From: AAAI Technical Report WS-02-10. Compilation copyright © 2002, AAAI (www.aaai.org). All rights reserved.

Information and Product Quality Dynamics in Tiered Supply Networks

Leif M. Johnson and **Dr. Peter R. Wurman**
lmjohns3@eos.ncsu.edu, wurman@csc.ncsu.edu
Department of Computer Science
North Carolina State University
Raleigh, NC 27695 USA

Abstract

Economic supply networks can be viewed as complex systems composed of multiple, independent, interacting agents. Agents (companies) form relationships with one another by trading goods for other goods, services, or money. This paper presents a model for supply networks composed of a fixed number of discrete tiers of companies. Each company has mutable parameters (product quality, information integrity, profit sharing) that determine specific aspects of the company's behavior. After describing the model we apply it to investigate the global behavior of multi–tiered supply networks. In particular we are interested in how information integrity and product quality change under various circumstances, and how these parameters influence and are influenced by the global structure of a supply network. Preliminary experiments concentrated on finding parameter values that encouraged stability and found complex relationships among available revenue, profit sharing by individual companies, and initial wealth of new companies. Further investigations have revealed a few subtle trends, but they have emphasized the difficulty in finding more precise, yet globally applicable, data analysis techniques.

Introduction

Economic supply networks can be viewed as complex systems composed of multiple, independent, interacting agents. Agents (companies) form relationships with one another by trading goods for money or for other goods. Many previous works have examined the behavior of interactions between or among two or more independent agents in various situations (e.g. Axelrod (1984) and Lindgren & Nordahl (1994) with the Iterated Prisoner's Dilemma (IPD), Von Neumann & Morganstern (1947) with other general game theory investigations, and Greenwald & Stone (2001) with agents for electronic auctions).

The importance of effective supply chain management is highlighted by Cisco's recent writeoff of $2.2 billion of unneeded parts due to supply chain misinformation (Kaihla 2002). This event underscores the importance of accurate and precise information in supply chain relationships.

A significant amount of recent research has investigated supply chain issues from an IT perspective, one focus of which is the design of agents and protocols as active components in the supply chain (Chen *et al.* 1999; Collins & Gini 2001; Malone *et al.* 1988). These systems address the technical problems associated with finding contractors on the network, announcing RFQs, and placing bids in response. This perspective is focused on the construction of a single link in the chain.

Agent-based modeling has also been used to study inventory control through the supply network. Chandra & Chilov (2001), for example, present a model similar to the one we present, but they analyze the model with respect to what types of information need to flow among supply chain elements. Swaminathan, Smith, & Sadeh (1998) use supply chain modeling to optimize an existing supply network as an integrated part of the supply chain formation and management processes.

However, the majority of work on supply chains focuses on the management of individual chains. With a few exceptions, very little attention has been paid to the formation and dynamic properties of economies of supply networks. One such exception is the work by Walsh & Wellman (1998; 1999) who study the dynamic properties of market-based supply chain formation.

Our paper proposes a model for analyzing supply networks composed of discrete tiers of companies and examines the behavior of the model with respect to how information and product quality change under various circumstances. In particular, we are interested in the configuration and reconfiguration of networks that entail many supply chains which may interact with one another. This approach resonates with Parunak, Savit, & Riolo (1998), who argue the benefits of using agent based modeling in the analysis of supply chain behavior. In particular, Parunak, Savit, & Riolo contrast the agent based approach to the systems based approach, which is traditionally based on mathematical flow modeling.

Our work examines the supply chain from a higher level of abstraction than these previous studies. The motivation for the model described in this article stems from several questions, including:

1. How do different levels of information integrity affect relationships in a complex social environment? Do individuals with a propensity for inaccuracies or for correct information tend to cluster together? Answers to this question

Copyright © 2002, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

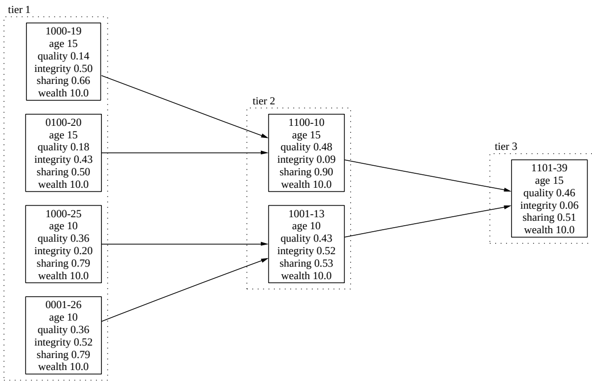

Figure 1: An example supply network with three tiers.

could potentially be related to the clusters of cooperators and defectors that one finds in the spatial iterated Prisoners' Dilemma (IPD), in which strategies are assigned locations in 2–dimensional space (e.g. Lindgren & Nordahl (1994)).

2. How do resources flow in an economic network where the individual agents in the network have some (limited) choice in their interaction partners?
3. How do bankruptcies, expansions, spin-offs, or start-ups affect the dynamics of economic supply networks?
4. Does a supply network that provides individual companies with a limited choice of interaction partners encourage the development of high or low information integrity on a global scale?

The bulk of this paper presents a model for supply networks that captures some, but certainly not all, of their interesting properties. Following the model description is a brief discussion of the experiments performed, followed by our results to date and directions for further investigation.

Model

The supply networks considered in this paper can be visualized as directed acyclic graphs in which companies (agents) are nodes and trade agreements are edges from a supplier to a customer. (The terms "network" and "graph" are used interchangeably in this paper.) Products flow along the graph edges, and money flows in the opposite direction. Figure 1 illustrates a simple example network configuration.

We restrict our model to tiered supply networks. Companies belong to one of a discrete set of T tiers that form a partition of the set of all companies. The "bottom tier" (tier 1) contains companies that might be viewed as raw material suppliers, while the "top tier" (tier T) contains companies that might be seen as retailers.

Companies act as individual agents in the supply network by changing certain parameters that govern the flow of money through the company; these parameters are described next. Money flow, on the other hand, is governed entirely by mathematical functions defined externally to any particular agent; these functions are described following the company parameters. It is these functions that form, in some sense, the economic landscape in which the agents exist. Agents may change certain aspects of their local landscape, however, during the economic cycle that is described as the last part of this section.

In the remainder of this paper, unless otherwise specified, we will be discussing a company i that belongs to tier t. The set of n companies that supply products to i is denoted by S, and the set of m companies that are customers of i is denoted K. When there is no ambiguity the subscripts on company parameters will be omitted.

Company Parameters

To keep the model simple, all companies are capable of selling only one type of product each. All products are represented in the model as binary strings of length $T+1$. Each unique binary 'word' represents a unique product, and the number of 1's in a product description is equal to the tier at which it is (possibly) produced. Thus, if company i produces product p, p must have exactly t 1's; if i is in tier 1 then p must have exactly one 1, and if i is in tier T then p must have exactly one 0. Companies are not allowed to change their product, so i will produce p as long as it is in the network.

To address the motivating questions, company i includes several parameters that may change over time. These are:

Product quality Company i produces p at a specific quality level q that is represented as a real number in the open interval $(0, 1)$. Product quality affects the flow of money in the network.

Information integrity Company i has an information system with a real valued integrity value l from the open interval $(0, 1)$. This parameter could be seen as the level of sophistication of i's information or records system, with values closer to 0 being less accurate.

Profit sharing Company i shares at most $\delta \in (0, 1)$ of its revenue r with each of its n suppliers each iteration. In this way, i retains at least $(1-\delta)$ of its revenue. In this version of the model, we assume that each supplier receives a maximum of an equal portion, $\frac{\delta r}{n}$, of the shared profit. This amount may be reduced for individual suppliers as described below.

To determine which companies are allowed to change their parameters (and which have gone bankrupt), each company i has wealth w which it can use to make investments in the company. Each iteration, i makes profits π (which might be negative if costs exceed revenue), which get added to the company's wealth at the end of the iteration. In addition, each company has a supply chain management parameter, $\alpha \in (0, 1)$, that determines to what extent it incorporates its trading partners' costs into its own optimization behavior, as described in the Experiments Section below.

Profit and Money Flow

Company i with quality q and information integrity l earns a profit each iteration given, in general, by the company's revenue minus the company's costs, $\pi = r - c(q, l)$ where the cost function is given by

$$c(q, l) = q + l + c_f + \sum_{j \in S \cup K} (1 - l_j)$$

for a fixed cost parameter c_f. The sum of information integrity values is over company i's trading partners j (either suppliers or customers). Thus the cost function is designed to be increasing with product quality and company information integrity, and decreasing with trading partner information integrity. The c_f could allow us to vary the cost model in different vertical markets or individual companies within the network. For experiments described in this paper, however, c_f was held at a constant value of 1.

The revenue function differs based on a company's tier; companies in the top tier of the supply network sell to abstract demand, while the demand for products below the top tier is created by the actions of companies in the tier above.

For top-tier product p, there is R_p total revenue available each iteration. The amount of revenue that company i will receive when producing p will depend upon the number and quality of the other producers of p. We let C be the set of c companies that participate in the market for p (note that $i \in C$). The revenue generated for i is given by

$$r(C) = R_p \left[(1 - k) \beta(C) + \frac{k}{c} \right]$$

where β is a weighting function and k is an exogenous parameter from the open interval $(0, 1)$ that guarantees at least $\frac{kR_p}{c}$ revenue for each company in C. For the experiments described here, $k = \frac{1}{3}$. Let $z_i = \sum_{j \in C} (q_i - q_j)^2$ be a metric of i's total quality distinction from its competitors. Then β is given by

$$\beta(C) = \frac{\frac{1}{n} + z_i}{1 + \sum_{j \in C} z_j}\,.$$

The weighting function is designed to encourage diversification in top tier markets by providing less revenue to companies whose quality values are more similar to their competitors' quality values.

All companies not in the top tier must earn profits through trade agreements. For a company i, recall that K is the set of i's customers. Then i's revenue is given by

$$r(K) = \sum_{j \in K} \delta_j \pi_j \frac{1 - (q_i - q_j)^2}{n_j}$$

where $\delta_j \in (0, 1)$ is customer j's profit sharing, and n_j is the number of companies supplying to j. Thus i will receive a maximum of $\frac{\pi_j \delta_j}{n_j}$ when $q_i = q_j$ and a minimum of 0 when q_i differs from q_j by 1. This revenue function is designed to capture the notion that companies that make high quality products need high quality inputs, and companies that produce low quality products don't need, and can't afford to use, high quality inputs.

Companies that do not acquire the necessary inputs do not earn any revenue; this revenue is lost and the market demand remains unsatisfied.

It should be noted that in the preceding discussion, as well as the rest of the paper, the modeling choices have been made primarily to set up the necessary tension in the marketplace between competing objectives. For example, the cost function creates an incentive to decrease quality and information integrity, but also to seek out trading partners with high information integrity. The revenue function repels top tier companies from each other, and pushes middle tier companies toward the quality levels of their customers.

Economic Cycle

Each iteration of the supply network consists of five steps: generation, negotiation, profit, investment, and removal. The steps proceed sequentially, and data are collected from companies at the end of the removal step. Each step is described in more detail below.

Generation. Initially the supply network starts with no companies. New companies are added to the top tier based on revenue available in the markets for top tier products. To add new companies, the following algorithm is used for each product p in the top tier:

1. Make a list C of companies that produce p.
2. Create a new company j that produces p.
3. Calculate j's profits $\hat{\pi}$ given the existing competitors in C and the existing total revenue available for p.
4. If $\hat{\pi} > \theta_{vc}$ then add j to C and go to step 2. Otherwise halt.

New companies start with a fixed initial wealth w_0, but their quality, integrity, and sharing parameters are each initialized to random values in $(0, 1)$. In this way, companies that do poorly go bankrupt and allow other companies with potentially better adapted parameter values to try to gain an economic foothold in the supply network. Because new companies begin with random parameter values, the supply network does not do any learning or adaptation on a global scale, though it would certainly be interesting to investigate how various learning algorithms could improve parameters for new companies. Note also that this process only generates companies in the top tier. Lower tier companies are created based on demand during the negotiation phase.

Negotiation. To exchange money and products, companies in a supply network negotiate trade agreements with one another. Companies in the model will arrange trade agreements only with companies that supply at least one of their input or output needs. Company i will be able to submit bids to all companies j in tier $t_i + 1$ whose product strings p_j share at least one 1 in the same position as p_i. Similarly, i will be able to receive bids only from companies j in tier $t_i - 1$ whose product strings p_j have at least one overlapping 1 with p_i. Companies are limited to supplying only one customer at a time, and they are not allowed to drop a current contract during the bidding process, but all companies that have supply needs are allowed to solicit bids, allowing companies to find new trading partners if their previous partners went bankrupt. Companies that have 1's in their product strings for which they do not have a supplier are not allowed to have customers. That is, a company must be adequately supplied before it may become a supplier itself.

The negotiation process begins with companies in the top tier and proceeds down to the bottom tier recursively. Consider, for example, a new company i in the top tier; it needs to have T product bits supplied to it, so i solicits bids from all companies in tier $T-1$ that have overlapping product strings. (Each of these potential suppliers has $T-1$ 1's in its product string, so i will in general need only two suppliers.) A bid is just a specification of the supplier j and the customer i. Company i then selects a combination of bids from this solicitation list that will satisfy its product needs. If, after this bidding process, i still requires a supplier, say for bit ϵ in its product string, then a new company j will be created in tier $T-1$ such that p_j and p_i share a 1 in bit position ϵ. To ensure that j does not supply i without itself being supplied, j solicits bids from companies in tier $T-2$ using the same process. This recursive company creation process continues until the bottom tier is reached. A similar recursive process takes place to create suppliers for any already existing company that still lacks product bits (either due to bad luck negotiating or to the loss of a trading partner).

Profit. Company i in the top tier earns revenue r in the market for p according to the equations presented earlier. Then i passes a portion of r on to its suppliers following the revenue equation for lower tier companies. In turn, each company j that receives revenue r_j from i passes a portion on to its suppliers. In this way, revenue trickles down to the bottom tier of the network, but only if the companies involved are adequately supplied.

Investment. After receiving profits from customers and sharing profits with suppliers, each company i adds its profits π to its wealth w. Then as long as i has wealth $w > w_0$, it is allowed to invest its profits by altering its company parameters; for each investment step, w decreases by w_0, so more profitable companies are allowed to invest more often than less profitable ones.

Clearly, this portion of the process is the focal point of our investigations to date. The company investment process is highly mutable and allows companies to perform a limited search of their economic situation to calculate potentially profit–increasing changes in company parameters. The specific changes allowed are described in the Experiments section below.

Removal. After investing, companies with negative wealth are considered bankrupt and are removed from the network. Since a bankrupt company no longer participates in the network, its trade agreements are voided. The broken contracts propagate to the tier T company that was downstream of the bankrupt company. In this way, a bankruptcy forces both downstream and upstream companies to renegotiate contracts and reevaluate their product supply needs. Although the bankruptcy forces recontracting within the particular supply chain in which a company was located, it may not immediately force other companies out of business: the other companies affected by the bankruptcy stay in business as long as their wealth exceeds their costs, giving them a chance to find new relationships in the network.

Analysis

Initial analysis concentrated on parameter values that encourage the development of stable supply networks. The parameters involved are the initial wealth w_0, the venture capital threshold θ_{vc}, the revenue R_p available for each top tier product p, and the individual companies' sharing parameters δ_i. These interactions are fairly straightforward mathematically, but the random nature of the initial companies' values throws some instability into the model.

The revenue R_p available for each top tier product clearly only has an impact on the total number of companies in the network. If, for example, we let C be the set of companies in the top tier that produce p, then C will be able to earn at least $\frac{kR_p}{c}$ and at most R_p each iteration, as mentioned earlier. This implies a relation for finding the maximum number of competitors $c_{\max}$ in any given market as

$$c_{\max} = \frac{kR_p}{\theta_{vc} + 1}$$

where the 1 in the denominator is the minimum possible cost for a company, as obtained from the cost function above.

Whereas the available revenue R_p determines mostly the number of top tier companies that will compete in the market for a product p, the venture capital threshold θ_{vc} and the sharing parameters δ_i have the most impact on the stability of the resulting networks (see the Results Section below). To find out what levels we need for a network with T tiers, we can calculate the amount of revenue that a top tier company i needs to pass on to its suppliers by creating a hypothetical supply tree below i. We assume companies need exactly two suppliers. Bottom tier companies y and z incur costs c_y and c_z, respectively, each iteration. So for a company v in tier $t_v = 2$ that is supplied by y and z in tier $t_y = t_z = 1$, $\delta_v r_v \geq c_y + c_z$ for y and z to avoid bankruptcy, so $r_v \geq c_v + \frac{c_y + c_z}{\delta_v}$ for a stable configuration. In turn, if u and v in tier 2 supply w in tier 3, then

$$r_s \geq c_s + \frac{r_u + r_v}{\delta_s} \geq c_s + \frac{c_u + c_v + \frac{c_w + c_x}{\delta_u} + \frac{c_y + c_z}{\delta_v}}{\delta_s} .$$

To guarantee the feasibility of a network, we assume the sharing parameters $\delta_i = \delta_{\min}$ and the costs $c_i = c_{\max}$ are at their minimum and maximum values, respectively. This gives a minimum necessary revenue

$$\theta_{vc} \geq c_{\max} \sum_{i=0}^{T-1} \left(\frac{2}{\delta_{\min}} \right)^i$$

for top tier companies. Given the cost function presented earlier, the maximum cost for any given company actually occurs when all companies have low information integrity, so we can set $c_{\max} = 1 + 0 + 1 + 3 = 5$. This is an important aspect of the supply networks in general, since individual companies face a conflict of interest in setting low information integrity for themselves but desiring high information integrity in their trading partners.

Experiments

The model described above was implemented using the Python programming language, and data were generated and collected on an Intel based PC. Runs of 100 iterations of the simulator generally took between five and sixty minutes to complete, depending primarily on the total amount of revenue available in the supply network.

Generation. We set $\delta_{\min}$ as discussed in our analysis to 0.5, so δ_i was chosen randomly for new companies from the open interval $(0.5, 1)$. Then based on the analysis, for a network with three tiers we set θ_{vc} slightly greater than 105, and for a network with four tiers, we set θ_{vc} slightly greater than 425.

Investment. During the investment phase of each iteration, companies with wealth $w > w_0$ are allowed to change two of their parameters (quality, q, and information integrity, l), as mentioned above. To do so, companies perform a sort of gradient search in the quality–integrity plane by calculating nine utility values: one value uses the current q and l values, and the other eight use the eight square grid points in a close neighborhood of q and l (the set of grid points used is $\{q-h, q, q+h\} \otimes \{l-h, l, l+h\}$). Here, h is the neighborhood size; it was fixed at 0.1 for these simulations.

Recall that i has m suppliers and n customers, and that α is i's supply chain management parameter. Then the calculated utility values are specified by

$$U_{(q',l')} = r\left(q', l'\right) - \alpha\left(m+n\right)\left(1-l'\right) - c\left(q', l'\right)$$

where r and c are the revenue and cost functions described previously. Values of α closer to 1 will tend to make i incorporate more of its $m+n$ trading partners' losses for l' being closer to 0. These utility values are designed to reflect a company's projected profits; after evaluating them, i sets its actual quality and information integrity parameters to those of the corresponding grid point with the highest utility.

Results

Experiments to date have focused on the stability of the networks given this limited company flexibility. In the course of these investigations we found a few general trends that are interesting but will likely need more in-depth research to clarify. Due to the large grain size of our results to date, we refrain for the most part from drawing too many conclusions about the model; very specific conclusions would be premature at this point. Nonetheless, there are some interesting trends worth noting.

The first set of experiments we performed concentrated simply on finding values for global network parameters (w_0, R_p, $theta_{vc}$, and δ_i for each company) that would encourage stable network formation. We found that such parameter combinations are extraordinarily difficult to locate, and it is likely that even parameter combinations that provide stable results for one seed of the random number generator will have a chance of providing unstable results for a different seed.

In general, as anticipated the initial wealth w_0 provides companies with more opportunity to survive in the absence of trading partners. An interesting consequence of this is values for this parameter that allow companies to survive even one turn without being fully supplied (i.e. $w_0 > 5$) tended to make the companies in the top tier of the network last the entire duration of the simulation. This is closely related to the recursive company generation process described above, since it is nearly impossible for a top tier company to lack product bits after a negotiation process.

The other two parameters generally followed the predictions given in the analysis section, with a few notable exceptions. First, we found that stable networks were capable of forming even with a venture capital threshold far below the theoretical values calculated in the Analysis Section above (e.g. stable 4–tier networks formed with $R_p = 475$ and $\theta_{vc} = 200$ formed stable configurations by iteration 75). This is clearly possible since we used in our calculations the minimum δ_i and the maximum cost possible; in simulations the actual values vary and allow for lower costs and higher profit sharing. Although it was fairly difficult to form a stable network, when stable networks did form, they remained stable until the end of the simulation. This is encouraging and hints at the possibility of creating shocks in specific markets to investigate the networks' reactive behavior.

Another interesting trend in the simulations we performed involved the companies' motions in q-l–space: in general, companies tended to move to lower product quality values, while the information integrity values reveal two attractors, one at $l = 0$ and the other at $l = 1$. Figure 2 shows web-type plots of the quality and information integrity values in the supply network over the course of the simulation. In these graphs, a company's parameter value at iteration x is plotted on the horizontal axis, and its value at iteration $x+1$ is plotted on the vertical axis. In this way we see that the positively sloped diagonal line through the center represents companies that retain the same parameter values over time, while those below the line reduce their parameter values, and those above the line increase their values. It is interesting to note that in the information integrity graph, companies with high information integrity tend to increase their information integrity (more points are above the diagonal line in the high value region), while those with already low information integrity tend to decrease their information integrity.

Future Work

Although the supply network model described in this paper is sophisticated enough to capture interesting behavior in economic supply networks, there are not many solid conclusions to be drawn as of yet. The results described above hint at some of the desired richness of behavior even though the space the agents are allowed to explore is quite limited. Further research in the coming months will explore the effects of changing the model in a variety of ways. In particular, we anticipate several axes of investigation that will be interesting.

First, we feel that a substantial amount of work on data analysis techniques could yield great rewards in terms of revealing local trends in these supply networks that the current global analysis techniques might tend to smooth out. In particular, as mentioned at the beginning of the paper, we are

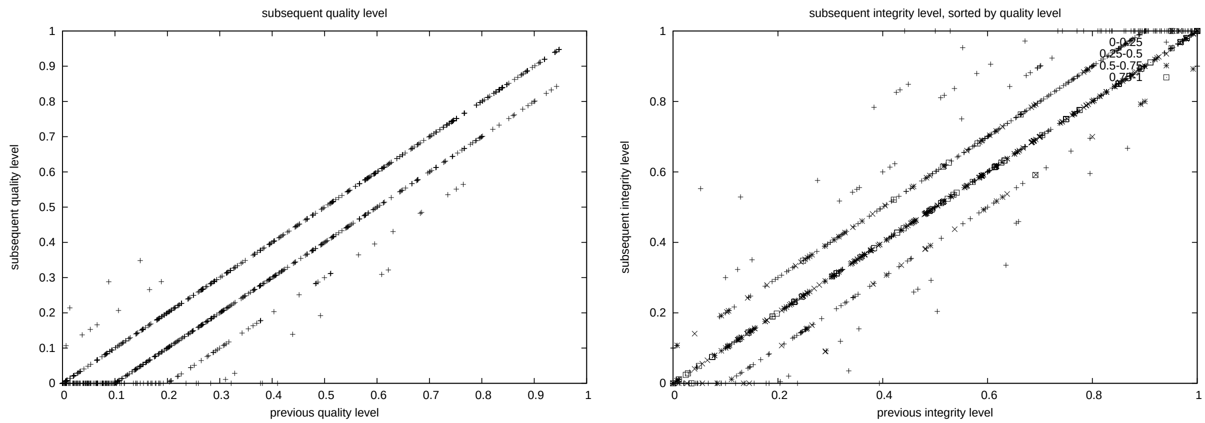

Figure 2: Web graphs of quality (left) and information integrity (right) for companies in successive iterations of the supply network.

interested in finding local areas of particular types of behavior (e.g. local pockets of high quality traders in a network composed primarily of low quality companies). Although our results to date are encouraging, we feel they are covering up more interesting local behavior.

Second, as mentioned previously, it will be interesting to investigate the results of creating shocks (e.g. sudden removal of one middle tier product, sudden introduction of many low or high quality competitors in a market, etc.) in an otherwise stable supply network.

Finally, the model presented here, although it is already rather complex and displays some rich behavior, could clearly be expanded to take into account other types of behavior. In particular, trade agreements could expire after a variable number of iterations (either random or known at negotiation time). Another possible expansion of the model could include a more complex gradient search method during investment, possibly including other company parameters.

Acknowledgements

We would like to thank our blind reviewers' comments for their valuable suggestions on ways to improve our paper.

References

Axelrod, R. M. 1984. *The Evolution of Cooperation*. New York: Basic Books, Inc.

Axelrod, R. M. 1997. *The Complexity of Cooperation: Agent-based Models of Competition and Collaboration*. Princeton, NJ: Princeton University Press.

Chandra, C., and Chilov, N. 2001. Simulation modeling for information management in a supply chain. In *Proceedings of the twelfth Annual Conference of the Production and Operations Management Society, POM–2001, March 30–April 2, 2001, Orlando Fl.* Production and Operations Management Society.

Chen, Y.; Peng, Y.; Finin, T.; Labrou, Y.; Cost, S.; Chu, B.; Yao, J.; Sun, R.; and Wilhelm, B. 1999. A negotiation-based multi-agent system for supply chain management. In *Agents '99 Workshop on Agents for Electronic Commerce and Managing the Internet-enabled Supply Chain*.

Collins, J., and Gini, M. 2001. A testbed for multi-agent contracting for supply-chain formation. In *Agents '01 Workshop on Agent-based Approaches to B2B*.

Greenwald, A., and Stone, P. 2001. Autonomous bidding agents in the trading agent competition. *IEEE Internet Computing*.

Kaihla, P. 2002. Inside cisco's $2 billion blunder. *Business 2.0*.

Lindgren, K., and Nordahl, M. G. 1994. Cooperation and community structure in artificial ecosystems. *Artificial Life* 1:15–39.

Malone, T. W.; Fikes, R. E.; Grant, K. R.; and Howard, M. T. 1988. Enterprise: A market-like task scheduler for distributed computing environments. In Huberman, B. A., ed., *The Ecology of Computation*. North Holland.

Parunak, H. V. D.; Savit, R.; and Riolo, R. L. 1998. Agent-based modeling vs. equation-based modeling: A case study and users' guide. In *Multi-agent Systems and Agent-based Simulation (MABS '98)*, volume LNAI 1534, 10–25. Springer.

Swaminathan, J. M.; Smith, S. F.; and Sadeh, N. M. 1998. Modeling supply chain dynamics: A multiagent approach. *Decision Sciences* 29(3):607–632.

Von Neumann, J., and Morganstern, O. 1947. *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press.

Walsh, W. E., and Wellman, M. P. 1998. A market protocol for distributed task allocation. In *Third International Conference on Multiagent Systems*, 325–332.

Walsh, W. E., and Wellman, M. P. 1999. Efficiency and equilibrium in task allocation economies with hierarchical dependencies. In *Sixteenth International Joint Conference on Artificial Intelligence*, 520–526.