which the trades can take place (defining the duration of each "game"), after which each agent calculates the payoff to itself.

The agents are not allowed to cooperate in any explicit way and they are also assumed to be self-interested. In particular, each agent i is trying to maximize its own utility function, which in most cases is assumed to be *quasilinear*. The combination of goods owned by an agent should include several complementary and substitutable goods in order for the game to be interesting; otherwise one might be able to find an equilibrium to the game analytically. The mechanism used in order to exchange commodities is several different auctions during which each unit of a certain commodity is traded in exchange for a monetary payment. We will assume that there is *no discriminatory pricing* in these auctions, which means that if two agents wish to buy the same good at the same time they will have to make the same payment. We will also assume that *similar goods are sold in auctions with similar rules*. Other than that, we allow the auctions to have a wide variety of rules, as far as the quantities and types of goods sold, the auction closing times and allocation rules and whether agents may act as buyers or sellers in the auction. The question we are interested in answering is what bids to place at each auction. There are therefore 3 main parameters to determine: the *quantity* of each good to be bought or sold, the *prices* offered for each individual unit and the *times* at which the bids are placed.

Due to space considerations we cannot present the rules of the TAC game here, which are necessary for the reader to fully comprehend the observations presented in latter sections, but one may find them on the web at $http : \backslash\backslash www.sics.se\backslash tac$. We selected the TAC game for our experiments because it encapsulates most of the issues of the general problem setting and is thus an appropriate test-bed for evaluating our agent design and our methodology.

2.2 Our proposed methodology

There are two main parts in our methodology. Initially we decompose the problem into subproblems and decide on the range of possible strategies that can be applied. Then we use rigorous experimentation to evaluate these strategies and to determine the best strategy overall.

2.2.1 Problem Decomposition

According to our methodology we must decompose the problem as much as possible. This involves two phases:

1. The quantities to be bought for each commodity are determined independently of the price and the placement time of the bid and are computed by maximizing the utility of the agent assuming that all the goods are bought at some predicted prices and that every unit will be bought instantly. A dedicated module called "the planner" (or optimizer) is doing this task. For more information on how this module of agent WhiteBear is implemented, see [5].
2. Once the quantity that the agent should buy has been computed we need to find the price and the placement time of each bid. We consider separately each different auction type, which corresponds to a different commodity type, and we compute the space of possible "partial strategies"[1] for this auction type. Initially we compute the extent of this space by setting boundary partial strategies, that is strategies that place bids at the lowest or highest prices that a rational agent would consider possible for this setting and at the earliest or latest time possible. This defines the boundaries of the strategy space for each auction type. Having accomplished this, it is necessary to generate intermediate strategies[2] to use in our experiments, since we cannot possibly consider all possible strategies in the strategy space. This is done systematically by
   - computing equilibrium strategies for each auction (making some relaxations if necessary) and then using our observations about the equilibrium to select a strategy,[3]
   - combining the boundary strategies, e.g. by bidding as low as possible in some cases and the highest possible in others, and
   - modifying the boundary strategies using empirical knowledge from the domain and strategies that work well in other domains, e.g. using the idea of strategic demand reduction to modify the "buy all tickets at the beginning of the game" strategy in the TAC game.

   For the more details and the specifics of how this part of the methodology is applied to the TAC domain, see section 3.

2.2.2 Experiment Setup

The number of possible combinations of partial strategies coupled with the fact that each strategy can behave differently based on the mixture of agents that participate in the game, leads to a prohibitive number of experiments that need to be performed in order to get the "overall best strategy". We propose one possible way of exploring this space: we determine how the behavior of the multi-agent system changes when exactly one of the partial strategies is varied (changed) between the participating agents, while all the rest are shared by all the agents and kept fixed. To do this in a systematic way we use several sets of experiments, each of which is designed to evaluate one particular partial strategy in different mixtures of agents.[4]

In each experiment set we keep a fixed number of agents who are using the intermediate strategies, while systematically changing the mixture of agents using the boundary cases. This will explore sufficiently the different multi-agent environments that the agents can participate in, since the behavior caused by the intermediate strategies is within the bounds of the behavior caused by the boundary strategies. Using statistical tests we can evaluate the performance of the agent and determine whether some agent performs significantly better than the others. Therefore each experiment set explores the strategy space in one dimension.

In order to explore the whole strategy space, we run several experiment sets, varying different partial strategies in each set; the partial strategies that are not varied are set to the best strategy (or strategies found so far). In case our experiments show that the fixed strategy we used for one particular subproblem was not the best one, then we vary this partial strategy in the next experiment. After a sufficient number of experiment sets, we reach an "optimal strategy"; if any one of its partial strategies is changed the agent using this modified strategy cannot do better than the optimal strategy.

In addition, if we want to explore a feature that would improve the performance of the optimizer, we set half the agents to use this feature and half not to, while all agents use the best partial strategy combinations that we have found thus far. An extra benefit of

[1] We called these strategies partial because they only deal with one particular type of auction. These are usually computed analytically or using domain knowledge.

[2] If the boundary strategies will place bids at price p_{low} and p_{high} in a certain case, then the intermediate strategy should place its bid at price $p : p_{low} \leq p \leq p_{high}$.

[3] In our previous work we had not considered this seriously because we believed that finding an analytical solution even to a relaxed problem was not feasible. However we have recently computed usable equilibria, like the one presented in theorem 3.1, that we use to generate partial strategies.

[4] This is important since the performance of each strategy depends on the strategies used by the other agents. In previous work, we did not present a clear methodology for what sets of experiments to run, but we clarify this process in this paper.

this methodology is that it allows us to derive general observations about the behavior of certain strategies in different domains. For a description of the experiments run in the case of TAC see section 5.

3. BIDDING STRATEGIES

Each different commodity in TAC is traded in a different type of auction. According to our methodology, we need to generate partial strategies for every different set of auctions. We will concentrate our discussion here towards the bids for plane tickets and hotel rooms, since these two present interesting tradeoffs for our agent.

3.1 Paying for adaptability

The purchase of plane tickets for the TAC 2003 game presents an interesting dilemma. Ticket prices are expected to increase significantly with time. Thus delaying the purchase of plane tickets increases the flexibility of the agent and hence provides the potential for a higher income, at the expense of paying more for these tickets.

According to our methodology the first step is to decide the boundary strategies. Since the only issue is the time of bid placement, two obvious strategies are to buy everything at the beginning or to defer all the tickets purchases at a much later time. Initially we set this later time to be right after 2 hotel auctions have closed. The reason for this is (i) that at that time the intentions of the other agents can be partially observed by their effect on the auctions' bid prices meaning that the room prices approximate sufficiently their potential closing prices and (ii) that the plane ticket price increases after this point tend to be quite substantial. However this is still not a very good boundary case; a further improvement is to buy some tickets at the start of the game. We buy the "almost certain to be used" tickets (about 50%) at the beginning (computed based on the client preferences and the ticket prices) and we have empirically observed that these tickets are almost never wasted. Another intermediate strategy comes from the idea of strategic demand reduction [7]: we compute the minimum number of tickets which, if left unpurchased, will allow the agent to complete its itineraries even if it fails to buy a hotel room on days during which it wishes a lot of rooms. Thus 80% to 100% of the tickets are now bought at the beginning.

Further improvements were obtained (i) by estimating the likelihood of price increases and using this information to bid earlier (resp. later) for tickets whose price is very likely to increase significantly (resp. little) and (ii) by using historical averages of the hotel prices in previous games to set the prices that the optimizer uses at the beginning of the game in order to calculate the agent's utility. We also implemented especially for this paper the price prediction technique presented by the Walverine team (see [8]), but we did not see much difference in the agent's performance.

Note that this tradeoff does not exist under the TAC 2004 rules (see section 4).

3.2 Bid Aggressiveness

The main issue when bidding for hotel rooms is how *aggressively each agent should bid* (the level of the prices it submits in its bids). If it bids low it might get outbid, while if it bids high (i.e. aggressively) it is likely to enter into price wars with the other agents. As far as the timing of the bids is concerned, there is little ambiguity about what the optimal strategy is. The agent should wait until the first hotel auction is about to close to place its first bids in order not to increase the prices earlier than necessary, nor to give away information to the other agents.

The first boundary strategy is to place low bids: the agent bids an increment higher than the current price. This is the lowest (`L`) possible aggressiveness since the agent will never wish to bid less. The other boundary strategy is that the agent bids progressively closer to the marginal utility δU (for a particular room this is the change in utility that occurs if the agent fails to acquire it) as time passes. Since the agent will likely lose money if it bids above the marginal utility, this is the highest (`H`) possible aggressiveness. We also combined these into intermediate strategies by selected the following compromise: the agent that bids like the aggressive (`H`) agent for rooms that have a high marginal utility δU, and bids like the low-aggressiveness (`L`) agent otherwise. This is the agent of medium or moderate (`M`) aggressiveness.

Another partial strategy for this tradeoff can be computed by using equilibria like the following:

THEOREM 3.1. *Assume that N risk-neutral agents wish to buy 1 unit of a certain good each. An independent seller sells m units of the desired good in an m^{th} price auction, i.e. the good is sold to the agents which submitted the m highest bids at a price equal to the lowest winning bid. The agents have valuations (utilities) u_i which are i.i.d. with probability distribution $F(u)$ in the first round. There can be a second round with probability $(1-p)$. If a second round does exist, the agents have new i.i.d. utilities $\widetilde{u}_i$ drawn from some distribution $H(u)$ and can submit new bids as long as they are greater or equal to the bid price from the end of the first round and the agents are not allowed to remove previously placed bids. Each agent i knows (more accurately for its own utility) that its own utility $\widetilde{u}_i$ in the second round is drawn from distribution $G(u)$. The equilibrium strategy is the solution of the differential equation*

$$\big(u_i - g(u_i) + \frac{1-p}{p} \cdot U_L(g(u_i))\big) \cdot \frac{\Phi'(u_i)}{g'(u_i)} = \big(\Phi(u_i) - Y(u_i)\big) \cdot \Psi(g(u_i)) \qquad (1)$$

where
$\Phi(x) = \sum_{i=0}^{m-1} C(N-1,i) \cdot \big(F(x)\big)^{N-1-i} \cdot \big(1-F(x)\big)^{i}$,
$Y(x) = \sum_{i=0}^{m-2} C(N-1,i) \cdot \big(F(x)\big)^{N-1-i} \cdot \big(1-F(x)\big)^{i}$,
$\widetilde{\Phi}(x) = \sum_{i=0}^{m-1} C(N-1,i) \cdot \big(H(x)\big)^{N-1-i} \cdot \big(1-H(x)\big)^{i}$,
$\widetilde{Y}(x) = \sum_{i=0}^{m-2} C(N-1,i) \cdot \big(H(x)\big)^{N-1-i} \cdot \big(1-H(x)\big)^{i}$,
$\Psi(Q) = 1 + \frac{1-p}{p} \cdot \Big\{ \int_Q^{+\infty} G'(z) \cdot \frac{\widetilde{\Phi}(Q)-\widetilde{Y}(Q)}{\widetilde{\Phi}(z)-\widetilde{Y}(z)} \cdot e^{\int_Q^z \frac{-\widetilde{Y}'(\omega)}{\widetilde{\Phi}(\omega)-\widetilde{Y}(\omega)} \cdot d\omega} \cdot (\widetilde{\Phi}(z) + \int_Q^z \frac{\widetilde{Y}(\omega) \cdot \widetilde{\Phi}'(\omega)}{\widetilde{\Phi}(\omega)-\widetilde{Y}(\omega)} \cdot d\omega) \cdot dz + \sum_{k=0}^{m-1} \frac{N \cdot (m-k)}{m \cdot (N-k)} \cdot C(N-1,k) \cdot \{(N-1-k) \cdot H'(Q) \cdot \big(H(Q)\big)^{N-2-k} \cdot \big(1-H(Q)\big)^{k} \cdot \int_0^Q G(\omega) \cdot d\omega - k \cdot H'(Q) \cdot \big(H(Q)\big)^{N-1-k} \cdot \big(1-H(Q)\big)^{(k-1)} \cdot \int_0^Q G(\omega) \cdot d\omega + \big(H(Q)\big)^{N-1-k} \cdot \big(1-H(Q)\big)^{k} \cdot G(Q)\}\Big\}$ *and*
$U_L(Q) = -\sum_{k=0}^{m-1} \Big\{ \frac{N \cdot (m-k)}{m \cdot (N-k)} \cdot C(N-1,k) \cdot \big(H(Q)\big)^{N-1-k} \cdot \big(1-H(Q)\big)^{k} \cdot \int_0^Q G(\omega) \cdot d\omega \Big\}$.
The boundary condition is $g(0) = 0$.

This equation has a unique solution; for the proof of this and other equilibria see [6]. We are currently experimenting with this strategy and will present experimental results in a forthcoming paper.

4. AGENT REDESIGN FOR TAC 2004

The rule changes in 2004 eliminated the tradeoff in the purchase of plane tickets. In fact it can be shown that it is a dominant strategy not to buy any tickets whatsoever at the beginning of the game. The reason for this is that until after the middle of the game there is no increase in the expected price of the plane tickets, so the agent can defer these purchases and thus has more flexibility in its planning.

The model of the random walk for the ticket changes is known, since the prices are perturbed every 10 seconds by a value drawn uniformly from

- $[-10, x(t)]$, if $x(t) > 0$
- $[x(t), 10]$, if $x(t) < 0$
- $[-10, 10]$, if $x(t) = 0$

where $x(t) = 10 + (t/540) \times (x - 10)$, and t is the elapsed time from the beginning of the game. The only unknown variable is the maximum change x which is drawn uniformly from $[-10, 30]$. But we can compute the probability of $x = z, z \in [-10, 30]$ as follows:

Let $\{y_i, t_i\}, \forall i \in \{1, \ldots, N\}$ be the pairs of the price changes y_i observed at times t_i and let N be the number of such pairs. $P[x = z] = \frac{1}{41}, \forall z \in \{-10, \ldots, 30\}$ and $P[x = z] = 0, \forall z \notin \{-10, \ldots, -30\}$. Since x is independent of the times $\{t_i = T_i\}$ when the changes occur, therefore $P[x = z] =$ $P[x = z | \{\bigwedge_{i=1}^{N} (t_i = T_i)\}]$. In addition the price change y_i only depends on time t_i and is independent of all $t_j, \forall j \neq i$, therefore $P[y_i = Y_i | (x = z) \wedge (t_i = T_i)] =$ $P[y_i = Y_i | (x = z) \wedge \{\bigwedge_{i=1}^{N} (t_i = T_i)\}]$. Given the observations made so far, the probability that $x = z$ for any $z \in \{-10, \ldots, 30\}$ is $P[x = z | \{\bigwedge_{i=1}^{N} (t_i = T_i)\} \wedge \{\bigwedge_{i=1}^{N} (y_i = Y_i)\}] =$ $\frac{P[(x=z) \wedge \{\bigwedge_{i=1}^{N} (y_i = Y_i)\} | \bigwedge_{i=1}^{N} (t_i = T_i)]}{P[\bigwedge_{i=1}^{N} (y_i = Y_i) | \bigwedge_{i=1}^{N} (t_i = T_i)]} =$ $\frac{P[x=z | \bigwedge_{i=1}^{N} (t_i = T_i)] \cdot p_z}{\sum_{\zeta=-10}^{30} \left\{ P[x=\zeta | \bigwedge_{i=1}^{N} (t_i = T_i)] \cdot p_\zeta \right\}} = \frac{\frac{1}{41} \cdot p_z}{\sum_{\zeta=-10}^{30} \frac{1}{41} \cdot p_\zeta} = \frac{p_z}{\sum_{\zeta=-10}^{30} p_\zeta}$, where $p_z = P[\bigwedge_{i=1}^{N} (y_i = Y_i) | (x = z) \wedge \{\bigwedge_{i=1}^{N} (t_i = T_i)\}]$. Given the value of x, $\{y_i = Y_i\}$ are independent of each other, thus $p_z = \prod_{i=1}^{N} P[y_i = Y_i | (x = z) \wedge \{\bigwedge_{i=1}^{N} (t_i = T_i)\}] =$ $\prod_{i=1}^{N} P[y_i = Y_i | (x = z) \wedge (t_i = T_i)]$. This probability is known; we compute $x(z, T_i) = 10 + \frac{(z-10) \cdot T_i}{540}$ and we have that

- if $x(z, T_i) > 0$ then
  $P[y_i = Y_i | (x = z) \wedge (t_i = T_i)] = \frac{1}{\lfloor \frac{(z-10) \cdot T_i}{540} \rfloor + 21}$
  if $-10 \leq Y_i \leq x(z, T_i)$,
  otherwise $P[y_i = Y_i | (x = z) \wedge (t_i = T_i)] = 0$.
- if $x(z, T_i) = 0$ then
  $P[y_i = Y_i | (x = z) \wedge (t_i = T_i)] = \frac{1}{21}$, if $-10 \leq Y_i \leq 10$,
  otherwise $P[y_i = Y_i | (x = z) \wedge (t_i = T_i)] = 0$.
- if $x(z, T_i) < 0$ then
  $P[y_i = Y_i | (x = z) \wedge (t_i = T_i)] = \frac{1}{1 - \lceil \frac{(z-10) \cdot T_i}{540} \rceil}$
  if $x(z, T_i) \leq Y_i \leq 10$,
  otherwise $P[y_i = Y_i | (x = z) \wedge (t_i = T_i)] = 0$.

Using these estimates we can compute the expected price increase at any point. Therefore we modified the strategy used in the agent as follows:

We decide how many plane tickets to buy depending on their expected price increase; the higher the increase is, the larger the percentage of desired tickets that the agent buys. A case of special interest is when the probability that the hidden parameter $x < 0$ for that flight becomes high, an event that only occurs in the second half of the game; in this case our agent buys all desired tickets for that particular flight. This is done not only because there is a definite price increase predicted in this case, but also because by that time the agent has secured most of the needed hotel rooms. Based on the performance of our agent in the competition it is clear that this strategy works extremely well.

It should be noted that we also tried other variations of this strategy, e.g. buying all desired quantity of a certain plane ticket when its price would reach the minimum price possible. However, based on experiments the we ran, it did not appear that such changes improved the performance of the agent, therefore they were not used in our agent that participated in TAC 2004.

	Average Scores			Statistically Significant Diff?		
#WB-M2H	WB-M2L	WB-M2M	WB-M2H	M2L/M	M2M/H	M2L/H
0 (178)	2466	2626	N/A	✓		
2 (242)	2251	2344	2359	✓	×	✓
4 (199)	2051	2101	2056	×	×	×
6 (100)	N/A	1138	865		✓	

Table 1: Scores for agents WB-M2L, WB-M2M and WB-M2H as the number of aggressive agents (WB-M2H) participating increases. In the last rows, ✓indicates statistically significant difference in the scores of the corresponding agents, while × indicates statistically similar scores. *The number inside the parentheses is the total number of games for each experiment and this will be the case for every table.*

		WB*xSM	WB*xSH	WB*x2M	WB*x2H
#	x=A	1941	1887	1744	1677
games	x=M	1729	1645	1686	1706
(206)	Difference?	✓	✓	×	×

Table 2: The effect of using historical averages in the optimizer. Early bidding agents benefit the most from this.

5. EXPERIMENTAL RESULTS

In this section we describe the controlled experiments we performed in order to determine the "best overall strategy" and the conclusions we drew from them concerning the tradeoffs described in section 3.

Notation To distinguish between the different strategies (or versions of the agent), we use the notation WB*xyz, where (i) x is M if the agent just models the plane ticket prices, A if, in addition to the modeling of plane ticket prices, historical averages are used in the optimizer, and W if we also use the price prediction technique of the Walverine team, (ii) y takes values 0, 2, which means that the agent buys its unpurchased tickets when the y^{th} hotel auction closes, or the value S, which means that the version based on strategic demand reduction is used and (iii) z characterizes the aggressiveness with which the agent bids for hotel rooms and takes values L,M and H for low, medium (moderate) and high degree of aggressiveness respectively.

Performance Evaluation To formally evaluate whether one version outperforms another, we use *paired t-tests*. Since we usually use multiple instances of a certain agent in each experiment set, we compute the t-test for all possible combinations of instances. We decide that the difference is statistically significant, if the probability (for any such combination) that the agent scores are drawn from the same distribution is less than $p = 5\%$.

5.1 Determining the strategy for TAC 2003

In the first experiment set we varied the hotel bidding strategy. We used agents WB-M2z (z=L,M,H), keeping all other partial strategies fixed, and we used a constant number of 2 instances of agent WB-M2M, while the number of agents WB-M2H was increased from 0 to 6. The rest of the slots were filled with WB-M2L. The results of this experiment are presented in table 1. Agent WB-M2M performs "reasonably well" in every situation, since it bids enough to maximize the probability that it is not outbid for critical rooms, and avoids "price wars" to a larger degree then WB-M2H. Based on these results we did not use low aggressiveness agents in the next experiments.

In the next experiment we evaluated the benefit of using histor-

#WB*AyH	Agent Scores 1	2	3	4	5	6	7	8	Average Scores WB*ASM	WB*A0M	WB*ASH	WB*A0H	Statistically Significant Difference? WB*ASM/WB*ASH	WB*A0M/WB*A0H
2 (428)	2458	2423	2412	*2394*	*2418*	*2400*	2474	*2377*	2431	2404	2474	*2377*	×	×
4 (421)	2577	2591	*2575*	*2560*	2545	2615	*2520*	*2515*	2584	*2568*	2580	*2518*	×	×
6 (318)	2503	*2330*	2387	2382	2353	*2355*	*2424*	*2366*	2503	*2330*	2374	*2382*	✓	×

Table 4: Scores for agents WB*AyM and WB*AyH (where y=S or 0) as the number of aggressive agents (WB*AyH) participating increases. The agents above the stair-step line are WB*AyM, while below are WB*AyH (the scores when y=0 are presented in *italic*).

#WB*A0z	Average Scores WB*ASM	WB*ASH	WB*A2M	WB*A2H	WB*A0M	WB*A0H	Statistically Significant Difference? A2z/ASz	ASz/A0z	A2z/A0z
0 (413)	*2175*	2213	*1936*	1904	*N/A*	N/A	M✓ H✓		
2 (650)	*2115*	2110	*2023*	1862	*2103*	2130	M✓ H✓	M× H×	M✓ H✓
4 (438)	*2404*	2419	*2261*	2207	*2347*	2385	M✓ H✓	M× H×	M✓ H✓
6 (1023)	*2430*	2442	*N/A*	N/A	*2367*	2382		M✓ H✓	

Table 3: Scores for agents WB*A2z, WB*ASz and WB*A0z (where z=M or H) as the number of early bidding agents (WB-A0z) participating increases. In the last rows, we compare the scores for the M and H aggressiveness of the two version; so M✓ in the A2z/Asz box means that the difference between A2M and ASM is significant.

ical prices for the price prediction. We introduced this feature and run an experiment in which we examined the benefit that agents WB*M2M, WB*M2H, WB*MSM and WB*MSH gain if historical prices are used. The results are presented in table 2. We observe that the agents which bid earlier are the ones who benefit from the use of this feature and that the increase of the price estimate causes the planner to generate itineraries which use slightly fewer rooms. This decreases the price wars between agents and improves their scores. Given that this feature improved (or did not deteriorate) the performance of all agents, we decided to use it in all our later experiments.

Having varied the hotel bidding strategy in the first experiment set, we then vary the plane ticket bidding strategy. We use 2 agents WB*ASz (z=M,H) in the mixture of agents and change the number of the other agents WB*A0z and WB*A2z (boundary strategies); half of these are of Medium and half of High aggressiveness, since these partial strategies were the best in the first experiments. The results are presented in table 3. Strategy $y = 2$, performs worse than the other two. The other two perform similarly overall. The only case, in which the WB*ASz's performance is statistically better than that of the early bidders (WB*A0z), is when there are lots of early bidders. From these results we determine that the strategic demand agent is performing most consistently and that is the reason we used it in our TAC agent.

Having determined that the strategic ($y = S$) and early bidding ($y = 0$) agents are the best strategies for the plane tickets, and given that we used partial strategy $y = 2$ in the first experiment set, we need to revisit our experiment for the hotel room bids. We vary the number of aggressive agents (H) from 2 to 6 while reducing the number of medium aggressiveness agents (M). Furthermore half of them use strategy $y = 0$ and half $y = S$. The results are presented in table 4. One can notice that in most cases the agents perform quite similarly to each other. Furthermore agent WB*ASM outperforms the other agents when the aggressive agents are in the majority and does statistically similar or better in all other cases. Another observation is the overall scores remain similar even when the number of aggressive agents increases (as opposed to the experiment presented in table 1), mainly due to the fact that these agents use historical prices in the optimizer. Since agent WB*ASM performs overall at least as good as the other agents, both when we vary the strategy for the hotels and for the plane tickets, *it follows that we have explored the strategy space that we had and we should use agent WB*ASM in TAC*, which we did.

One of the directions in which we had not put much effort as far as our agent design is concerned, was the prediction of the closing prices for the hotel auctions. Most other agents have spent a considerable effort in this area. Our agent discounts the effect of the lack of a good prediction thanks to the bidding strategies that it uses. It was however our intention to implement some better prediction in order to examine the result on the performance of the agent. To this effect we chose to implement the price prediction algorithm of the Walverine team (it is uses tatonment to compute the prices; for more details see [8]), since it is shown to generate some of the best price predictions among all TAC agent implementations. We have tested this method by running four instances of agent WB*ASM using this prediction to set the prices used by the optimizer at the beginning of the game (let us call this agent WB*WSM) and four others that did not use this. We run 235 games and the performance of the two agents was actually quite similar. The average scores for agents WB*ASM and WB*WSM were respectively 2263 and 2251. This seems to indicate that the strategy we selected deals successfully with the errors (which are not catastrophic though) in the price predictions that the optimizer uses.

5.2 Determining the strategy for TAC 2004

The previous experiments determine why both in TAC 2002 and 2003 we used agent WB*ASM. The rule changes in 2004 leave only one tradeoff to be examined: the hotel bidding strategy. Therefore only one experiment set is necessary in this case, one in which the strategy for the hotel room bidding is varied. We use 2 agents of moderate (M) bidding aggressiveness and vary the number of agents of low (L) and high (H) aggressiveness. The results are presented in table 5. The differences between the different versions of the agent are now much smaller; this is a result of the fact that the agents no longer buy the flight tickets early and thus are able to change their plans much more easily, which in turn reduces the performance difference between the various strategies. From this experiment set, we can conclude that the low aggressiveness strategy is still dominated by the other two, but less so than in the case of TAC 2003. It is unclear which strategy (moderate or high aggressiveness) is best, as they perform quite similarly and the differences are not statistically significant; in fact in most cases there is a high probability of similarity in the score distributions. In TAC 2004, we used the moderate strategy, but based on these experiments the aggressive strategy would have performed quite well. In two cases, more games are needed in order to determine whether the difference is statistically significant (denoted by "?" on table 5).

The biggest drawback in our methodology is probably the fact that in order to get results with a reasonable amount of confidence for each experiment we need to run a large number of games. If one also considers the fact that quite a few experiments are needed in order to explore the strategy space, this leads to a huge number of games and that takes a very long time to run. The large number of games needed is mostly due to the fact that the TAC game has an overwhelming number of randomness involved, so that even when two instances of the same agent are running they can get considerably different results in the same experiment depending on the number of games played. This is more of a problem during the finals of the Trading Agent Competition, when a very small number of games are played. To partially alleviate this problem the University of Michigan group decided to perform a linear regression of the scores that an agent gets against a number of parameters that

	Agent Scores								Average Scores			Stat. Significant Difference?		
#H agents	1	2	3	4	5	6	7	8	Low	Moderate	High	Low/Mod	Mod/High	Low/High
0 (259)	3310-15	3275-17	3264+15	3295-21	3260 -8	3276 +1	3258+21	3223+24	3263	3293	N/A	$\times$		
2 (320)	3263-12	3293-15	3209 -7	3229 +4	3227 -7	3187 +5	3260+21	3234+11	3213	3278	3247	$\checkmark$	$\times$	?
4 (385)	3191 +2	3191+12	3163 -6	3205-12	3214+22	3268-11	3240-16	3193 +9	3184	3191	3229	$\times$	$\times$	?
6 (413)	2899-31	2855 -9	2818+15	2842+23	2815+16	2934-20	2854 -1	2847 +6	N/A	2877	2852		$\times$	

Table 5: Scores for agents of low (L), moderate (M) and high (H) bidding aggressiveness as the number of aggressive agents participating increases. In each experiment agents 1 and 2 are of moderate (M) aggressiveness. The agents above the stair-step line are of low aggressiveness (L), while the ones below are the most aggressive (H). The averages scores for each agent type are presented in the next rows. In the last rows, $\checkmark$ indicates statistically significant difference in the scores of the corresponding agents, while $\times$ indicates statistically similar scores.

affect these scores. This was an interesting idea and a good starting point. However we did not want to assume that two different games can be characterized just by these random parameters. In the case of our experiments we have enough data to do more. Given that in our experiments we run pairs of the same agent we can run a linear regression in order to minimize the difference between the scores of two instances of the same agent in the same game. This is the "least" assumption that one could make about the scores in any game, that is that agents using the same exact strategy should perform similarly in the same game. We used this idea and we calculated the regressions and the adjustments for the experiments presented in table 5 (the average adjustments are presented next to the corresponding scores). In both cases the adjustments made were rather small and the scores of the agents were not equalized, however an examination of the results demonstrates that in most cases the statistical confidence (both in the case when there is a difference and when the scores are similar) after adjusting the scores has increased, which indicates that we can get more accurate results with fewer experiments. This is the case for example in the first row of table 5; before the score adjustment, it was unclear that the scores were statistically similar. Given that the adjustments that this process makes to the scores are rather small, we can reasonably argue that we are justified in using this improvement, since the results of each experiment are not altered. However the correlation of the linear regression predictions of score differences against the actual differences was between $R = 0.28$ and $R = 0.36$ depending on the case (row), which means that this adjustment cannot account for some of the randomness in the game. This means that the number of experiments needed to gain statistical confidence in the results is indeed reduced, but not by a large percentage.

For the record, the parameters found in each of the four experiments to be the most relevant to the score (among the 10 parameters that we used) were the sum of the number of days an agent's clients want their trips to last, the sum of the rooms per day that an agent wants weighted by the total demand for each day among all agents, the sum of the n highest valuations of its clients' entertainment preferences, where n depends on the length of the client's trip, and the sum of the hotel bonuses (for the expensive hotel).

6. CONCLUSIONS

To summarize we clarified the procedure that determines what experiments to run in order to explore the strategy space and we presented a "complete" set of experiments for determining an overall "best" strategy in the Trading Agent Competition (TAC) both under the 2003 and 2004 rules. We made this process more efficient by improving one of the biggest problems of our methodology; the large and time-consuming number of experiments needed is reduced, by accounting to some degree for the effects of random parameters. We also extended our methodology by generating partial strategies that are based on equilibria and in particular we briefly examined the Bayes-Nash equilibria for an m^{th} price auction with multiple closing times. In addition we presented the changes that were necessary for our agent as a result of the 2004 TAC rule changes. Finally we implemented a better algorithm for price prediction and the results showed that this improvement is made redundant by the flexible strategy that our agent uses.

We used our observations to determine our entries into the TAC competition and our agents performed consistently well, including winning TAC 2004 with a statistically significant difference in score from the runner up. We continue our work by examining and evaluating the performance of the agent when it uses a bidding strategy based on the computation of Bayes-Nash equilibria.

Acknowledgements We would like to thank the organizers of both competitions for their technical support during the running of our experiments and the competition (well over 10000 games!).

7. REFERENCES

[1] P. Anthony, W. Hall, V. Dang, and N. R. Jennings. Autonomous agents for participating in multiple on-line auctions. In *Proc IJCAI Workshop on E-Business and Intelligent Web*, pages 54–64, 2001.

[2] 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 MultiAgent Systems*, pages 399 – 400, July 2000.

[3] C. Preist, C. Bartolini, and I. Phillips. Algorithm design for agents which participate in multiple simultaneous auctions. *In Agent Mediated Electronic Commerce III (LNAI), Springer-Verlag, Berlin*, pages 139–154, 2001.

[4] P. Stone, R. Schapire, M. Littman, J. Csirik, and D. McAllester. Attac-2001: A learning, autonomous bidding agent. In *Agent Mediated Electronic Commerce IV. LNCS, volume 2531. Springer Verlag, Berlin.*, 2002.

[5] I. A. Vetsikas and B. Selman. A principled study of the design tradeoffs for autonomous trading agents. In *Proceedings of the 2nd International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 473–480, 2003.

[6] I. A. Vetsikas and B. Selman. Bayes-nash equilibria for m-th price auctions with multiple closing times. *ACM SIGecom Exchanges (to appear)*, 2005.

[7] R. Weber. Making more from less: Strategic demand reduction in the fcc spectrum auctions. *Journal of Economics and Management Strategy Vol6(3)*, pages 529–548, 1997.

[8] M. P. Wellman, D. M. Reeves, K. M. Lochner, and Y. Vorobeychik. Price prediction in a trading agent competition. In *Journal of Artificial Intelligence Research*, 2004.

[9] M. P. Wellman, P. R. Wurman, K. O'Malley, R. Bangera, S. de Lin, D. Reeves, and W. E. Walsh. Designing the market game for TAC. *IEEE Internet Computing, April*, March/April 2001.