The 2007 Procurement Challenge: A Competition to Evaluate Mixed Procurement Strategies

Alberto Sardinha, Michael Benisch, Norman Sadeh, Ramprasad Ravichandran, Vedran Podobnik[a], Mihai Stan[b]

August 2008
CMU-ISRI-07-123R[c]

[a]University of Zagreb, Croatia
[b]University of "Politehnica" of Bucharest, Romania
[c]This report supercedes CMU-ISRI-07-123

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

The platform for the procurement challenge is based on the baseline Supply Chain Trading Agent Competition game originally designed by Carnegie Mellon's e-Supply Chain Management Laboratory in collaboration with the Swedish Institute of Computer Science (SICS). The game was later refined with the help of the University of Minnesota. The first four authors, who are with the e-Supply Chain Management Laboratory, would like to thank SICS for sharing its code and allowing them to customize it for the Procurement Challenge described in this paper. This work has been supported by the National Science Foundation under ITR Grant 0205435 as well as by a grant from SAP Research.

Keywords: Autonomous Agent, Electronic Commerce, Trading Agent

Abstract

Global competition is putting a premium on the ability to manage risk through flexible and agile Web-enabled procurement practices. This article discusses the design of the 2007 "Supply Chain Management - Procurement Challenge" (SCM-PC), a competition designed by the first three authors to evaluate the performance of mixed procurement strategies that balance risk through combinations of long-term, quantity-flexible contracts and one-off contracts. Specifically, the SCM-PC Challenge revolves around a PC assembly scenario, where Web-enabled trading agents developed by different teams compete for components required to assemble different types of PCs. Collectively the authors represent the top three entries in the 2007 Procurement Challenge. They present the strategies their teams developed for the competition, compare their performances, and discuss lessons learned from the competition.

1 Introduction

The Web is enabling manufacturing enterprises to explore more flexible and agile procurement practices. As product life cycles become shorter and demand becomes more difficult to predict, both manufacturers and suppliers are looking for new ways of sharing risk. Manufacturers are trying to move away from static, long-term contracts that would require them to take unacceptable levels of risk, while trying to secure price and availability guarantees from their suppliers. Simultaneously, suppliers are seeking demand and payment guarantees from their customers. This in turn translates into long-term procurement contracts, where suppliers and manufacturers often agree on some levels of flexibility (e.g. quantity, price, or service levels). These contracts are often supplemented with one-off procurement arrangements to accommodate surges in demand or address significant price disruptions.

1.1 The Supply Chain Trading Agent Competition and the Procurement Challenge

The International Supply Chain Trading Agent Competition ("TAC-SCM") was established to stimulate the development and evaluation of novel Web-enabled supply chain trading strategies. The competition, which was introduced in 2002 [2], has been held as an annual event since the summer of 2003, attracting over 120 entries from 60 teams in 21 countries in its first five years. It revolves around a personal computer (PC) assembly scenario, where agents developed by different teams compete against one another for customer orders for different types of PCs and for the different components required to assemble these PCs. Each game brings together six agents developed by six different teams and simulates 220 days of operation, with each day being simulated in 15 seconds. The tournament requires agents to compete in hundreds of games, each simulating different market conditions and bringing together different combinations of competitors.

In 2006, encouraged by the success of the Supply Chain Trading Competition, the first three authors embarked on the design of a variation of the TAC-SCM scenario that focuses solely on procurement decisions. The motivations for the design of this new game were several:

- **Evaluate Procurement Decisions in Isolation**: By requiring supply chain trading agents to simultaneously compete in component and finished-goods markets, the original TAC-SCM competition (now referred to as the "TAC-SCM baseline game") requires agents to continuously manage a complex set of decisions (i.e. customer bidding, procurement and coordination between both sets of decisions), making it sometimes difficult to determine why one agent performs better than another. By designing a variation of the game that requires agents to focus solely on the procurement side of the problem, the authors wanted to make it easier to evaluate the performance of competing procurement strategies. Insight gained from such a game is expected to also be easier to transfer to industry, given the typical decoupling between procurement and marketing decisions imposed by today's Enterprise Resource Planning (ERP) architectures.

- **Introduce Quantity-Flexible Procurement Contracts**: The TAC-SCM baseline game uses a uniform mechanism for agents to negotiate procurement contracts with suppliers. This mechanism does not explicitly distinguish between long-term and one-off procurement contracts, though it has been shown that successful TAC-SCM agents modulate the horizons and quantities of their procurement contracts to effectively implement hybrid strategies that combine both long-term and short-term contracting practices [1]. By explicitly introducing quantity-flexible long-term contracts, the new procurement challenge (SCM-PC) requires agents to more explicitly decide how to balance one-off and long-term procurement, making it easier to analyze and compare the behavior of different agents.

- **Lower the Barrier to Entry**: By introducing a variation of the TAC-SCM game that requires teams to only worry about procurement, it was felt that the Procurement Challenge would lower the barrier to entry for prospective competitors and possibly entice new teams to join the TAC-SCM community.

The specifications of the "Supply Chain Management - Procurement Challenge" (SCM-PC), as it is now known, were published in late 2006 and the web server and agentware required for the competition were released in early 2007. The remainder of this article provides a detailed description of the procurement challenge (Section 3), including the motivations behind key design decisions, descriptions of the top three agents in the first edition of the tournament held at AAAI-2007 in Vancouver in July 2007 (Section 4), and a detailed account and analysis of the results of the tournament (Section 5). A summary of lessons learned and concluding remarks are provided in Section 6.

2 Related Work

The Procurement Challenge builds on three complementary lines of research:

1. Sourcing and Procurement
2. Autonomous Bidding
3. Supply Chain Management Games

This is further discussed below.

2.1 Sourcing and Procurement

An important line of research in supply chain management has been concerned with the trade offs associated with different sourcing and procurement strategies. For instance, Pyke and Johnson have argued that different types of sourcing strategies are better suited for different situations, with critical, high-value added components better handled through strategic partnerships and commoditized components available from multiple sources more effectively handled through dynamic e-procurement [27]. Peleg et al. compared such pure strategies with a mixed strategy combining both short-term and long-term elements, showing that the superiority of one strategy over the others depends on contract terms [23]. Bensaou helped debunk the myth that Japanese car manufacturers rely solely on long-term strategic partnerships with suppliers, instead advocating the management of portfolios of buyer-supplier relationships covering a wide spectrum of possible arrangements [6]. A review of models for constructing short-term and long-term contracts in business-to-business markets has been conducted by Kleindorfer and Wu [16]. Elmaghraby also provides an excellent review of trade offs between different sourcing strategies [11]. Albeinitz et al. have shown that portfolios of quantity flexible procurement contracts used in combination with spot market procurement can contribute to reducing the manufacturer's expected profit and financial risk [10]. More recently, Nagali et al. have reported using new risk management techniques to support the development of portfolios of procurement contracts to save hundreds of millions of dollars in procurement costs at Hewlett-Packard [19]. Collectively, this body of work suggests that decisions relating to balancing one-off and long-term procurement contracts are becoming increasingly important and that their complexity warrants additional research.

2.2 Autonomous Bidding

The Artificial Intelligence and Electronic Commerce research communities have developed autonomous bidding agents that combine machine learning and stochastic optimization techniques. Work in this area has revolved around the development and benchmarking of bidding agents in the context of the Trading Agent Competition. A first game known as TAC-Travel was introduced by Wellman et al. in 1999 [34] with the first tournament held in July 2000. This game requires travel agents to compete against one another for air tickets, hotel reservations and event tickets to best match the requirements of their different customers. Techniques developed in the context of the TAC-Travel scenario are discussed in [33].

The Supply Chain Trading Competition, which was introduced in 2002 by the third author, was developed for the same research community but focuses on a supply chain trading scenario, where PC assembly agents concurrently compete for customers for different types of PCs and for the components required to assemble these PCs [2]. During its first five years, the TAC-SCM scenario has attracted over 120 entries from 21 countries and resulted in over 50 publications [32]. These include overall analysis of different tournaments (e.g. [2, 1]), novel supply chain trading techniques (e.g. [3, 14]) and descriptions of entire agents. In the latter category, Kiekintveld et. al. discuss how their DeepMaize agent dynamically coordinates sales, procurement and production strategies while maintaining profitability [15]. Fuzzy reasoning techniques used in the SouthamptonSCM agent are discussed in [12]. Details about the TacTex agent are provided in [21] and [22], including a discussion of machine learning techniques developed to predict bid prices of other agents [20]. [4] provides an in depth description of the different modules composing the CMieux agent. Podobnik, Petric and Jezic describe their CrocodileAgent agent in [26], [24] and [25]; Stan et. al. discuss their PhantAgent entry in [29]. Mathematical programming models and heuristic algorithms for bidding and scheduling developed for the Botticelli agent are detailed in [5]. In [13], the RedAgent team presents an internal market architecture to coordinate micro-agents that individually handle different subsets of decisions. Descriptions of many other techniques and agents can be found at [32].

2.3 Supply Chain Management Games

Supply chain management games are not new. The *Beer Distribution Game* originally developed at the Massachusetts Institute of Technology in the sixties has been played by several generations of students and executives and is an excellent illustration of what is now known as the *bullwhip effect* [17, 7]. Recent efforts to develop more comprehensive supply chain management games include our own work in the context of the Supply Chain Trading Agent Competition [2] as well as the work of Corsi et al. [9]. The MIT Procurement Game can be viewed as a precursor of the Procurement Challenge game discussed in this paper, as it also looks at the negotiation of portfolios of long-term quantity-flexible contracts [18]. In contrast to this game, the Procurement Challenge combines two phases: one involving the negotiation of long-term quantity flexible contracts and the other requiring agents to dynamically manage the resulting portfolios of contracts in combination with opportunities for one-off procurement contracts. The end result is a richer set of strategic interactions between the agents and scenarios where agents are evaluated over a significantly larger set of decisions. Like the TAC-SCM baseline game, the procurement challenge builds on multi-agent supply chain modeling and simulation concepts first introduced by Swaminathan et al. [30, 31].

3 TAC-SCM Procurement Challenge

The specification of the TAC-SCM Procurement Challenge was first released in 2007 [28] and was designed to provide a competition platform that isolates the procurement decisions faced by manufacturer agents in the baseline TAC-SCM game [2, 8]. This challenge game also extends the space of procurement options available to the manufacturer agents by allowing them to enter long-term contracts with suppliers agents.

The TAC-SCM Procurement Challenge (or "SCM-PC") was designed to promote the development of supply chain trading agents that are capable of effectively coordinating their procurement decisions. The game revolves around a personal computer (PC) assembly supply chain consisting of competing PC manufacturer agents (or agents), their component supplier agents (or suppliers) and their customer agents (or customers). This challenge requires agents to manage supply chain risk by negotiating long-term, quantity-flexible procurement contracts (or long-term contracts) and supplementing these contracts with one-off procurement contracts.

The SCM-PC game features three manufacturer agents competing for supply contracts from ten different suppliers. Each game simulates one hundred days, with each lasting ten seconds of real time. A server simulates the customers and suppliers, and provides banking, production, and warehousing services to the individual agents. The agents receive messages from the server on a daily basis informing them of the state of the game, such as the current inventory of components. They must send responses to the same server before the end of the day indicating their decisions, such as the orders they would place with the suppliers. At the end of the game, the agent with the most money in the bank is declared the winner.

Long-term contracts are negotiated at the start of a game. Each contract stipulates a minimum and maximum weekly quantity that an agent commits to purchasing for the entire duration of the game. Each day, the agents may also decide to procure additional components by negotiating one-off contracts. Each agent has the same amount of assembly capacity that it can use to produce any type of PC (different PC types consume different amounts of assembly capacity depending on their configuration). The assembly process is simulated by the game server, and also provides a warehouse where agents can store components and finished PCs. A daily production and delivery schedule are also generated by the game server for each agent. Customer orders are produced and delivered as soon as the required components have been purchased from suppliers.

3.1 Customer Demand, Production and Delivery

Throughout the course of a game the simulated customer demand varies according to a stochastic process whose actual parameter values are unknown to the agents. However, on any given day, the agents receive the exact same set of orders from customers representing one third of the total demand (this allows agents to focus solely on the procurement aspect of the game). Each order consists of a product type, a quantity, a due date and a price per unit. As in the baseline TAC-SCM game, the customer demand is generated according to a Poisson distribution. The distribution has an average (unknown to agents) that changes on a daily basis based on a random walk (for details see the latest specifications of the baseline game [8]).

Production and delivery are handled by the server in a greedy fashion, orders with higher revenue are placed at the top of a priority queue. When an order reaches the top of the queue the server checks whether or not each agent has enough components and available assembly capacity to produce it. Those agents with enough components and capacity exchange them for the revenue associated with the order.

3.2 Long-term Contracts

Long-term contracts are used to distribute risk between suppliers and manufacturer agents. Each contract consists of a minimum (Q^{lts}_{min}) and maximum (Q^{lts}_{max}) weekly quantity. A manufacturer agent that enters into such a contract commits to purchasing a quantity between Q^{lts}_{min} and Q^{lts}_{max} each week at a pre-negotiated unit price (p_{exec}). The agent also commits to pay at least a weekly reservation price (p_{res}) regardless of how many units it actually orders, which effectively raises the actual unit price on weeks when the agent purchases less than Q^{lts}_{max}. To ensure that each game presents a mix of long-term and one-off contract options, we assume that each component is available from a supplier that offers long-term contracts exclusively and another one that only offers one-off contracts.

When the game starts, the manufacturer agents have the option of negotiating long-term contracts for each component, and these contracts are awarded based on second price auctions. Each long-term supplier first announces a reserve price, ρ_j, in its auction for component j. The supplier then waits ten seconds for the agents to submit their bids, which must have prices greater than or equal to ρ_j. A bid consists of a requested maximum weekly quantity and an execution price that the agent is willing to pay for each component. The minimum weekly quantity and the reservation price are not specified by the agents' bids, they are calculated by the suppliers as follows:

- $Q^{lts}_{min} = Q^{lts}_{max}/(1 + \alpha)$, where α changes from one game to another and is announced at the start of each game. Note that larger values of α lead to larger gaps between Q^{lts}_{min} and Q^{lts}_{max}. Thus, the α parameter determines the flexibility of the contract.

- $p_{res}/(p_{res} + p_{exec}) = \beta$, where β also changes from one game to another and is announced at the start of each game. The β parameter can be seen as a measure of the risk taken on by the manufacturer agent that enters into the contract. Larger values of β lead to agents paying a larger fraction of the maximum contract price, regardless of how many units they actually request.

The long-term contract supplier allocates up to 100% of its weekly nominal capacity (C^{nom}_{week}) to the bidding agents. Quantities are allocated based on the requested maximum weekly quantities, starting with the bid that has the highest execution price. Each agent's long-term contract has an execution price that is computed as the next highest price below its own bid ("second highest price" rule). The allocation proceeds until there are no bids left or until the long-term supplier has run out of capacity (based on its weekly nominal capacity). In the latter situation, the last manufacturer agent to receive a contract may end up with a maximum weekly quantity that is less than what it had requested. Each supplier's actual capacity follows a stochastic process similar to that of the customer demand. This process is identical to the way supplier capacity is handled in the baseline TAC-SCM game (for details see that game's latest specification [8]).

Agent	**Requested maximum weekly quantity**	**Bid price**
1	1000	850
2	800	950
3	1200	870

Table 1: Long-term Contract Bids

Table 1 presents an example of bids sent to a long-term contract supplier. In this example, let us assume that the weekly nominal capacity of the supplier is 2695 and the auction's reserve price is 800. Agent 2 is the first to receive a long-term contract with a maximum weekly quantity of 800 units and an execution price of 870, followed by Agent 3 that gets a long-term contract with 1200 units/week for 850 per unit and Agent 1 that only gets a contract with 695 units/week for 800 per unit, namely the supplier's reserve price.

At the beginning of any given week, each agent decides how much to actually order under its long-term contracts. If the total quantity of a given component requested by the agents is less than the actual supplier's capacity, all agents get the full quantities they requested. If the total quantity requested by the agents exceeds this capacity, the supplier computes the ratio of demand it can satisfy based on its actual capacity. Each agent then receives a quantity that is proportional to this ratio, so that all agents with long-term contracts are treated equally and receive the same fraction of their actual demand that week.

3.3 One-off Contracts

The one-off contract negotiations proceed exactly as in the baseline TAC-SCM game. The following provides a brief overview - for further details, the reader is referred to the latest specification of the baseline game [8].

Every day, manufacturer agents can send up to 5 requests for quotes (RFQs) to each one-off supplier with a specified reserve price, quantity and delivery date (we refer to the length of time between the day the RFQ is sent and the delivery date in the future as the "lead time" of the RFQ). A supplier receives all RFQs on a given day, and processes them together at the end of the day to find a combination of offers that approximately maximizes its revenue. On the following day, the suppliers send back to each agent an offer for each RFQ with a price that is guaranteed to below the specified reserve price, a quantity that is less than or equal to the specified quantity, and an estimated delivery date. Due to capacity restrictions, the supplier may not be able to supply the entire quantity requested in the RFQ by the due date. Thus, it responds by issuing up to two modified offers, each of which relaxes one of the two constraints:

- **Quantity**, in which case offers are referred to as *partial offers*.
- **Due date**, in which case offers are referred to as *earliest offers*.

Both long-term and one-off suppliers have a limited capacity for producing a component, and this limit varies throughout the game according to a mean reverting random walk. Moreover, suppliers limit their long-term commitments by reserving some capacity for future business. The pricing of components is based on the ratio of demand to supply, and higher ratios result in higher prices. Each day the suppliers estimate their free capacity by scheduling production of all current and pending orders as close to their due dates as possible. The price offered in response to an RFQ is equal to the requested components base price discounted by a function proportionate to the supplier's free capacity before the RFQ due date. The manufacturer agents normally face an important trade-off in the procurement process: pre-order components for the future where customer demand is difficult to predict, or wait to purchase components and risk being unsuccessful due to high prices or availability.

A reputation rating is also used by the suppliers to discourage agents from driving up prices by sending RFQs with no intention of buying. Each supplier keeps track of its interaction with each agent, and calculates that agent's reputation rating based on the ratio of the quantity purchased to quantity offered by the agent. If an agent's reputation falls bellow a minimum value, then the prices and availability of components are affected for that specific agent. Therefore, agents must carefully plan the RFQs sent to suppliers.

4 TAC SCM-PC Agents

This section describes the approaches of the top three SCM-PC agents: PhantAgent (University "Politehnica" of Bucharest), CMieux (Carnegie Mellon University) and CrocodileAgent (University of Zagreb). Due to the complexity of the SCM-PC game there are many differences between each of these three agents. However, we will focus on the key differences that affect their long-term and one-off contract strategies. For example, one of these differences is the way that future demand is predicted and how this prediction is used to create long-term procurement strategies. Another difference is how one-off contracts are handled, with PhantAgent and CrocodileAgent relying on strategies that generate requests for quotes with fixed lead times and CMieux implementing a dynamic strategy that adapts these lead times to identify and selectively take advantage of component price fluctuations.

4.1 PhantAgent

PhantAgent's decision process can be thought of as consisting of three different sub-problems: forecasting needed components, handling long-term contract procurement, and generating one-off contract orders.

4.1.1 Calculating Needed Components

At the beginning of each day, PhantAgent estimates the number of components it will need for the remainder of the game. In order to determine this number, PhantAgent first estimates the number of components it expects to have in inventory on each of the remaining days. The expected inventory levels are estimated by iterating through each day, adding component arrivals that are due and subtracting the estimated component usage. The main difficulty in this process involves determining a good estimate of daily component usage. We will refer to the daily usage of component j as Q_j. To estimate Q_j PhantAgent combines two heuristic values:

- The first heuristic value, $E[Q_j]$, assumes that component usage will be exactly equal to the expectation of the stochastic process described in the game specification. This can be calculated offline by assuming that the changing means of the Poisson distributions governing customer demand will remain in the middle of their range (i.e. that the demand will not be too high or too low).

- The second heuristic value, $\bar{Q}_j$, is a moving average of component usage from the past 10 days. This value is computed based on the observed draws from the Poisson distributions governing customer demand.

Both of these heuristic values have certain weaknesses. The problem with using $E[Q_j]$ is that it is not flexible to demand variations and fails to account for fluctuations in demand throughout the game. By ignoring such fluctuations the agent will often either run out of components or be left with excess inventory when the game nears its end. The problem with using $\bar{Q}_j$ is that demand at the beginning of the game can be significantly different from the demand at the end. Thus, the long lead time orders placed at the beginning of the game based on the demand at that time may not match the demand when the components arrive. To avoid these problems PhantAgent uses a weighted average of both heuristic values in its prediction of daily component usage, $\widehat{Q}_j$. The average is weighted towards $E[Q_j]$ during the beginning phase of the game and $\bar{Q}_j$ is weighted more heavily during the end phase of the game when there is too little time left for the demand to change much. The formula is given in Equation 1.

$$\widehat{Q}_j = \frac{D-d}{D} E[Q_j] + \frac{d}{D} \bar{Q}_j \tag{1}$$

where:
D - the total number of days in the game.

Towards the end of the game $\widehat{Q}_j$ is slightly scaled down to avoid excess inventory due to changes in demand from the time when long lead time requests were made.

4.1.2 Handling Long-term Contract Procurement

Long-term contracts have the potential to provide lower prices and higher guarantees on availability. PhantAgent prioritizes availability over price in its long-term contracts. To take advantage of high selling prices during times of high demand and low supply, it was empirically determined that bidding the average one-off contract prices from the past several games enabled the agent to reliably procure the quantity it desired. The quantity requested from each long-term supplier for each component j was chosen empirically to be a fixed fraction of the expected weekly demand for each component based on the game specification.

Throughout the game PhantAgent exercises the option to increase weekly order quantities if there is a need for components and the one-off contract suppliers are charging more than the long-term contract prices.

4.1.3 Generating One-off Contract Orders

For one-off contract requests, PhantAgent typically uses all five available RFQs each day to request components with fixed lead times (for SCM-PC in 2007, values used were $\{2, 3, 10, 25, 45\}$ days).

The agent adjusts its requested quantities according to current market conditions.

Short Lead Time Requests with One-off Contracts

Requests with very short lead times (such as lead times of 2 and 3 days) are treated independently of the other RFQs and are used primarily to maintain a steady stock of components. The price at which components are sold over such short lead times tend to vary widely from one day to the next, depending on current demand. Handling these daily variations is the main concern here. By using low reserve prices in its short-term procurement RFQs, PhantAgent ensures that it only purchases components in the short-term markets when prices are particularly attractive while maintaining a high reputation (the impact of an agent's reputation is described in Section 3.3). Specifically, the reserve price is set to a value that is slightly over the *minimum* price obtained in the last 5 days (using the *average* price was also considered, but results were observed to be worse).

Long Lead Time Requests with One-off Contracts

In addition to the short lead time requests described above, PhantAgent uses one-off contracts with long lead times (e.g. 10, 25 and 45 days) to complement its long-term contracts. Since PhantAgent uses fixed lead times for all one-off contract requests, the main decisions regarding long lead time requests are choosing appropriate order quantities and reserve prices. The general principle governing these choices is to make long lead time orders only if they are expected to be better or equal to orders with short lead times. In most cases the orders with longer lead times have lower prices than those with shorter lead times. However, orders with shorter lead times are important towards the end of the game due to the need for more precise inventory management to prevent the agent from finishing the game with unused components.

Order quantities are chosen simply based on the difference between the expected usage and the components already expected from prior procurement. The reserve price is calculated the same way as the reserve prices in the very short lead time requests (i.e. the minimum price obtained in the last 5 days). This ensures that long lead time orders only result in offers that are at least as good as current short lead time ones.

4.2 CMieux

The strategy used by CMieux for the long-term contract negotiation and procurement is described in section 4.2.1. The strategy used by CMieux for negotiating one-off contracts is essentially the same as the technique used by the CMieux agent for the baseline TAC-SCM game. Therefore, we will only present a summarized version of this strategy in section 4.2.2 (for a more complete description the reader is referred to [4]).

4.2.1 Long-term Contract Negotiation and Procurement

CMieux's strategy for long-term contract negotiation is similar to that of PhantAgent. The bid price chosen by CMieux is the average one-off contract price computed over the past several games (unless this price is lower than the supplier's stated reserve price, ρ_j, in which case CMieux bids slightly above itρ_j). Thus, when reserve prices are low the agent is willing to buy components from the long-term contract suppliers for a price equal to what it expects to be the average one-off contract price over the course of the game.

The requested weekly quantity for each component j, Q_j^{lts}, is considered a parameter in the model that must be adjusted to reflect the amount of risk the agent is willing to take with the long-term contracts. Empirically, it was determined that on average the loss from unsold inventory in low customer demand games outweighed the profits from sales in higher demand games. This problem occurred due to the very low flexibility (approximately 15% on average) between the minimum and maximum weekly quantities, which was not large enough to cover all the different customer demand scenarios. Thus, a somewhat conservative strategy was adopted with the requested maximum weekly quantity adjusted to suit low demand games.

4.2.2 One-off Contract Procurement

CMieux's one-off procurement strategy uses dynamically chosen lead times for its RFQs, which makes it significantly different from that of the other two agents' fixed lead time strategies. CMieux uses the same procurement module as in the baseline TAC-SCM game to handle all aspects of requesting and purchasing components from the one-off contract suppliers. The module is designed to rapidly adapt to changing market conditions. This includes monitoring component price fluctuations by probing suppliers and adjusting RFQ lead times to exploit periods when different components are relatively inexpensive. Each day, the procurement module performs two tasks: i) given that agents are limited to sending five RFQs per day to each supplier, CMieux optimizes the lead times and quantities of its RFQs to exploit periods when each component is inexpensive, ii) when suppliers respond to its RFQs with offers, it attempts to identify a particularly promising subset of current of the offers to accept. We provide a brief summary of the procurement module below. Additional details are given in a sister publication describing the CMieux baseline agent [4].

Sending Supplier Requests

To determine the amount of components needed, the procurement module computes the difference between the inventory required to satisfy customer orders (both current and expected) and inventory (both current and projected from both long-term and one-off contract suppliers for the remainder of the game). Clearly, CMieux does not *need* to procure this entire difference each day, since most components are not needed immediately. Instead, the agent can spread required component quantities across many days. This not only enables CMieux to more aggressively procure components it needs in the short-term but also gives it the flexibility to identify the cheapest times when to procure the remaining components.

The process of computing what specific requests to send to suppliers is decomposed by component type. For each component type, the procurement module generates several sets of lead times and searches for the set most likely to minimize procurement costs. Uncertainty about future need for different components as well as their future availability and price is handled by computing a "utility" for each component. The utility of a component is considered to be its estimated value to the agent minus its forecast price. The value of a component reflects the price at which the agent expects to be able to sell PCs it can build with it, as well as the expected price and availability of other components the agent might need to procure to be able to assemble these PCs. The agent searches through the space of possible RFQs for a set that provides the greatest overall utility. The reserve price of each RFQ is set to be the average utility of the components it includes.

4.3 CrocodileAgent

In contrast to PhantAgent and CMieux, CrocodileAgent uses a more dynamic strategy for choosing the component quantities to request from long-term suppliers. Specifically it adjusts the quantity it requests based on the supplier's reserve price. For one-off contracts, the agent is similar to PhantAgent in that it also uses static lead times for one-off RFQs. CrocodileAgent further adds a large one-off procurement request at the beginning of the game to jump start its production.

4.3.1 Negotiating Long-term Contracts

After the long-term suppliers have announced reserve prices for each component, CrocodileAgent determines a quantity to request from each supplier based on its reserve price. The requested quantity Q_j^{lts} for a component j is linearly adjusted between two empirically chosen parameters Q_j^{max} and Q_j^{min}. Specifically, given a long-term supplier's stated reserve price ranging between ρ_{min} and ρ_{max}, CrocodileAgent modulates its long-term component bid quantity Q_j^{lts} according to Equation 2.

$$Q_j^{lts} = \left(\frac{\rho_{max} - \rho_j}{\rho_{max} - \rho_{min}}\right) Q_j^{max} + \left(1 - \frac{\rho_{max} - \rho_j}{\rho_{max} - \rho_{min}}\right) Q_j^{min} \tag{2}$$

This leads to a strategy that requests more from long-term suppliers with lower reserve prices and less from long-term suppliers with higher reserve prices. Prices in long-term contract bids are simply set to the reserve price for the corresponding component ($p_{exec} = \rho_j$).

Each week CrocodileAgent adjusts the actual quantity to exercise (Q_j^{order}) on each long-term contract based on the amount of components it has in inventory (N_j). The quantity that the agent actually procures each week from each of its long-term contracts is modulated to maintain the corresponding component's inventory level between two empirically determined levels N_j^{min} and N_j^{max}. In other words CrocodileAgent scales back its weekly long-term contract orders when it has an ample supply of components and scales them up when it is running short.

4.3.2 Negotiating One-off Contracts

CrocodileAgent breaks up the procurement problem with one-off contracts into two different subproblems: a strategy for the first day of the game (*day 0*) and a replenishment strategy throughout the rest of the game. A close examination of the baseline TAC-SCM Game rules [8] (which also defines the SCM-PC one-off contract supplier model) suggests that ordering components at the very beginning of the game (*day 0* procurement) results in lower prices since one-off suppliers have not yet committed any of their capacity. Although the concept of the *day 0* procurement strategy has some similarities with long-term contract negotiation, these two procurement strategies are totally independent because each component is available from two different suppliers: one that only offers long-term contracts and one that only sells components with one-off contracts.

One-off Procurement on *Day 0*

By implementing a special *day 0* one-off strategy at the beginning of the game, the agent has the opportunity to possibly buy cheap components with longer lead times. Specifically, CrocodileAgent sends a fixed set of five RFQs on *day 0* (the specific values of the lead times, quantities and reserve prices are given in Table 2). These *day 0* lead times and quantities were empirically fine-tuned over a large number of games.

Lead time	7 days	14 days	21 days	52 days	77 days
Reserve price (% of nominal)	107%	97%	92%	77%	69%
Quantity for CPUs	300	350	400	450	450
Quantity for all other types	600	700	800	900	900

Table 2: The actual parameter values used by CrocodileAgent for *day 0* RFQs

One-off Procurement During the Game

During the course of the game CrocodileAgent uses one-off contracts with fixed lead times to fill in gaps in its inventory not covered by its long-term contracts or its *day 0* strategy. At the beginning of each day, the agent calculates the quantity of each component that has been previously ordered but not yet delivered. This quantity is multiplied by a distance factor, so that orders with longer lead times have a smaller weight than orders with shorter lead times. If the quantity in inventory plus the outstanding quantity for a component is less than a threshold, the agent shifts to a more aggressive strategy where all five available RFQs are sent to the corresponding one-off supplier with short lead times and relaxed reserve prices.

CrocodileAgent uses the following additional heuristics to adjust reserve prices and request quantities for one-off procurement:

- When the inventory level of a certain component is low the agent begins to procure components with higher reserve prices than usual to ensure a sufficient supply.
- When the agent experiences a rapid rise in demand it increases the quantities it is requesting from one-off suppliers. This helps to ensure that the agent does not run out of components and consequently lose potentially profitable PC orders.
- Special attention is also given to the end of the game. One-off contract requests at the end of the game are adjusted to help ensure that the agent is not left with excess inventory.

5 2007 TAC SCM-PC Results and Analysis

This section presents the results of the final rounds of the 2007 TAC-SCM Procurement Challenge. The final rounds were held at the Twenty-Second Conference on Artificial Intelligence (AAAI-07). They featured twelve games, three games for each combination of three agents out of the four finalists. The final standings are presented in Table 3, the value in the fourth column is the average profit accumulated by each agent over the course of the nine games it played in.

In addition to the overall competition results we also performed a finer pairwise comparison of the top three agents to account for the fact that they did not all participate in the same games.

Agent	Games Played (out of 12)	Number of Games Won	Average Score
PhantAgent	9	4	8,731,513
CMieux	9	6	7,405,743
CrocodileAgent	9	2	6,399,115
Warrior	9	0	4,200,440

Table 3: Final Standings for the 2007 SCM-PC

Agent	Games Played	Number of Games Won	Average Score
CMieux	6	4	7,149,838
PhantAgent	6	2	6,788,197

Table 4: "Pairwise" Comparison Between CMieux and PhantAgent for the 2007 SCM-PC

Tables 4, 5 and 6 present the performance of pairs of agents in all of the games involving them both. For the top three agents there are three distinct pairs and each pair participated in six common games. As can be seen, the pairwise results provide a different ranking with CMieux ahead of the other two agents and PhantAgent ahead of CrocodileAgent. It is also worth noting that these results are consistent with the number of games won by each agent throughout the finals, with CMieux winning 6 out 9 games, PhantAgent 4 out of 9 games and CrocodileAgent 2 out of 9 games.

The discrepancy between the overall rankings and pairwise rankings can be explained by the varying demand conditions faced by the agents in different games. PhantAgent achieved a higher overall score than CMieux because it participated in one game with a high customer demand that CMieux did not play in.

We will now present analysis of several important aspects of the game as well as graphs that illustrate the effect of the strategies adopted by the top three agents. Section 5.1 describes the sales volume of each agent, section 5.2 presents their one-off and long-term contract mixes, and section 5.3 presents their average procurement costs.

Agent	Games Played	Number of Games Won	Average Score
CMieux	6	4	8,286,761
CrocodileAgent	6	1	4,385,217

Table 5: "Pairwise" Comparison Between CMieux and CrocodileAgent for the 2007 SCM-PC

Agent	Games Played	Number of Games Won	Average Score
PhantAgent	6	3	10,027,071
CrocodileAgent	6	1	7,096,601

Table 6: "Pairwise" Comparison Between PhantAgent and CrocodileAgent for the 2007 SCM-PC

5.1 Customer Orders and Deliveries

To measure the sales volume of the top three SCM-PC agents, we calculated their *realized demand percentage*, or the fraction of the total possible demand that they were able to satisfy. Figure 1 presents a pairwise comparison of the average realized demand (with 95% confidence intervals) of the top three agents. As in Tables 4, 5 and 6, the values shown for each pair in Figure 1 are calculated using only the games that pair participated in. CrocodileAgent had the highest average realized demand amongst all agents. However, the overlapping confidence intervals show that there was no statistically significant difference between any of the agents.

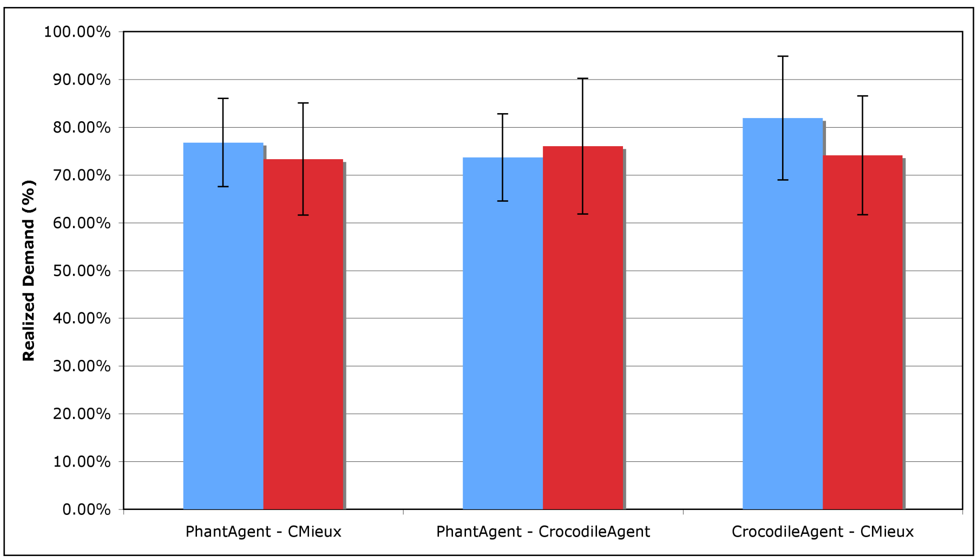

Figure 1: Average Realized Demand

5.2 Quantity Ordered from the Suppliers

The average number of components ordered by each of the top three agents from the one-off and long-term contract suppliers with 95% confidence intervals is presented in Figures 2 and 3. These graphs show that all three of the top agents procure a substantial amount of components from the more stable long-term market, but tend to buy significantly more from the one-off contract market. While long-term contracts provide some amount of flexibility in the weekly orders, they are negotiated when the agents have very limited information about customer demand: only information about the general parameters used to generate random walks as provided in the game specification.

It is not surprising therefore that all agents chose to be conservative in their long-term contract bids and relied more heavily on one-off contracts: one-off contracts can be negotiated on a daily basis, giving agents more flexibility to adapt to actual market conditions during the course of the game. Accordingly, higher demand games also resulted in a higher proportion of components being procured from one-off suppliers. Among the three agents, CrocodileAgent was the least conservative in its long-term contract strategy, which is consistent with its strategy of offering lower prices in its long-term bids (bidding the reserve price, as discussed in section 4.3.1). CMieux comes next, with PhantAgent relying the least on long-term contracts (and the most on one-off contracts).

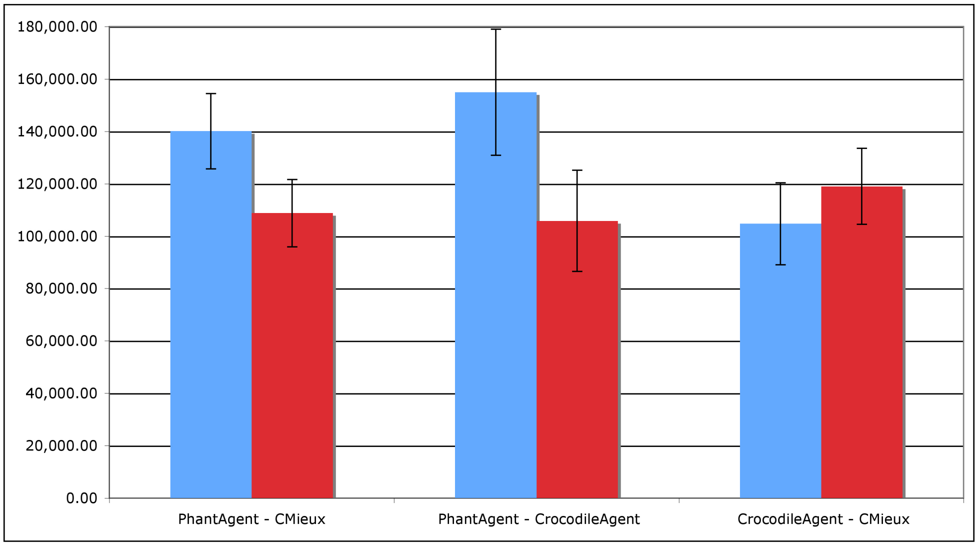

Figure 2: Average Number of Components Ordered from the One-off Contract Suppliers

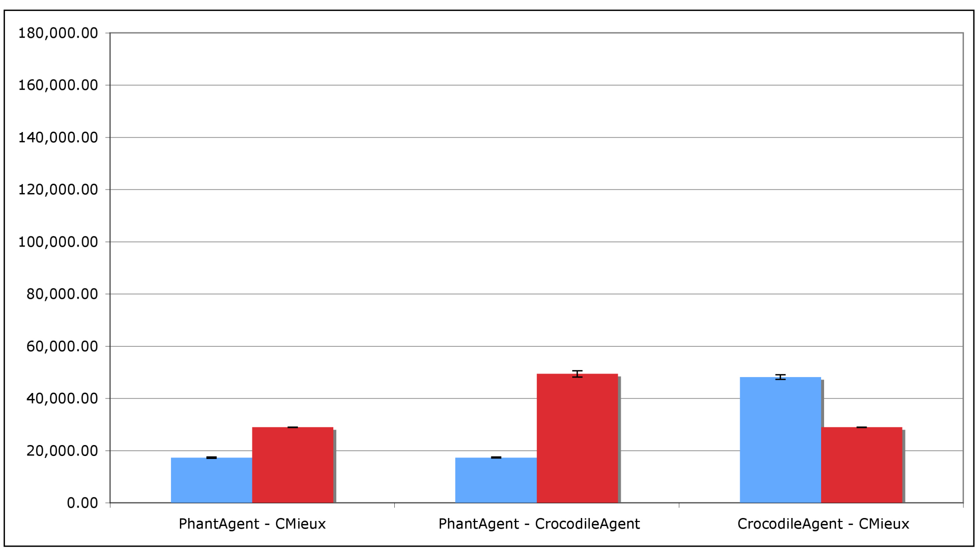

Figure 3: Average Number of Components Ordered from the Long-term Contract Suppliers

5.3 Component Prices

Figure 4 displays the average weighted prices of components purchased by each agent in both the long-term and one-off contract markets (with 95% confidence intervals). The graph shows that CMieux's procurement prices were significantly better than the other two agents when compared across both markets.

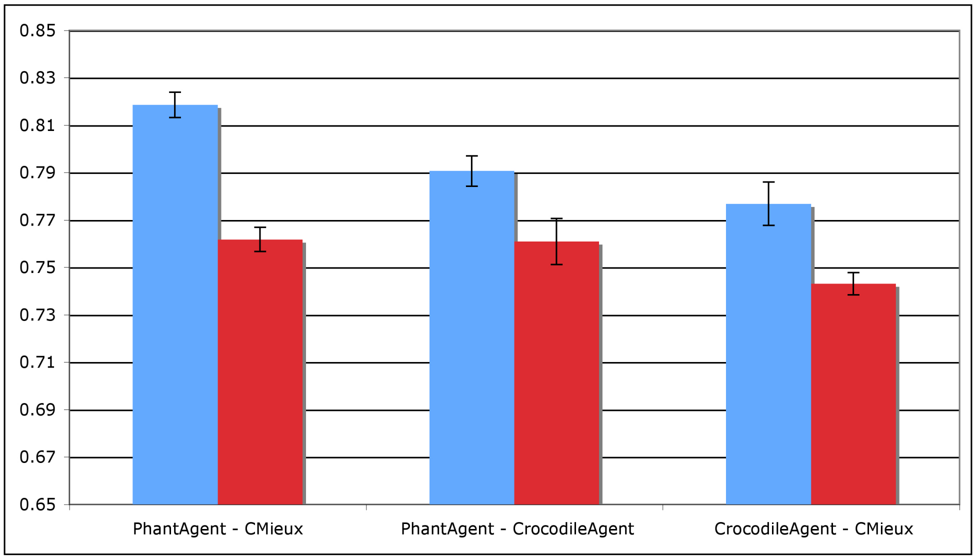

Figure 4: Average Weighted Prices of Components from Long-term and One-off Contracts

A finer analysis that looks separately at prices paid in the one-off and long-term contract markets shows that CrocodileAgent was able to get significantly lower prices for long-term contracts than the other two agents (Figure 5). CMieux's long-term prices were a close second to CrocodileAgent with an average difference of 1.69%. In the one-off contract market, it can be seen that CMieux has a significant edge over the other two agents (Figure 6).

As already indicated CrocodileAgent's low long-term procurement costs can be attributed to its strategy of bidding the reserve price on long-term contracts (see section 4.3.1). The fact that CMieux had the lowest overall procurement costs suggests that CrocodileAgent was not able to procure enough from the long-term markets to overcome CMieux's price advantage in the one-off contract markets. CMieux's advantage in the one-off markets can likely be attributed to its dynamic one-off contract strategy, enabling the agent to optimize the lead times of its daily one-off RFQs, while the other agents rely on more static strategies.

6 Conclusions and Future Work

With the Web opening the door to more agile procurement strategies where manufacturers manage risk through a mix of long-term and one-off contracts, it becomes increasingly important to develop a better understanding of the risk and behaviors associated with these practices. In this paper, we provided an overview of the the Supply Chain Trading Agent Competition Procurement Challenge (SCM-PC), a game developed by the first three authors to capture many of the strategic interactions

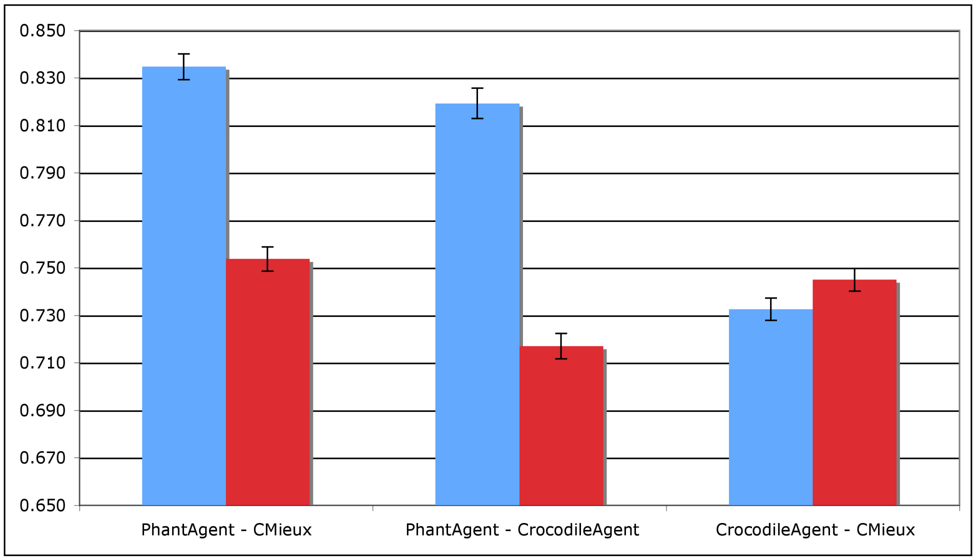

Figure 5: Average Prices of Components from Long-term Contracts

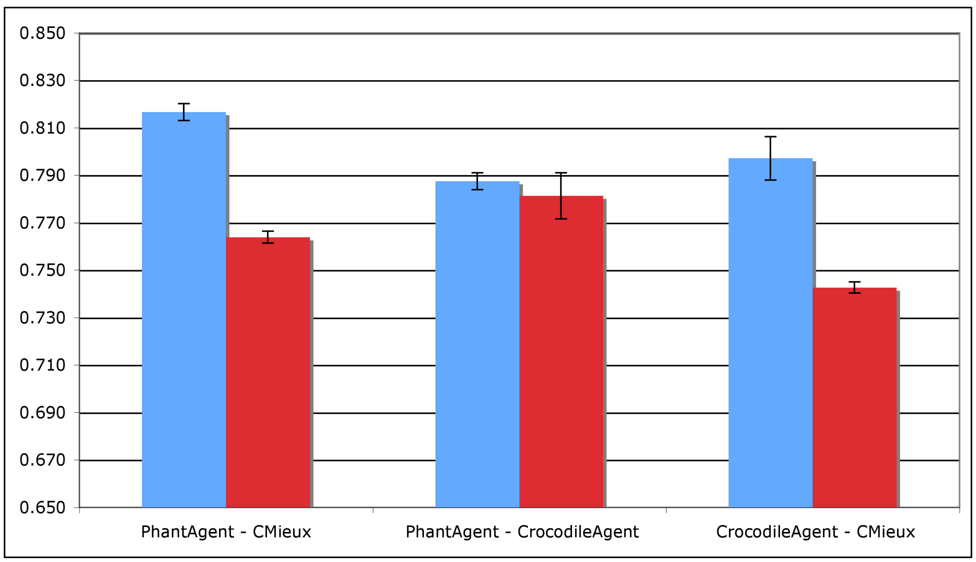

Figure 6: Average Prices of Components from One-off Contracts

entailed by these new practices. The game builds on the success of a scenario introduced by the third author in 2002 in the context of an annual tournament, now known as the TAC-SCM baseline game, a somewhat more complex game that requires competing agents to coordinate both component procurement and customer bidding activities while managing their finite assembly capacity. In contrast, the SCM-PC scenario was designed to (1) focus solely on procurement decisions, and (2) introduce an environment where manufacturing agents need to manage risk through a combination of long-term quantity flexible contracts and one-off contracts. This scenario has the merit of capturing decisions that more closely resemble those faced by many companies today, given the decoupling between procurement and sales traditionally imposed by prevailing ERP architectures.

We proceeded to provide an overview of the top three entries in the first edition of the Procurement Challenge, which was collocated with the AAAI-07 conference in Vancouver. Results from the competition indicate that all three agents developed strategies that combined both long-term and one-off procurement opportunities, suggesting that the overall design of the game was successful. A finer analysis shows however that, in its initial form, the competition was still biased towards one-off contracts in part due to a dearth of information available to agents at the beginning of the game, when they need to negotiate long-term contracts. An additional incentive for being conservative with long-term contracts in this first edition of the competition had to do with the separation between long-term and one-off suppliers. Agents that do not bid aggressively on long-term contracts are still guaranteed a chance to compete for component capacity set aside by one-off suppliers. In 2008, the first three authors decided to lift this restriction and allow each supplier to sell its capacity under both long-term and one-off contracts, a decision expected to lead to more competitive behaviors in the long-term contracting phase of the game. Future versions of the game could also allow agents to enter long-term contracts in mid course, rather than limit long-term contract negotiation to the initial phase of the game.

References

[1] J. Andrews, M. Benisch, A. Sardinha, N. Sadeh, What differentiates a winning agent: An information gain based analysis of TAC-SCM, in: Proceedings of Trading Agent Design and Analysis Workshop, 2007.

[2] R. Arunachalam, N. Sadeh, The supply chain trading agent competition, Electronic Commerce Research and Applications 4 (1) (2005) 66–84.

[3] M. Benisch, J. Andrews, N. Sadeh, Pricing for customers with probabilistic valuations as a continuous knapsack problem, in: Proceedings of ICEC'06, 2006.

[4] M. Benisch, J. Andrews, A. Sardinha, N. Sadeh, CMieux: Adaptive strategies for competitive supply chain trading, Electronic Commerce Research ApplicationsTo appear.

[5] M. Benisch, A. Greenwald, I. Grypari, R. Lederman, V. Naroditsky, M. Tschantz, Botticelli: A supply chain management agent, in: Third International Joint Conference on Autonomous Agents and Multi-Agent Systems, 2004.

[6] M. Bensaou, Portfolios of buyer-supplier relationships, Sloan Management Review (1999) 35–44.

[7] S. Chopra, P. Meindl, Supply Chain Management, Pearson Prentice Hall, New Jersey, 2004.

[8] J. Collins, R. Arunachalam, N. Sadeh, J. Eriksson, N. Finne, S. Janson, The supply chain management game for the 2007 trading agent competition, CMU-ISRI-07-100, Carnegie Mellon University (2006).

[9] T. Corsi, S. Boyson, A. Verbraeck, S. van Houten, C. Han, J. MacDonald, The real-time global supply chain game: New educational tool for developing supply chain management professionals, Transportation Journal 45 (2006) 61–73.

[10] V. M. de Albeniz, D. Simchi-Levi, A portfolio approach to procurement contracts, Production and Operations Management 14 (1) (2005) 90–114.

[11] W. J. Elmaghraby, Supply contract competition and sourcing policies, Manufacturing and Service Operations Management 2 (4) (2000) 350–371.

[12] M. He, A. Rogers, X. Luo, N. R. Jennings, Designing a successful trading agent for supply chain management, in: Proceedings of AAMAS'06, 2006.

[13] P. W. Keller, F.-O. Duguay, D. Precup, RedAgent-2003: An autonomous, market-based supply-chain management agent, in: Third International Joint Conference on Autonomous Agents and Multi-Agent Systems, 2004.

[14] C. Kiekintveld, J. Miller, P. Jordan, M. P. Wellman, Forecasting market prices in a supply chain game, in: International Joint Conference on Autonomous Agents and Multiagent Systems, 2007.

[15] C. Kiekintveld, M. P. Wellman, S. Singh, J. Estelle, Y. Vorobeychik, V. Soni, M. Rudary, Distributed feedback control for decision making on supply chains, in: Fourteenth International Conference on Automated Planning and Scheduling, 2004.

[16] P. Kleindorfer, D. Wu, Integrating long- and short-term contracting via business-to-business exchanges for capital-intensive industries, Management Science 49 (11) (2003) 1597–1615.

[17] H. L. Lee, V. Padmanabhan, S. Whang, The bullwhip effect in supply chains, Sloan Management Review (1997) 93–102.

[18] MIT Forum for Supply Chain Innovation, The MIT procurement game, Website, http://supplychain.mit.edu/ProcurementPortal/ (2003).

[19] V. Nagali, J. Hwang, D. Sanghera, M. Gaskins, M. Pridgen, T. Thurston, P. Mackenroth, D. Branvold, P. Scholler, G. Shoemaker, Procurement risk management (PRM) at Hewlett-Packard company, Interfaces 38 (1) (2008) 61–75.

[20] D. Pardoe, P. Stone, Bidding for customer orders in TAC SCM, in: AAMAS-04 Workshop on Agent-Mediated Electronic Commerce, 2004.

[21] D. Pardoe, P. Stone, Predictive planning for supply chain management, in: Proceedings of Automated Planning and Scheduling'06, 2006.

[22] D. Pardoe, P. Stone, An autonomous agent for supply chain management, in: G. Adomavicius, A. Gupta (eds.), Handbooks in Information Systems Series: Business Computing, Elsevier, 2007.

[23] B. Peleg, H. Lee, W. Hausman, Short-term e-procurement strategies versus long-term contracts, Production and Operations Management (2002) 458–479.

[24] A. Petric, V. Podobnik, G. Jezic, The CrocodileAgent 2005: An overview of TAC SCM agent, Lecture Notes in Artificial Intelligence, Subseries of Lecture Notes in Computer Science 4452 (2007) 219–233.

[25] A. Petric, V. Podobnik, G. Jezic, The CrocodileAgent: Designing a robust trading agent for volatile e-market conditions, Lecture Notes in Artificial Intelligence, Subseries of Lecture Notes in Computer Science 4496 (2007) 597–606.

[26] V. Podobnik, A. Petric, G. Jezic, The CrocodileAgent: Research for efficient agent-based cross-enterprise processes, Lecture Notes in Computer Science 4277 (2006) 752–762.

[27] D. Pyke, M. Johnson, Sourcing strategy and supplier relationships: Alliances vs. eprocurement, in: C. Billington, H. Lee, J. Neale, T. Harrison (eds.), The Practice of Supply Chain Management, Kluwer Publishers, 2003, pp. 77–89.

[28] A. Sardinha, M. Benish, J. Andrews, N. Sadeh, The 2007 supply chain trading agent competition - procurement challenge, CMU-ISRI-07-106, Carnegie Mellon University (2007).

[29] M. Stan, B. Stan, A. M. Florea, A dynamic strategy agent for supply chain management, in: SYNASC, 2006.

[30] J. M. Swaminathan, N. M. Sadeh, S. F. Smith, Information exchange in the supply chain, CMU-RI-95-36, Carnegie Mellon University (1995).

[31] J. M. Swaminathan, S. Smith, N. M. Sadeh, Modeling supply chain dynamics: A multiagent approach, Decision Sciences 29 (1998) 607–632.

[32] Trading Agent Competition Website, TAC research publications, Website, http://tac.eecs.umich.edu/researchreport.html (2007).

[33] M. P. Wellman, A. Greenwald, P. Stone, Autonomous Bidding Agents: Strategies and Lessons from the Trading Agent Competition, MIT Press, 2007.

[34] M. P. Wellman, P. R. Wurman, A trading agent comeptition for the research community, in: IJCAI-99 Workshop on Agent-Mediated Electronic Trading, 1999.