CS886: Multi-Agent Systems

Project Two: Trading Agent Competition

Analysis of SouthamptonTAC in TAC'01

Michael Jarrett (99318764)

June 18, 2004

Abstract

The Trading Agent Competition is a yearly competition held to inspire practical research in the area of agents acting within marketplaces. Agents are required to make travel arrangements for clients by competing in auctions for hotels, flights, and entertainments.

SouthamptonTAC is an agent that first competed in the 2001 competition, placing fifth overall. Despite not achieving top ranks, it had the highest average score, and the most consistent performance.

The performance of SouthamptonTAC demonstrated that a risk-averse agent could indeed perform well in the competition, while demonstrating how important an assessment of market competitiveness is. Changes made to give SouthamptonTAC the ability to make such assessments improved SouthamptonTAC in subsequent years.

1 Introduction

The first International Trading Agent Competition[1] (TAC) was an agent competition held during the International Conference on Multi Agent Systems in July 2000 in Boston. In the TAC, agents play games in groups of eight, biding against each other in simultaneous auctions to secure travel packages. At the end of the bidding period, each agents allocates the goods won to their clients, and reports their allocation to the server. The server then awards a score based on how effectively the agent satisfied each client's preferences.

With the success of the first competition, a second TAC[2] was held in Tampa in 2001. Using knowledge gained from the first competition, the rules were adjusted to increase agent competitiveness. This paper investigates the agent SouthamptonTAC[3], a very successful agent during this competition.

Section 2 will present a summary of the trading agents competition, including a summary of improvements from the first TAC, and general game trends, while section 3 will mention some of the general issues faced by competitors. Section 4 will describe the approach of the SouthamptonTAC agent in the trading agent competition, section 5 will discuss the performance of the agent, and section 6 will examine the effectiveness of the agent, considering its weak and strong points. Finally, section 7 will present conclusions and summary.

2 Trading Agent Competition

The second Trading Agent Competion (TAC) simulates a marketplace for goods related to travel plans from the fictional town of 'TACTown' to Tampa over a five day period. The goal of each agent in the competition is to win flights, hotels, and entertainment tickets, and allocate these to their clients based on their preferences.

The agent must secure both the departure flight and return flight, as well as hotel rooms for each client's stay in one of the two hotels available. Furthermore, entertainment tickets are available on each night for three events, and each client has a certain value attached to having a ticket for each event.

2.1 Clients

Each agent in a game is assigned eight randomly-generated clients to satisfy. Each client has an ideal arrival day (IAD) and ideal departure day (IDD), which represent the client's preferred days of travel. The score for a client is penalized for travelling on days other than their ideal travel days.

$$\text{travel penalty} = 100[|\text{IAD} - \text{AD}| + |\text{IDD} - \text{DD}|]$$

Hotel Value (HV) represents how much a client values a high-quality hotel, between 50 and 150. The score for the client is increased by HV if the client stays in that hotel.

Bonuses are also given for each type of entertainment given, between 0 and 200 for each client. If a client has a ticket to that form of entertainment for a night of their stay, they gain the respective bonus.

The score given for a client is zero if a feasible travel package is not available. Otherwise it is:

$$Client\ score = 1000 - travel\ penalty + hotel\ bonus + entertainment\ bonus$$

2.2 Auctions

All purchasing and sales are run as simultaneous auctions. There are a total of twenty-eight auctions. These auctions run for twelve minutes, with the exception of the hotel auctions, which start randomly closing after four minutes.

2.2.1 Flights

Eight auctions are run for flights, one for each of four arrival and departure days. The price for these flights are initially between \$250 and \$400, and are adjusted every 30-40 seconds. The prices are constrained to stay within the range of \$150 to \$800. The flights may only be sold by the game auctioneer, which provides an effectively unlimited supply.

The prices for flights vary by a random number in the range from -\$10 to $x(t) = 10 + (x_a - 10)(t/720)$, where t is the game time in seconds, and x_a is a random value (unknown to the participating agents) in the range of 10 to 90. At $x_a = 10$, the prices represent a uniform random walk, which was the case in the first TAC. In most cases however, the price will show an increasingly upward bias as game time progresses.

2.2.2 Hotels

There are two hotels available in Tampa: the Tampa Towers (TT), and the Shoreline Shanties (SS). A stay in the Tampa Towers hotel will increase a client's score by their HV.

Eight auctions are run for hotels, one per night for each hotel for nights one through four. These are open-cry sixteenth-price auctions, with the top sixteen bidders acquiring a room.

To encourage early competition, after four minutes, one randomly chosen hotel has the auction closed. Each subsequent minute, another hotel auction is closed until game end.

2.2.3 Entertainment

Three forms of entertainment are available in Tampa: an amusement park (AP), alligator wrestling (AW), and a museum (MU).

Twelve entertainment auctions are run, one per event per night for nights one through four. These are continuous double auctions, where any agent may place a buy order or sell order. Twelve event tickets are assigned to each agent at the start of the game. In each auction, the agent finds out the current ask price when they place a bid, and will have their bids instantly satisfied if there is a lower ask.

Tickets can be 'sold short', meaning that an agent can sell tickets that it does not currently own, and take profit on them. The agent can wait for the same kind of ticket to be sold and then purchase it, hopefully at a lower price than it sold it for. However at the end of the game, if the agent has entertainment tickets outstanding, it is assessed a 200 point penalty per ticket.

2.3 Allocation

After the auctions close, the auction server calculates the optimal allocation of goods to clients for each agent, given the goods it has acquired. The score assigned an agent is the sum of the scores obtained by each of the eight clients, minus the amount spent on acquiring goods. The amount spent on acquiring goods can be decreased in the case of entertainment auctions, where agents can actually make profit by selling tickets.

2.4 Changes from TAC'00

A few changes were made to the rules between TAC 2000 and TAC 2001, based on the market trends observed during the former.

The auction time was reduced to twelve minutes from fifteen during the bidding phase, and the allocation phase was removed completely from TAC'01. TAC'00 demonstrated that optimal allocations could be calculated in a short time period, so optimal allocations were generated by the server.

To encourage hotel bidding, the first TAC would close hotel auctions after a random period of inactivity. It was found that this was insufficient to encourage bidding, as agents would just hold the auctions open with small incremental bids until the end, where prices would suddenly spike. The random closure per minute after the fourth minute encouraged the agents to bid earlier for hotels.

Flights in the first TAC had a uniform price distribution, meaning that there was no advantage to buying flights until the last possible instant, where uncertainty about hotels was at a minimum. With the upward bias, risk-seeking agents interested in low flight prices would buy most of their flights early on in the auction.

The initial entertainment distribution was also changed. Entertainment trading was not a large factor in the first game since tickets were uniformly distributed. In the new version, two ticket types are given on two days, totalling exactly twelve tickets for each agent. These lumped allocations encouraged the agent to trade for tickets for days and events they didn't have.

2.5 General Trends

As a welcome change from the first TAC, bids for flights came quite early in the game. Simpler agents simply bought the flights for their clients' ideal days right at the start, often within a minute of game start. Smarter agents observed the variation in price to predict the random parameter, then would bid based on the relative cost of delaying a bid versus uncertainty in acquiring hotels. Even in this case, bids for flights were concentrated in the first four minutes of the game.

The hotel bids did not show the skyrocketing prices that were shown in the first game, as agents were forced to bid early on hotels. As hotel priced were updated every minute, bids tended to be issued in bursts around this time. This led to other auctions showing similar bidding bursts, as agents would change their travel plans in response to hotel closures.

3 Issues to Consider in TAC'01

3.1 Calculating Optimal Allocations and Completions

While it is not required that agents calculate the optimal allocation of acquired resources to their clients at the end of the auction period, most agents rely on some base target allocation, which it uses as a guide to determine what other resources to acquire.

The problem of optimally allocating resources to clients is equivalent to the 'weighted set packing problem'[4], which is NP-complete. However, since there are only eight clients, and only a small number of resources, in the vast majority of cases this can be calculated using standard linear programming techniques within seconds, and was the method of choice for TAC'01. In the first TAC competition, several other methods were experimented with. Optimal AI search techniques are searches that are guaranteed to find an optimal solution before suboptimal solutions. These were able to find solutions quickly, and tended to scale better with larger numbers of clients. Most TAC'00 competitors used various forms of greedy algorithms, which were able to get consistently within a few percent of optimal solutions.

3.2 Bidding Timing

Agents need to make strategic decisions about when they will place their bids. At the highest level, a later bid will lead to less uncertainty, but can result in missed opportunity or higher prices. An early bid can take advantage of lower prices and early hotel closures, but risks acquiring goods that it will later not have a use for.

These behaviours tend to roughly correspond to the concept of risk-averseness. An agent that acquires goods early is generally called risk-seeking. These agents normally get better prices on their goods, but lock themselves into specific travel plans. To protect the investment in their previous goods, they often buy high priced hotels later to complete their travel packages. A risk-averse agent delays purchases when it can, allowing it to change travel plans with minimal penalty. These agents often don't get the best prices on flights, but have the flexibility to back down from hotels that are bid on too aggressively.

The advantage of each approach is demonstrated by the change of strategy between TAC'00 and TAC'01. In TAC'00, almost all the bidding was conducted very late in the game, when uncertainty was lowest. In TAC'01, there was much more incentive to bid early, leading to most agents acquiring their goods earlier, despite the risk.

While still an issue, the time required to make a bid itself was not as big a factor in the second TAC, since bidding was done earlier. Bid timing was still done by some agents, to ensure they could have bids accepted just before one of the hotel auctions was closed at random.

3.3 Price Prediction

With agents acquiring goods earlier, the expected utility of the packages they construct will depend on what the agent predicts the final sale price of goods it currently does not hold will be.

A variety of approaches were used here. Some use historic data from previous games it has played; this can be quite effective, but often suffers in later rounds as agents' behaviours are tuned by their creators. Others fitted curves or used basic projections based on the amount of time remaining. Still others incorporated behavioural models based on qualitative observations of the competition made by their creators.

4 SouthamptonTAC

SouthamptonTAC was an agent that competed in the second TAC, and competed in a revised state in the third TAC. It was implemented by researchers at the University of Southampton in the United Kingdom.

The architecture of the agent is that of a control loop as shown in figure 1. In a cycle lasting between six and thirty seconds depending on load, the agent queries the server for the latest prices in the auctions. This is processed to extract the relevant information from the bids, and is then forwarded to the allocator and a hotel price predictor, which also forwards its results to the allocator.

The allocator takes the current holdings of the agent and the predicted prices of goods yet to be obtained, and calculates an optimal allocation to clients. While this is in general an NP-complete problem, it was demonstrated during TAC'00 that it is reasonable to calculate this using standard integer programming approaches within seconds for the current problem size of eight clients. SouthamptonTAC learned from this lesson, and uses integer programming to find an optimal allocation.

The allocator feeds its results to three bidding units; one for flights, one for hotels, and one for entertainment bidding. These units place updated bids with the auction server.

Bidding is divided into three phases, which affect the behaviour of the agent. The *probing*

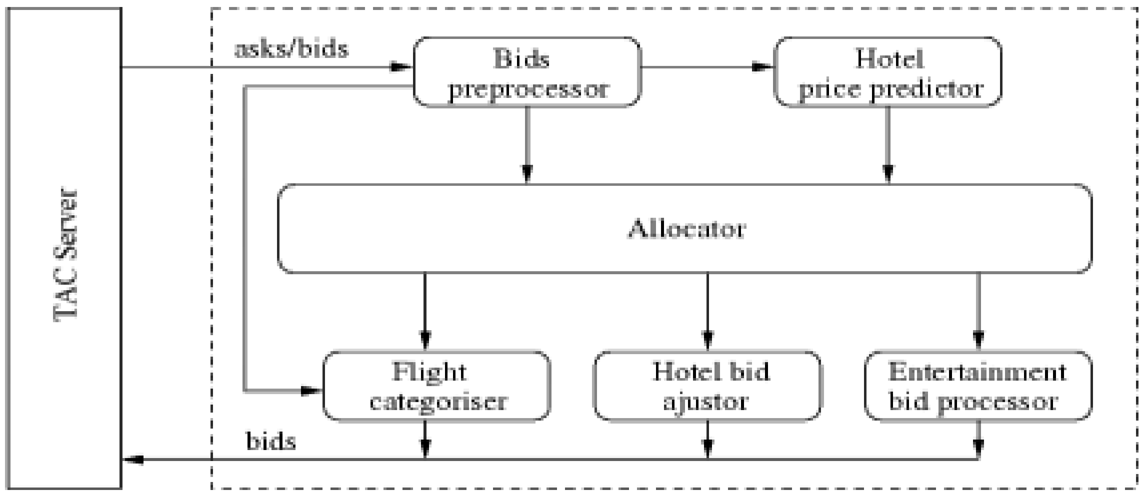

Figure 1: Control loop of SouthamptonTAC[3]

stage lasts up to minute four, during which the agent places initial hotel bids, and acquires a few of the most important flights. The *decisive* stage lasts from minutes four through eleven, and is where the agent makes strategic decisions about which resources to obtain. Finally, the twelfth minute is the *finalization* stage, where the agent places the bids necessary to obtain any goods not currently in its possession.

4.1 Flights

A critical aspect of flights is determining the rate at which the prices are slated to grow. However, this is difficult to determine accurately until late in the game, since the positive bias is increasingly applied as time progresses.

SouthamptonTAC classifies each flight auction into a numbered category between 0 and 3, representing the belief about the value of x_a. Category 0 represents $[10, 15]$, 1 to $[15, 30]$, 2 to $[30, 60]$ and 3 to $[60, 90]$. The exact equations for calculating this category in each bidding cycle is given in detail in [3]. A flight is promoted to the next category should any one-time price increment be seen as inconsistent with the current category.

A total of eight flights in each direction are required, assuming each client has a legitimate travel package. Eight to ten of the most critical flights are obtained early during the *probing* stage. The rest of them are delayed based on their categorization. Required category 3 flights are purchased immediately, category 2 flights at beginning of the *decisive* stage, category 1 flights once the hotel rooms are secured, and category 0 flights during *finalization*.

This takes a calculated approach to risk, where flights that will likely not be much more expensive are not purchased until the agent is sure it needs them. While taking a moderate approach to risk, compared to most of the other agents in the competition, this is widely risk-averse.

4.2 Hotels

The optimal hotel allocation is determined by the allocator, which needs an estimate of the closing price for each room. This is done based on a fuzzy inference system, with a set of rules derived from basic economic theory and general observations from previous games.

Fuzzy inference systems work similarly to an expert system, with rules given as if-then implications using basic boolean logic. However, the variables, rather than being boolean, are in fact members of 'fuzzy sets'. These are functions that map a variable to a value between zero and one, where zero is equivalent to a false, and one is equivalent to a true. The if-then rules are applied using these fuzzy values, and the output is then converted using another fuzzy set to a value in the output domain. More information on fuzzy inference systems can be found at [5].

Five factors were observed to affect hotel price, and are represented in the rules as fuzzy logic variables.

- Current hotel ask price.
- Counterpart hotel ask price
- Counterpart hotel closing time (if closed)
- Current game time
- Rate of change of hotel price.

A 'counterpart' hotel room is a room that is tightly tied to the results of another. For example, on day three, if the Tampa Towers closes early, it's likely that the Shoreline Shanties auction for day three will see a sudden increase as agents that didn't get their desired hotel adjust their bids.

The rules are designed to set an output fuzzy variable for the predicted rate of change for the hotel in the near future. This is used to predict the closing price, and allows the allocator to guess the approximate value choosing the hotel.

Sometimes, a hotel will show a sudden increase, for example, if its counterpart hotel closes early. In this case, the agent will 'withdraw' its bid to avoid overpaying for the room. It bids for a single room at ask+1. If a room is rapidly increasing in price, this will be quickly outbid, and therefore replaced. Since the rules are designed not to allow bid to be withdrawn, it is unclear if this is actually a legitimate tactic, or an exploit of the implementation.

4.3 Entertainment

Entertainment tickets are sold by the simplistic formula of bidding on tickets a client values over a certain threshold (currently set to 100), and offering all available tickets for prices that guarantee it reasonable profit. In both cases, the agent sets the price approaching its own utility for the ticket as game time progresses. Unallocated tickets are valued between 50 and 80 based on the number of tickets remaining, where 80 is the historical average for entertainment from the first TAC.

When evaluating open bids from other agents, fuzzy sets are used to opportunistically accept deals. The fuzzy set is 1 at the agent's current price, and 0 at the agent's utility for the ticket. The threshold for accepting an offer is set to 0.9, but could easily be tuned.

5 Agent Performance

For a newcomer, SouthamptonTAC performed admirably, having the largest average score and smallest standard deviation in score throughout the competition. It competed well, surviving even among highly risk-seeking agents.

During seeding rounds, SouthamptonTAC had the largest average score of all agents, at 3162. This score trumps many of the agents that later moved to become finalists, and thoroughly defeats some of the contenders from the first TAC such as RoxyBot and ATTac.

As the top competitors started playing amongst themselves, SouthamptonTAC started falling. During its semi-final heat, it placed second with 3615, with the third highest score overall. Livingagents was the first place winner at 3660, only slightly higher than SouthamptonTAC, which is to be expected, as discussed in section 6. Both scores were much higher than the third place agent, Urlaub01, at 3485.

During the final round, SouthamptonTAC placed fifth, with a score of 3254, however this was due to a crash during one game costing the agent several thousand points. With the score included, the agent would have placed third with 3530 average. The first and second place agents were livingagents and ATTac respectively.

The full competition results can be found in [2].

6 Analysis

SouthamptonTAC seemed to be a very strong competitor overall, obtaining the highest mean score. However, it seemed to struggle against some of the more aggressive agents, such as the ones encountered in the final round.

Livingagents is a genuine surprise in the second TAC; its success even against hardened competitors like ATTac runs counter to its simple strategy. The agent is essentially open-loop, deciding on an allocation, placing high bids on the items it wants, and then just sits mostly idle for the rest of the game. Compared to the learning strategy of ATTac and fuzzy logic control loop of SouthamptonTAC, this is quite primitive. The strategy will either succeeed spectacularly or fail miserably; if it wins all the auctions it wants, its score will be high. However, especially in competitive rounds (normally caused by risk-seeking agents), it may not achieve its desired allocation, and therefore have incomplete travel plans. This is an extreme in the area of risk-seeking behaviour.

ATTac was nowhere near as risk-seeking as livingagents, though it was still fairly aggressive with flights and hotels. It benefitted heavily from its learning abilities in the final round. Once ATTac was able to predict the behaviour of the auctions in the final round, it quicky started to catch up, to reach a second-place finish.

Relative to many agents in the finals, SouthamptonTAC is considered risk-averse, since it delays many decisions, especially on flights, until later in the game. This leads to flights being more expensive for it than a risk-seeking agent like livingagents, which buys all its flights immediately. However, a risk-averse agent will perform more consistently, as demonstrated from the standard deviation in score averages: SouthamptonTAC has the lowest at 867.3, while livingagents has an incredible 6506.0.

One observation made by the creators of SouthamptonTAC is that risk-seeking agents have a large effect on its ability to score well in games. If there are many risk-seeking agents, prices will skyrocket on hotels, and the risk-seeking agents will get into a price war that drains any potential utility gains. The risk-averse agents will detect this and switch to a less-contested hotel. SouthamptonTAC in particular is well-equipped to deal with this, since its fuzzy logic rules can quickly detect a situation where a price will grow quickly.

In competitions with mostly risk-averse agents, there are more benefits to act as a risk-seeking agent. Since prices are unlikely to grow very high, it is better to buy hotels and flights early where prices are still low. In this case, SouthamptonTAC does not perform as well, since its settings for the competition are fixed, and result in risk-averse behaviour. The ability to take advantage of a non-competitive game would allow SouthamptonTAC to raise its average score during these instances.

One feature present in ATTac that was not present in SouthamptonTAC is basic learning behaviour. In particular, ATTac learns from price changes in hotels, and uses that information to recognize similar situations in other games. Especially in later rounds, where there is only a few agents participating, this was shown to have a major competitive advantage. ATTac's scores in the finals, for example, quickly increased as it started learning trends from previous games. This feature in particular would be extremely useful in SouthamptonTAC, which while employing fuzzy logic, relies on a large number of constants chosen (and not always well) based on observations and reasoning of the creators.

SouthamptonTAC, like many agents, did not put much emphasis on optimizing profit from entertainment auctions. The prices offered for entertainment that was not valued by clients was rather low, with the upper limit being the average price from the first competition. With a proper fuzzy logic and time decay system, the minimum value of 100 before trying to acquire tickets serves only to limit the ability to make small additional gains from entertainment, while selling tickets cheaply that could greatly benefit other agents. With the lumped allocation of entertainment tickets, and the fact that they do not affect travel feasibility, a much more aggressive entertainment stance would be beneficial.

6.1 SouthamptonTAC-02

SouthamptonTAC was improved after the competition, and entered into the third TAC the next year, where it placed second in the finals[6]. Several lessons were taken from TAC'01 and used to produce the improved SouthamptonTAC-02[7]. One improvement that likely led to dramatic improvement in performance was the classification of each game as competitive, semi-competitive, or non-competitive. This classification affected many of the game constants, and how aggressively the agent would bid. Which classification to choose is based on past games, and can change during the game.

A further improvement was that of entertainment auctions. Rather than just bidding on tickets which present a utility over a fixed constant, the entertainment bidder bids on whatever entertainment is needed for an optimal allocation.

7 Conclusions

The Trading Agent Competition has served as an effective means to bring together researchers from a variety of academic and industry research settings to apply AI techniques to real-world agent environments. The competition itself has evolved as the relative strengths and weaknesses of agents in general have been shown. Agent techniques for dealing with complementary and substitutable auctions have developed enormously during this time.

SouthamptonTAC proved to be a formidable competitor within the second TAC. The use of fuzzy logic to represent general patterns within hotel trading gave the agent an insight into closing prices, allowing for an allocation-completion to be very effective. Its risk-averse nature allowed it to perform more consistently than the other finalists, with only a minimal sacrifice in scores.

The educational value of SouthamptonTAC has been demonstrated by the improvements made for SouthamptonTAC-02. The varying of behaviour based on the competitiveness of the environment gave it the flexibility to succeed in a larger range of situations.

Potential improvements could include advancing the learning capabilities of the agent, allowing it to tune many of the parameters that are currently chosen by the designers. A more aggressive entertainment strategy should also be adopted.

References

[1] P.Stone, A. Greenwald, "The First International Trading Agent Competition: Autonomous Bidding Agents," to appear *Electronic Commerce Research*, vol. 5, no. 1, 2005.

[2] M. Wellman et al., "The 2001 Trading Agent Competition," *proc. 14th Conf. on Innovative Applications of Artificial Intelligence*, pp. 935-941, Edmonton, 2002.

[3] M. He, N. Jennings, "SouthamptonTAC: Designing a Successful Trading Agent," *proc. 15th European Conference on Artificial Intelligence,* pp. 8-12, Lyon, France, 2002.

[4] K. Hoffman, "Set Covering, Packing and Partitioning Problems," http://iris.gmu.edu/~khoffman/papers/set_covering.html (current June 2004).

[5] M. Kantrowitz, E. Horstkotte, and Cliff Joslyn, "What is a fuzzy expert system", *Answers to Questions about Fuzzy Logic and Fuzzy Expert Systems,* http://www-2.cs.cmu.edu/Groups/AI/html/faqs/ai/fuzzy/part1/faq-doc-4.html (current June 2004).

[6] A. Greenwald, "The 2002 Trading Agent Competition: An Overview of Agent Strategies," *AI Magazine,* vol.24, no.1, 2002.

[7] M. He, N. Jennings, "Strategy for SouthamptonTAC-02," *Trading Agent Competition,* http://tac.eecs.umich.edu/research_reports/SouthamptonTAC-02.pdf (current June 2004).