An Analysis of the 2005 TAC SCM Finals

Jan Ulrich
Department of Computer Science
University of Texas at Austin, Austin TX 78712, USA
janulrich@mail.utexas.edu
Supervisor: Dr. Peter Stone

Abstract

Supply chains are the basis of the world economy. With a computer simulation of a supply chain management scenario, the Trading Agents Competition in Supply Chain Management (TAC SCM) promotes research in this field to try and automate the process. By letting computer agents optimize the manufacturing process, consumer goods become cheaper. Presented is an analysis of the 2005 Finals of the Trading Agent Competition in Supply Chain Management. Statistical methods identified the most important agent attributes to success. The different agents in the finals were then evaluated on those attributes.

Introduction

A supply chain is the process of producing a finished product from raw materials and components and then selling it to the final consumer. It is the interaction between all the different entities working in the production and sales process. Consequently supply chain management is the planning and implementation of the supply chains. In most supply chain management problems there are suppliers, manufacturers, and customers which all have to be coordinated to work in unison.

Supply chain management is being undertaken in every production industry ranging from automobiles to children's toys. Therefore, there is much interest in supply chain management research. In 2003, TAC SCM was created in addition to the original Trading Agents Competition which focused on the travel agent problem. This new competition, developed by Carnegie Mellon University and the Swedish Institute of Computer Science, realized the need for a medium to do supply chain research.

The game is a scaled down version of supply chain management in the real world. By limiting the number of resources and products that can be built, the game becomes feasible while still being based on reality. The core supply chain problems need to be dealt with while different issues such as new product design and depreciation are not considered. Market changes such as inflation and trade policies are not dealt with in this competition.

Shortages, market behavior, and monopolies (as seen in 2003) are all problems addressed in the competition. This competition then provides the incentive to solve these problems through new research ideas. Two years ago a team came up with a preemptive strategy to counter a self-destructive mutually unprofitable equilibrium (Wellman *et al.* 2005). These situations happen in all the time in what are called prisoner's dilemmas. In cases like this, market problems in the game are overcome.

Just like in the business world, there is limited information available during the game. The manufacturers interact only with the suppliers and customers and don't get direct information from each other. There is a market report and daily price report which give some information of how the market is doing at different points during the game.

In this paper I will first explain the Trading Agent Competition in Supply Chain Management game. Then I will identify different measurable attributes of agents that affect game performance. By analyzing these attributes I will show what is important to winning the game. Finally I will analyze TacTex-05's performance and set up guidelines to building a successful agent.

TAC SCM[1]

In the Trading Agent Competition in Supply Chain Management, six computer manufacturers who procure components, schedule factory production, and sell finished computers compete to maximize profits. The simulation consists of 220 days where manufacturers compete in the supplier market and customer market. The sale of computers takes place in a first-price sealed bid auction. Computer components are bought from profit-seeking suppliers that create competition among manufacturers. There are 10 different components consisting of cpus, motherboards, memory, and hard drives that are used to make 16 different computers. Each day the manufacturers interact with the suppliers by requesting supplies, accepting offers, or receiving deliveries. The manufacturers also interact with the customers by offering prices for the demanded computers and delivering finished computers.

The supplier capacity and the customer demand vary according to a mean reverting random walk. The computer sales price is determined by the auction, while the component costs are determined by both the manufacturer demand and the supplier capacity. Computers and components are

[1] http://www.sics.se/tac/tac05scmspec_v157.pdf

ordered with a given due date and with a daily penalty for late delivery. If a manufacturer wants a certain component, he sends a RFQ (request for quantity) to the supplier. The supplier then sends an offer back to the customer. This could be the earliest possible delivery or a partial quantity if the supplier cannot meet the original request. The manufacturer then either accepts or rejects the order and the order gets delivered by the due date if accepted.

A new aspect of the game that was added this year is player reputation. Each supplier keeps track of a reputation for each of the manufacturers. The reputation is the percentage of requests that were actually accepted. As long as this is in the reasonable range, the reputation is considered perfect. When placing a request, a reserve price is specified which is the highest amount you are willing to pay, so it is possible to avoid unwanted offers that would need to be rejected.

The competition consists of a qualifying round, then a seeding round, quarter-finals, semi-finals, and eventually the finals. These different rounds last several months with teams altering their agents as the competition progresses. This year the competition has 32 entrants from around the world. The teams play each other in different match-ups until the finals bracket where the participants play each other several times (8 for the quarter-finals and 16 for the semi-finals and finals).

The 2005 Finals

Position	Agent	Average Score
1	TacTex-05	4.741 M
2	SouthamptonSCM	1.604 M
3	Mertacor	546 272
4	Deep Maize	-220 503
5	MinneTAC	-311 844
6	Maxon	-1.985 M

This supply chain management game can be viewed as a set of actions taken each day depending on the daily variables. There is one optimal strategy for each agent given what the others did that can be found at the end of the game. The goal is to get as close as possible to this optimal strategy, which becomes very difficult due to the vast number of variables to consider.

I worked on developing the winning agent as it is discussed in the paper of (Pardoe & Stone 2006). We ended up decisively winning the 2005 Finals, but I wanted to find out why we won by such a big margin. I helped improve the agent during the tournament, but it was hard to understand the performance improvement of the agent, especially not knowing the changes of the other agents. I therefore decided to analyze this year's Finals. The Finals are the best part of the competition to analyze. By then the teams have found their best strategies and the best teams are playing each other, making the market more efficient (Jordan *et al.* 2006).

Attribute Analysis

After each game is played, a log file is generated which contains all the information of the game. Thus after the game is done, what was limited information before is now public information. Log parsers exist with graphical analysis that let people interpret the games. Using the log files from the Finals, I was able to extract data on several attributes. I then did a T-test to see the significance of each of them. There were sixteen games in the Finals, but most attributes are specific to each manufacturer, so for those there are 96 data points available.

There are two different kinds of attributes I analyzed in this paper: those that are particular to each game and those that are particular to each agent and each game. The attributes that were the same for all the manufacturers but different in each game were Customer Demand, Supplier Capacity, Average Supply Price, Game Number, Interest Rate, and Storage. The attributes that varied across manufacturers and games were Component Purchase Price, Computer Sales Price, Factory Utilization, Average Component Stock, and Lead Times.

Game Dependent Attributes

Customer Demand The customer demand fluctuates according to a mean reverting random walk. Some games are high demand games where the customer demand is generally above average, whereas others are low demand games. The sixteen different types of computers are divided into three categories: low-end, mid range, and high-end. The number of customer RFQs (request for quotation) in the low-end and high-end markets is determined by the same random variable. The mid range market has its own independent random variable. There are generally slightly more computer requests in the mid range. The demand in these categories is averaged for the whole game.

Supplier Capacity The supplier capacity also varies according to a mean reverting random walk. Each of the eight suppliers has a different production capacity for both of their products. In the beginning suppliers start with either high or low capacity (nominal capacity $\pm$ 35%) and gradually return to the nominal capacity. This capacity determines how many products are produced, which therefore determines a surplus or shortage. The supplier capacity is averaged over all the days and all 10 different supplier components.

Average Supply Price The average supply price is based on the day's orders, the production capacity, product commitments, and reputation. The day's orders are added to the previous commitments, which are then subtracted from the production capacity to get the available capacity. When planning future production, some of the available capacity is reserved for short term orders. This prevents a manufacturer from placing a huge order and thus causing a shortage. Agents with low reputations receive higher prices because their requests are taken into consideration later, after more commitments have been made. The average supply price is the average price of all different components over the whole game.

Game Number The game number is the order in which the game was played. During the tournament people could not give any sort of input to the computer agent, but the agent could modify itself all it wanted as long as it was automated. Since the semifinals and finals were played against the same opponents, adaptation could be useful. The finals were played on two different servers which meant there were two games being played at the same time. This attribute keeps track of the order of the game being played to see any effects of adaptation.

Interest Rate The bank will supply the agents with a unlimited loan, but the interest rate (which is set at the beginning of the game) determines how much interest is paid on the loan. The amount of interest paid on deposits in the bank is always half of the interest rate for loans.

Storage Cost There is a storage cost associated with any inventory of components or computers that the agents hold. The storage cost is different in every game but is somewhere between 25% and 50% of the base price of components, which is an arbitrary price given to each component as stated in the game specifications (Collins *et al.* 2004).

Manufacturer and Game Depended Attributes

Component Purchase Price and Computer Sales Price The prices of different components can be compared as a percentage of the respective base price. The component purchase price is the average percent of the base price paid for components. Similarly, the computer sales price is the average percent of the base price of a computer, which is the sum of the base price of its components.

To get the average component cost I decided to use a calculation that more accurately reflects the realities of the game than a simple average. Since computers have to be built from the four types of components: cpus, motherboards, memory, and hard drives, I decided to average the component cost of each of these types. Within each type I did a weighted average based on the amount of each component purchased. I did each of these calculations with the percentage of the base price.

The average sales were computed by calculating the weighted average based on the amount of computers sold. This was also done using the percentage of the base price of the computers components.

These calculations are more realistic to the game, yet they remain close to the straight forward, averaged values.

Factory Utilization Each day there are only 2000 assembly cycles in the manufacturer's factories. The manufacturer allocates these cycles to produce the different computers. However, the manufacturer needs to have the required components in order to build the computers. If the manufacturer runs out of even one of the required components, the computer cannot be built, and the factory sits idle, wasting cycles. The factory utilization is the percentage of the total cycles that are used. The results are averaged over the whole game. This is very closely related to the total number of computers produced. Some computers require a different number of cycles to built, but there is not much relative difference between factory utilization and the total number of computers produced.

Average Component Stock Each manufacturer keeps a component stock so they don't run out of components and the factory sits idle. There are several considerations when deciding how many components to keep on hand. There is a storage cost which varies from game to game on all inventory including components. Buying inventory now means you pay the current price, so if components get more expensive, you are better off, but if components get cheaper, you lose profits. The capacity of suppliers changes randomly, so a higher component stock helps avoid component shortages.

Lead Times This attribute measures the average number of days in advance of delivery components are ordered. This average is weighed by the number of components in each order.

Hypothesis

Hypothesis: TacTex-05's superior performance depended on the flexible short-term ordering, the predictive models of the component and computer prices, and the adaptation of initial component ordering and endgame sales in the Finals.

TacTex-05 was a very strong participant in the 2005 TAC SCM competition. Before doing this analysis, I was not convinced I knew the actual reasons for its proficient performance. However, I had seen the agent improve throughout the competition, so I had an idea of what was important.

This year, we created a whole new supply manager for our agent. The reason for this was that the change in the supplier rules, altered the entire game. In 2003 the best strategy was first day long-term orders. The rules were changed in 2004 to try and make it a better game, which worked to some extend, but there was still a "Zero Day" singularity (Eriksson, Finne, & Janson 2006). In 2005 reputation was added to the supplier model, and the supplier behavior was redone to try and remove this singularity.

Since the supplier side was changed much more than the customer side, we decided to concentrate on the supplier side. We switched to a short-term ordering strategy as opposed to the long-term orders we had used previously. Short-term ordering was important to winning in the 2005 competition, since it provided much more flexibility, as long as supplies were still available. With the new improvements to the game model, it turned out that a shortage in components as the game went on was much less likely than before, therefore TacTex-05's short-term ordering was most likely consequential in its success.

TacTex-05 used a predictive approach to play this game. Supplier's capacity was predicted so that the cheapest components could be bought. This predictive approach was used on the customer side as well (Pardoe & Stone 2005). For the customer side, the Offer Acceptance Predictor was used in planning the production and sale of computers. The Offer Acceptance Predictor was a function that returned the chances of selling a computer at the given price. This function was created by learning from previous games.

The last attribute that most likely contributed to TacTex-05's success was adaptation in the Finals. During the Semi-Finals, Deep Maize scored higher than TacTex-05, but during the Finals, TacTex-05 did better, even accounting for

the games where Deep Maize did not respond to the server. Since TacTex-05 used adaptation mainly in the Finals, I think this is the advantage that TacTex-05 had.
A simple solution to finding the most profit should be to maximize the following equation:

$$Max\{(SalesPrice - Cost) * \#ofComputers\}$$

That is, profit is just the sales price minus all the costs times the number of computers that can be sold. So in order to increase profit, an agent would have to increase their sales price, decrease their component costs and increase the number of computers to sell. It becomes more complicated because the question is how can all these things be achieved? One of my goals was to find the best ways to achieve all these objectives.
The following is the analysis I performed to test my hypothesis and to see which aspects of the game are important to winning.

Statistical Analysis

After collecting the data from the finals through analysis of the logfiles, I statistically analyzed the data to find out which attributes are significant. However some of the agents had network problems during the finals and did not respond for several days or they could also have taken too long computationally. I used a filter similar to the one described in (Wellman *et al.* 2006) to remove games where agents missed more than 7 days. The filter identified three games (3718 on tac3 and 4254, 4259 on tac4) that were bad, so I was left with 13 games to do my analysis.
After applying the filter, Deep Maize would move into second place. It was the agent that was most affected by applying the filter. In all the following figures the teams are ranked according to the 2005 results. However, the data being used is filtered, so Deep Maize should be in second place in the ordering.
I regressed the players scores vs the attributes listed in the section above. By using the actual score, I was able to see which attributes contributed to a high scoring game. These included ones that player could control and others that were not controlled by the players such as storage cost or interest rate. To see what was important to an agent's performance, I then calculated the relative score which was the difference of an agent's score to the average score of the game. By using the relative score, I was able to eliminate uncontrollable factors in an agent's performance.
I then performed a T-test on each of the attributes with the following null hypothesis: $H_0 : \beta_j = 0$. The T-Statistic was considered Significant at the 95% level of a two-tailed test if the absolute value was above 2.000 for 64 degrees of freedom. The results below are discussed in order of significance starting with the most significant.

Game Attributes	Actual Score (T-Stat)	Relative Score (T-Stat)	Sig?
Computer Sales Price	7.90	2.18	Y/Y
Component Purchase Price	-5.97	-3.42	Y/Y
Factory Utilization	4.17	5.73	Y/Y
Storage Cost	-3.61	-1.63	Y/N
High/Low Customer Demand	3.15	-2.55	Y/Y
Medium Customer Demand	3.12	-1.94	Y/N
Game Number	-2.60	-0.75	Y/N
Average Supply Price	-2.56	-0.51	Y/N
Average Component Stock	-1.37	-2.96	N/Y
Lead Times	-1.20	0.34	N/N
Interest Rate	1.09	0.19	N/N
Supplier Capacity	-0.76	0.98	N/N
Average Computer Stock	0.69	0.92	N/N

N = 78 $\bar{R}^2$(actual) = 0.86 $\bar{R}^2$(relative) = 0.66

Computer Sales Price

The price that the computers are sold at has a very high significance. The higher the price you can sell computers, the more profit you can make. It is therefore most important to place the most efficient bids possible. With less competition from competitors, a larger share of the higher priced customer requests can be fulfilled. The significance of the computer sales price went down when testing the relative score, which means this attribute is much more game dependent and higher scoring games have more opportunities for high sales.

Component Purchase Price

The cheaper components can be bought, the higher the manufacturer's score. Since all the agents are competing for the same components, those that buy components when they are cheapest will have the greatest advantage. This requires supplier price prediction or flexibility to adjust to changing prices. When comparing the effect on the actual scores and the relative scores, it can be seen that just as in the computer sales price, the component purchase price is important to both scores, but a high scoring game depends both on an agent's strategy and on the availability of lower component purchase prices.

Factory Utilization

Since factory utilization is the bottleneck in production due to the limited number of production cycles, a high utilization is a significant factor in success. The only time it would be a smart choice to stop production is when there is no possibility of profit since the components are more expensive than the market price of computers. As long as making computers is profitable, the more computers that are produced, the more profit is made. Factory utilization becomes much more important when comparing relative scores. Out of the most significant attributes, factory utilization is the one that can be totally controlled by the agent, so it becomes the most significant in determining relative scores.

Storage Cost

Storage cost takes away directly from profits. Therefore games with a higher storage cost will generally have lower profits. Since all manufacturers have the same storage cost, this attribute becomes less significant when regressing against relative scores.

High/Low Customer Demand

A higher customer demand will increase scores, since there will be more computer requests and therefore more chances to make profits. The High/Low demand has a greater impact than the Medium demand because there are 10 computers in the High/Low category, as opposed to 6 computers in the mid range.
The interesting result is that the sign is flipped for relative scores. When there is high demand, there is less competition between the different agents and thus the relative scores are more similar and thus smaller. When demand is lower, the competition between manufacturers increases and there is a much bigger difference in relative scores.

Medium Customer Demand

Similarly to the High/Low Customer Demand, the higher the demand for computers, the higher the profits of the manufacturers. The sign is also flipped for relative scores for the same reason listed above.

Game Number

As the game went on the scores generally got lower. With all other attributes held constant, this shows the effect of adaptation. Using actual scores, the significance was larger than with relative scores which meant some of the trend towards lower scores was due to game related factors. The average game scores can be seen in Figure 3. As can be seen from Figure 2, Deep Maize performed the best. Since the average trend is negative, the agents that did better took advantage of the non-adaptive agents. However the relative scores are not significant so an effect cannot be distinguished.
The way individual agents fared can be seen in Figure 1. The player's relative scores are plotted against the set of 7 pairs (8 pairs of games with 3 removed by the filter). A regression line is added to see the trend in each of their scores.
Figure 2 showed the average change in score for each of them. I did a regression on the relative game scores vs the game number, each of agents as a binary variable, and all the interaction variables. I then took the derivative with respect to the game number to get the adaptation effect. This left me with a constant and the interaction variables. I could now get each agent's specific derivative by filling in the binary variables accordingly. To find the average change in score throughout the finals, I multiplied the slope by 7, the number of game pairs - 1. I was not able to get the variances because I did not have the correlations between the interaction variables and the game number.

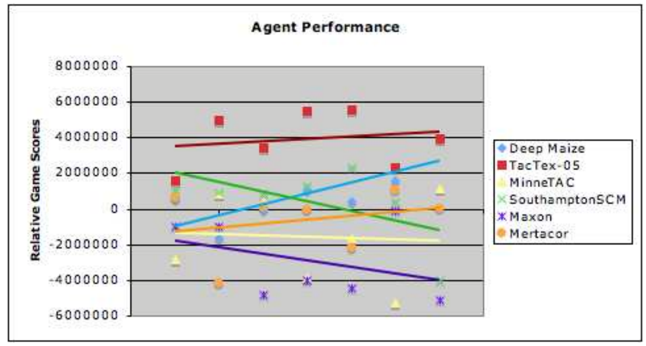

Figure 1: Relative Performance of the agents through the 7 valid game pairs. Relative performance is the difference between the average game score and the agent's score. Each pair of relative scores (one from each server) is again averaged.

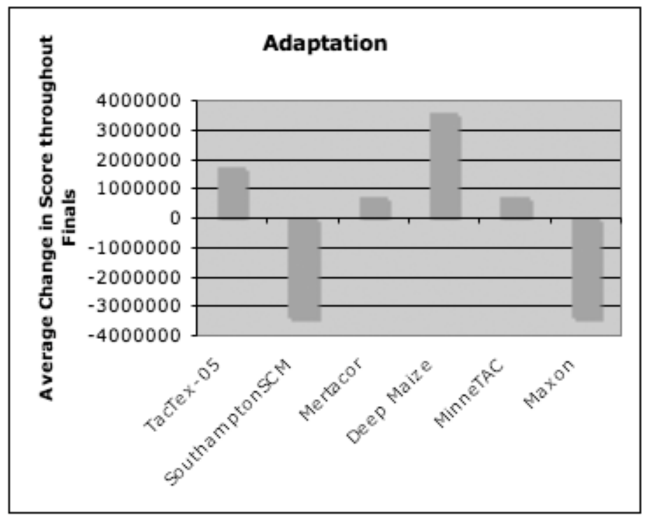

Figure 2: Calculated average change in score from the beginning of the finals to the end of the finals using filtered data.

$$Score = \beta_0 + \beta_1 GameNum + \beta_2 DM + \beta_2 TT + \beta_3 MT + \beta_4 SH + \beta_5 MA + \beta_6 GameNum * DM + \beta_7 GameNum * TT + \beta_8 GameNum * MT + \beta_9 GameNum * SH + \beta_{10} GameNum * MA$$

$$\frac{d}{dGameNum} Score = \beta_1 + \beta_6 DM + \beta_7 TT + \beta_8 MT + \beta_9 SH + \beta_{10} MA$$

$$\frac{d}{dGameNum} Score[TT] = \beta_1 + \beta 7$$

Note: DM = Deep Maize, TT = TacTex-05, MT = MinneTac, SH = SouthamptonSCM, MA = Maxon

Mertacor was the base value, thus all the binary variables are 0 for Mertacor.

As can be seen in Figure 2, there are big differences in the adaptations of different players. The two teams with the best

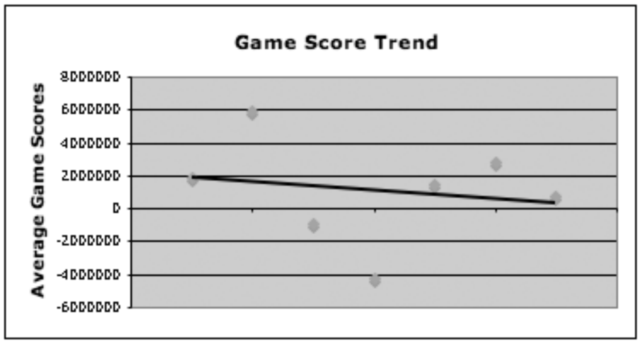

Figure 3: The average score trend as the tournament games progressed for the 7 valid game pairs.

adaptation were Deep Maize and TacTex-05. Deep Maize updated their prediction of customer market prices based on previous game information. TacTex-05 adapts mainly at the beginning and end of the game (Pardoe, Stone, & Van-Middlesworth 2006). According to feedback from the Deep Maize team, the inclusion of these market prices in the predictions should have had a minimal effect on Deep Maize's actions. This leads me to conclude that Deep Maize benefited more from TacTex-05's adaptation than TacTex-05. As the finals continued, TacTex-05 placed higher and higher orders on the first day. Due to the limited amount of RFQs per day, the game ordering was shifted to longer horizons, which benefitted Deep Maize how was very aggressive in early game ordering and production.

SouthamptonSCM and Maxon did not seem to adapt at all. In fact they did much worse as the competition continued. Mertacor and MinneTAC either adapted just enough for their performance not to be affected by the increased performance of TacTex-05 and Deep Maize, or they did not adapt at all and were unaffected by TacTex-05's and Deep Maize's changing strategies.

The agent's performance also has to be put into perspective with respect to their ranking. An agent such as TacTex-05 that is doing much better than the other agents might not be able to adapt was well since it is already performing much better. The adaptation curve is much steeper for an agent that is doing worse.

Average Supply Price

The lower the price of components, the higher the scores will be in the game. This result is to be expected, but because it is an important factor in determining the scores of the game, it has to be included in the statistical regression. Its significance also shows how important it is compared to the other attributes. The prices of components depend on many things including production capacity, manufacturer demand, and the point in the game. When regressing relative scores against this attribute, average supply price becomes insignificant because it is a variable that is solely dependent on the game conditions.

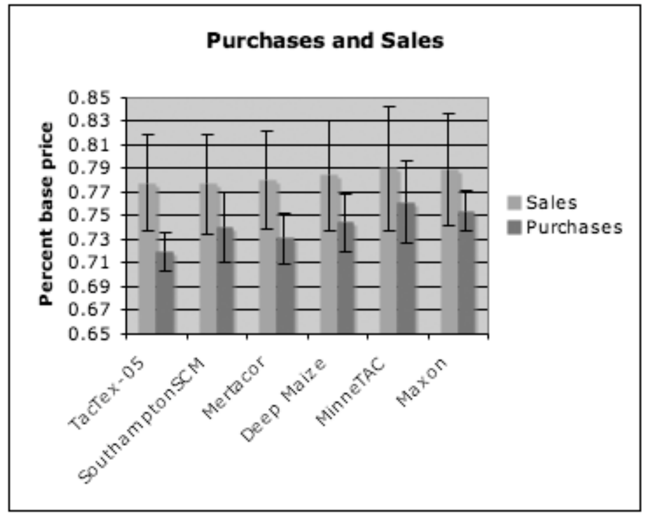

Figure 4: Average purchase price of components and average sales price of computers for each of the agents.

Average Component Stock

The smaller an agent's average component stock, the higher the score of the agent will be. This is because there a less storage costs to pay. As long as this stock is equally distributed across the different components, there won't be a drop in computer production. The hard part is keeping your component stock balanced. This attribute becomes much more important when looking at relative prices as it is completely controlled by an agent's strategy.

Insignificant Attributes

The rest of the attributes have a general trend that can be seen by the sign of the T-Statistic. However, the trend is not strong enough to distinguish it from no influence at all with 95% certainty.

The Different Agents

TacTex-05 was the TAC SCM champion in 2005. As stated in my hypothesis, this game is so complex that it is hard to figure out exactly what aspects of an agent are important in winning the competition. Was my hypothesis correct in that short-term component ordering, the offer acceptance predictor in planning and sales, and adaptation in the Finals, were the key attributes of TacTex-05 that led to winning the competition? In the last section I described the aspects of the game that were important. The following section compares the 2005 finalists by all the different agent-dependent attributes.

Component Costs and Computer Sales

The profit made by each sale is the price the computer was sold at minus the purchase price of the components. The standard error on the profit assumes that the average component purchase price and the average computer sales price are uncorrelated. As can be seen in Figures 4 and 5, TacTex-05 had the lowest purchase price and the highest profit. The

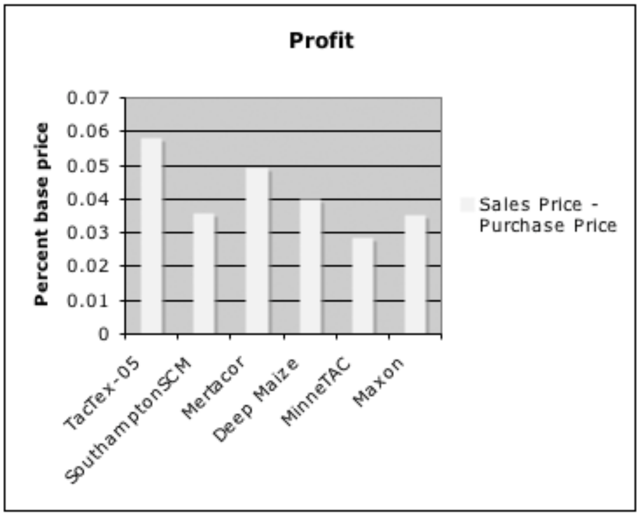

Figure 5: Profit per computer calculated by subtracting purchase cost from sales cost.

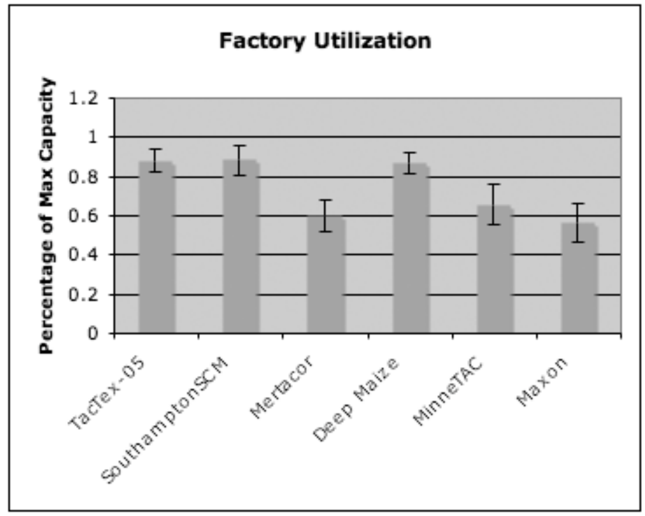

Figure 6: Average factory utilization throughout the game for each of the agents.

low purchase price was due to the price prediction that was used and the 3-day extremely short orders (Pardoe & Stone 2006). The agent also had the lowest sales price, but the sales price had higher variance across the different agents than the purchase price. The advantage in getting cheap components more than made up for the lower sales price.

Mertacor had a higher profit than both SouthamptonSCM and DeepMaize (which would have scored higher according to this filtered data), but scored lower than both of these competitors. The reason for this was that Mertacor produced fewer computers as can be seen in both the Factory Utilization and the Total Computers Produced. Even though the margins were higher, without enough volume, Mertacor could not make more profit than either of the other two agents.

Factory Utilization

The relevancy of factory utilization can be seen by the fact that the downward sloping trend of the percentage of maximum capacity matches the final rankings. Players with higher factory utilization were able to produce more computers and therefore make more profits. Higher factory utilization also meant less shortages where there were not enough components to produce. Factory utilization is a good measure of how well the agent was able to manage its inventory. A slowdown in production is almost always due to a shortage in components. The only time this is not the case is when there is no possibility to make a profit. This was never the case in the 2005 Finals.

No one can have 100% factory utilization because the agents have no inventory at the start of the game, so they cannot start producing. 97% is the highest possible factory utilization. However, none of the agents were anywhere near this maximum since it would not be wise to produce that much. The minimum number of days to wait for production to start is four, because that is the soonest any components can arrive at the factories. The sooner a manufacturer can get components, the higher the average factory utilization will be. Manufacturers also tend to slow down production near the end of the game, as they do not want to be left with unsold excess inventory. The importance of factory utilization can be seen by the results of the tournament. Mertacor had the second highest profit margin per computer sold, but they were outscored by both SouthamptonSCM and Deep Maize (after the filter), because both of these agents had higher factory utilization and thus volume.

Computer and Component Inventory

As can be seen in the Figures 7 and 8, the teams vary greatly in the amount of inventory they keep on hand. There does not seem to be much significance in the amount of inventory held. It depends on the agent's strategy whether to keep a high inventory of components to produce computers or try and lower storage costs through made-to-order computers. A short-term order strategy would favor a lower component inventory while a long-term order strategy would require periods of high inventory. The question of whether manufacturers keep a high computer inventory or not depends on whether the agent thinks computers will get more expensive and thus more profitable in the future. Low inventory had its advantages in the 2005 tournament, as long as factory utilization was kept high. Deep Maize outperformed SouthamptonSCM (after the filter) because they had a lower average component inventory, and thus storage cost, and slightly more profit per computer.

Component Lead Times

Figure 9 shows the strategy each of the agents used to procure components. Agents with shorter lead times ordered components short-term as opposed to long-term ordering. In the 2004 tournament long-term orders were the preferred strategy. The supplier behavior led the agents to order all there components on the first day and a degradation of the

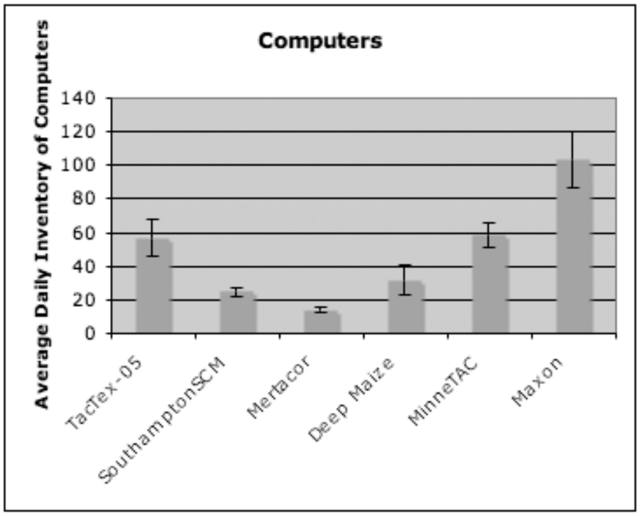

Figure 7: Average number of computers in inventory for each computer type.

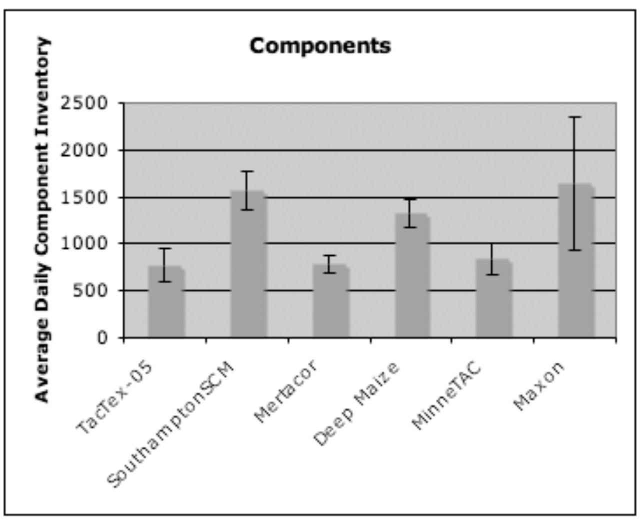

Figure 8: Average number of components in inventory per component type.

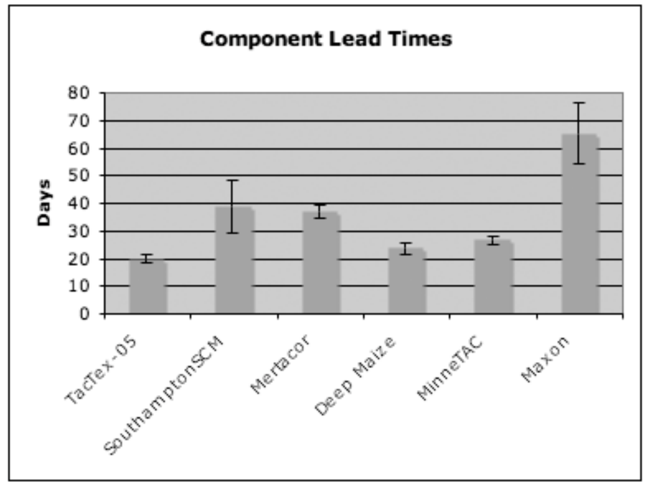

Figure 9: The average number of days between the placement of a component order and delivery.

game (Wellman *et al.* 2005). This year it can be seen that the strategies were much more mixed and probably short-term if they were to be categorized.

TacTex-05 Analysis

I was correct in the fact that TacTex-05's supply-side manager led to its success in this year's competition. The component purchase price was the second most significant attribute in determining an agent's actual and relative score. TacTex-05 had the lowest purchase price which is the lowest average buying price as a percentage of component's base prices. This was due to the fact that TacTex-05 switched to a short-term purchasing system. It gave TacTex-05 more flexibility to buy components on the cheapest days. The cheapest purchasing days were found through supplier capacity prediction. As can be seen in Figures 10 and 11, these short-term orders proved to be very effective for TacTex-05. Figure 10 shows the price of component orders vs the difference in order date and delivery date for all the manufacturers in the game. As can be seen in Figure 10, prices varied a lot and there was generally no price difference between short-term orders and long-term orders. What is not included in this graph is component cost. Generally, long term orders need to be made with higher quantities, because it is impossible to quickly react to inventory changes. Therefore long-term orders generally instill higher storage costs as the high number of components purchased are built into computers. Therefore even though the slope of the regression line is almost 0, the cost of storage would make it positive, thus making short-term purchases cheaper.

Figure 11 shows the specific purchases of TacTex-05. The agent focused on short-term orders, and was able to receive these as cheap or cheaper than any other agent. Since these short-term orders allowed more flexibility and therefore had a lower storage cost, this was important to winning the 2005 Finals.

The second part of my prediction was partly correct.

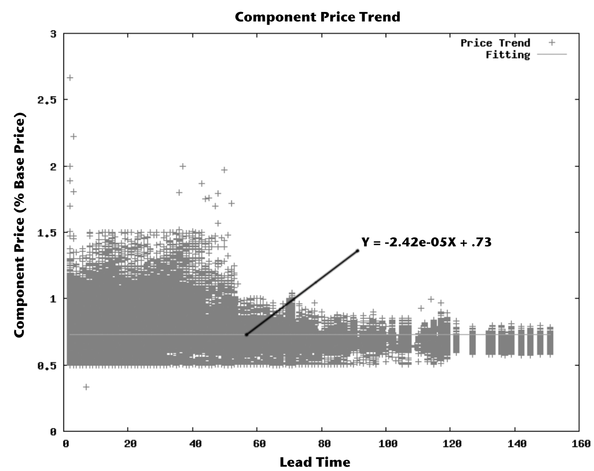

Figure 10: The cost of component orders based on the time in days between order and delivery for all agents.

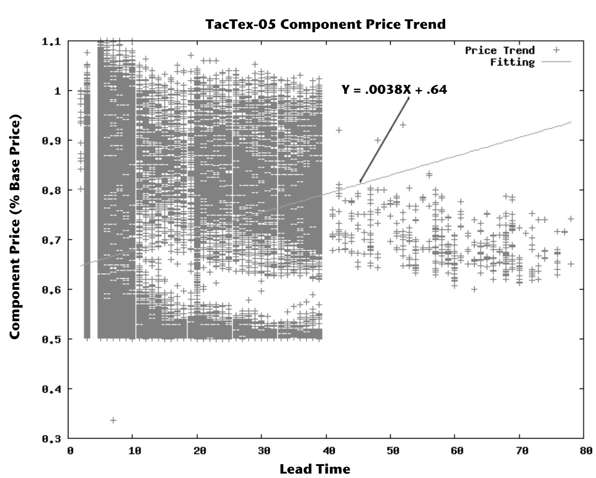

Figure 11: The cost of component orders based on the time in days between order and delivery for all TacTex-05.

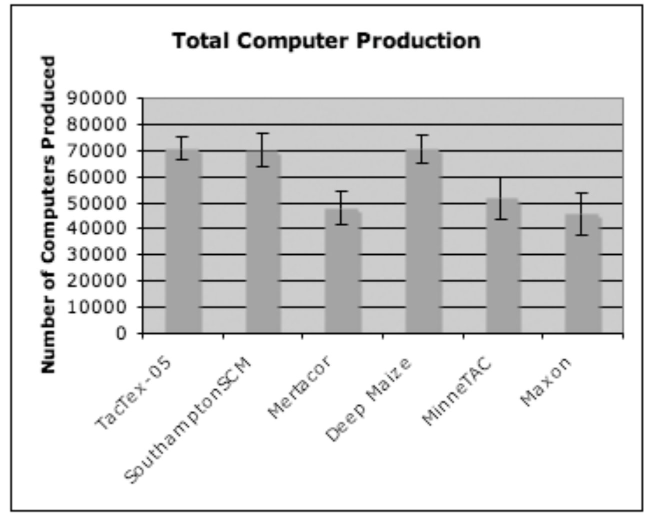

Figure 12: The total number of computers produced by each agent throughout the entire game.

The predictions in the Supply Manager did help in procuring TacTex-05 the lowest component prices. However the Offer Acceptance Predictor was not as effective as I had thought. TacTex-05 had the lowest average computer sales price, I was wrong about my hypothesis that the customer side predictions gave the winning advantage. The close range in customer sales prices between the different competitors shows just how competitive the finals were. The different agents all had excellent customer sales mechanisms. In another analysis of the 2005 finals, (Jordan *et al.* 2006) found that all the different agents had a 97.1% bid efficiency, which was the ratio of actual profit made per bid and the possible profit that could have been made if a particular agent had perfect information about his opponent's bids. With such an incredibly high percentage, all the cutomer demand managers were good.

Even with the lowest computer sales average, TacTex-05 still had the highest profit per computer sold. The basic profit equation is maximized if with the highest profit per computer and the highest number of computers sold. It can be seen from Figure 12 below that TacTex-05, SouthamptonSCM, and Deep Maize all produced the highest total number of computers. TacTex-05 therefore had the highest profit per computer and the highest total number of computers produced and therefore won the competition.

The agent that had the biggest improvement as the finals progressed was Deep Maize. However, after seven games TacTex-05 still scored higher than Deep Maize. TacTex-05 had the most documented adaptation that I am aware of as described in the paper by (Pardoe, Stone, & VanMiddlesworth 2006), which is why it ranked second according to adaptation. TacTex-05's adaptation decreased SouthamptonSCM's (the second place finisher) score significantly, while Deep Maize benefited the most as the finals progressed. Even though Deep Maize beat TacTex-05 in the Semi-Finals, the Finals were more competitive and TacTex-

05 won due to the factors mentioned above. The adaptation had trade offs for TacTex-05: some of its opponents did better and some of them did worse. Overall TacTex-05 improved its score, so the adaptation was a good idea.

Conclusion

When comparing the relative performance of the different agents, the most significant attribute is factory utilization. Since the number of factory cycles is the only constraint in production, this variable needs to be maximized. While maximizing factory utilization, other attributes such as component and computer prices and average inventory need to be considered as well. Therefore when considering a concept for a new agent, the first step would be to make sure that enough components are always available to keep the factory working. A good way to do this is to keep an inventory threshold. Short-term orders are a very effective way to keep the inventory low and above the threshold, while at the same time getting cheap component costs.

Opponent modeling and agent adaptation are also important aspects in the game. In the 2005 Finals, the agents that adapted over time did much better than the agents that did not. This will become even more important in the 2006 tournament where the opponents names will be available during the game.

The best computer sales strategy would be one used by MinneTAC, the agent with the highest average sales price (Ketter *et al.* 2004). This strategy differentiates between high demand games and low demand games. In high demand games an optimistic approach is best where prices are determined by maximizing profit. Expected profit is calculated based on the current game variables such as current customer demand, inventory level, production rate, and the time left in the game. Offer acceptance probability is calculated from past data collected over all the games played. During low demand games, a more practical approach is used where the objective is to minimize leftover or unprofitable computer production. The market is tracked with a distribution that will tell you the chances of individual computer type offers.

Future Work

This work was done on the 2005 Finals of the Trading Agents Competition in Supply Chain Management. These competitions happen every year and this kind of analysis can be done on each year's games. Such continued analysis would show the trends in the improvement of the agents. The competitiveness of the agents increases every year as more time is spent working on them. I only analyzed the finals of this years competition. Since there were 32 teams participating, the performance of the different teams throughout the entire competition could be studied. There is almost an endless amount of analysis that can be done on such a complex game with so many game logs available.

More analysis on the agent's adaptation could be done with more data available. Deep Maize had the best adaptation, but would the agent's performance increases level off as its performance came close to TacTex-05's performance? Short-term ordering was shown to be the best strategy during this year's finals, but if more players play an aggressive short-term strategy would the supplier pricing trend stay indifferent between short-term and long-term orders, or would it become advantageous to make long-term orders?

References

Collins, J.; Arunachalam, R.; Sadeh, N.; Eriksson, J.; Finne, N.; and Janson, S. 2004. The supply chain management game for the 2005 trading agent competition.

Eriksson, J.; Finne, N.; and Janson, S. 2006. Evolution of a supply chain management game for the trading agent competition. *AI Communications*.

Jordan, P. R.; Kiekintveld, C.; Miller, J.; and Wellman, M. P. 2006. Market efficiency, sales competition, and the bullwhip effect in the tac scm tournaments. *AAMAS-06 Workshop on Trading Agent Design and Analysis(TADA/AMEC)*.

Ketter, W.; Kryzhnyaya, E.; Damer, S.; McMillen, C.; Agovic, A.; Collins, J.; and Gini, M. 2004. Minnetac sales strategies for supply chain tac. *Third Int'l Conf. on Autonomous Agents and Multi-Agent Systems*.

Pardoe, D., and Stone, P. 2005. Bidding for customer orders in TAC SCM. In Faratin, P., and Rodriguez-Aguilar, J., eds., *Agent Mediated Electronic Commerce VI: Theories for and Engineering of Distributed Mechanisms and Systems (AMEC 2004)*, volume 3435 of *Lecture Notes in Artificial Intelligence*. Berlin: Springer Verlag. 143–157.

Pardoe, D., and Stone, P. 2006. Predictive planning for supply chain management. In *Proceedings of the International Conference on Automated Planning and Scheduling*.

Pardoe, D.; Stone, P.; and VanMiddlesworth, M. 2006. Tactex-05: An adaptive agent for tac scm. *AAMAS-06 Workshop on Trading Agent Design and Analysis(TADA/AMEC)*.

Wellman, M. P.; Estelle, J.; Sing, S.; Vorobeychick, Y.; Kiekintveld, C.; and Soni, V. 2005. Strategic interactions in a supply chain game. *Computational Intelligence*.

Wellman, M. P.; Jordan, P. R.; Kiekintveld, C.; Miller, J.; and Reeves, D. M. 2006. Empirical game-theoretic analysis of the tac market games. *AAMAS-06 Workshop on Game-Theoretic and Decision-Theoretic Agents*.