THE SUPPLY CHAIN TRADING AGENT COMPETITION

Raghu Arunachalam, Norman M. Sadeh*
CMU-CS-04-164
September 27, 2004

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213-3891

Also appears as ISRI Technical Report CMU-ISRI-04-133
And
To Appear in "Electronic Commerce Research & Applications Journal"

Keywords. Agent-based systems, market game, multi-agent simulation, procurement, intelligent agents, supply chain management, trading agents.

Abstract

Supply chain management deals with the planning and coordination of bidding, production, sourcing and procurement activities associated with one or more products. It is central to today's global economy, leading to trillions of dollars in annual transactions worldwide. With the emergence of electronic marketplaces, it is only natural to seek automated solutions that are capable of rapidly evaluating a large number of bidding, sourcing and procurement options. In this paper, we detail a game we have designed to promote the research and evaluation of such solutions under realistic conditions. The game requires agents to manage the assembly of PCs, while competing with one another both for customer orders and for key components. We discuss how the game captures the complexity, stochasticity and competitive nature inherent to supply chain environments. A Web-based multi-agent simulation platform developed for the game was implemented in 2003 and validated in the context of the first *Supply Chain Management Trading Agent Competition* (TAC-SCM). A total of 20 teams from around the world competed with one another. We review agent strategies developed by different teams and discuss the merits of competition-based research over more traditional research methodologies in this area.

1. INTRODUCTION

Supply chain management is concerned with planning and coordinating bidding, production, sourcing and procurement activities across the multiple organizations involved in the delivery of one or more products. It is central to today's global economy, leading to trillions of dollars in annual transactions worldwide. Supply chains are highly dynamic environments that are subject to:

- *market fluctuations*, such as surges in customer demand or drops in supply availability;
- *operational contingencies*, such as delays in supply delivery, losses of capacity, or quality problems; and,
- *changes in strategies* employed by competitors, customers or suppliers

Accordingly, supply chain performance can significantly benefit from decision making processes that constantly monitor changing conditions and dynamically evaluate available trading and operational options in light of these conditions. With the emergence of electronic marketplaces, automated programs or "intelligent agents" offer the promise of significantly increasing the number of options one can consider and of substantially improving supply chain performance. Simple versions of such programs have been demonstrated in other domains, though the prospect of delegating routine supply chain decisions to software agents still makes many managers nervous. How will the agents react under changing conditions? Could competitors develop strategies that exploit some of their potential weaknesses? While routine planning and control decisions in static supply chains are now relatively well understood, this is not the case of dynamic supply chain trading environments, where companies more openly compete for customer orders and components. Evaluating the benefits and possible limitations of intelligent agent functionality in these more challenging environments can not be convincingly done by just relying on traditional methodologies, where a given technique is evaluated under a set of predefined conditions. Instead, supply chain trading environments require evaluation methodologies that capture their inherently competitive nature.

In this paper, we argue that a well designed and well publicized suite of open competitions, where teams from around the world are pitted against one another, can go a long way in helping identify promising automated (or semi-automated) supply chain trading techniques. Specifically, this paper introduces the *Supply Chain Management Trading Agent Competition* ("TAC-SCM"), which we designed in 2003 in collaboration with members of the Swedish Institute of Computer Science [3, 21]. We show how the competition captures the inherent complexity, stochasticity and competitive nature of supply chain trading environments. The first edition of the competition was held in August 2003 and attracted a total of 20 teams from nine different countries. We discuss the approaches and strategies used by the competing teams, results from the competition and lessons learned. At the time this paper was brought to press, the next edition of the competition was already under way with over 30 entries, a sign of the rapidly growing popularity of the competition and its perceived research value.

The rest of this paper is structured as follows. In the next section we present a brief review of the literature. Section 3 and 4 discuss our objectives in designing TAC-SCM. They also provide an overview of the game. Challenges posed by the game are addressed in Section 5, including a discussion of the complexity associated with different game abstractions. We also cover the challenges associated with the inherently stochastic and competitive nature of the game and with the incomplete information with which each agent has to operate. Results of the trading agent competition are presented in Section 6. The following section looks at some of the approaches and strategies employed by the various agents that took part in the 2003 TAC-SCM Tournament. We conclude with a discussion of lessons learnt and future work.

2. LITERATURE REVIEW

To the best of our knowledge, little research has been reported on coordinating sourcing ("who to order from"), procurement ("how much to order and when"), finite capacity production and customer bidding. One notable exception is the work of Kjenstad and Sadeh on multi-tier, finite-capacity capable-to-promise and profitable-to-promise supply chain functionality, where different techniques are compared in the context of an agent-based simulation platform [15, 22]. In general, work reported to date has looked at somewhat simpler abstractions, such as problems focusing on production decisions [8, 30] or procurement decisions [12, 16]. Although it is valuable, much of the research conducted so far has ignored or considerably simplified key temporal and capacity constraints [4, 5, 27].

Thomas and Griffin provide an excellent overview of coordinated supply chain management models [26]. Bassok and Akella have developed a single-period model that considers both procurement and production costs [6]. They assume a single-machine production scenario with a single critical raw material and multiple products with stochastic demand. The authors demonstrate the benefit of jointly optimizing production and procurement decisions. Sun and Sadeh consider the integrated sourcing/procurement and production problem faced by a manufacturer that has to satisfy a set of customer orders, each with its own due date, tardiness penalty and component requirements [23]. Supplier procurement bids for each component may differ in both price and delivery dates. The manufacturer's objective is to identify a production schedule and a selection of procurement bids that minimize the sum of its procurement costs and tardiness costs.

Lau and Lau [16] have looked at the procurement and sourcing problem faced by a buyer that has to choose between a low cost/long leadtime supplier and a high cost/short leadtime one. Demand is deterministic while procurement leadtimes are stochastic. The authors derive an analytical expression for a total cost function that takes into account optimal order quantity, optimal reorder point, and the ratio of orders placed with each of the suppliers. Gurnani, et.al. [12] consider a procurement scenario where a single manufacturer produces a single product composed of two components. Three suppliers are available, one for each component and a third that sells both components in pairs. They investigate the problem of how much to buy from each supplier with a view to minimizing procurement, storage, and backlog costs.

While the production planning literature has largely ignored the bidding and bid selection dimensions of supply chain trading, the e-business research community has primarily focused on these issues, while ignoring capacity and temporal constraints. Babaioff and Nissan [4] have presented bidding protocols that result in the efficient allocation of goods among partners in a linear supply chain. Their approach organizes auctions in terms of a series of production markets and derives a double auction mechanism that has the attractive feature of being Incentive Compatible (IC), Individually Rational (IR), and of resulting in a Balanced Budget (BB). Babaioff and Walsh [5] have extended these results to accommodate a broader class of supply chain topologies while still retaining the IC, IR, and BB properties and while yielding high allocation efficiencies. Walsh and Wellman [27] have proposed a family of decentralized protocols, based on the task dependency network model, to negotiate supply contracts. All these models operate under a single period, and assume only single unit transactions. They ignore capacity and temporal constraints. Negotiation between supply chain entities is restricted to a single attribute, namely price. Experience has shown that suppliers generally prefer to compete on more than just price – competing solely based on price just yields lower profit margins. Accommodating more general and more realistic forms of supply chain negotiation requires models that also capture the responsiveness of supply chain entities, their ability to meet specific delivery dates, to accommodate surges or drops in demand, to meet different sets of quality requirements, etc. Ultimately, negotiating within these richer frameworks benefits both suppliers and customers. Suppliers get to differentiate themselves across a broader range of attributes and customers can select suppliers based on a richer set of considerations.

An early example of such a model is the approach taken by Sadeh, et al. [21] in MASCOT, where bids from supply chain entities are evaluated based on both cost and delivery dates. Zeng and Sycara [29] also discuss a real time supply chain formation model where supply chain entities negotiate based on costs and responsiveness. Also relevant is the work of Collins, et al. in MAGNET [9], which provides a framework where agents negotiate the coordination of tasks constrained by temporal and capacity considerations.

There have been a number of simulation tools developed to analyze supply chain performance. They include both research prototypes and commercial software aimed at supporting supply chain re-engineering efforts (e.g. [25]). Other tools have been developed to study more specific research issues such as the impact of different information exchange protocols (e.g. [24]) or the well-known "bullwhip effect" [17, 18]. Most prominent among these latter tools are software simulations of the "Beer Game" [7], which has been an integral part of many operations management curricula. Originally developed at MIT in the mid 1960s as a board game, it allows players to experience the well known 'bullwhip effect', namely the amplification of downstream demand fluctuations as they propagate upstream to distributors, manufacturers and their suppliers. The Beer Game supply chain consists of four stages (producer, distributor, wholesaler, and retailer) with a single player operating at each stage. To the best of our knowledge, none of these many supply chain simulation testbeds have been designed to capture the competitive nature of supply chain trading environments, where multiple companies vie with one another for both customer orders and supplies. On the other hand, earlier computerized trading games, such as the Santa Fe Double Auction Tournament conducted in 1990 by Friedman and Rust [19], have focused on significantly simpler scenarios than those likely to be found in practical supply chain trading environments (e.g. single good, no production/capacity constraints, etc.).

3. TAC-SCM: MOTIVATIONS AND DESIGN OBJECTIVES

The Trading Agent Competition (TAC) has been an annual event since its inception in 2000 [28]. TAC provides a forum for researchers studying trading agents and focuses their energies on a common problem. Between 2000 and 2002, TAC revolved around a Travel Shopping Game, now referred to as "TAC Classic" [28]. In this game, each trader is a travel agent responsible for organizing trips to Boston for eight clients. The travel agent's objective is to maximize the total satisfaction of its clients relative to the money spent procuring plane tickets, hotel rooms and tickets to entertainment events in different markets. Each client has a different set of preferences (e.g. travel dates, room requirements, etc.).

In Summer 2002, encouraged by the success of their competition, TAC organizers solicited proposals for the addition of a new game that would introduce new challenges to the trading agent community and be representative of a somewhat more strategic segment of the economy. The Supply Chain Management Trading Agent Competition, also known as TAC-SCM, was conceived by the authors in response to this call for proposals. It was eventually selected from several other proposals as the new TAC game to be introduced in the Summer 2003. It was further refined during the Fall 2002, in collaboration with members of the Swedish Institute for Computer Science (SICS). As researchers who had been working for over ten years in supply chain management [22,23, 24, 25], our motivation for developing a supply chain trading game was twofold. First, we wanted to tap the talent and insight of the trading agent community and see how its technologies could be brought to bear in the context of supply chain management scenarios. Second, through our own research (e.g. [15,22]), we had come to the realization that traditional evaluation methodologies, where one compares a technique with a pre-specified set of competing approaches, were insufficient to fully validate supply chain trading techniques. A game, we believe, is an ideal setting to capture the competitive nature inherent to realistic supply chain trading scenarios.

More specifically, a good candidate game needed to address the following issues:

- *Strategizing.* The game scenario had to leave room for supply chain trading agents to strategize about their own choices and those of the competition.
- *Uncertainty and incomplete information.* Supply chain management is in great part about managing uncertainty and risks and working with incomplete information. Uncertainty comes in many shapes and forms and includes dealing with changes in both upstream and downstream market conditions as well as internal contingencies. As they make decisions, supply chain entities only have a partial view of the state of the world (e.g., information about competitors' inventory positions and order books, available supplier capacity, etc.) The inclusion of these elements was thus essential.
- *Realism*. The applied nature of supply chain management dictated that any effort in this area would need to be made with an eye on actual practice. If the lessons learned from

designing successful agents were to find wider acceptance, scenarios depicted in the game would need to be plausible. We use the term "plausible" deliberately, as it was not our intention to restrict the game to depicting current supply chain scenarios. To the contrary, our objective has been to stretch the state of the art and examine future possible supply chain trading practices.

- *Generality*. The challenges introduced in the game needed to be representative of a broad class of supply chain situations.
- *Simplicity*. To be successful and encourage participation, it was critical that the game be simple enough to enable a good number of competitors to submit entries. This, of course, to a degree, runs counter to the aforementioned realism requirement. Getting the right balance was one of the challenges in designing the game.

4. TAC-SCM: GAME OVERVIEW

The TAC-SCM game was designed to promote the development of supply chain trading agents that are capable of effectively coordinating their sourcing, procurement, production, and customer bidding decisions. The game revolves around a personal computer (PC) assembly supply chain consisting of six competing PC assemblers (or "manufacturers"), their component suppliers and their customers [21]. (See Figure 1.) The PC assemblers bid on *requests for quotes* (RFQs) for PCs of different types from customers. When they win a bid, they are responsible for procuring the necessary components, assembling the PCs and delivering them in a timely fashion. If they fail to meet the delivery date requested by their customer, they incur a penalty proportional to the delay. Each type of PC requires four key components: CPU, motherboard, memory and hard drive. CPUs and motherboards are available in two different product families, Pintel and IMD. A Pintel CPU only works with a Pintel motherboard, while an IMD CPU can only be incorporated in an IMD motherboard.

Figure 1. Overview of the TAC-SCM Game

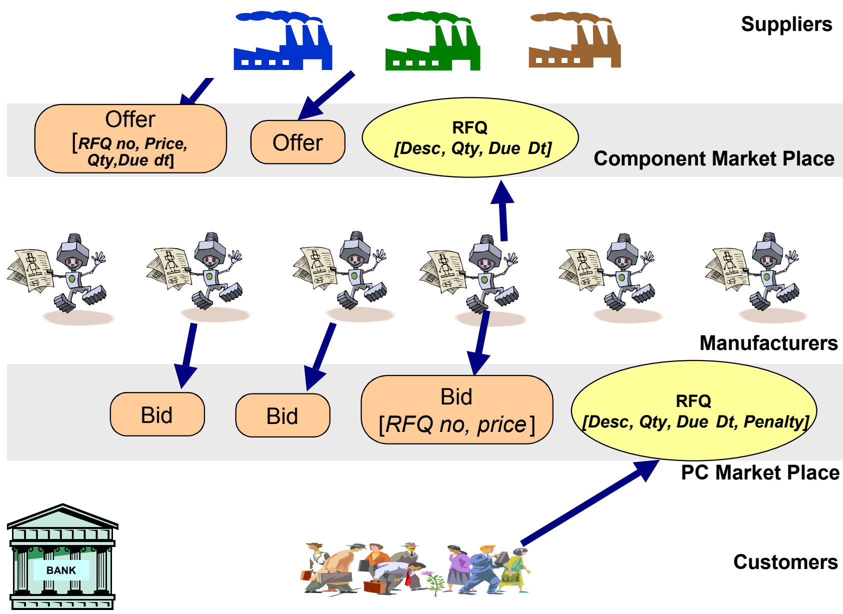

Each component is available in two different specifications and can be procured from two potential suppliers. (See Table 1.)

Table 1. Components in TAC-SCM

Components	Suppliers	Component specification
CPU	Pintel	2 Ghz
		5Ghz
	IMD	2 Ghz
		5Ghz
Motherboard	Basus	For Pintel CPUs
		For IMD CPUs
	Macrostar	For Pintel CPUs
		For IMD CPUs
Memory	MEC	1GB
		2GB
	Queenmax	1GB
		2GB
Hard Drive	Watergate	300 GB
		500 GB
	Mintor	300 GB
		500GB

Each PC assembly agent is endowed with an identical assembly cell capable of assembling any type of PCs. The cell operates a fixed number of hours per day (its capacity) and different PCs have different assembly times. PC assembly agents can store both components and finished PCs, enabling them to procure components and assemble PCs ahead of time, whether for orders they have already secured or in anticipation of future demand.

Each agent has a bank account from which it draws money when it purchases components and where it receives money for products it delivers to customers. Penalties for missing delivery dates are also taken from the agent's bank account. Agents are allowed to borrow money from the bank. They are charged interest when they owe money and credited interest when they have a positive balance. The aim of the competition is to end up with as much money as possible in one's bank account. This, in turn, requires securing a good number of customer orders at a high enough price and components at a low enough price to make a profit, while meeting one's delivery commitments.

Each competitor enters an agent that is responsible for the following tasks:

- Negotiate supply contracts
- Bid for customer orders
- Manage daily assembly and delivery activities

The game is played over a period of 220 simulated days (each day lasting 15 seconds) and the above three tasks are performed by the agents daily. At the end of the game the agent with the most money in the bank is declared the winner. Additional details about the game specifications can be found in [1].

Figure 2 illustrates the timing of the various events that occur during the game. At the start of each day, an agent receives:

Figure 2. Key Daily Events

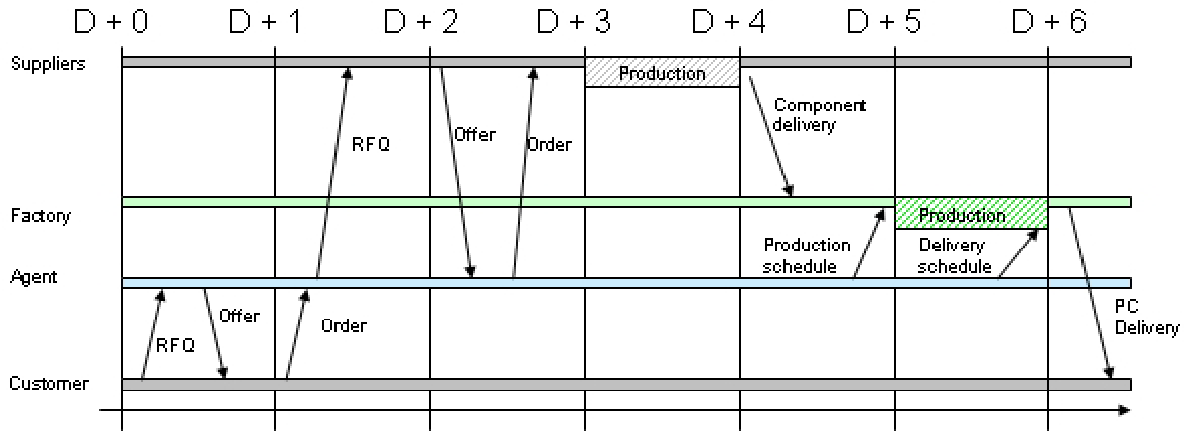

- <u>From Customers:</u>
  - Requests For Quotes (RFQs) for different types of PCs
  - A list of orders won by the agent following bids (or offers) it submitted to customers the previous day
- <u>From Suppliers</u>:
  - Quotes (or offers) for the delivery of components in response to RFQs the agent had sent the suppliers the previous day
  - Delivery of supplies it had ordered earlier. The supplies (components) can be used for production the day after the delivery.

Note that there is at least a two-day lag between the time an agent submits an RFQ to a supplier and the time the supplies can be used to assemble a PC. Lags such as this, along with the competitive nature of the game, force agents to plan ahead of time and create incentives for taking risks (e.g., incentives to order components in anticipation of future demand and to hedge against possible supply shortages).

- <u>From the Bank</u>:
  - Statement with the agent's account balance.

- From its Factory:
    - Inventory report (quantity of components and finished PCs available)

During the course of the day, the agent has to determine (1) which customer RFQs to bid on, if any, and the particulars of its bids; (2) which components to procure and the specifics of RFQs to be sent to different suppliers (including requested quantities and delivery dates); (3) which supplier offers to accept, if any; (4) which PC orders to assemble for, subject to availability of supplies – the list of PCs to be assembled on a given day is referred to as the *production schedule*; and, (5) which assembled PCs to ship to which customers – the list of PCs to be shipped to customers on a given day is referred to as the *delivery schedule*

4.1. Negotiating Supply Contracts

In order to procure supplies an agent issues RFQs to potential suppliers. An RFQ specifies the type of component required, quantity and due date. (See Figure 3.)

Figure 3. Format of Procurement RFQs and Supplier Offers

Format of an RFQ to a Supplier
RFQ::= <RFQ-Id, Component-type, Quantity, Due-Date>

Format of a Supplier offer:
Offer::= <Offer-Id,RFQ-Id,Component-type,Quantity,Price>

Based on its existing and projected inventory (available-to-promise quantities), a supplier replies to the RFQ by issuing one or more bids (or offers). Bids include the price at which the supplier offers the components, a quantity, and promised delivery date. Specifically, if the supplier can satisfy the RFQ in its entirety, a single bid is returned with the full RFQ quantity, the requested delivery date and a price. On the other hand, if available-to-promise inventory on the requested delivery date is insufficient, the supplier responds by issuing up to two amended offers:

- A *partial offer* is generated if the supplier can deliver only part of the requested quantity on the due date. The offer is for the fraction of the RFQ's quantity that the supplier can deliver in time.
- An *earliest complete offer* is generated to reflect the earliest day (if any) on which the supplier can deliver the entire RFQ quantity.

Amended offers are mutually exclusive. In other words, an agent can only select one or the other (or neither). All offers made by suppliers are only valid for a day, after which quantities are freed for other offers to be made. Available-to-promise quantities are constantly updated by suppliers to reflect the bids they have submitted, the orders they have received as well as their actual capacities, which randomly fluctuate from one day to the next. Because a supplier's capacity fluctuates, it may not always be able to meet its delivery commitments. This, in turn, requires agents to monitor the performance of their suppliers and possibly adjust their supply delivery expectations. Finally, suppliers adjust unit prices in their offers subject to how much available-to-promise capacity they have left. As remaining supplier capacity goes down, component prices go up. Additional details on the finite capacity supplier model used in the game, including the use of a reputation mechanism to discourage agents from swamping suppliers with RFQs they do not need, can be found in [1].

4.2. Bidding on Customer Orders

Customers exhibit demand by issuing RFQs to the agents. Each customer RFQ includes a type of PC, a quantity and a due date. (See Figure 4.)

Figure 4. Format of a Customer RFQ

Format of a Customer RFQ

```
RFQ ::= <RFQ-Id, PC-type, Quantity, Due-Date, Penalty,
Reserve-Price>
```

The penalty specified in the RFQ is the daily price that the agent has to pay for missing its delivery commitment - up to a maximum of 5 days, at which point the order is canceled and the maximum five-day penalty levied against the agent. The reserve price is the maximum unit price that the customer is willing to pay for the item specified in the RFQ. The customer discards all bids that are priced above the reserve price, or that fail to meet the RFQ's quantity or due date. The agent with the lowest bid price wins the order. In case of a tie the customer makes a random choice among the tied bids.

On any given day, a number of RFQs are issued by customers for a number of different types of PCs with delivery dates distributed over a period ranging between *current_date + 3 days* and *current_date + 12 days*. Details on the probability distributions used in the game also can be found in [1].

5. CHALLENGES POSED BY THE GAME

The objective of each agent is to have accumulated as much money as possible in the bank by the end of the game subject to constraints of finite capacity and component availability. Each day, agents are faced with an exponential number of bidding, sourcing, procurement and scheduling options. The game is further complicated by uncertainty in both supply and demand and by strategizing by competitor agents. To maximize their overall profits, agents generally need to submit competitive bids to customers and to secure supplies in a timely and cost effective manner. Their fixed capacity limits the number of customer bids they can handle without defaulting on delivery commitments. The daily choices faced by an agent can be viewed as four interacting sets of decisions (or "Problems"):

- *Problem 1*: Which customer RFQs to respond to and with what bids
- *Problem 2*: Which combinations of supplier bids to accept

- *Problem 3*: How to schedule assembly operations
- *Problem 4*: Which finished PCs to ship to which customers

Each of these problems is challenging in its own right:

- *Problem 1*. Even under deterministic conditions and with full knowledge of what the competition might do, Problem 1 requires examining an exponential number of RFQ combinations one could reply to, determining the one that will maximize the agent's overall profit while respecting capacity and supply availability constraints. In practice, an agent does not know what its competitors will do. Submitting high-price bids could result in higher profits but also lowers the chance that the bid will be accepted by the customer, as competing agents become more likely to submit lower price bids. Submitting too many low-price bids could lead to a situation where the agent wins more orders than it can satisfy, given its limited assembly capacity and available supplies. In general, to be competitive, agents can not afford to just look at their own constraints, they need to maintain models of what the competition is likely to do and attempt to forecast future supply and customer market conditions.
- *Problem 2*. This problem is similarly challenging. It too presents agents with an exponentially large number of options to consider. Procurement has to be well synchronized. Often, there is no point in acquiring a subset of components if it will take several more weeks to acquire matching components required to fully assemble some PCs. For each component type associated with an order, an agent may have to combine multiple supplier bids, as the quantities of individual bids may not be sufficient. Like Problem 1, Problem 2 requires taking into account uncertainty about future market fluctuations both at the supplier end and at the customer end. This is in part because

agents may not want to limit their procurement activities to requirements associated with their current order book. Instead, they may want to stock up on components to be able to offer shorter delivery times to prospective customers. In addition, they may want to hedge against future supply shortages, whether due to production delays at one or more suppliers or due to strategic decisions made by competitors (e.g., a competitor could try to pre-empt others from acquiring some components).

- *Problem 3.* Under deterministic conditions, Problem 3 can be shown to be a generalized version of the single machine total weighted tardiness scheduling problem, a well-known NP-hard problem [10]. In practice, uncertainty about supply deliveries and future demand adds to the challenges associated with this problem. An agent may want to secure supplies and schedule assembly activities a little ahead of time to hedge against possible delays in supply deliveries. It may also want to do so to free future capacity and give itself extra flexibility to accommodate future possible customer RFQs.
- *Problem 4.* While agents can be expected to assemble many of their PCs in response to specific orders ("make-to-order" or "assemble-to-order" practices), uncertainty in the game creates an incentive for producing a little extra ahead of time (e.g., to hedge against delays in supply deliveries or to be competitive on customer RFQs with particularly short leadtimes). Each day an agent has the flexibility of reallocating PCs to different customers prior to shipping them. Determining the optimal allocation of PCs to customers requires looking at the due dates and penalties associated with each order and deciding which orders, if any, to sacrifice or to delay based on available finished goods inventory and projections of when additional PCs of different types will be ready for shipping.

Beyond their individual difficulty, the above problems interact with one another. Ultimately, a competitive agent is not one that just does a good job at solving each of these problems in isolation. It also must be good at closely coordinating each of these decisions. Bidding on customer RFQs has to be well coordinated with scheduling and procurement decisions. If scheduling falls behind, bidding may need to become more selective (e.g., by increasing bid prices or reducing the number of RFQs the agent responds to). If customer bidding is more successful than expected, procurement may need to be ramped up and scheduling may be faced with tough choices. In [23], Sun and Sadeh look at a deterministic variation of Problems 2 and 3 and show that a technique that concurrently optimizes decisions in Problem 2 and 3 will do significantly better than approaches that take a more decoupled view of these problems. The same generally holds for TAC-SCM as a whole. Heuristic solutions that do a good job at coordinating the four types of problems identified above can be expected to do significantly better than solutions that rely on more simplistic views of the interactions between these problems.

Through the complexity of the sub-problems it entails, the uncertainty associated with both customer and supplier markets and the opportunities for strategizing, TAC-SCM encapsulates many of the tradeoffs one can expect to find in typical supply chain trading environments. As in these environments, the size of the problems faced by an agent, the pace at which decisions have to be made (15-second days) and the multiple sources of uncertainty preclude the development of any type of “optimal solution”. By requiring agents to compete in a number of games with randomly generated market conditions, the tournament ensures that agents are extensively evaluated before being allowed to move to the next round. Simple-minded agent strategies that might do well under particular situations will generally fail miserably under others. To be competitive, agents have to exhibit strategies that dynamically adapt to the situation at hand, making the game one that provides exciting opportunities for developing innovative solutions that dynamically adjust agent planning and trading behavior.

6. THE 2003 TAC-SCM TOURNAMENT

Twenty teams registered to participate in TAC SCM 2003. The participants featured organizations from nine countries, as shown in Table 2.

Table 2. TAC SCM 2003 Participants

AGENT	TEAM LEADER	AFFILIATION
TAC-o-matic	Jim Holmström	Uppsala Universitet
UMBCTAC	Rong Pan	University of Maryland, Baltimore County
PSUTAC	John Yen	Pennsylvania State University
PackaTAC	Peter Wurman	North Carolina State University
Botticelli	Amy Greenwald	Brown University
Deep Maize	Satinder Singh	AI Lab, CSE Division, U. Michigan, Ann Arbor
Jackaroo	Dongmo Zhang	University of Western Sydney
Sirish	Sirish K. Somanchi	North Carolina State University
Mertacor	Kyriakos Chatzidimitriou	Aristotle University of Thessaloniki
Argos	Taner Bilgiç	Bogazici University
Zepp	Ovidiu Trascu	Teamnet - Politehnica University of Bucharest
HarTAC	Wilfred Yeung	Harvard University
RonaX	Wolfram	Xonar GmbH
DummerAgent	Yilanya John-Alex	University of Texas
RedAgent	Doina Precup	McGill University
MinneTAC	John Collins & Maria Gini	University of Minnesota
TacTex	Peter Stone	The University of Texas at Austin
DAI_hard	Arjita & Sabyasachi	University of Tulsa
Socrates	Maria Fasli	University of Essex
Whitebear	Ioannis A. Vetsikas	Cornell University

On average, each team typically involved somewhere between three and five members. The 2003 TAC-SCM research community totaled around 80 people. While we have no hard figures, we estimate that each team typically spent about six months preparing for the competition with team members devoting on average 25% of their time to the effort.

The TAC team at the Swedish Institute of Computer Science, which co-organized the event along with the authors, set up and administered the game servers (www.sics.se/tac). Two servers were used to play the entire competition. Teams ran their agents from their home facilities by connecting to the SICS TAC servers. The preliminary rounds (one qualification and two seeding rounds) of the competition took place between July 7 and 18, 2003. The final rounds were held on August 11 to 13 as an exhibition at the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI 2003) in Acapulco, Mexico.

The final rounds took place over three consecutive days, featuring quarter-finals on Day 1, semi-finals on Day 2 and finals on Day 3. The motivation for organizing the competition around multiple rounds was to allow for a sufficiently large number of games to be played in each round and to slowly weed out less competitive agents with the objective of allowing the six best teams into the finals. The semi-finals featured twelve agents, which were broken into two groups that each played a total of nine games, each using one of the two SICS TAC servers. The top three teams in each group then proceeded to the finals, where sixteen games were played, using both servers.

Standings at the end of the semifinals and finals are shown in Table 3a and 3b, respectively. The value in the third column is the average profit accumulated by an agent over the course of a game. The presence of many negative scores reflects the high level of competition among agents in the final rounds. With the exception of TAC-o-matic, each of these agents had positive average scores in the quarter-finals. However, as the better agents were brought to compete against one another in the final rounds, interactions between their strategies caused a number of them to start losing money. Also, while teams were not allowed to modify their agents during a given round, changes were allowed as teams moved from one round to the next. This placed some agents in competitive situations they had never faced in earlier rounds.

Table 3a. TAC-SCM Semi-Finals Results

GROUP 1 RANKINGS	AGENT	AVERAGE PROFIT IN MILLIONS $
1	Red Agent	25.09
2	DeepMaize	15.28
3	PackaTAC	8.697
4	PSUTAC	-1.555
5	TAC-o-matic	-13.50
6	HarTAC	-32.95
GROUP 2 RANKINGS	**AGENT**	**AVERAGE PROFIT IN MILLIONS $**
1	Botticelli	-4.831
2	whitebear	-9.579
3	TacTex	-15.54
4	Sirish	-20.21
5	MinneTAC	-24.98
6	UMBCTAC	-29.91

Table 3b. TAC-SCM 2003 Final Results

RANKINGS	AGENT	AVERAGE PROFIT IN MILLIONS $
1	Red Agent	11.610
2	DeepMaize	9.473
3	TacTex	5.017
4	Botticelli	3.330
5	PackaTac	-1.680
6	Whitebear	-3.453

7. AGENT DESIGNS IN THE 2003 TAC-SCM TOURNAMENT

During the preliminary rounds, a number of TAC-SCM 2003 agent designers opted for an early procurement strategy intended to exploit the lower prices of supplies at the start of the game. As indicated in Section 4.1, component unit prices quoted by suppliers are a function of remaining available-to-promise capacity. Because at the start of the game suppliers still have all their capacity available, they offer components at a discount – possibly as much as a 50% discount, creating a strong incentive for agents to procure early. (See [1] for details.) On the flip side, early procurement comes with the burden of higher inventory costs (in the form of interests) and inflexibility with respect to customer demand. Early experimentation by a few agents with this strategy during the preliminary rounds proved rather effective, leading more agents to adopt it over time. While initial adoption of this strategy (termed *Day 0 procurement*) by a few agents had little impact on the overall dynamics of the game, adoption by an increasing number of agents resulted in a race condition of sorts, akin to hoarding. This is because agents that did not procure early were left with little hope of procuring supplies later in the game, as supplier capacity was committed very early on. This equilibrium is not very desirable and can become quite destructive, as agents pay hefty inventory costs and gamble on the materialization of future customer demand. Two teams, *DeepMaize* and *RedAgent*, managed to counter the destructive influence of this Day 0 procurement strategy, using very different approaches.

7.1. DeepMaize's Day 0 Fake Strategy

The performance of *DeepMaize* in the early rounds was modest (10th in the second seeding rounds and 4th in a group of nine in the quarter-finals). *DeepMaize* by design ignored the singularities caused by the start and end of the game and performed best once the game had settled in a steady state. Much effort by the design team had gone into effective demand forecasting and consequently it was a key competency of the agent [14]. Widespread Day 0 procurement meant that the game took longer to achieve a steady state, thus affecting the performance of *DeepMaize*. To overcome this, the *DeepMaize* team devised a radical strategy aimed at neutralizing the effectiveness of Day 0 procurers.

Day 0 procurers handled the singularity of Day 0 by hardwiring their large procurement behavior on Day 0, with agents reverting back to steady state operation on subsequent days. Starting in the semifinals, the *Deepmaize* team exploited this lack of flexibility by sending very large Day 0 RFQs to suppliers for each of the sixteen components for delivery on Day 30. The quantity requested in each RFQ was large enough to require a supplier to dedicate its entire capacity for 170 days to satisfy it. Accordingly, supplier responses to *DeepMaize*'s Day 0 RFQs could only be in one of three possible forms:

- Two amended offers: one *partial offer* indicating how much the supplier could deliver by Day 30 and one *earliest complete offer* indicating by when the supplier could deliver the balance of the RFQ's quantity.
- If the supplier had no available-to-promise capacity left prior to Day 30, a single amended offer in the form of an *earliest complete offer* with the date by which the full RFQ quantity could be delivered.
- Possibly no offer, if the supplier did not have enough capacity left.

DeepMaize had no intention of ever purchasing the entire quantities requested in its RFQs – hence, the term "Day 0 fake strategy." Instead, it would ignore *earliest complete offers* and only accept *partial offers* for Day 30, thereby keeping enough flexibility to adjust its procurement strategy during the course of the game – using its demand monitoring and forecasting functionality. However, the effect of *DeepMaize*'s Day 0 fake strategy was rather disruptive for Day 0 procurers. This is because suppliers reserve capacity when they respond to RFQs. This capacity is held for a day. If the supplier does not hear back from the agent the capacity is subsequently released. This meant that all the components offered to *DeepMaize*, including large quantities in *earliest complete offers*, were unavailable to Day 0 procurers and remaining quantities were priced higher. This simple counter worked because Day 0 procurers relied on hardwired Day 0 strategies and automatically reverted back to their default procurement strategy on Day 1. Had Day 0 procurers come back on Day 1 and placed similar orders, they often would have found that supply capacity had increased again and that prices had dropped, as *DeepMaize* had only ordered a small fraction of the quantities requested in its Day 0 RFQs. This strategy allowed *DeepMaize* to ensure a more steady stream of supplies throughout the game and played to its strength of demand forecasting. In the semifinals where this strategy was first deployed, *DeepMaize* ended up with second place, behind only *RedAgent.*

7.2. RedAgent

RedAgent approached the problem of coordinating sourcing, procurement, production and bidding under uncertainty by developing a novel multi-agent market architecture [13]. The architecture features internal markets where local micro-agents compete for raw material, capacity, and customer orders. The micro-agents use different parametric heuristics to evaluate trading options [18]. Parameters are set based on a combination of data learned from earlier games and data collected in the current game (e.g., winning customer bid prices for different types of PCs on the previous day). This architecture provides for a particularly tight coordination between procurement and bid submission activities.

Through their interconnections, the micro-markets enable *RedAgent* to make external procurement decisions that reflect the latest customer market developments (e.g., success in capturing new customer orders). Similarly, these internal markets ensure that *RedAgent's* bidding strategies are responsive to the agent's current procurement situation (e.g., success in procuring different sets of components will lead the agent to bid more aggressively on customer RFQs for PCs that use these components). This tight coupling proved crucial to the success of the agent in the finals. In contrast to Day 0 procurers, *RedAgent*'s architecture provides for a much more gradual approach to procurement, enabling it to adjust over time. In the quarter-finals, this strategy was impeded by Day 0 procurers, leading *RedAgent* to finish third in its group after *PackaTAC* and *PSUTAC*.

As *DeepMaize* introduced its Day 0 fake strategy in the semifinals, it freed inventory and enabled *RedAgent*'s strategy to perform much more effectively, as illustrated by a number of performance metrics. Specifically, Figure 5 shows the aggressiveness of agents in terms of bidding on customer RFQs.

Figure 5. Percentage of Customer RFQs Bid Upon by Each Agent and Percentage of Bids Won in the Finals

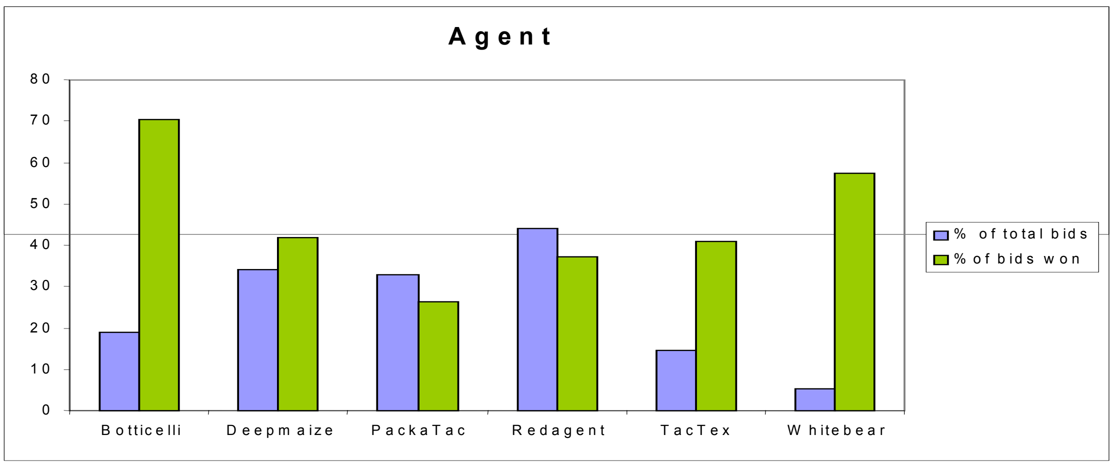

RedAgent dominates this statistic, bidding on 44% of customer RFQs with *DeepMaize* coming in second with 34%. At the bottom is *Whitebear*, an agent that relied on a Day 0 procurement strategy, procuring almost all its components on Day 0. Because of this strategy, *WhiteBear* was most affected by *DeepMaize*'s Day 0 fake strategy, and ended up spending most of the games waiting for supplies. The high percentage of bids won by *Botticelli* and *Whitebear* simply reflects the fact that these agents generally submitted bids with lower prices with Whitebear ultimately losing money (See Figures 6 and 7). Overall, *RedAgent*'s market share was the highest, slightly ahead of *DeepMaize*'s. (See Figure 8). While on average *TacTex* was able to make a larger per unit profit than *RedAgent* and

DeepMaize, its marketshare was less than half that of these two agents, which led it to finish third, with about half the money accumulated by either *RedAgent* or *DeepMaize*.

Figure 6. Average Revenue Per Game Day During the Finals

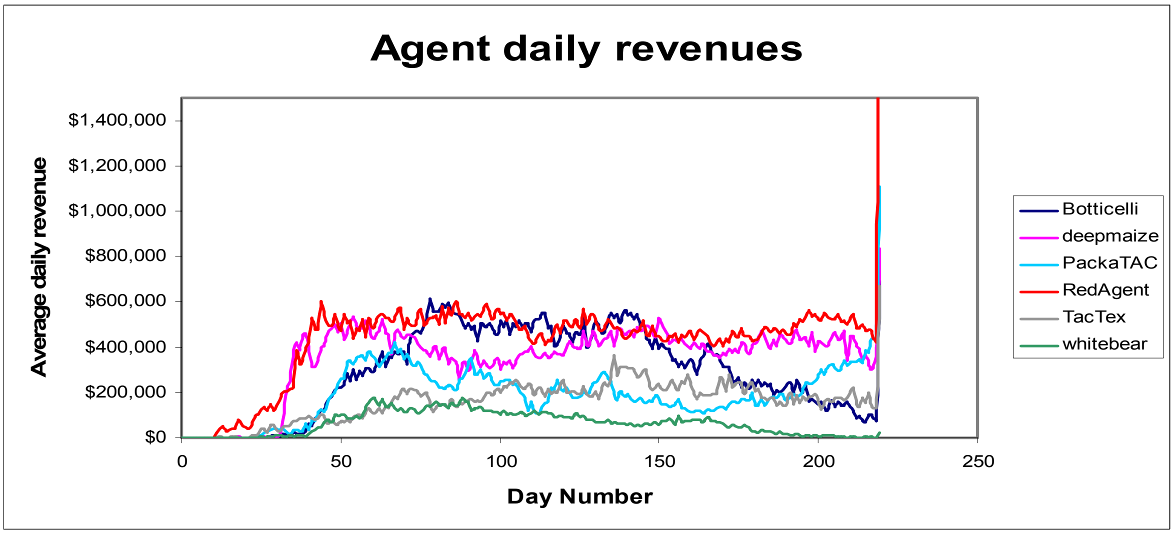

Figure 7. Average Per Unit Profit Per PC Sold During the Finals

Figure 8. Average Selling Price of PCs Sold by Agents in the Finals

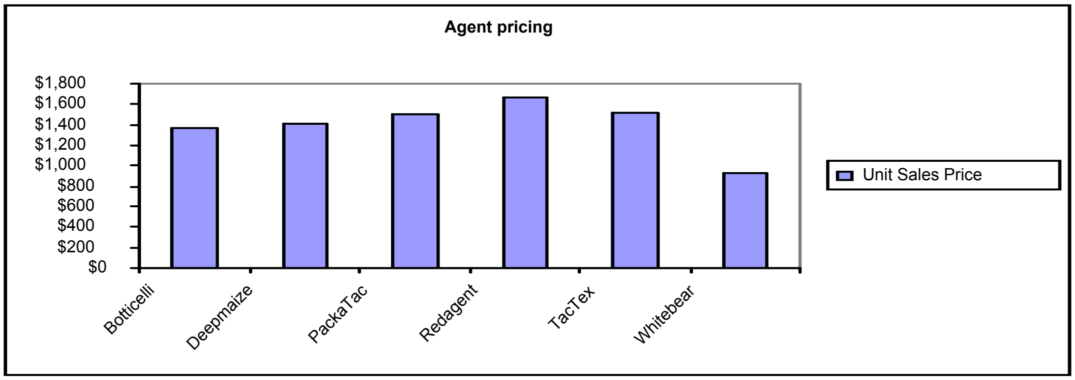

Figure 9. Agent Market Share in the Finals

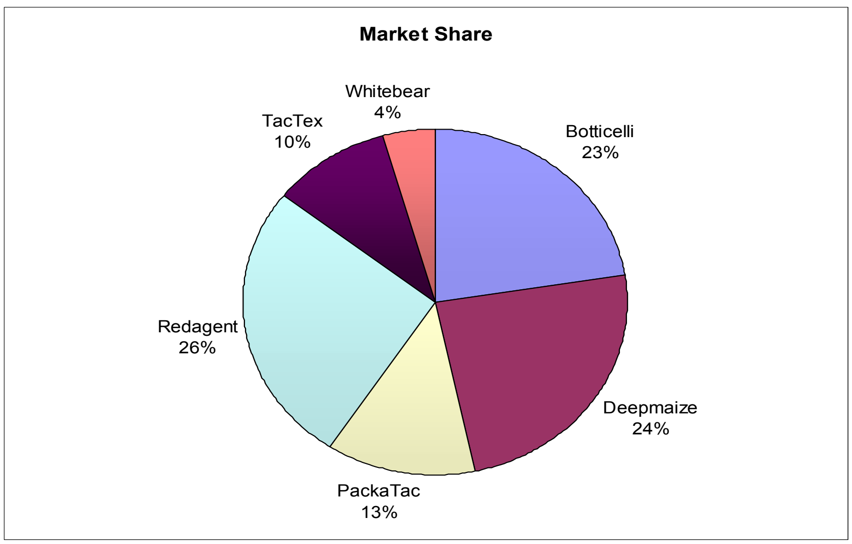

Figure 10. Average Component Inventory Levels During the Finals

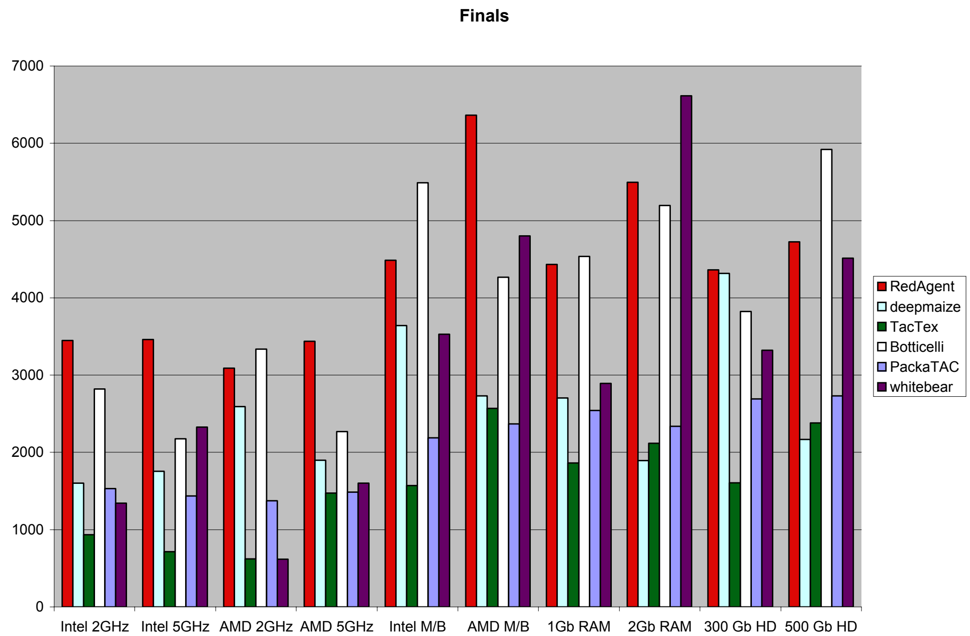

Figure 11. Average PC Inventory Levels During the finals

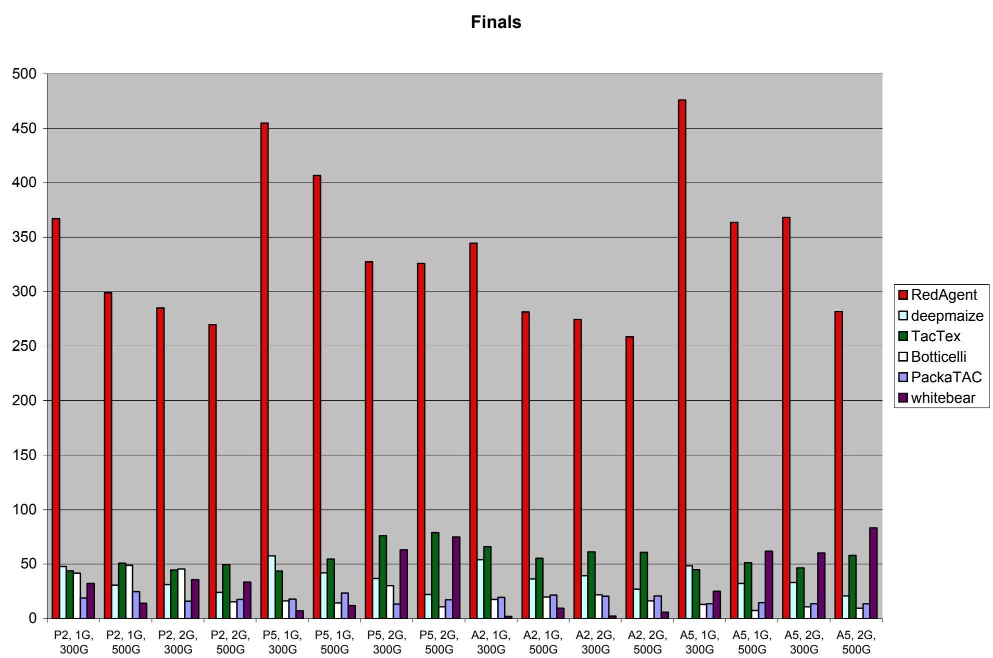

The large number of customer RFQs that both *RedAgent* and *Deepmaize* were able to bid on indicates that the two agents that dominated the finals were also the ones that were the most responsive. The two agents accomplished this in different ways. *RedAgent* carried the highest inventory among the six finalists, allowing it to bid on a number of orders with short deadlines (See Figure 10 and 11) [13] . *Deepmaize* on the other hand used its superior demand forecasting to achieve similar responsiveness with somewhat less inventory [11].

RedAgent's tight co-ordination between procurement and order bidding and the adaptability of its micro-markets to prevailing market trends was particularly dramatic during the start and end phases of the game. With multiple agents still using variations of Day 0 procuring strategies, most agents spent the early period of the games waiting for components, unable to produce PCs. The wait was often quite long as an agent requires all four types of components before it can start assembling PCs. Table 4 shows when agents first began production during each of the sixteen games in the finals.

Table 4. First Day of Production During the Finals

Game	RED AGENT	DEEPMAIZE	PACKATAC	BOTTICELLI	WHITEBEAR	TACTEX
1423	5	34	22	32	59	11
1424	5	35	32	44	72	10
1425	6	23	36	49	220	36
1426	5	24	21	21	39	9
1427	8	22	33	40	220	45
1428	6	129	27	36	55	9
1429	6	39	22	40	112	26
1430	6	32	37	54	220	40
1264	7	33	25	56	220	8
1265	6	23	37	49	220	43
1266	6	35	38	40	78	11
1267	6	33	42	50	69	9
1268	6	32	40	25	19	24
1269	6	220	23	57	37	9
1270	6	10	25	54	38	41
1271	6	27	37	56	38	24

Thanks to relatively modest component orders spread throughout the game, *RedAgent* had a significant head start, reliably beginning assembly by Day 7 - the latest start being on Day 8 and the earliest on Day 5 . The average revenue earned on each of the game days during the finals is shown in Figure 6. On average, *RedAgent* generated $632,000 in daily sales with virtually no competition during the early part of the game (the first 27 days). Figures 7, 8 and 9 depict the average unit profit margin, unit selling price and market share further illustrating the success of the *RedAgent* architecture. A similar, though more dramatic, cornering of the market can be seen at the very end of the game, as shown in Figure 8. *RedAgent* picks up on average about $4,000,000 in revenue at the end, $2,000,000 more than the nearest competition. This is simply accomplished by progressively lowering its target finished goods inventory so that it reaches zero at the end of the game. *RedAgent's* start game and end game performance, as well as its steady state performance, attest to its ability to adapt and exploit profitable segments of the market. Further, they point out an important distinction between the strengths of *DeepMaize* and *RedAgent*. While *DeepMaize* performed best during steady state operations and had to employ a special procurement strategy (Day 0 fake strategy) to minimize the singularity of the start of the game, *RedAgent's* architecture proved more responsive to start and end game singularities.

8. SUMMARY AND CONCLUDING REMARKS

Supply chain trading environments present companies with the challenge of evaluating exponential numbers of sourcing, procurement, scheduling and customer bidding options under uncertain market and operational conditions. Intelligent agent functionality offers the promise of significantly increasing supply chain trading performance by automatically evaluating a much larger number of options than a human manager could. At the same time, these technologies are largely untested, making managers nervous about their performance.

How are these technologies going to behave under changing market conditions or in the face of competitors looking for strategies aimed at defeating them? Traditional research methodologies that evaluate techniques by comparing them against predefined sets of solutions are not sufficient to answer these and related questions, as they fail to capture the inherently competitive and strategic nature of supply chain trading. In this paper, we have argued that a more promising approach, or at the very least a complementary one, involves the development of open competitions that pit alternative solutions against one another.

To the best of our knowledge, TAC-SCM is the first competition to successfully capture the combinatorial complexity, uncertainty and strategic dimensions associated with realistic supply chain trading scenarios. It does so, while retaining a sufficient level of simplicity to allow teams to develop competitive solutions in a matter of a few months. We believe that this balance between realism and simplicity has been key to the early success of the competition with 20 teams from nine different countries participating in the first edition of the tournament. At the time of writing, the second edition of the competition is already under way, featuring over 30 entries, a reflection of the success of the competition and its perceived research value.

The success of TAC-SCM goes beyond the high number of competitors it managed to attract. The game proved too complex for any simple-minded strategy. Its sophisticated model of supply chain negotiation and its multiple sources of uncertainty seem to elude the design of any form of "optimal" solution, requiring instead that agents closely monitor changing conditions and adjust their behavior accordingly. This was illustrated by *RedAgent*, the winner of the 2003 tournament. By forcing agents to compete in a number of games before moving to the next round, the competition also ensures that agents are evaluated across a number of different market conditions.

The 2003 tournament also revealed areas of the game that could be further improved. In hindsight, component discounts offered on Day 0 were probably excessive, placing too much emphasis on start game strategies. This has been corrected in the 2004 competition [3]. Nevertheless, Day 0 procurers were unable to effectively capitalize on this particular singularity, as their strategies were eventually countered by *DeepMaize*'s Day 0 fake strategy. In the end, this counter cleared the way for *RedAgent* to win the competition, thanks to an adaptive architecture that dynamically coordinates procurement, sourcing, scheduling and customer bidding activities via internal micro-markets.

Finally, TAC-SCM also offers new insights into the maturity of automated supply chain trading technologies. No agent fully dominated the competition, including *RedAgent*, which only won with the help of *DeepMaize*'s Day 0 fake strategy. A quick look at the inventory levels carried by *RedAgent* (both component inventory and finished goods inventory) suggests that one should be able to eventually develop agents that perform significantly better (See Figure 12 and 13). In fact, *DeepMaize*, which finished second, managed to nearly match *RedAgent*'s profitability with significantly less inventory.

Overall, the 2003 tournament showed that supply chain trading technology can already deliver solutions capable of effectively evaluating very large numbers of sourcing, procurement, scheduling and customer bidding options under routine conditions. Even if these techniques appear to still have room for improvement, there is no question that they are far better than solutions any human could ever hope to develop manually. However, when it comes to strategic decisions, today's solutions still seem to fall short and be too brittle. *DeepMaize*'s Day 0 fake counter was not discovered by the agent itself but rather by its developers, who modified their agent over night. In the short to medium term, this would generally argue for the development of mixed-initiative supply chain trading solutions, where managers remain in charge of key strategic decisions, controlling key parameters of their trading agents, while relying on the agents' speed to effectively operationalize these decisions.

Acknowledgements. The research reported in this paper has been funded by the National Science Foundation under ITR Grant 0205435. TAC-SCM is the product of a collaboration between Carnegie Mellon University's e-Supply Chain Management Laboratory and the Swedish Institute of Computer Science. In particular, the authors would like to acknowledge the contributions of Joakim Eriksson, Niclas Finne and Sverker Janson. They also wish to thank all members of the new TAC-SCM community for their feedback on the design of the game and their participation. The authors also thank the *RedAgent Team* for sharing Figure 10 and 11 with them.

REFERENCES

[1] R. Arunachalam, N.M. Sadeh, J. Eriksson, N. Finne, and S. Janson. The TAC Supply Chain Game. Technical Report CMU-CS-03-184, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, August 2003.

[2] R. Arunachalam, N.M. Sadeh, J. Eriksson, N. Finne, and S. Janson. The Supply Chain Management Game for the Trading Agent Competition 2004. Technical report CMU-CS-04-107, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, July 2004.

[3] R. Arunachalam, N.M. Sadeh, J. Eriksson, N. Finne, and S. Janson. Design of the Supply Chain Trading Competition. Workshop on Trading Agent Design and Analysis, 18th International Joint Conference on Artificial Intelligence, Acapulco, Mexico, August 9-15, 2003.

[4] M. Babaioff and N. Nissan. Concurrent Auctions across the Supply Chain. M. Wellman and Y. Shoham (Eds.), *Proceedings of 3rd ACM Conference on Electronic Commerce*, Tampa, FL, October 14-17, 2001, 1-10.

[5] M. Babaioff and W.E. Walsh. Incentive-Compatible, Budget-Balanced, Yet Highly Efficient Auctions for Supply Chain Formation. In N. Nisan (Ed.), *Proceedings of the 4th ACM Conference on Electronic Commerce*, San Diego, CA, June 9-12, 2003, 64-75.

[6] Y. Bassok and R. Akella. Multiplan Co-ordination: Combined Production and Ordering with Demand and Supply Uncertainity, *Management Science*, 37, 12: 15, 1991.

[7] S. Chopra and P. Meindl. *Supply Chain Management-Strategy, Planning, and Operation.* Prentice Hall, Englewood Cliffs, NJ, 2001.

[8] F.W. Ciarallo, R. Akella, and T.E. Morton. A Periodic Review Production Planning Model with Uncertain Capacity and Uncertain Demand: Optimality of Extended Myopic Policies, *Management Science*, 40, 3, 320-332, 1994.

[9] J. Collins, W. Ketter, and M. Gini. A Multi-Agent Negotiating Test Bed for Contracting Tasks with Temporal and Precedence Constraints. *International Journal of Electronic Commerce*, 7, 1, 35-57, 2002.

[10] J. Du and J.Y.T. Leung. Minimizing Total Tardiness on One Processor is NP-Hard. *Mathematics of Operations Research*, 15, 3, 1990, 483-495.

[11] J. Estelle, Y. Vorbeychik, M. Wellman, S. Singh, C. Kiekintveld, and V. Soni. Strategic Interactions in a Supply Chain Management Game. Technical Report, Computer Science and Engineering, University of Michigan, April 2004.

[12] H. Gurnani, R. Akella, and J. Lehoczky. Optimal Order Policies in Assembly Systems with Random Demand and Random Supplier Delivery, *IIE Transactions*, 32, 8, 701-714, 2000, 865-878.

[13] P.W. Keller, F.O. Dugay, D. Precup. RedAgent 2003: An Autonomous, Market-Based Supply Chain Management Agent. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multi Agent Systems, Volume 3*, July 19-23, 2004, IEEE Computer Society Press, 2004, 1182-1189.

[14] C. Kiekintveld, M.P. Wellman, S. Singh, J. Estelle, Y. Vorobeychik, V. Soni, and M. Rudary. Distributed Feedback Control for Decision Making on Supply Chains. In *Proceedings of the 14th International Conference on Automated Planning and Scheduling*, 2004, 384-392.

[15] D. Kjenstad. Coordinated Supply Chain Scheduling. Ph.D, thesis and NTNU Report 1998:24, Department of Production and Quality Engineering, Norwegian University of Science and Technology (NTNU), Trondheim, November 1998.

[16] H.S. Lau and A.H.L. Lau. Coordinating Two Suppliers with Offsetting Lead Time and Price Performance. *Journal of Operations Management*, 11, 1994, 327-337.

[17] Lee, H.L., V. Padmanabhan, and S. Whang, S. The Bullwhip Effect in Supply Chains. *Sloan Management Review*, Spring 1997, 93-102.

[18] D. Simchi-Levi, P. Kaminsky, E. Simchi-Levi. *Designing and Managing the Supply Chain*. McGraw-Hill, New York, NY, 1999.

[19] R.W. Rosenthal. Rules of Thumb in Games. *Journal of Economic Behavior and Organization*, 22, 1, 1993, 1-13.

[20] D. Friedman and J. Rust. *The Double Auction Market: Institutions, Theories, and Evidence.* Addison Wesley, Reading, MA, 1992.

[21] N. M. Sadeh, R. Arunachalam, J. Eriksson, N. Finne and S. Janson, TAC'03: A Supply Chain Trading Competition, *AI Magazine*, 24, 1, 2003, 83-91.

[22] N.M. Sadeh, D.W. Hildum, D. Kjenstad and A. Tseng. MASCOT: An Agent Based Architecture for Dynamic Supply Chain Creation and Coordination in the Internet Economy. *Journal of Production, Planning and Control*, 12, 3, 2001, 211-223.

[23] J. Sun and N.M. Sadeh. Coordinating Multi-Attribute Reverse Auctions Subject to Finite Capacity Considerations. Technical Report CMU-ISRI-03-105, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, May 2003.

[24] Swaminathan, J.M., N.M. Sadeh, and S.F. Smith. Information Exchange in the Supply Chain. Technical Report CMU-RI-TR-95-36, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, 1995.

[25] Swaminathan, J.M., S.F. Smith, and N.M. Sadeh. Modeling Supply Chain Dynamics: A Multiagent Approach, *Decision Sciences*, 29, 30, 1998, 607-632.

[26] D.J. Thomas and P.M. Griffin. Coordinated Supply Chain Management. *European Journal of Operational Research*, 94, 1, 1996, 1-15.

[27] W.E. Walsh and M.P. Wellman. Decentralized Supply Chain Formation: A Market Protocol and Competitive Equilibrium Analysis. *Journal of Artificial Intelligence Research*, 19, 2003, 513-567.

[28] M.P. Wellman, P.R. Wurman, K O'Malley, R. Bangera, S-D. Lin, D. Reeves, and W.E. Walsh. Designing the Market Game for a Trading Agent Competition. *IEEE Internet Computing*, 5, 2, 2001, 43-51.

[29] D. Zeng and K. Sycara. Dynamic Supply Chain Structuring for Electronic Commerce among Agents. In M. Klusch (Ed.), Chapter 10, *Intelligent Information Agents*, Springer, Heidelberg, Germany, 1999.

[30] P. Zipkin. Models for Design and Control of Stochastic Multi-Item Batch Production Systems, *Operations Research*, 34, 11, 1986, 91-104.