Predictive Planning for Supply Chain Management: Adapting to Competitor Behavior*

David Pardoe
Department of Computer Sciences
The University of Texas at Austin
dpardoe@cs.utexas.edu

Introduction

In today's industrial world, supply chains are ubiquitous in the manufacturing of many complex products. Traditionally, supply chains have been created through the interactions of human representatives of the various companies involved. However, recent advances in planning, scheduling, and autonomous agent technologies have sparked an interest, both in academia and in industry, in automating the process (Kumar 2001).

From a planning and scheduling perspective, supply chain management simultaneously requires long-range inventory management, mid-range customer negotiations, and short-term factory scheduling, all of which interact closely.

One barrier to supply chain management research is that it can be difficult to benchmark automated strategies in a live business environment, both due to the proprietary nature of the systems and due to the high cost of errors. The Trading Agent Competition Supply Chain Management (TAC SCM) scenario provides a unique testbed for studying and prototyping supply chain management agents by providing a competitive environment in which independently created agents can be tested against each other over the course of many simulations in an open academic setting (Arunachalam & Sadeh 2005). In a TAC SCM game, each agent acts as an independent computer manufacturer in a simulated economy. The agent must procure components such as CPUs and memory; decide what types of computers to manufacture from these components as constrained by its factory resources; bid for sales contracts with customers; and decide which computers to deliver to whom and by when.

One crucial challenge in supply chain management is that decisions must often be made in the face of considerable uncertainty. For instance, purchases of production resources may need to be negotiated long before accurate information about customer preferences becomes available. This challenge is particularly evident in TAC SCM, where sources of uncertainty include the capacity of suppliers to deliver components, the nature of customer demand, and the actions of other agents as they compete for components and customers.

I have designed an agent to compete in TAC SCM, TacTex-05 (winner of the 2005 competition), that addresses this uncertainty by taking a predictive approach to its many planning and scheduling decisions. In particular, TacTex-05 makes predictions concerning the types and quantities of computers that will be requested by customers, the capacities of component suppliers and the prices they are likely to offer, and the probability that an offer to a customer will be accepted at a particular price. Planning and scheduling takes place using these predictions.

In this abstract, I will first provide details on the TAC SCM scenario and give an overview of the design of TacTex-05. Then I will describe my current work, which focuses on learning to adapt to the behavior of competing agents.

The TAC Supply Chain Management Scenario

In this section, I provide a brief summary of the TAC SCM scenario. Full details are available in the official specification document (Collins *et al.* 2005).

In a TAC SCM game, six agents act as computer manufacturers in a simulated economy managed by a game server. The length of a game is 220 simulated days, with each day lasting 15 seconds of real time. The game can be divided into three parts: i) component procurement, ii) computer sales, and iii) production and delivery, as expanded on in the remainder of this section and illustrated in Figure 1.

Component Procurement

The computers are made from four components: CPUs, motherboards, memory, and hard drives, each of which come in multiple varieties. From these components, 16 different computer configurations can be made. Agents must purchase these components from a set of suppliers managed by the game server.

Agents wanting to purchase components send requests for quotes (RFQs) to suppliers indicating the type and quan-

*This abstract is largely based on a paper in the ICAPS 2006 techical program (Pardoe & Stone 2006) that contains both a complete description of the TacTex-05 agent and a number of experimental results. I would like to thank my advisor, Peter Stone, for assistance with this work.

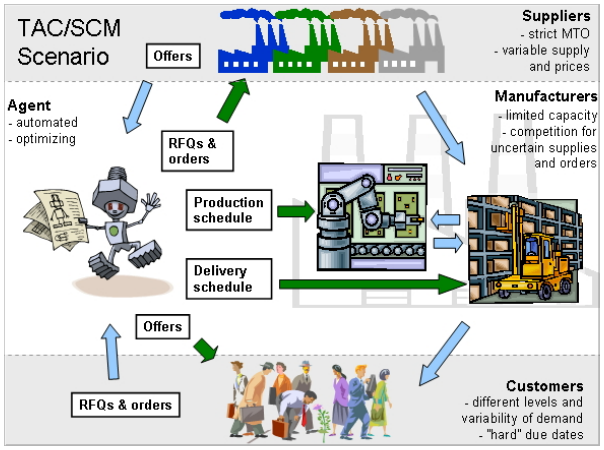

Figure 1: The TAC SCM Scenario (figure taken from (Collins *et al.* 2005)).

tity of components desired, the date on which they should be delivered, and a reserve price stating the maximum amount the agent is willing to pay. Suppliers respond to RFQs the next day by offering a price for the requested components if the request can be satisfied. Agents may then accept or reject the offers.

Suppliers have a limited capacity for producing components; this capacity varies throughout the game according to a random walk. The price offered in response to an RFQ depends on the fraction of the supplier's capacity that is free before the requested due date.

Computer Sales

Customers wishing to buy computers send the agents RFQs consisting of the type and quantity of computer desired, the due date, a reserve price indicating the maximum amount the customer is willing to pay per computer, and a penalty that must be paid for each day the delivery is late. Agents respond to the RFQs by bidding in a first-price auction: the agent offering the lowest price on each RFQ wins the order. The number of RFQs sent by customers each day depends on the level of customer demand, which fluctuates throughout the game.

Production and Delivery

Each agent manages a factory where computers are assembled. Factory operation is constrained by both the components in inventory and assembly cycles. Each day an agent must send a production schedule and a delivery schedule to the server indicating its actions for the next day. The production schedule specifies how many of each computer will be assembled by the factory, while the delivery schedule indicates which customer orders will be filled from the completed computers in inventory. Agents are required to pay a small daily storage fee for all components in inventory at the factory.

Overview of TacTex-05

Given the detail and complexity of the TAC SCM scenario, creating an effective agent requires the development of tightly coupled modules for interacting with suppliers, customers, and the factory. TacTex-05 is a fully implemented agent that operates within the TAC SCM scenario. In this section, I present a high-level overview of the agent.

Agent Components

Figure 2 illustrates the basic components of TacTex-05 and their interaction. There are five basic tasks a TAC SCM agent must perform:

1. Sending RFQs to suppliers to request components
2. Deciding which offers from suppliers to accept
3. Bidding on RFQs from customers requesting computers
4. Sending the daily production schedule to the factory
5. Delivering completed computers

The first two tasks are assigned to a *Supply Manager* module, and the last three to a *Demand Manager* module. The Supply Manager handles all planning related to component inventories and purchases, and requires no information about computer production except for a projection of future component use, which is provided by the Demand Manager. The Demand Manager, in turn, handles all planning related to computer sales and production. The only information about components required by the Demand Manager is a projection of the current inventory and future component deliveries, along with an estimated replacement cost for each component used. This information is provided by the Supply Manager.

The tasks to be performed by these two managers can be viewed as optimization tasks: the Supply Manager tries to minimize the cost of obtaining the components required by the Demand Manager, while the Demand Manager seeks to maximize the profits from computer sales subject to the information provided by the Supply Manager. In order to perform these tasks, the two managers need to be able to make predictions about the results of their actions and the future of the economy. TacTex-05 uses three predictive models to assist the managers with these predictions: a predictive *Supplier Model*, a predictive *Demand Model*, and an *Offer Acceptance Predictor*.

The Supplier Model keeps track of all information available about each supplier, such as TacTex-05 's outstanding orders and the prices that have been offered in response to RFQs. Using this information, the Supplier Model can assist the Supply Manager by making predictions concerning future component availability and prices.

The Demand Model tracks the customer demand in each of the three market segments, and tries to estimate the underlying demand parameters in each segment. With these estimates, it is possible to predict the number of RFQs that will be received on any future day. The Demand Manager can then use these predictions to plan for future production.

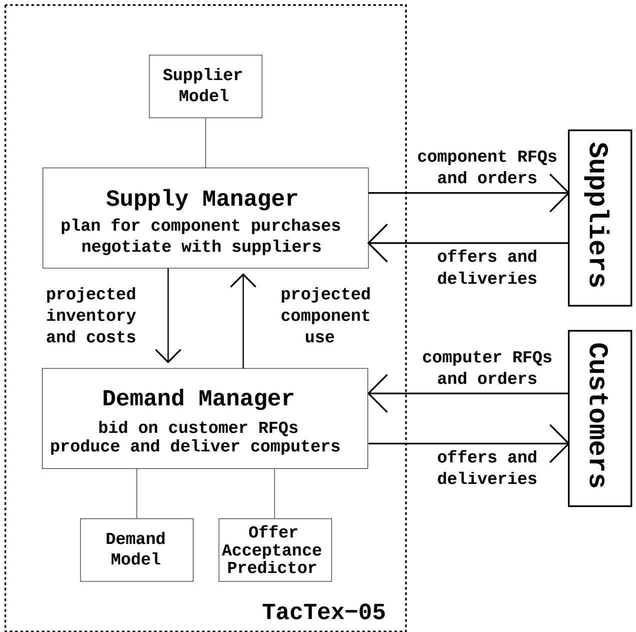

Figure 2: An overview of the main agent components

When deciding what bids to make in response to customer RFQs, the Demand Manager needs to be able to estimate the probability of a particular bid being accepted (which depends on the bidding behavior of the other agents). This prediction is handled by the Offer Acceptance Predictor. Based on past bidding results, the Offer Acceptance Predictor produces a function for each RFQ that maps bid prices to the predicted probability of winning the order.

Current Work: Adapting to Competing Agents

The TAC SCM competition consists of a series of rounds. During each round an agent faces the same five opponents in a number of games. When analyzing competition results, it quickly becomes apparent that the nature of the economy within a game depends heavily on the agents participating. An agent that consistently achieves a high profit against one set of opponents may lose a large amount of money against a different set of opponents in a different round. This fact suggests the potential value of designing an agent that can adapt to the behavior of whatever opponents it happens to be facing during a particular round. Enabling TacTex-05 to adapt in such a fashion is the primary focus of my current work. (The general development of such adaptive agents in agent-based economies will be the focus of my thesis; in addition to studying the TAC SCM domain, I have explored auction domains in which seller agents adapt the parameters of auction mechanisms in response to the observed behavior of bidding agents (Pardoe *et al.* 2005).)

The primary means by which TacTex-05 can be made more adaptive is through improvements to the predictive modules described previously. In particular, I would like to improve long-term predictions of computer prices and component prices, both of which can vary considerably based on opponent behavior. Currently, the predictions made by the predictive modules are based primarily on observations from the current game. Another source of information that could be useful in making predictions is the events of past games, made available in log files kept by the game server.

The potential benefit from basing predictions on the results of these past games is illustrated by the one form of adaptation used by TacTex-05 during the 2005 TAC SCM competition. At the beginning of each game, many agents place relatively large component orders (when compared to the rest of the game) to ensure that they will be able to produce computers during the early part of the game. Prices for some components may also be lower on the first day than they will be afterwards, depending on the due date requested. Determining the optimal initial orders to place is difficult, because no information is made available on the first day of the game. As a result, many agents use the same hard-coded initial orders in each game. TacTex-05 takes advantage of this fact by basing its predictions of early-game component prices on the prices observed in past games. An analysis of the final round of competition (Pardoe, Stone, & VanMiddlesworth 2006) showed that first-day prices were unusually attractive due to the purchasing patterns of the agents. As a result of its adaptivity, TacTex-05 recognized this opportunity and purchased significantly more components on the first day of each game than its competitors. The savings on component costs accounted for much of TacTex-05's winning margin. Although this example illustrates the value of adaptation, it is admittedly ad hoc. One goal of my current work is to identify additional opportunities for adaptation automatically, through techniques that will generalize to other domains.

One possible approach is the use of machine learning techniques to develop more accurate predictive models. In fact, I explored this possibility in past work (Pardoe & Stone 2004), finding that learned predictors could indeed improve agent performance. There is one primary drawback to this approach, however: it requires that an agent be able to draw training data from a large number of games against the same opponents. A single round of competition consists of a relatively small number of games, at most 16, raising the question of how a machine learning approach could successfully be applied. In particular, during the first game of a round, there would be no data from which the agent could learn.

Thus, I am currently exploring means by which TacTex-05 can begin a round of competition with fairly general predictive modules, and then revise them based on data as it becomes available, rather than starting tabula rasa. One valuable resource in my work is the TAC Agent Repository,[1] a collection of agents made available by competition participants for research purposes. By simulating rounds of competition with various combinations of these agents,

[1] `http://www.sics.se/tac/showagents.php`

along with variations of TacTex-05 designed to exhibit particular behaviors, I can observe a wide range of different economies. Using these simulations I hope to answer the following questions:

- What properties remain the same from one set of opponents to another? (e.g., component prices tend to decrease over the course of a game)
- What properties are highly dependent on the set of opponents? (e.g., how quickly computer prices rise when customer demand increases)
- What fixed predictive models result in the best performance across a wide range of opponent sets?
- As additional data becomes available, how can these predictive models successfully be revised?

Answering the last question presents the largest challenge from a learning perspective. Approaches I am currently investigating include the use of online learning methods for combining expert advice (where each expert represents a predictive model learned for a particular opponent set) and metalearning methods (in which the performance of a learning system is improved through experience with a family of related tasks – in this case various opponent sets).

References

Arunachalam, R., and Sadeh, N. 2005. The supply chain trading agent competition. *Electronic Commerce Research and Applications* 4:63–81.

Collins, J.; Arunachalam, R.; Sadeh, N.; Eriksson, J.; Finne, N.; and Janson, S. 2005. The supply chain management game for the 2006 trading agent competition. Technical report. Available from http://www.sics.se/tac/tac06scmspec_v16.pdf.

Kumar, K. 2001. Technology for supporting supply-chain management. *Communications of the ACM* 44(6):58–61.

Pardoe, D., and Stone, P. 2004. Bidding for customer orders in TAC SCM. In *AAMAS 2004 Workshop on Agent Mediated Electronic Commerce VI: Theories for and Engineering of Distributed Mechanisms and Systems*.

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

Pardoe, D.; Stone, P.; Saar-Tsechansky, M.; and Tomak, K. 2005. Adaptive auctions: Learning to adjust to bidders. In *The Fifteenth Annual Workshop on Information Technologies and Systems*.

Pardoe, D.; Stone, P.; and VanMiddlesworth, M. 2006. TacTex-05: An adaptive agent for TAC SCM. In *AAMAS 2006 Workshop on Trading Agent Design and Analysis / Agent Mediated Electronic Commerce*.