

**Pushing the limits of rational agents: the
Trading Agent Competition for Supply Chain
Management**

John Collins

Computer Science Department, University of Minnesota

Wolfgang Ketter

Rotterdam School of Management, Erasmus University

Norman Sadeh

School of Computer Science, Carnegie Mellon University

December 2009

CMU-ISR-09-129

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

Over the years, competitions have been important catalysts for progress in Artificial Intelligence. We describe one such competition, the Trading Agent Competition for Supply Chain Management (TAC SCM). We discuss its significance in the context of today's global market economy as well as AI research, the ways in which it breaks away from limiting assumptions made in prior work, and some of the advances it has engendered over the past six years. TAC SCM requires autonomous supply chain entities, modeled as agents, to coordinate their internal operations while concurrently trading in multiple dynamic and highly competitive markets. Since its introduction in 2003, the competition has attracted over 150 entries and brought together researchers from AI and beyond in the form of 75 competing teams from 25 different countries.

Keywords: Autonomous Agent, Electronic Commerce, Trading Agent

Introduction

Many important developments in Artificial Intelligence have been stimulated by organized competitions that tackle interesting, difficult “challenge” problems. Examples include chess, robot soccer, poker, robot navigation, stock trading, and others. Economics and artificial intelligence share a strong focus on rational behavior. Yet the real time demands of many domains do not lend themselves to traditional assumptions of rationality (Simon, 1979; Wellman, 1996). This is the case in trading environments, where self-interested entities need to operate subject to limited time and information. With the Web mediating an ever broader range of transactions and opening the door for participants to concurrently trade across a number of markets, there is a growing need for technologies that empower participants to rapidly evaluate very large numbers of alternatives in the face of constantly changing market conditions. AI techniques such as neural networks and genetic algorithms are already routinely used in support of automated stock trading scenarios. Yet, the deployment of these technologies remains limited, and their proprietary nature precludes the type of open benchmarking that is critical for further scientific progress.

The Trading Agent Competition for Supply Chain Management (TAC SCM)¹ was conceived by the third author in 2002 as a way of focusing the attention of researchers in AI and beyond on the increasingly complex problem of managing supply chains in today’s global economy. More specifically, it was designed to foster the development of new techniques to manage risk and adapt to changing conditions while concurrently trading in multiple market places (Sadeh et al., 2003a). The initial version of the game was designed and implemented through a collaboration between Carnegie Mellon University and the Swedish Institute of Computer Science (Arunachalam and Sadeh, 2005; Eriksson et al., 2006), with subsequent refinements introduced through a collaboration with the first author (Collins et al., 2005). Over the years, all three authors have also contributed successful entries in the competition.

Supply chains are the foundation of today’s global economy, with annual flows worth tens of trillions of dollars. As companies continue to focus on core competencies and outsource functions ranging from procurement of raw materials and components to logistics, after sales support, and recycling/remanufacturing operations, they weave increasingly complex networks of interdependent organizations often spanning multiple continents. Pressure to shorten product lifecycles, reduce costs and offer higher levels of customization is simultaneously forcing organizations to explore increasingly flexible contractual relationships (e.g. price, volume or service-level flexibility) aimed at reducing inventory risks while providing protection against shortages and price fluctuations. By their very nature, these more flexible relationships place a premium on the ability of supply chain entities to rapidly adapt to changing market conditions. Those capable of doing so reap significant benefits in the form of more efficient operations and higher profit margins. Yet failures under these less forgiving scenarios can also be catastrophic, ranging from companies going out of business because they made the wrong bets, to critical supplies failing to be delivered in time in the face of disruptive events such as hurricanes, strikes or pandemics.

TAC SCM builds on the observation that supply chains should not be viewed as monolithic entities that can be centrally optimized, but instead consist of multiple self-interested entities each operating according to its own objectives and policies (Swaminathan et al., 1998). Whereas each real-world supply chain exhibits its own peculiarities, TAC SCM is designed to capture major sources of complexity common to large numbers of supply chains, while shielding researchers from

¹See <http://www.tradingagents.org> for more information

less relevant idiosyncracies. Supply chain entities are modeled as autonomous agents that concurrently compete with one another in both end-product and component market places subject to both exogenous and endogenous sources of uncertainty. In artificial intelligence terms, these agents must act autonomously to maximize their expected utilities in an environment that is highly dynamic, partially observable, and strongly affected by the actions of competing agents.

Starting with the first open competition in 2003, TAC SCM has attracted over 150 entries from 75 teams distributed across 25 different countries. These teams represent a variety of research interests, including supply-chain management (Sardinha et al., 2009), agent architectures (Collins et al., 2009a; Benisch et al., 2009), economic decision-making (Kiekintveld et al., 2006; Ketter et al., 2009), empirical game theory (Jordan et al., 2007), dynamic pricing (Benisch et al., 2006a; Ketter et al., 2007), machine learning (Pardoe and Stone, 2007), economic market modeling (Ketter et al., 2009), fuzzy logic (He et al., 2006), stochastic optimization (Benisch et al., 2004), and a large variety of other areas².

In the following sections we present the TAC SCM game scenario and review the decisions competing supply chain trading agents have to make in the game. Along the way, we discuss some of the challenges associated with the design of the game and of successful agents. We highlight the game’s relationship to current and future supply chain practices. We also review how TAC SCM has been used in education and what is required for new teams to enter the competition, including the availability of software that can help to develop and fine-tune agents.

The TAC SCM scenario

The TAC SCM game captures key features of a multi-tier supply chain with multiple actors competing in each tier. The game models 220 days (or 44 five-day weeks, which we informally refer to as “a year”) of operation in 55 minutes of real time. Participating agents must operate through three conceptual phases that are characteristic of the launch, steady production, and eventual phasing out of multiple end-products. The end products are assumed to be different types of PCs, though the simulation model is in no way restricted to this particular type of product.

Each PC model requires a different combination of components. To promote lean supply chain management practices, TAC SCM assumes that PC models are phased out by the end of the one-year period, and their residual value and that of the major components they require is effectively zero. In other words, the scenario penalizes agents that hoard more components than they need or assemble more PCs than they are able to sell. Teams compete by entering their supply chain trading agents in the competition and playing a large number of games against different combinations of competitors. Each trading agent is responsible for procuring components from multiple suppliers, manufacturing finished products (different types of PCs), and selling those products to customers. All the while it competes with other trading agents entered by other teams, who are also trying to purchase the same components and sell the same products, as component and end product market conditions change (e.g. because of the actions of other trading agents or because of exogenous conditions such as suppliers losing some capacity or demand for some particular types of PCs dropping over time). By requiring agents to compete in hundreds of games, the competition is able to evaluate agent performance across a large number of market conditions and competitive settings

²A Google Scholar search on “TAC-SCM” returns over 350 hits. While some of these links may be duplicates or erroneous, they give a sense of the intensity and breadth of research activity the competition has stimulated over the past several years.

- from markets with low product demand and high component availability, to markets with high product demand and low component availability, to anything in between including markets that transition between these extremes for different end products and components.

The schematic overview of the TAC SCM scenario shown in Figure 1 will help explain the specifics of the game. Several agents (e.g., MinneTAC, TacTex, Botticelli) compete with each other as manufacturers of personal computers, purchasing computer components (CPU, motherboard, hard drive, and memory) from suppliers (IMD, Pintel, etc.) in a procurement market, and selling their finished products to customers in a sales market. Each agent has identical production and warehouse facilities, each sees the same customer demand, and each initially has equal access to suppliers, although reputation effects can result in preferential treatment as the game progresses. Customer demand and supplier capacity and prices are highly variable both within the course of a game and across different games. Each agent starts with no inventory and an empty bank account, and must borrow (and pay interest) to build up an inventory of computer components before it can begin assembling and shipping computers. The agent with the largest bank account at the end of the game is the winner of that particular game; performance of agents is averaged across many games in the competition setting. Actions of other agents are visible only through their effects on the customer and supplier markets.

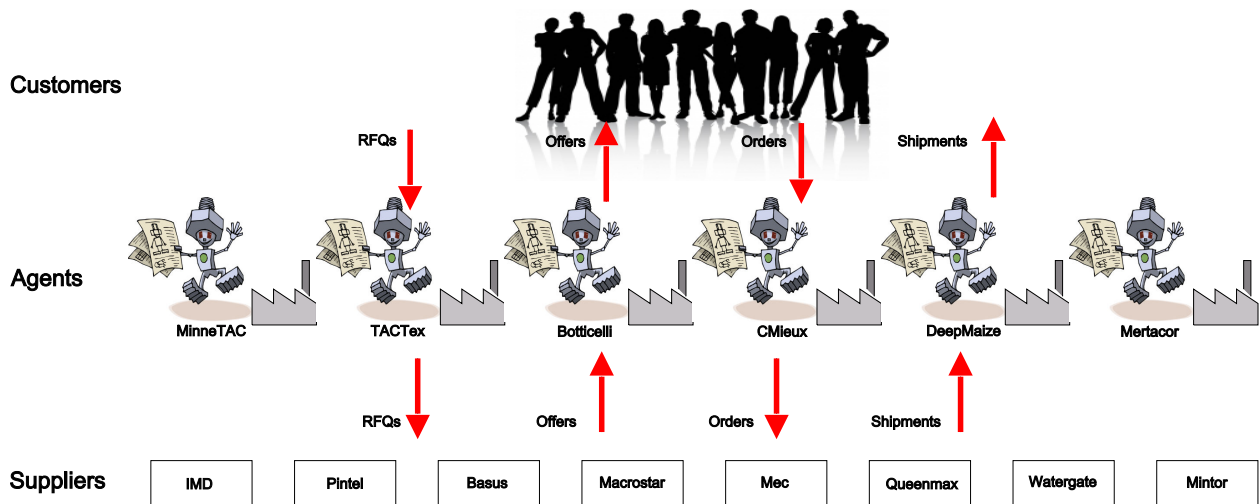


Figure 1: Schematic view of TAC SCM simulation scenario. Agents assemble and sell 16 different products, using 10 component types purchased from suppliers.

Figure 2 shows a typical sequence of interactions between a TAC SCM agent and its environment over the course of a simulated day. Each day lasts 15 seconds, which limits the reasoning time available to agents. In the real world, supply chains require the management of many more components with many events accruing during the course of the day; a 15-second day in the TAC SCM game is therefore not unrealistic. At the beginning of each cycle, each agent receives a bundle of messages from the server, representing customer demand, customer orders arising from bids placed during the previous cycle, status updates from the bank and warehouse, and supplier offers arising from supplier RFQs issued during the previous cycle. Before the end of the cycle, the agent must decide how to bid on new customer RFQs, which supplier offers to accept, what parts to request from suppliers, what finished goods to ship to customers, and how to allocate its limited factory

and inventory resources to production of new finished products. The actual delivery of components from suppliers to the agent’s warehouse, and delivery of finished products from the agent’s warehouse to customers, along with associated payments from and to the agent’s bank account, are handled by the server.

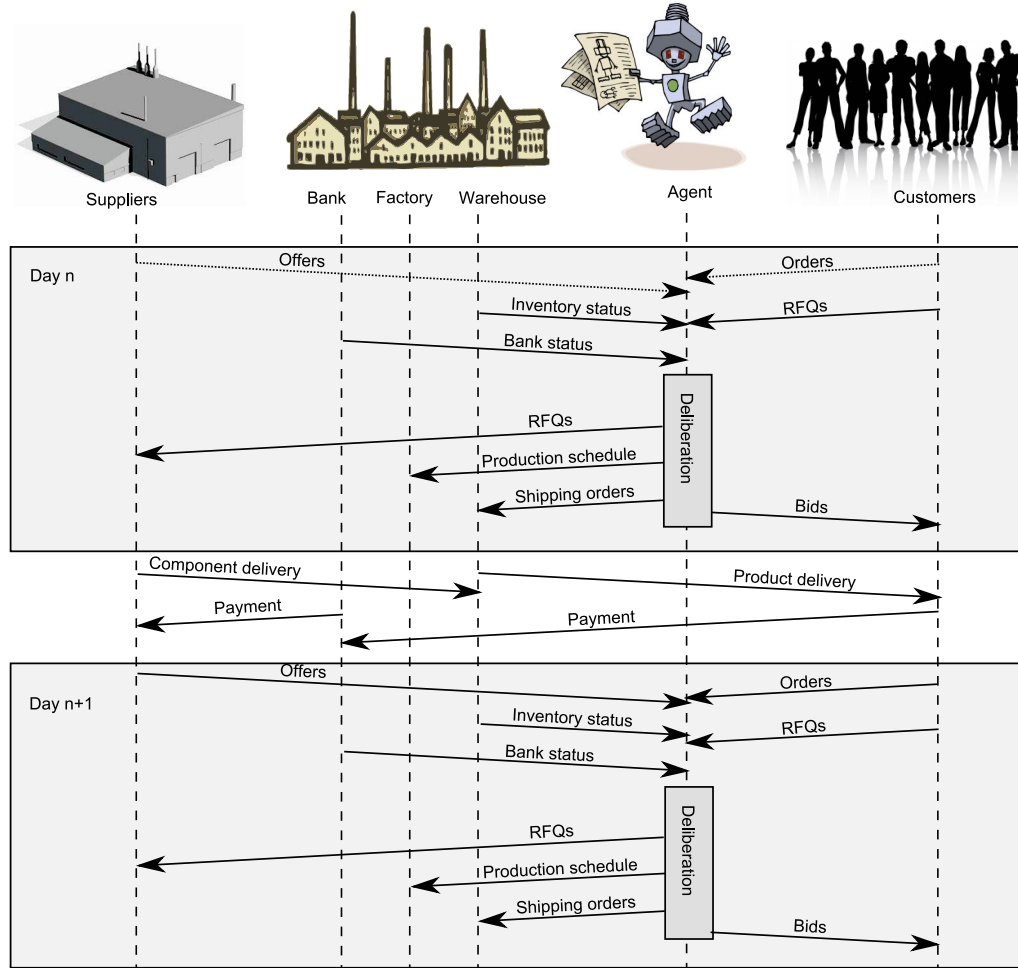


Figure 2: Timeline of an agent’s interactions with its environment.

From an AI perspective, the TAC SCM game requires agents to concurrently compete in multiple markets (the procurement and sales markets) subject to numerous sources of uncertainty, while simultaneously managing their internal production and logistics operations. The markets also exhibit some level of interdependence. For instance, different end products require different but overlapping combinations of components, and hence demand for different components will have some correlation over time. The agents also have to operate with incomplete information. For instance, agents do not know how market conditions will evolve and they do not see the private data of other agents such as their inventory positions, their order books, the prices at which they procure components from suppliers, the component purchases they have made, the price at which they are selling different PC models on a given day, etc. Yet, each day as they place requests for components and offer end products for sale, they receive feedback from the environment. This

feedback comes in multiple forms, including whether they manage to sell different types of PCs and at what price, how many components they manage to procure from different suppliers, at what price and leadtime. Agents can leverage this feedback to update probabilistic models of the markets and of what other agents are doing. These models in turn can help them dynamically adapt their strategies. Strategies can include adjusting one's product mix, stocking up on components that are expected to be in low supply, and increasing the price of products that seem to be in high demand or for which there seems to be little competition.

The following three subsections further detail elements of the game's procurement market, production and logistics, and the sales market. These subsections are intended to convey the level of sophistication involved in the design of the TAC-SCM Competition. Some readers may just want to jump to the section on "Agent decision problems."

Procurement market

The procurement market consists of eight suppliers, each of which carries two product lines. Each supplier operates according to a lean, make-to-order policy. In other words, production is driven by actual demand rather than demand forecast. The capacity of each supplier production line varies from day to day according to a mean-reverting random walk that captures effects such as loss of capacity (e.g. due to maintenance) and exogenous demand (e.g. components used in products not modeled in this supply chain). Agents may request price quotes from suppliers, specifying a particular component, quantity, delivery date, and reserve price. Suppliers respond with quotes that reflect how busy they are, with higher quotes being returned when they are running near capacity and lower quotes when they have lower order books. Supplier commitments are based on estimates of their future capacity and, as such, are not entirely reliable. As a result, capacity variations can lead to shipment delays.

To build a finished product, an agent needs one each of four different component types: a CPU, a motherboard, a disk drive, and a memory card. There are two CPU suppliers, Pintel and IMD. Their processors are not interchangeable; Pintel CPUs must be assembled with Pintel motherboards, and IMD CPUs must be assembled with IMD motherboards. Both Pintel and IMD make "fast" and "slow" CPUs. Disk drives and memory cards each come in two different sizes. The result is that the CPUs are single-sourced, while all other component types are dual-sourced. The lack of substitutability between CPU sources means that the CPU market is more volatile and unpredictable than the markets for other component types.

Individual suppliers are approximately revenue-maximizing entities, and they manage risk in two ways. First, they will not commit their entire capacity at any one time; instead, they reserve a portion (approximately half) of future capacity for future business. Second, they keep track of whether agents follow through with orders when offers are made. The result is that agents must manage their individual "reputations" with respect to each supplier, by keeping their ratio of orders to offers above a threshold. Failure to do so results in higher prices and lower availability of parts in comparison with competing agents.

The procurement market generally yields lower prices for longer lead-times, but at times of oversupply, prices can be lower for very short-term requests. Order lead-times can extend to the end of the game, which is 220 days at the beginning of a game. The longest customer lead-time is 12 days, and supplier prices tend to peak in the range of 8-15 days lead-time. Supplier orders require a 10% down-payment, and so the cost of funds can be a significant factor for long lead-time orders.

Production and logistics

Once an agent has acquired the necessary parts to assemble computers, it must schedule production in its finite-capacity factory. Each computer model requires a set of parts, and a specified number of assembly cycles. Assembled computers are added to the agent's finished-goods inventory, and may be shipped to customers to satisfy outstanding orders.

Warehouse capacity for both components and finished products must be purchased at a price that is a function of the value of the stored materials. This effectively places a premium on keeping inventories under control. In addition, each individual component type is used in multiple types of finished goods, across multiple market segments. The result is that there is potentially some opportunity cost to converting parts to finished goods without having sales commitments for the finished goods. Further incentive to keep inventory under control arises from the fact that at the end of the simulation, unsold inventory has no residual value.

Sales market

The sales market uses a reverse, first-price, sealed-bid auction mechanism. Each day, customers issue requests for quotes (RFQs) for the products they wish to buy. Each RFQ specifies a computer model, quantity, delivery date, and a maximum or *reserve* price, as well as a daily penalty amount that the agent must pay if it fails to meet its sales commitments. Penalties can make failure to ship on time quite expensive, and customer orders are canceled if they are more than five days late, which also eliminates income from the sale.

The sales market trades in 16 product types, segmented into high-end, medium, and low-end products. Customer demand varies from day to day independently in each segment, controlled by a trend that changes daily through a bounded random walk. Figure 3 shows a qualitative view of the supply and demand behavior of one segment of the customer market. The shape of the demand curve is controlled by the current overall demand, and by the uniform distribution of reserve prices between 0.75-1.25 of the nominal cost of components. The shape of the supply curve is limited at the high-quantity end by the inventory status and aggregate production capacities of the competing agents, and at the low-price end by the minimum cost of components. The detailed shape of the supply curve is a function of the combined bidding strategies of the competing agents. It is not directly observable within a game, but segments of it can be deduced through post-game examination of data.




Figure 3: Supply and demand in the sales market.

The game design challenge³

Designing an economic game that serves a research agenda and provides an interesting and reasonably accessible challenge to researchers is a difficult undertaking. The original goals for TAC SCM, described by Arunachalam and Sadeh (2005), included realism and generality, uncertainty and incomplete information, strategic interactions, and simplicity. The scenario must have enough realism to be relevant to practitioners, enough generality to be representative of a broad range of supply-chain situations, and enough subtlety to require new ideas, without unnecessary complexity that would make agent design and data analysis more difficult than it needs to be. Supply chain management must deal with uncertainty, risk, and limited information. An interesting simulation scenario should include significant variability in prices, availability, and demand, and should limit visibility of markets and competitors to approximate the view of a real-world supply-chain manager. The scenario should reward careful management of risk. The game and its market mechanisms must allow and encourage strategic behavior, and yet be free of unintended exploitable weaknesses. In addition, in order to attract enough interest to provide good competition, game designers must provide the simulation infrastructure and a basic agent framework that encapsulates the interactions between the infrastructure and the agent. Serviceable visualization and analysis tools are needed in order to demonstrate the concepts of the game, and to support researchers as they develop and analyze their agents.

The design of TAC SCM was carefully tuned over the first three years to make the competition interesting and challenging. Empirical evidence based on the many games played over the years by agents developed by 75 different teams suggests that opportunities for strategic manipulation that have the potential to subvert the purpose of the game have been eliminated (Ketter et al., 2004; Wellman et al., 2005). Agents must manage their reputations with respect to each supplier; this discourages agents from creating false demand by making large requests and then turning down the resulting offers, thereby inflating prices. Because suppliers reserve some capacity for future demand, it is very difficult to “corner” the market for some component type. It is still possible and indeed common for agents to manipulate prices in both the sales and procurement markets. For example, knowing that agents can see only the highest and lowest order prices in the sales market, agents can make isolated offers well below current prices, thereby seeding opponent price models with misleading information. This technique can be used to drive prices down with minimal impact on profits, in situations where prices are above the “knee” in the customer demand curve, or when an agent believes it has a lower procurement cost basis for a specific product than its competitors.

The parameters of the scenario are set to ensure that decision coordination between procurement and sales is reasonably challenging. Figure 4 is a histogram of the daily customer RFQ count over 200 games, approximately 44,000 observations, which shows the overall balance between supply and demand. Superimposed on the histogram are the mean customer demand, the aggregate capacity of six agent factories, and the expected supplier capacity. The key message from this balance is that an agent can expect to buy enough parts to keep its factory busy, but a strategy that simply tries to keep the factory busy is likely to result in a large unsold inventory at the end of an average game because the expected customer demand cannot absorb the production of a set of agents that behave in this way⁴. On the other hand, there are some games in which the agents cannot supply all the demand, and the variability inherent in the simulation can lead to serious imbalances between

³Note to the reviewers: This section could be a sidebar

⁴This balance was first introduced in the 2005 competition. Destructive price wars were a common problem in the early rounds of that competition until the full-production agents were eliminated

customer demand for specific products and the availability of the parts required to build them. The best agents are able to adapt their behaviors to exploit these imbalances.

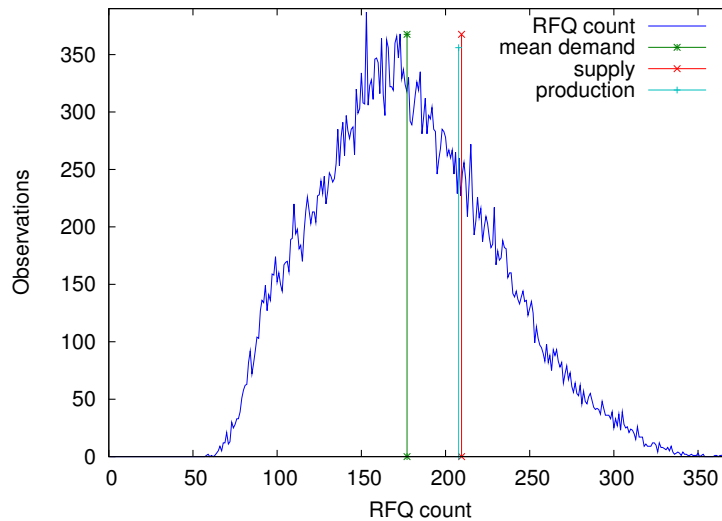


Figure 4: Distribution of demand in the TAC SCM customer market, compared with the expected availability of parts, and the aggregate production capacity of six agent factories.

The game platform consists of a server that simulates not only the suppliers and customers, but also the agent factories and warehouses, along with a bank. Agents join a simulation through standard Internet connections. This allows research teams to work with their own tools and hardware, and greatly simplifies the operation of open competitions. Since many agents are essentially compute-bound, this structure arguably gives some advantage to teams who have more or better hardware at their disposal. So far, there is little evidence that this is a factor; agent design appears to be a significantly greater predictor of performance than hardware capability.

Agent decision problems

To be competitive, an agent must purchase components and manufacture products it can sell, and, to the extent possible, it must sell what it has built at a profit. After seven years of competition and publication, the competition in the final stages of the annual tournament is quite intense. Profit margins are very slim, and prices in the customer market are seldom very far above the component costs seen by the agent with the most effective procurement strategy. Prices in both the customer and supplier markets can be quite volatile as agents continually adjust their behaviors to take advantage of the markets and of any weaknesses in their opponents. Kiekintveld et al. (2006) identify three key issues that a successful TAC SCM agent must address: dealing with substantial *uncertainty* in a highly *dynamic* economic environment, in competition with other self-interested agents whose behavior is naturally *strategic*.

A successful agent design for TAC SCM must make a large number of decisions every 15 seconds. Attempts to construct and maintain opponent models must account for the fact that the sales and procurement markets are “oligopoly” markets, which means that one must model and anticipate the effects of one’s own actions. For example, any attempt by an agent to increase its market share for

a highly profitable product will typically depress prices and hence the profitability of that product. It also means that daily variations in the behavior of a competitor can introduce significant “noise” into the price signals that are observed by an agent.

In the remainder of this section we describe the basic decision problems a successful TAC SCM agent must address, and review some of the successful approaches that have been developed as a result of the ongoing competition.

Procurement

Agents seek to purchase raw materials (computer components) at a price that will allow finished goods to be sold at a profit. But purchasing low-cost components is not enough. There is a limit to the rate at which the agent’s factory can absorb components, and components that cannot be assembled into usable products are worse than useless – they represent sunk cost that may not be recoverable, and the agent must pay to store them.

Most agents use some sort of projected inventory model to determine their procurement needs, along with a supplier pricing model. For example, Kiekintveld et al. (2004) describe an inventory model that projects future inventories by adding supplier shipping commitments and subtracting existing and expected future consumption. Whenever a future shortage is discovered, the agent forms a goal of replenishing its inventory to cover the shortage. If the expected shortage is more than a few days in the future, the agent must decide whether to place an order immediately, or wait until some later date and place an order with a shorter lead-time. This decision is driven by a supplier pricing model that observes supplier price quotes to build up a history of prices at various lead-times; the agent then uses this model to estimate bounds on each supplier’s uncommitted capacity (this works because supplier pricing is a deterministic function of uncommitted capacity, but the signal is sparse and extremely noisy), and it uses these capacity estimates to predict current and future prices.

Benisch et al. (2006b, 2009) describe a procurement process that spreads its requests to each supplier over time in a way that attempts to minimize overall cost. The difference between target and predicted inventory is projected out to the end of the game, and each projected shortfall initiates a process that distributes purchasing across available suppliers and across time, attempting to exploit minima in the prices predicted by its supplier-price model.

Procurement is arguably the most difficult decision problem in the TAC SCM scenario. Possibly because of this difficulty, there is strong evidence from work by Andrews et al. (2008) that procurement performance is the best predictor of overall agent performance. This observation was also one of the motivations behind the launch in 2007 of a separate *Procurement Challenge* focusing specifically on benchmarking agents’ procurement strategies (Sardinha et al., 2009).

Production scheduling

In the early competitions, some agents experimented with various methods for building near-optimal production schedules. For example, Benisch et al. (2004) describe a stochastic programming formulation for production scheduling. The goal was to maximize the probability that products produced would actually be sold. At any given time, an agent has a set of outstanding customer orders, with due dates spread out over up to 11 days in the future. It also has a set of customer RFQs on which it can place bids, and it has some expectation of future demand. In the approach described by Benisch et al., bids are placed before the production schedule is generated. The production

schedule attempts to satisfy all outstanding commitments first, ordered by due date, then as many outstanding sales offers as possible, ordered by expected profitability and due date. The schedule is projected into the future by several days, and the availability of uncommitted capacity is used to control sales volume targets.

Most agents simply keep track of uncommitted production capacity and use it to constrain sales quotas, and schedule production using a straightforward greedy method. Evidence from Andrews et al. (2008) suggests that production scheduling performance is not a strong differentiating factor among existing agents.

Sales

During each simulation cycle, each agent sees the full customer demand in the form of a bundle of RFQs. It must decide which requests to bid on, and what the bid prices should be. The agent must consider a number of factors in making this decision, including its own current and expected inventory situation, expected cost of inventory replenishment, existing sales commitments, available factory capacity, and its own models of future demand and market prices. A number of approaches have been tried for controlling this decision, ranging from fuzzy logic to linear programming. All of them must somehow solve a constrained optimization problem with some degree of approximation, in limited time.

The MinneTAC agent (Ketter et al., 2007) controls its bidding in the customer market on three principles. First, market prices are tracked and projected using a model that is trained with a large body of historical game data, and fine-tuned using in-game market monitoring. A Gaussian mixture model classifies market situations or “economic regimes” (Ketter et al., 2009). Projection of future price trends uses a recursive Markov model. Given sales and procurement price predictions over a planning horizon, a linear program sets daily sales quotas for each product over the horizon, subject to constraints arising from inventory, factory capacity, and (projected) customer demand.

The SouthamptonSCM agent (He et al., 2006) uses fuzzy reasoning to decide which customer requests to respond to, and what the bid prices should be. Its overall goal is to maximize its factory utilization, as long as it can do so profitably. Demand and inventory are represented as fuzzy sets, and a set of rules convert those into control variables which are combined into a “price adjustment factor” that is applied to the prices observed in the market on the previous day. SouthamptonSCM uses a separate fuzzy rule base to control pricing near the end of a game, where the criterion for profitability requires moving as much of its remaining inventory as possible.

Benisch et al. (2006a) describe a very different approach to pricing in the sales market. Their CMieux agent treats sales pricing as a *continuous knapsack problem*. The goal is to offer a price to every potential customer (at least those for which a profitable sale is possible) that maximizes total revenue, while avoiding overcommitment of its inventory and production capacity in expectation. They show that the optimum price sets the demand fraction for the different products such that the first derivative of a “reward function” is equal across all products.

The small number of players in the TAC SCM markets has been a tempting target for opponent-modeling techniques. Unfortunately, the specific actions of individual opponents are not visible during the simulation – only their aggregate effects on the markets can be observed. However, the simulation server keeps a detailed log of agent interactions, and this data is available after the completion of each tournament round. Pardoe and Stone (2009) have experimented with a variety of opponent modeling schemes, using these detailed simulation records. They were able to show that there is some advantage to learning from a body of training data containing exactly the mix

of opponents that one is competing against, but it is not a large effect.

One way in which progress is evident over the years of the competition is through the distribution of sales prices. We see this in two different ways. First, agents have become more consistently profitable, although they do not always earn back their initial investment. This means, for example, that price wars have been rare in the more recent tournaments. Second, the daily sales price ranges for given products averaged around 3% in the 2008 tournament, and 1.5% in the 2009 tournament. In contrast, the width of the daily price distributions during the first two years of competition averaged over 7% and frequently went well over 10% in tournament games.

Decision coordination

The TAC SCM scenario places a premium on effective coordination of decisions affecting multiple markets and internal resources. Inventory planning is complicated by the fact that a given part may be used in multiple products, and a shortage of a particular part can prevent an agent from participating in significant segments of the customer market. Because demand in the three customer market segments can vary independently over a wide range, a strategy that strongly decouples procurement from sales is unlikely to meet customer demand effectively without carrying excess inventories of the parts that are not currently in demand. In the first years of the competition, before the customer market was segmented, this decoupling was a very common strategy. Ketter et al. (2010) identify the problem of decision coordination as a crucial element in the design of an agent for TAC SCM, and review the published literature on agent design to discover a wide variety of approaches to this problem. The authors believe that the existence of such variety is an indication that much is yet to be learned about designing such agents.

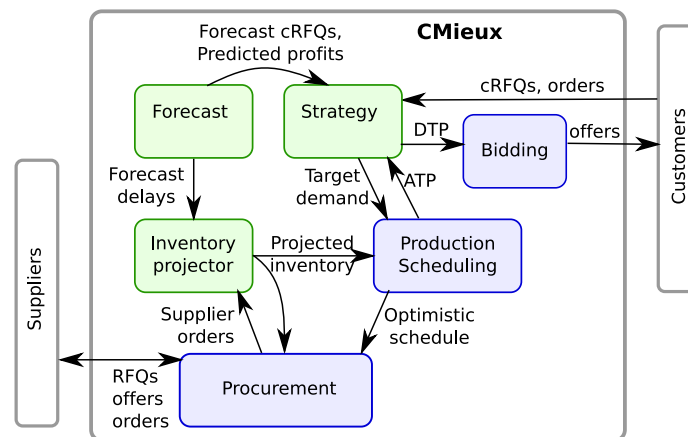


Figure 5: CMieux coordinates decisions using a separate Strategy module.

The CMieux agent (Benisch et al., 2009) from Carnegie Mellon is an example of an agent that clearly separates decision coordination from details of procurement, sales, and production scheduling. A schematic diagram of the CMieux design is shown in Figure 5. The Strategy module sets overall goals for the remainder of the system, such as the portion of expected demand to target, and the portion of the production schedule (*ATP*, or *Available to Promise* quantities for each end product) that should be sold to customers (*DTP*, or *Desired to Promise* quantities for each end product). The Forecast module observes the markets and makes predictions about demand, prices,

and delays in supplier shipments. The Inventory Projector combines that with current inventories and expected supplier deliveries to generate inventory projections over time. Procurement uses the projected inventory along with an optimistic version of the production schedule (what Production would expect to build if there were no inventory constraints) to decide what to order from suppliers, and supplies the Inventory Projector with actual supplier orders. CMieux reached the finals in 2007, 2008 and 2009.

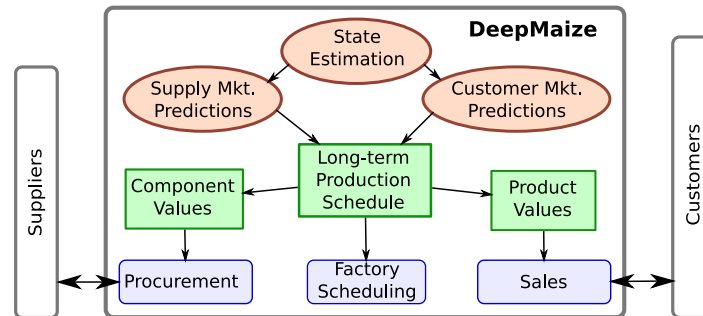


Figure 6: DeepMaize coordinates decisions through a long-term production schedule, using value-based decomposition.

The DeepMaize agent (Kiekintveld et al., 2006) from the University of Michigan coordinates its decisions through “value-based decomposition”. In this approach, a long-term production schedule is constructed by incrementally adding the products that are expected to return the highest marginal profits at multiple points in the future. The general scheme is summarized in Figure 6. This approach depends on reasonably accurate pricing models for both the customer and supplier markets that effectively capture price-quantity tradeoffs. The two prediction components shown in the diagram, along with an off-line machine-learning process, are responsible for producing those models. Given the resulting long-term production schedule, the Procurement module attempts to provide the necessary components to fill it, and Sales uses it to set prices in the customer market. DeepMaize has been a finalist in all of the TAC SCM tournaments. It placed third in 2006 and 2007, and first in 2008 and 2009.

Real-world agent-enabled SCM

Since the mid nineties, artificial intelligence techniques have contributed to the development of new supply chain decision support techniques, starting with the work of Swaminathan et al. on agent-based supply chain modeling and simulation, which was applied in the context of business process re-engineering efforts at major electronics and grocery firms (Swaminathan et al., 1998). Around the same time, Sadeh et al. also reported on the initial deployment and evaluation of AI-based decision support tools for supply chain coordination at Raytheon (Sadeh et al., 1998). This work led to the development of the MASCOT multi-agent architecture for coordinated decision support within and across multiple supply-chain entities (Sadeh et al., 2003b) and demonstrated how agent-assisted lateral coordination of manufacturing operations across organizations can increase both profitability and customer responsiveness in the face of high load factors and a variety of contingencies. A third line of work strongly influenced by research in artificial intelligence has seen the development and fielding of technologies for reverse supply chain auctions. This includes work by Wurman et al.

(1998) on configurable auction technology eventually commercialized by Ariba. It also includes work on more expressive mechanisms for reverse auctions conducted by Sandholm et al. (2005) and commercialized by CombineNet, a company that conducts large-scale procurement auctions where buyers and sellers can express a wide range of constraints and preferences beyond price.

Among other objectives, TAC-SCM was designed to promote the development and benchmarking of adaptive supply chain trading technologies required to better manage risk in supply chain environments characterized by increasingly flexible contractual relationships, such as those resulting from reverse auctions organized by companies like CombineNet. These long-term contractual relationships are typically characterized by flexibility in price, quantities and service levels and often entail arrangements where supply chain entities need to dynamically manage complex portfolios of flexible supply chain contracts (Martinez de Albeniz and Simchi-Levi, 2005). This work itself was strongly influenced by studies conducted in the mid to late nineties, showing that, contrary to popular belief, a number of manufacturers did not rely solely on long-term strategic partnerships with suppliers and that more research was needed on how to effectively manage portfolios of buyer-supplier relationships covering a wide spectrum of possible arrangements (e.g. study of Japanese car manufacturers by Bensaou (1999)). A review of models for constructing short-term and long-term contracts in business-to-business markets has been conducted by Kleindorfer and Wu (2003). Elmaghraby (2000) also provides an excellent review of tradeoffs between different sourcing strategies. Martinez de Albeniz and Simchi-Levi (2005) have shown that portfolios of quantity flexible procurement contracts used in combination with spot market procurement can contribute to higher expected profits and lower financial risk. Nagali et al. (2008) have reported using a similar risk management model to support the development of portfolios of procurement contracts, achieving savings of hundreds of millions of dollars in the procurement of flash memory used in printers assembled by Hewlett-Packard. In 2007, the third author and his colleagues launched a variation of the supply chain trading competition focusing solely on the management of long-term, quantity flexible procurement contracts and one-off procurement contracts (Sardinha et al., 2009).

Teaching

Beyond its research impact, the TAC-SCM scenario has also contributed to classroom education around the world (US, Canada, the Netherlands, Brazil, UK, Australia, etc), both at the undergraduate and graduate levels. Typically, students are required to either develop new entries from scratch or develop and evaluate alternative designs for modules of a given trading agent. In the process, they gain hands-on experience with online learning and stochastic optimization techniques. They also learn to better appreciate the complexity associated with competitive environments, where the success of one's strategy depends on the strategies of others and how quickly one can adapt to changes in these strategies. Having students work on TAC-SCM agents or modules has also proven to be an excellent way to expose students to software engineering concepts, especially when they work in teams, with different members each in charge of developing and evaluating modules or functionality whose performance is often dependent on modules or functionality developed by other team members. This type of work fits naturally in a computer science curriculum, but is less appropriate for management science and economics students.

As a result, games such as the MIT beer game⁵ (Sterman, 1989, 1992), are still the prevalent teaching tools in business and management programs, despite their simplistic setup. TAC SCM has

⁵<http://beergame.mit.edu>

the potential to replace these games and expose students (and managers) to much more dynamic and complex supply chain scenarios. Getting there however will require building an additional layer of mixed initiative functionality, where human decision makers (students or managers) have the ability to tweak high-level, human-oriented parameters, while relying on underlying agent-based decision support to process the very large number of options and decisions necessary to implement resulting strategies. Initial work on a mixed-initiative version of TAC SCM and the MinneTAC trading agent is currently under way (Nelson et al., 2009).

Getting involved

A working trading agent is a complex piece of software. TAC SCM agents must not only make coordinated decisions; they must also interact correctly with the game server, and typically they must produce data needed for empirical research. The organizers of the competition have worked to keep the game interesting and to minimize the barrier to entry, by providing the game server along with a software infrastructure that handles the agent-to-server interface and does the basic data management tasks. A simple “dummy” agent is included with this infrastructure, forming a foundation for more sophisticated agents. Teams whose agents do well in competition are strongly encouraged to make their agents available to the community⁶. As a result, a number of teams have provided both binary and source for working agents that are significantly more competitive than the dummy agent. The availability of top-performing agents enables novel types of research, such as the empirical game theory work of Jordan and Wellman (2007).

Research requires data, and both research and agent development depend on basic analysis tools. The game infrastructure includes a tool for parsing the logs produced by the game server. This tool provides a basic user interface that shows day-by-day activity in the procurement and sales markets, along with bids, offers, orders, inventory levels, factory utilization, bank account balances, and other data. For example, Figure 7 shows the production and inventory display. The log analysis tool is programmable with simple Scheme scripts to enable data extraction, and is distributed with sample code that dumps game data into a database for further analysis.

Figure 8 shows an example of a set of market-oriented analysis tools built by students at CMU on top of the basic logfile parser. This tool is useful for understanding market interactions among agents, such as market share and bidding behavior.

The MinneTAC agent (Collins et al., 2008, 2009a), shown schematically in Figure 9, is a complete, easily configured agent available in source form⁷. Building a working agent on the MinneTAC foundation is much less work than building a competitive agent on the lower-level framework that is distributed with the TAC SCM server. As we can see from Figure 9, MinneTAC uses a very different design approach from the other agents we have examined. The Repository acts as a “blackboard,” and the various components interact only through the Repository. The Oracle component is a wrapper for a large number of small modules, called “Evaluators”, that can be strung together as specified in a configuration file to do the necessary analysis and prediction tasks requested by the decision components. The actual coordination among decision components happens because they share some of those Evaluators. Specifically, both the Sales Manager and the Supplier Manager use sales quotas produced by one of the Evaluators.

⁶See <http://www.sics.se/tac/showagents.php> for the current list

⁷at <http://tac.cs.umn.edu>



Figure 7: Logfile analysis example.

Ultimately, the process of generating and analyzing data from TAC SCM simulations requires that many games be run. The simulation environment is designed to separate the server from the agents over network connections, and many agents are nearly compute-bound while they make their daily decisions. This means that running games requires either a large cluster configuration, or coordination of processes across multiple machines in a network environment. In addition, the high variability of the game scenario coupled with random behavior on the part of some agents may require analyzing large numbers of games to reach statistically interesting conclusions. For example, the game theory analysis described by Jordan et al. (2007) required over 12,000 games. The process of manually setting up and running such experiments is daunting even for the most dedicated graduate student. To address this problem, Collins et al. (2009b) describe a framework for managing multi-game experiments through a simple web-based user interface.

A primary factor that makes the TAC SCM scenario interesting and challenging is the high variability of market conditions within and across games. This variability drives up the number of simulations required to achieve statistical significance when comparing agents or agent configurations. Because the game server generates its random behavior using pseudo-random sequences, it is possible to “re-run” games when evaluating alternate agent configurations, as described by Sodomka et al. (2007). The result is a dramatic reduction in the number of games needed in an experiment design. The game server distributed through tac.cs.umn.edu supports this level of control.

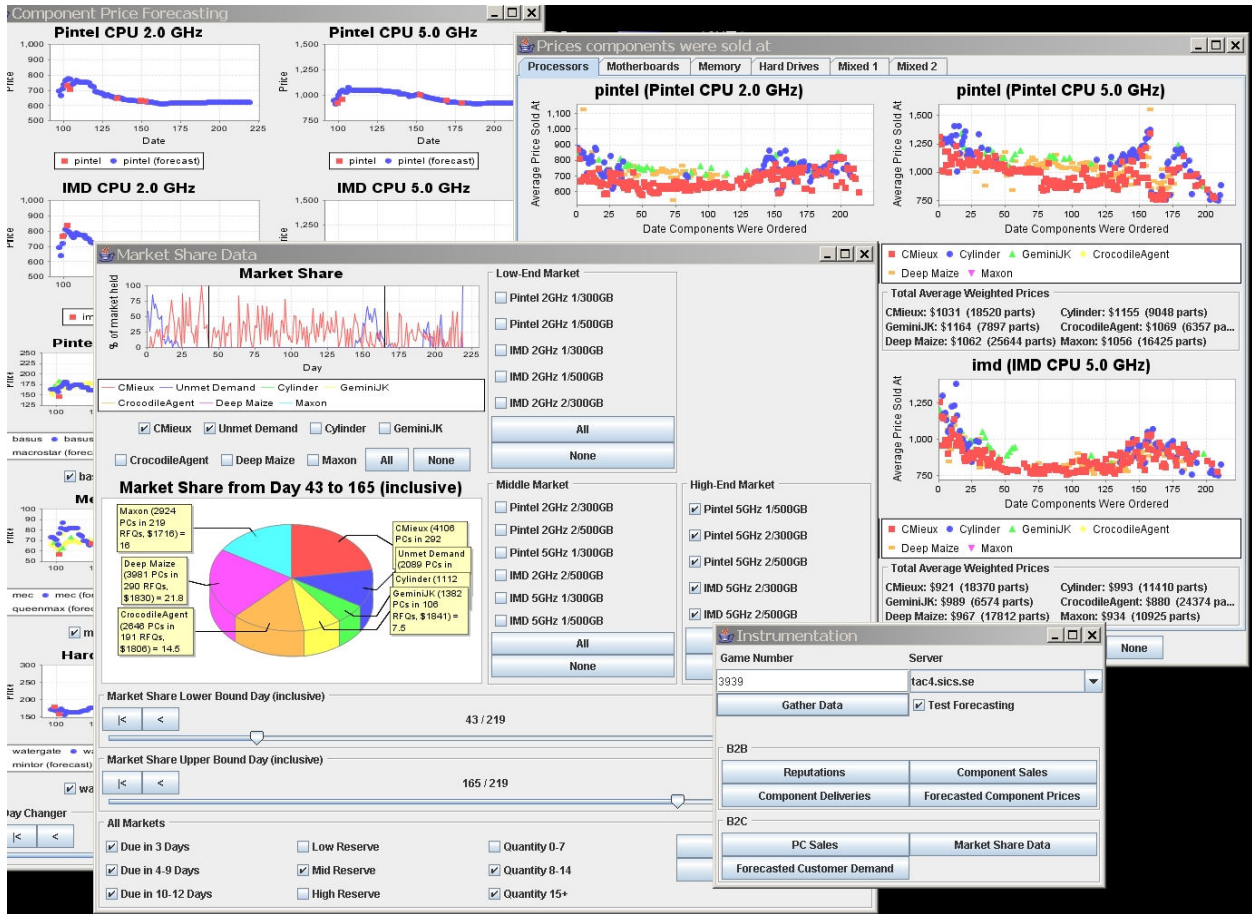


Figure 8: Game analysis tools made available to competitors by Carnegie Mellon.

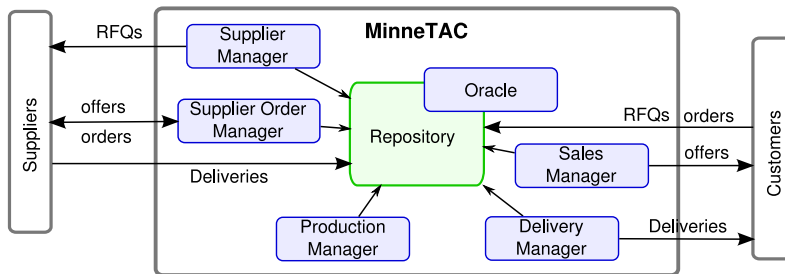


Figure 9: MinneTAC: Coordination through the Repository, details depend on configuration.

Conclusion and future work

Organized competitions such as TAC SCM have been effective tools for driving AI research into a range of interesting, complex domains that are both socially and economically important, and difficult for a single research team to address. The rapid rise of internet-enabled business interactions makes the supply-chain management domain, like many real-world problem areas, increasingly challenging for human decision making. At the same time, the complexity of such interactions is

beyond the analytic scope of formal game theory. The result is that evaluation of new approaches to decision-making is very difficult in isolation. The multi-year competition format, with active participation of motivated teams of researchers and regular publication of new techniques and results, makes such evaluation possible. At the same time, the need to work with management science and economics practitioners expands the scope of contributions from AI research and stimulates the AI community with new and interesting challenges.

Agents that operate effectively in the TAC SCM environment must be able to sense and model their environment and predict their own impacts on that environment. They must be able to deal with substantial uncertainty and limited visibility of the important features of their environment. They must maximize their utilities in expectation, while carefully managing risk. They must make a number of coordinated decisions within strict time constraints. Agents may engage in strategic interactions with their competitors, through manipulation of the shared environment. After seven years of competition, there is no clearly dominant approach to agent design and decision processes. This is evident in the results of the 2009 competition, in which the two top agents very nearly tied, using very different approaches. The DeepMaize agent from the University of Michigan uses very careful control of inventory and profitability over a wide range of market conditions, while the TacTex agent from the University of Texas is more aggressive and strategic, taking large risks in building up inventory during low-demand periods when procurement prices are low, and exploiting these inventories when demand recovers and procurement prices rise.

Over the past seven years, considerable progress has been made in developing effective techniques and architectures to manage risk in dynamic supply chain environments, with models from the TAC-SCM competition influencing ongoing work at large companies such as HP. While research in this area is far from over and there is still room for very significant advances, the TAC-SCM community will also have to turn its attention to packaging many of the technologies it has already developed into human-oriented decision support tools. These tools will have to allow supply chain managers to remain in control of key strategic decisions while delegating many more minute, real-time optimization decisions to agent-oriented functionality. To be effective, this line of work will have to identify a meaningful balance between the frequency and level of details in supply chain updates and decisions it exposes users to. This balance will have to be sufficient for supply chain managers to feel that they remain in control of key sensitive decisions where their own insight is critical. Yet it should not overwhelm users with information and decisions. Ultimately, developing mixed initiative functionality that meets these requirements will be critical to the broad uptake of TAC-SCM technology. Ongoing efforts in this area include work on a mixed-initiative version of the MinneTAC agent and the TAC SCM competition (Nelson et al., 2009). Early work on developing mixed initiative supply chain decision support functionality was also detailed by Sadeh et al. (2003b).

Acknowledgements

We would like to thank all research teams who have participated in TAC SCM over the years and made it a stimulating and highly research driven competition. Special thanks also to Sverker Janson and his group at the Swedish Institute for Computer Science for their contributions to the initial design and implementation of the TAC-SCM game. Another big thanks goes to the TAC community as a whole, which provides a platform for vivid exchange and feedback among the various researchers through an annual workshop on Trading Agent Design and Analysis (TADA).

The third author would also like to acknowledge the support provided for this work by the National Science Foundation under ITR Grant 0205435 as well as by a grant from SAP Research.

References

- Andrews, J., Benisch, M., Sardinha, A., and Sadeh, N. M. 2008. Using information gain to analyze and fine tune the performance of supply chain trading agents. In Collins, J., Faratin, P., Parsons, S., Rodríguez-Aguilar, J. A., Sadeh, N. M., Shehory, O., and Sklar, E., editors, *Agent-Mediated Electronic Commerce and Trading Agent Design and Analysis, 2007 Workshops, Selected and Revised Papers*, volume 13 of *Lecture Notes in Business Information Processing*, pages 182–199. Springer.
- Arunachalam, R. and Sadeh, N. 2005. The supply chain trading agent competition. *Electronic Commerce Research and Applications*, 4:63–81.
- Benisch, M., Andrews, J., and Sadeh, N. 2006a. Pricing for customers with probabilistic valuations as a continuous knapsack problem. In *Proc. of the 8th Int’l Conf. on Electronic Commerce*.
- Benisch, M., Greenwald, A., Naroditskiy, V., and Tschantz, M. 2004. A stochastic programming approach to scheduling in TAC SCM. In *Proc. of 5th ACM Conf. on Electronic Commerce*, pages 152–159, New York NY USA. ACM SIGecom, ACM Press.
- Benisch, M., Sardinha, A., Andrews, J., Ravichandran, R., and Sadeh, N. 2009. CMieux: adaptive strategies for competitive supply chain trading. *Electronic Commerce Research and Applications*, 8(2):78–90.
- Benisch, M., Sardinha, A., Andrews, J., and Sadeh, N. 2006b. CMieux: Adaptive strategies for competitive supply chain trading. In *Proc. of the 8th Int’l Conf. on Electronic Commerce*.
- Bensaou, M. 1999. Portfolios of buyer-supplier relationships. *Sloan Management Review*, pages 35–44.
- Collins, J., Arunachalam, R., Sadeh, N., Ericsson, J., Finne, N., and Janson, S. 2005. The supply chain management game for the 2006 trading agent competition. Technical Report CMU-ISRI-05-132, Carnegie Mellon University, Pittsburgh, PA 15213.
- Collins, J., Ketter, W., and Gini, M. 2008. Flexible decision support in a dynamic business network. In Vervest, P., van Liere, D., and Zheng, L., editors, *The Network Experience – New Value from Smart Business Networks*, pages 233–246. Springer Verlag.
- Collins, J., Ketter, W., and Gini, M. 2009a. Flexible decision control in an autonomous trading agent. *Electronic Commerce Research and Applications*, 8(2):91–105.
- Collins, J., Ketter, W., and Pakanati, A. 2009b. An experiment management framework for TAC SCM agent evaluation. In *Proceedings of the 2009 Workshop on Trading Agent Design and Analysis*, pages 9–13. AAAI, AAAI Press.
- Elmaghraby, W. J. 2000. Supply contract competition and sourcing policies. *Manufacturing and Service Operations Management*, 2(4):350–371.

- Eriksson, J., Finne, N., and Janson, S. 2006. Evolution of a supply chain management game for the trading agent competition. *AI Communications*, 19:1–12.
- He, M., Rogers, A., Luo, X., and Jennings, N. R. 2006. Designing a successful trading agent for supply chain management. In *Proc. of the Fifth Int’l Conf. on Autonomous Agents and Multi-Agent Systems*, pages 1159–1166.
- Jordan, P. R., Kiekintveld, C., and Wellman, M. P. 2007. Empirical game-theoretic analysis of the tac supply chain game. In *Proc. of the Sixth Int’l Conf. on Autonomous Agents and Multi-Agent Systems*, pages 1188–1195.
- Jordan, P. R. and Wellman, M. P. 2007. Best-first search for approximate equilibria in empirical games. In *Proceedings of the 2007 Workshop on Trading Agent Design and Analysis*, pages 11–18. AAAI.
- Ketter, W., Collins, J., and Gini, M. 2010. Coordinating decisions in a supply-chain trading agent. In Ketter, W., Poutre, H. L., Sadeh, N. M., Shehory, O., and Walsh, W., editors, *Agent-Mediated Electronic Commerce X and Trading Agent Design and Analysis VI*, Lecture Notes in Business Information Processing. Springer.
- Ketter, W., Collins, J., Gini, M., Gupta, A., and Schrater, P. 2007. A predictive empirical model for pricing and resource allocation decisions. In *Proc. of the 9th Int’l Conf. on Electronic Commerce*, pages 449–458, Minneapolis, Minnesota, USA.
- Ketter, W., Collins, J., Gini, M., Gupta, A., and Schrater, P. 2009. Detecting and Forecasting Economic Regimes in Multi-Agent Automated Exchanges. *Decision Support Systems*, 47(4):307–318.
- Ketter, W., Kryzhnyaya, E., Damer, S., McMillen, C., Agovic, A., Collins, J., and Gini, M. 2004. Analysis and design of supply-driven strategies in TAC-SCM. In *Trading Agent Design and Analysis Workshop TADA04*, pages 44–51.
- Kiekintveld, C., Miller, J., Jordan, P. R., and Wellman, M. P. 2006. Controlling a supply chain agent using value-based decomposition. In *Proc. of 7th ACM Conf. on Electronic Commerce*, pages 208–217.
- Kiekintveld, C., Wellman, M. P., Singh, S., and Soni, V. 2004. Value-driven procurement in the TAC supply chain game. *SIGecom Exchanges*, 4(3):9–18.
- Kleindorfer, P. and Wu, D. 2003. Integrating long- and short-term contracting via business-to-business exchanges for capital-intensive industries. *Management Science*, 49(11):1597–1615.
- Martinez de Albeniz, V. and Simchi-Levi, D. 2005. A portfolio approach to procurement contracts. *Production and Operations Management*, 14(1):90–114.
- Nagali, V., Hwang, J., Gaskins, D. S. M., Pridgen, M., Thurston, T., Mackenroth, P., Branvold, D., Scholler, P., and Shoemaker, G. 2008. Procurement risk management (PRM) at Hewlett-Packard company. *Interfaces*, 38(1):51.

- Nelson, A., Nyabuti, D., Collins, J., Ketter, W., and Gini, M. 2009. Toward human-agent competition in tac scm. In *Proceedings of the 2009 Workshop on Trading Agent Design and Analysis*, pages 63–72. AAAI Press.
- Pardoe, D. and Stone, P. 2007. Adapting in agent-based markets: A study from tac scm. In *Proc. of the Sixth Int’l Conf. on Autonomous Agents and Multi-Agent Systems*, pages 677–679.
- Pardoe, D. and Stone, P. 2009. Adapting price predictions in TAC SCM. In Collins, J., Faratin, P., Parsons, S., Rodriguez-Aguilar, J. A., Sadeh, N. M., Shehory, O., and Sklar, E., editors, *Agent-Mediated Electronic Commerce and Trading Agent Design and Analysis*, Lecture Notes in Business Information Processing, pages 30–45. Springer.
- Sadeh, N., Arunachalam, R., Eriksson, J., Finne, N., and Janson, S. 2003a. TAC-03 – a supply-chain trading competition. *AI Magazine*, 24(1):92–94.
- Sadeh, N., Hildum, D., and Kjenstad, D. 2003b. Agent-based e-supply chain decision support. *Journal of Organizational Computing and Electronic Commerce*, 13(3).
- Sadeh, N. M., Hildum, D. W., Laliberty, T. J., McA’Nulty, J., Kjenstad, D., and Tseng, A. 1998. A blackboard architecture for integrating process planning and production scheduling. *Concurrent Engineering: Research and Applications*, 6(2):88.
- Sandholm, T., Suri, S., Gilpin, A., and Levine, D. 2005. Cabob: A fast optimal algorithm for winner determination in combinatorial auctions. *Management Science*, 51(3):374–390.
- Sardinha, A., Benisch, M., Sadeh, N., Ravichandran, R., Podobnik, V., and Stan, M. 2009. The 2007 Procurement Challenge: a competition to evaluate mixed procurement strategies. *Electronic Commerce Research and Applications*, 8(2):106–114.
- Simon, H. A. 1979. Rational decision making in business organizations. *The American Economic Review*, 69(4):493–513.
- Sodomka, E., Collins, J., and Gini, M. 2007. Efficient statistical methods for evaluating trading agent performance. In *Proc. of the Twenty-Second National Conference on Artificial Intelligence*, pages 770–775.
- Sterman, J. 1989. Modeling managerial behavior: misperceptions of feedback in a dynamic decision making experiment. *Management Science*, pages 321–339.
- Sterman, J. D. 1992. Teaching takes off: Flight simulators for management education. *OR/MS Today*, 19(5):40–44.
- Swaminathan, J. M., Smith, S. F., and Sadeh, N. M. 1998. Modeling supply chain dynamics: A multiagent approach. *Decision Sciences*, 29(3):607–632.
- Wellman, M. P. 1996. The economic approach to artificial intelligence. *ACM Computing Surveys (CSUR)*, 28(4es).
- Wellman, M. P., Estelle, J., Singh, S., Vorobeychik, Y., Kiekintveld, C., and Soni, V. 2005. Strategic interactions in a supply chain game. *Computational Intelligence*, 21(1):1–26.

Wurman, P. R., Wellman, M. P., and Walsh, W. E. 1998. The Michigan Internet AuctionBot: A configurable auction server for human and software agents. In *Second Int'l Conf. on Autonomous Agents*, pages 301–308.

About the authors

John Collins is a lecturer in Computer Science at the University of Minnesota. He has over 30 years of industry experience in product development, research, and consulting. For the last ten years, he has been in academia, focused on trading agents and economic reasoning. He has been involved in TAC SCM from its inception. He led a major revision of the game specification in 2004 and 2005, served as Game Master for two years, chaired the 2007 TADA workshop, and is currently Operations Chair for TAC SCM, managing the game infrastructure and user support. He serves on the Board of Directors of the Association for Trading Agent Research, and on the editorial board of Electronic Commerce Research and Applications.

Wolfgang Ketter is Assistant Professor at the Department of Decision and Information Sciences at the Rotterdam School of Management of the Erasmus University. He received his Ph.D. in Computer Science from the University of Minnesota in 2007. He founded and runs the Learning Agents Research Group at Erasmus (LARGE). The primary objective of LARGE is to research, develop, and apply autonomous and mixed-initiative intelligent agent systems to support human decision making in the area of business networks, electronic markets, and supply-chain management. He has been involved in TAC SCM from its inception, was co-chairing the TADA workshop at AAAI 2008, the General Chair of TAC 2009, and is member of the Board of Directors of the Association for Trading Agent Research since 2009. His research has been published in various information systems, and computer science journals such as Decision Support Systems, Electronic Commerce Research and Applications, European Journal of Information Systems, INFORMS OR/MS Today and International Journal of Electronic Commerce. He serves on the editorial board of Electronic Commerce Research and Applications.

Norman Sadeh is a Professor in the School of Computer Science at Carnegie Mellon University. He is Director of the School's e-Supply Chain Management Laboratory and co-Director of its PhD Program in Computation, Organizations and Society. Norman's work in scheduling and supply chain management has resulted in tools and technologies fielded by organizations such as IBM, Raytheon, and the US Army as well as solutions commercialized by several major vendors (e.g. IBM, CACI, Numetrix - eventually acquired by JD Edwards, PeopleSoft and Oracle). He conceived and designed the TAC-SCM Competition in collaboration with the Swedish Institute of Computer Science and later worked with the first author to further refine the game's scenario. Norman received his PhD in Computer Science at Carnegie Mellon University with a major in Artificial Intelligence and a minor in Operations Research. He has authored over 160 publications and is also a Fellow of the Association for Operations Management. Other research interests include Web security and privacy, mobile and pervasive computing, the Semantic Web as well as social and policy issues associated with the emerging information society.