Decision Support Systems 57 (2014) 274–284

Contents lists available at ScienceDirect

Decision Support Systems

journal homepage: www.elsevier.com/locate/dss

Agent-assisted supply chain management: Analysis and lessons learned

William Groves [a,*], John Collins [a], Maria Gini [a], Wolfgang Ketter [b]

[a] Department of Computer Science and Engineering, University of Minnesota, United States
[b] Rotterdam School of Management, Erasmus University, The Netherlands

ARTICLE INFO

Article history:
Received 24 September 2012
Received in revised form 29 August 2013
Accepted 16 September 2013
Available online 25 September 2013

Keywords:
Decision support systems
Supply chain management
Key performance indicators
Economic simulation
Software agents
Trading Agent Competition

ABSTRACT

This work explores "big data" analysis in the context of supply chain management. Specifically we propose the use of agent-based competitive simulation as a tool to develop complex decision making strategies and to stress test them under a variety of market conditions. We propose an extensive set of business key performance indicators (KPIs) and apply them to analyze market dynamics. We present these results through statistics and visualizations. Our testbed is a competitive simulation, the Trading Agent Competition for Supply-Chain Management (TAC SCM), which simulates a one-year product life-cycle where six autonomous agents compete to procure component parts and sell finished products to customers. The paper provides analysis techniques and insights applicable to other supply chain environments.

1. Introduction

Supply-chain management is becoming increasingly sophisticated, driven by expectations of greater business agility and more closely coupled processes within as well as across organizations. For example, in *smart markets* participants have to use computational tools to understand the market characteristics and to anticipate market needs [1]. Central to smart markets are software agents with *adjustable autonomy*, i.e. programs capable of autonomous or semi-autonomous decision making, which are used to assist humans.

While many supply-chain business transactions are regulated by long-term contracts, it is critical to develop strategies that can anticipate and adapt to changes in the underlying market conditions [2]. Studies based on analytical methods simplify the problems (e.g., a single supplier and a single retailer) for tractability [3] and do not model how adaptive strategies affect market dynamics. Testing adaptive strategies in real-world markets can be risky, making it impossible to do the rigorous experiments needed to understand the interactions between competitors' strategies and the market. Competitive simulation environments, such as the Trading Agent Competition for Supply-Chain Management (TAC SCM), described in Section 3, overcome these obstacles and enable users to perform rigorous evaluation of strategies under varied market conditions. Agent modeling, which is at the core of these simulation environments, is more robust and flexible than centralized strategies based on traditional optimization methods in scenarios that have high complexity and uncertainty [4,5].

This work makes two major contributions. The first is to highlight the advantages of using a competitive simulation environment, where software agents compete in a market with other agents, for developing complex decision making strategies. Simulation and serious games are routinely used as educational support for students and professionals. Software agents are capable of playing such games (e.g., [6]) and can find good, often optimal, policies, but their use so far has been limited, despite their advantages.

The adoption of agent-based design forces the practitioners to specify precisely the decision processes and enables a thorough exploration of the interactions between competitors' strategies and market conditions. This kind of experimentation can elucidate unexpected, unintended, and undesirable high-level side effects of low-level business decisions. The phenomena we observe in TAC SCM emerge from the interactions among the competing agents and are not inherent in the design of the simulation. Running many simulations with different agents and market conditions is the only way to observe the effects of such interactions.

In TAC SCM, the software agents developed by different research groups are available for others to use. Since those agents use a variety of decision making methods, their availability makes TAC SCM a particularly rich testbed for exploration and assessment of decision strategies. For these reasons we think agent-based simulation will provide great dividends.

The second contribution is an extensive set of automated supply-chain key performance indicators (KPIs). Using TAC SCM as a testbed, we show how to compute the KPIs relevant to autonomous supply chains from data for guiding strategic, tactical, and operational decision making.

* Corresponding author at: 4-192 Keller Hall, 200 Union St SE, Minneapolis, MN 55455, United States.
E-mail addresses: groves@cs.umn.edu (W. Groves), jcollins@cs.umn.edu (J. Collins), gini@cs.umn.edu (M. Gini), WKetter@rsm.nl (W. Ketter).

0167-9236/$ – see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.dss.2013.09.006

The paper is organized as follows. We present related literature (Section 2) and introduce TAC SCM (Section 3). We describe briefly the KPIs (Section 4), followed by the KPIs for procurement (Section 5), production (Section 6), and sales (Section 7). Section 8 provides lessons learned from TAC SCM applicable to real-world supply chains. We conclude with thoughts on how these analysis tools could be applied to other markets.

2. Related work

Agent-based systems have long been proposed for managing business complexity (e.g., [7]). They are valuable for Smart Business Networks [8], which extend traditional business processes to systems that are integrated across organizations, where agents are used to support complex managerial decision making. Agent-based experimentation can find strategies for traditional operations research situations [4,9]. The advantages over non-agent methods are that agents' strategies are easier to understand due to the use of first principles in the design of the decision processes, do not assume a static environment, and are more robust to modeling errors.

The Association for Trading Agent Research (http://tradingagents.org) over the years has promoted the development of competitive agents for a variety of trading environments, including supply-chain management (TAC SCM) [10], price prediction [11], advertising auctions [12], and more. As a result, many teams have developed agents that have competed in the annual competitions and have shared them with the community. For a survey of decision strategies of many agents in the competition, consult [13].

In TAC SCM significant effort has gone into developing methods for agents to predict prices. Several agents use highly tuned prediction models, often based on specific features of the market, since price prediction is needed to drive a profitable sales strategy (e.g., [13,14]). Economic regimes have been used to characterize economic conditions, such as scarcity or oversupply, and to inform both pricing and resource allocation decisions [15,2].

Literature on trading agents in continuous double auction markets [16] provided inspiration on visualizing supply and demand. Visualization is valuable to support human decision making. For example, tools to generate dynamically a graphical presentation of decision processes improve agent-assisted decision-making in the overall supply-chain [17].

This work differs from previous work on TAC SCM because we focus on TAC SCM as a testbed for business automation instead of concentrating on how to develop an agent to win the competition.

While there is much existing literature on supply chain KPIs [18–20] and their analysis [21,22], this paper presents KPIs specifically relevant to autonomous supply chain processes. Some conventional KPIs are focused on measuring efficiencies of real-world operations that may involve process failures. These kinds of failures (i.e. data entry errors or equipment failures), while important in real-world processes, are already well studied in the literature. This work focuses on aspects of performance due to the oligopoly market, price risk, and supply and demand uncertainty.

3. Background on TAC SCM

TAC SCM [10] simulates a one-year product life-cycle in a three-tier supply-chain, with eight parts suppliers, six competing manufacturing agents that produce different models of computers, and many end customers. These agents compete to purchase parts in the *procurement market* and to sell the finished goods in the *sales market*. TAC SCM was inspired by companies such as Dell and Gateway Computer, both using a direct sales model [10].

Several aspects of TAC SCM make it relevant to the study of real supply-chain management systems. The market is an oligopoly of six manufacturers, so agent actions generate strong effects on the market. The supply and demand stochasticity is tuned to create situations of both excess supply as well as acute shortages. Supply and demand are only indirectly observable by the agents during the simulation, but agents affect each other indirectly through competition in the procurement and sales markets.

An agent must plan proactively: component parts must be ordered far in advance of delivery to minimize costs, but advanced purchases can force later sales that create losses. Sales commitments can only be made close to the delivery date, but the parts for upcoming sales should already be pre-ordered to take advantage of lower prices for long-term orders.

In both sales and procurement, there is uncertainty about which requests will actually be fulfilled due to the competitive actions of the other agents as well as market stochasticity. To handle times when demand exceeds factory capacity, agents can build inventories, but they incur an opportunity cost as the mixture of products built may not be optimal at sales time.

In TAC SCM, all communications are done via explicit market requests and responses. This enables the simulator to capture every interaction and provides a wealth of data that can be used after the simulation to visualize and analyze the activities and decisions of each agent.

Since multiple experiments can use the same set of agents, the same group of agents can be observed operating under different market conditions (e.g., levels of supply and demand), something not possible in real systems. To simplify experimental work and achieve statistical significance with fewer runs, we use a version of the simulation server [23] which supports repeatable pseudo-random sequences of the stochastic market factors. This enables testing of different strategies under repeatable market conditions.

While TAC SCM is not an exact model of any real world market, even the design decisions and hard limits imposed by the simulation are valuable aspects that are present in real markets. Market design properties of the simulation include an oligopoly, no direct observation of competitor behavior, a fixed component to product mapping, and no secondary market for goods. Hard limits include an oligopoly of exactly six agents, a preset simulation time limit, no availability of initial supplies, and lack of value of goods in the inventory at the end of the simulation. Removing those hard limits by randomizing experiments using a range of values would make numerical comparisons across experiments more difficult.

These hard limits are inspired by real-world markets and focus the competition on challenging aspects. For example, the oligopoly aspect causes agents to have greater market power when compared to scenarios with many agents, which is a less interesting case. When TAC SCM is adapted to match new market scenarios, these properties can be changed as well. The environmental responses to agent actions were designed to emulate emergent behaviors of real-world markets (i.e. greater demand for a component will drive up its price). A key difference that should not be discounted is the additional transparency made possible by the simulation: real-world market experiments are often impossible to repeat, and being able to measure emergent market phenomena through repeated observation is invaluable.

The TAC SCM simulation server is open source and can be adapted as needed. Most agents are also freely available, as source or binary code, so they can be used as competitors when experimenting with new decision making processes. Also, logs are publicly available from the simulation server.[1]

The main tools for this analysis are based on a SQL database which is populated by the simulation logs.[2] Each experiment run contains over 500,000 messages requiring uncompressed storage of around 100 MB. We use SQL queries against the database to generate plots of data of interest. Our analysis toolkit provides a collection of over fifteen modules

[1] http://tac.sics.se/.

[2] http://www.cs.umn.edu/~groves/tac/.