

Agent-Facilitated Real-Time Flexible Supply Chain Structuring

Daniel Dajun Zeng
Department of Management
Information Systems
University of Arizona
Tucson, AZ 85721
+1 520-621-4614
zeng@bpa.arizona.edu

Katia P. Sycara
The Robotics Institute
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213
+1 412 268 8825
katia@cs.cmu.edu

Abstract

Electronic commerce and the vast amounts of real-time information available through means of EDI and the Internet are reshaping the way enterprises conduct business. A new computational infrastructure and models are needed for a business to gain a competitive edge through effective use of this information base. In this paper, we develop a model of inter-organizational electronic commerce that explores various new choices and opportunities that the electronic marketplace offers. We explicitly address two major performance measures of supply chains—time and cost—in a unified fashion. In our work, we model different business entities as autonomous software agents interconnected via the Internet. The main research focus of our efforts is how to coordinate software agents in supply chains dynamically and flexibly such that goods and services can be delivered at the right time in a cost-effective manner.

1 Introduction

Supply chain management is very different from managing one isolated site. Business entities operating in various supply chains usually have complex interdependencies and interrelationships. Lee and Billington described many pitfalls of managing supply chains, one of which is directly related with “inefficient information systems” [18].

Thanks to the recent spectacular and fast-paced development of computing and network technology such as Internet and Software Agents, these deficiencies in enterprise information systems are quickly disappearing. Electronic commerce and the vast amounts of real-time information readily available from internal and external data sources are reshaping the way that enterprises conduct business [7, 16, 21]. From the business perspective, electronic commerce provides many new opportunities:

- Electronic commerce provides a new media for marketing and selling products or services. Web-based marketing and on-line shopping are two prominent examples.
- Locating product and service providers is much easier and more efficient in the electronic marketplace. Many online businesses specialize in delivering electronic catalog services and performing other intermediary functions such as business and product yellow-pages and matchmaking. Emerging computing paradigms such as Internet-based software agents also makes locating and accessing information increasingly easier. These software agents are able to proactively search various information sources according to the user’s or business’ specifications or preferences, to monitor these sources for new changes and developments, and to integrate information from various sources into a consistent set of results to be reported to the user [24]. This ease and efficiency in information search and retrieval contribute to the decrease in transaction costs.
- The electronic marketplace enables fast, easy, and reliable communication of data between different sites. It provides information infrastructure for business-to-business cooperation. In addition to support for simple transactional data exchange, the electronic marketplace and its associated network infrastructure and software specifications make it possible for businesses to form strategic alliance and share information and collaborate on a real-time basis [15]. With the emergence of the software agent based system development paradigm, the interoperability among distributed decision support systems (these systems may be developed and deployed by different businesses) becomes less of a problem.

How to take full advantage of these new opportunities and compete effectively in the electronic marketplace poses significant challenges to practitioners and researchers. In this

paper, we focus on inter-organizational electronic commerce in the context of supply chain management. In our study, we explicitly consider multiple decision criteria including both leadtime and cost. Our model is partially motivated by the observation made by many researchers and practitioners that the nature of competition in electronic commerce does not resemble undifferentiated Bertrand competition. By leadtime, we mean the amount of time that elapses from the instant that an order (or service request) is placed until it arrives. By cost, we mean the sum of the costs of all activities required to satisfy the order (or deliver the service). Some examples of these supply chain activities and decisions are supplier selection, subcontracting, transportation mode selection, production rate decision, etc.

In our study, we model different business entities as autonomous software agents interconnected via the Internet. These agents act on behalf of their human users/organizations in order to perform laborious information gathering tasks, such as locating and accessing information from various on-line information sources, filter away irrelevant or unwanted information, and provide decision support. Section 1.1 discusses some of the related literature. Section 2 presents a brief description of an AND/OR network-based supply chain model. This model needs to be integrated with other operational level decision making models such as inventory management to enable intelligent agents to make the full range of supply chain decisions. We present in Section 3 this integrated model and a computational study of making inventory decisions that take advantage of multiple leadtime/cost options. We conclude the paper in Section 4 by summarizing the results and pointing out other extensions to our agent-based multi-issue supply chain model.

1.1 Related Literature

Effective use of the Internet by individual users, organizations, or decision support machine systems has been hampered by some dominant characteristics of the Infosphere. Information available from the net is unorganized, error-prone, multi-modal, and distributed on server sites all over the world. The notion of Intelligent Software Agents (e.g., [26, 24]) has been proposed to address this challenge. In this paper, we model a supply chain as a multi-agent system where different business entities interact with one another through intelligent software agents that act on their behalf. In general, multi-agent systems can compartmentalize specialized task knowledge, organize themselves to avoid processing bottlenecks, and can be built expressly to deal with dynamic changes in the agent and information-source landscape. In addition, multiple intelligent agents are ideally suited to the predominant characteristics of the Infosphere (and in particular supply chain management), such

as the heterogeneity of the information sources, the diversity of information gathering and decision support tasks that the gathered information supports, and the presence of multiple users/organizations with related information and decision aiding needs.

In order for autonomous software agents to make sensible decisions in any nontrivial domain such as supply chain management, they need to have access to domain-specific decision making models and related computational mechanisms [19, 23, 6, 25, 27]. We briefly survey some of these models in literature that are most relevant to multi-issue (time and cost) supply chain management.

In manufacturing, the current focus on time-based competition strongly suggests the importance of and substantial interest in tradeoffs between time and cost [3, 4]. Time-based competitors focus on the entire value-delivery system. Porter offers insightful conceptualization of value chains [20]. He argues that competitive advantage lies not only in activities themselves but in the way activities relate to each other, to supplier activities, and to customer activities. Time responsiveness is an essential part of companies' core competitive strategies.

The first models that consider the possibility of purchasing shorter leadtimes at a premium cost appeared in [5] and [9], among others. The main objective of these papers is to find the optimal ordering policy that minimizes ordering, holding and penalty costs when subject to random demand. Structural results regarding the optimal replenishment policy were established when there are only two options and the leadtimes of the two options differ by one time unit.

Song and others in [22] studied the impact of stochastic *leadtimes* on the optimal inventory decisions and the optimal cost in a base-stock inventory model. The focus there is to evaluate the impact of the variability of leadtimes but not to derive an inventory policy which makes use of the availability of multiple leadtime/cost options. In [17] Lau and Zhao considered the order splitting between two suppliers that offer different leadtime with uncertainty. The authors assumed a constant splitting ratio among two suppliers and developed computational methods to compute the optimal ratio, ordering quantities and reordering point in a continuous review inventory setting. Several papers (e.g., [2]) deal with situations where leadtime is one of the decision variables. Their assumption is that by paying "leadtime crashing cost" leadtime reduction can be achieved. The goal of these papers is to find the single best leadtime option under single sourcing.

2 An AND/OR Network-based Supply Chain Model

We have developed a supply chain model, called LCT, based on an AND/OR network representation [13, 14]. This model is capable of capturing a variety of supply chain activities and decisions.

In LCT, a supply chain is modeled as a directed acyclic graph with parallel arcs. The model follows an activity-on-arc representation where each arc corresponds to a particular supply chain activity (production, transportation, subcontracting, etc.). Each activity/arc has two performance measures: leadtime and cost. In this supply chain network, nodes represent completion of activities and may be used to establish precedent constraints among activities. The graph is directed towards one particular “root node”. The root node corresponds to the retailer of the product that the supply chain produces. End customers interact with the root node only. Two types of nodes are defined to specify conditions for satisfying prior activities: *conjunction* and *disjunction* nodes. Conjunction nodes or AND nodes are nodes for which *all* the activities that correspond to the incoming arcs must be accomplished before the outgoing activities can begin; whereas disjunction nodes or OR nodes requires that *at least one* of the incoming activities must be finished before the outgoing activities can begin.

Based on LCT we have developed efficient computation methods to identify the entire efficient frontier between leadtime and cost in supply chains. This efficient frontier at the “root node” of the supply chain, i.e., the retailer point, compactly represents all the undominated, feasible combinations of supply chain activities. By a feasible combination of supply chain activities, we mean the set of activities that guarantee the availability of goods or services at the root node. We say a combination dominates the other when the former offers cheaper cost and shorter leadtime than the latter. Given the efficient frontier coupled with the market demand profile and pricing strategy at the root node, management can converge on the optimal tradeoff point specifying a particular supply chain configuration.

One of the limitations of the basic LCT model is that the model does not explicitly consider inventory. The LCT model suffices for “one-shot” scenarios in which at each decision-making episode only single-period demand is taken into account. In cases where multi-period or stochastic demand is considered, holding inventory at one or more places in the supply chain clearly has the potential of improving the performance of the whole system. We discuss the integration of the LCT model with inventory management in Section 3.

3 Leadtime-Cost Tradeoffs in A Stochastic Inventory Model

For simplicity, we assume that inventory can be held only at the root node of the supply chain network. In other words, we are concerned with managing inventory at a single echelon within the context of the LCT model. Since we only add the inventory capacity at the root node, the entire leadtime cost efficient frontier for the supply chain remains the same. We are interested in ways through which the end product retailer (and his agent) can take advantage of the availability of multiple efficient supply chain configurations offering various leadtimes and costs.

Integrating LCT with inventory models with constant demand rates such as EOQ types of models can be easily done. When we consider inventory management with uncertain demand, however, the situation changes dramatically. In this section, we first present a formal formulation of the problem.

As is standard practice in the inventory literature, we start with the finite horizon case. We study an N -period stochastic inventory problem in which there are m different ordering options. These options represent different leadtime and cost tradeoffs which can be computed using the LCT model given the network topology of a supply chain and time/cost information for the supply chain activities. We ignore the setup cost in this study¹. We use the following notation in our study.

N = the number of periods in the planning horizon.

m = the number of ordering/delivery options.

λ_i = the leadtime associated with option i , $i = 1, 2, \dots, m$.

We assume that these leadtimes are deterministic. Without loss of generality, we assume that $\lambda_i < \lambda_j$ when $i < j$.

τ = the maximum leadtime from all possible ordering options, $\tau = \lambda_m$.

c_i = the unit ordering cost with option i , $i = 1, 2, \dots, m$.

Without loss of generality, We assume that $c_i > c_j$ when $i < j$.

T = the random demand for the item during each time period. We assume that the demand is stationary and follows a known probability distribution $f(\cdot)$. Further, we assume that the demand is nonnegative and the demand distribution has a finite mean. We denote by t a generic value sampled from this demand distribution.

α = the one-period discount factor.

¹This is not entirely arbitrary since electronic commerce contributes to the setup cost reduction.

x_0 = the current stock level (on-hand inventory).

$x_1, x_2, \dots, x_{\tau-1}$ = the pipeline inventory from orders already placed. They are the outstanding orders such that x_1 is due at the start of the next period, x_2 is to be delivered two periods hence, etc.

z_i = the amount of goods to be ordered at the start of the present period using option i , $i = 1, 2, \dots, m$. These are the inventory decision variables.

$L(x)$ = the expected operational costs during the period, exclusive of ordering costs, with respect to the stock on hand before the demand occurs.

$$L(x) = \begin{cases} \int_0^x h(x-t)f(t)dt + \int_x^\infty p(t-x)f(t)dt & x > 0, \\ \int_0^\infty p(t-x)f(t)dt & x \leq 0 \end{cases} \quad (1)$$

We assume that the holding cost $h(\cdot)$ and penalty cost $p(\cdot)$ are nondecreasing and convex. The unfulfilled orders are backlogged. Since integration preserves the convexity, the convexity of $L(x)$ is easily seen.

$C_n(x_0, x_1, \dots, x_{\tau-1})$ = the minimum expected cost following an optimal policy, given that only n future periods are to be taken into account. The state vector $(x_0, x_1, \dots, x_{\tau-1})$ represents all the information about the current stock level as well as the amounts of goods whose orders have been submitted and are to be delivered during the following $\tau - 1$ periods.

The state space for this inventory problem with multiple ordering options is

$$\{(x_0, x_1, \dots, x_{\tau-1}) \mid x_0 \in R, x_i \in R^+ \text{ for } i = 1, \dots, \tau - 1\}$$

where the set of real numbers is denoted by R and the set of nonnegative real numbers is denoted by R^+ . The control space is

$$\{(z_1, z_2, \dots, z_m) \mid z_i \in R^+ \text{ for } i = 1, \dots, m\}$$

Given this system dynamics, the functional equation for $C_n(\cdot)$ is:

$$C_n(x_0, x_1, \dots, x_{\tau-1}) = \min_{z_i \geq 0, \text{ for } 1 \leq i \leq m} \left\{ \sum_{i=1}^m c_i z_i + L(x_0 + y_0) + \right.$$

$$\left. \alpha \int_0^\infty C_{n-1}(x_0 + y_0 - t + x_1 + y_1, x_2 + y_2, \dots, x_{\tau-1} + y_{\tau-1}, y_\tau) f(t) dt \right\}$$

where y_j is defined as follows:

$$y_j = \begin{cases} z_i & \text{if } j = \lambda_i, \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

We use y_j to indicate the amount of goods to be delivered j periods into the future which can be positive only if there exists an ordering mode i that offers that leadtime option j .

For this stochastic inventory model with multiple lead-time/cost replenishment options, we can prove the convexity of the value function [29]. Nevertheless, it seems extremely difficult to identify any simple structures from optimal replenishment policies. The complexity of the problem stems from the fact that although the convexity of the value function holds, the optimal controls appear to depend on all the x_i for $i = 0, \dots, \tau - 1$ rather than depend on some aggregated state such as the echelon inventory level $\sum_{i=0}^{\tau-1} x_i$.

Using the value iteration approaches in dynamic programming, we can compute the optimal control policies regardless of whether they follow the order-up-to structure or not. However, these value iterations approaches are almost impossible to scale up since the size of state space itself is exponential with respect to the maximum leadtime. For practical purposes, we need to find other more efficient algorithms.

To address these computational issues, we have developed several suboptimal polices that are easy to compute. Our primary interest is to optimize infinite-horizon average cost per stage. In Section 3.1, we report these policies and an experimental evaluation of their performances.

3.1 Inventory Policies with Multiple Lead-time/Cost Options

One way of avoiding representing the exponentially large state space is to apply the idea of *parameterizing the control policies*. We consider a class of control policies that can be characterized by a small number of parameters. Based on these parameters, we can easily deduce the corresponding control for any given arbitrary state. For instance, consider a one ordering option inventory model, the optimal policy—order-up-to policy—is characterized by one parameter—order-up-to level S . For our general m -mode inventory problems, the order-up-to policy or its variants have been proved nonoptimal. The simplicity of the order-up-to policy and its wide acceptance in practice, however, have rendered itself as a good candidate suboptimal policy. In our study, we consider the following policy fashioned after the order-up-to policy.

$$z_i = (S_i - \sum_{j=0}^{\lambda_i} x_j - \sum_{j=1}^{i-1} z_j)^+ \quad \text{for } i = 1, \dots, m - 1 \quad (4)$$

$$z_m = (S_m - \sum_{j=0}^{\tau-1} x_j - \sum_{j=1}^{m-1} z_j)^+ \quad (5)$$

We call this an “echelon order-up-to policy” due to the similarity between this policy and the policies developed in multi-echelon inventory research [8]. Intuitively, we can imagine that for each leadtime option, there is an “order-up-to” level that guarantees that the sum of future arrivals and committed orders using the current mode and quicker options reaches a predetermined level. The rationale behind this is similar to that behind Fukuda’s policy for two adjacent (in terms of leadtimes) options [9].

Although the echelon order-up-to policy is easy to compute and can be used to solve large-size problems, the resulting solution quality measured by the average per stage cost is not entirely satisfactory. To achieve lower and closer-to-optimal costs, we develop the following policy which is a generalization of the echelon order-up-to policy.

$$z_i = (\beta_i - \gamma_i (\sum_{j=0}^{\lambda_i} x_j + \sum_{j=1}^{i-1} z_j))^+ \quad \text{for } i = 1, \dots, m-1 \quad (6)$$

$$z_m = (\beta_m - \gamma_m (\sum_{j=0}^{\tau-1} x_j + \sum_{j=1}^{m-1} z_j))^+ \quad (7)$$

Recall that when we follow an echelon order-up-to policy, we weigh all the x_j and z_j equally for all the options. We analyzed a number of the optimal policies (for small-sized problems) produced by policy iteration or value iteration and observed that equally weighing x_j and z_j for all the modes clearly violates optimality. In the meantime, to a large extent, the points in the high-dimensional space of $(x_1, x_2, \dots, x_m, z_j)$ for $j = 1, 2, \dots, m$ can be separated by hyperplanes. These hyperplanes are not necessarily parallel. This motivates the development of the separating plane policy. This policy, like the echelon order-up-to policy, does not require explicit representation of the state space and therefore can be used in solving large-size problem instances. The separating plane policy involves more parameters than the echelon order-up-to policy and therefore are more expensive to compute. Note that any echelon order-up-to policy is a special case of a separating plane policy when all γ_i are equal to 1. The performance of the best separating plane policy will be at least as good as the best echelon order-up-to policy.

3.2 Experimental Comparisons

We use the echelon order-up-to policy computation to illustrate the major computational steps in our experiments. The separating-plane policy computation essentially follow the same steps.

The objective is to minimize the infinite horizon average cost per stage. We denote this cost by V . Given a particular

set of echelon order-up-to levels (S_1, S_2, \dots, S_m) , we know precisely how much to order using which mode for any arbitrary state. Under mild assumptions, V does not depend on the initial state. In addition, any meaningful order-up-to levels have to satisfy non-negativity and monotonicity conditions. Therefore, for the purpose of finding the best echelon order-up-to policy, we are solving the following multi-dimensional optimization problem:

$$\min_{S_1, S_2, \dots, S_m} V(S_1, S_2, \dots, S_m)$$

$$\text{subject to: } 0 \leq S_1 \leq S_2 \leq \dots \leq S_m \quad (8)$$

Since the analytic form of V in terms of echelon order-up-to levels is unknown, we use simulation to evaluate candidate order-up-to policies. Within simulation itself, we use infinitesimal perturbation analysis[10, 12] to obtain estimated sensitivities of V with respect to (S_1, S_2, \dots, S_m) .

A modified gradient projection method [11, 1] is used to carry out the outer loop optimization over (S_1, S_2, \dots, S_m) which in turn uses simulation to obtain estimated $V(S_1, S_2, \dots, S_m)$ and $(\partial V/\partial S_1, \partial V/\partial S_2, \dots, \partial V/\partial S_m)$. It turns out that $V(S_1, S_2, \dots, S_m)$ is neither concave nor convex with respect to (S_1, S_2, \dots, S_m) . Our modified gradient projection method makes use of some simple observations about the behavior of order-up-to levels to avoid the numerical difficulties imposed by the nonconvexity of the cost function. Numerical experience indicates that in most situations, our method is able to find the best echelon order-up-to policy.

The major focus of our experimental study is to evaluate the impact of the leadtime-cost efficient frontier on stochastic inventory management and the resulted savings. This study presents the supply chain managers with managerial insights as to when it is important to consider flexible supply chain structuring to realize the potential gains of the flexibility offered by the conjunctive/disjunctive supply chain network commonly seen in the electronic marketplace.

We simulated four normal demand distributions with the same mean demand 20 but different standard deviations including 1, 2, 6, and 10. The major steps of the simulation are as follows:

1. For each of the demand profiles, we first computed a one-mode “iso-utility” curve. This curve is a two-dimensional curve. The x -axis is leadtime; the y -axis is the unit ordering cost. A point on this curve (x, y) represents that there is one ordering option available to the inventory manager, offering leadtime x and costing y . All the points on this curve result in exactly the same infinite horizon average cost if they are used alone

and the optimal inventory policy is followed. This infinite horizon average cost includes both inventory operating cost (holding and penalty) and ordering cost. In other words, if the inventory manager is allowed to use only one ordering option, then all the points on this one-mode iso-utility curve are equivalent to him. This one-mode iso-utility curve establishes our comparison baseline.

2. Given this one-mode iso-utility curve, we explored how much improvement can be made if the inventory manager is given the choice of using two modes simultaneously. We fixed one of the two modes to be the mode that offers leadtime 0 and then we varied the other mode along the one-mode iso-utility curve. We examined what two-mode combinations would yield best performance and how much gain was realized by using two modes in comparison with the one-mode cases.

We assumed the following cost structure. The penalty cost p is set to 30. The holding cost h is set to 5. Furthermore, we fixed the ordering cost of the mode that delivers instantaneously to 15. To avoid confusion, we use the following notation: when we refer to “mode $[i]$ ”, we mean the ordering mode that offers leadtime i .

Figure 1 shows the one-mode iso-utility curves for four normal demand distributions that we examined. We observe that the curves corresponding to high-variance demands are below the ones corresponding to low-variance demands. This is quite intuitive since as ordering leadtime increases, high-variance demand results in relatively high inventory operating cost which has to be offset by the decrease in unit ordering cost. If we are allowed to use only one ordering mode, any mode on this curve delivers the same performance therefore we are indifferent. If we do have a choice of using more than one modes, we face the decision scenarios modeled by the LCT inventory model. To gain insights, we concentrate on using two modes. Since there are too many mode combinations to consider, we only focus on a subset of these combinations where we always use mode $[0]$.

The cost savings of using two modes simultaneously are quite dramatic as depicted in Figure 2. For instance, using mode $[0]$ and mode $[20]$ cuts the cost by more than a half! The experimental results also suggest that (a) it is better to couple two modes when their leadtime difference is substantial; (b) higher variance in demand allows more savings. In Table 1, we list the performance of echelon order-up-to policies versus separating plane policies for demand profile $N(20, 10^2)$. We note that separating plane policies performed better when the gap between the two leadtimes increases. We have experimented with other types of demand profiles and cost structures with similar results.

4 Concluding Remarks

Inter-organizational electronic commerce is reshaping the way enterprises conduct business. In this paper, we present a model of supply chain that explicitly addresses time and cost issues. We have coupled this model with inventory management and performed a computational study to find effective control policies. These models and computational mechanisms are essential for software agents to take advantage of abundant choices that come with the increasingly accessible worldwide information infrastructure, and to gain competitive edge in highly dynamic business environments.

We conclude this paper by presenting some of the extensions to our agent-based supply chain model.

- Other performance measures of supply chains such as product quality, service levels, etc., are not addressed in our current model. We plan to enrich our model to address these additional measures.
- In this paper, we discussed how to integrate inventory decisions into the LCT model. We assumed that inventory is held at the root of the supply chain. We are currently working on integrating LCT with multi-echelon production/inventory models.
- In addition to the operational decisions we have discussed in this paper, we are also interested in strategic aspects of supply chain management. In the research reported in this paper, we assumed that all leadtime/cost parameters are given. In cases where supply chain agents are responsible for making these pricing decisions, strategic interactions among business entities become critical. We are currently working on extending our work on learning and negotiation[28] to address issues arisen from adaptive multi-issue negotiation in supply chain management and electronic commerce.

Mode paired with mode 0	Order-up-to Cost	Separating Plane Cost
0	376.004	376.004
1	357.291	357.291
2	343.160	341.901
3	330.469	326.879
4	318.175	313.888
5	307.952	302.115
6	297.621	290.526
7	288.249	279.898
8	279.510	269.475
9	270.828	259.631
10	260.568	250.214
11	253.627	241.212
12	244.494	232.395
13	236.939	223.957
14	230.940	215.826
15	224.214	207.857
16	218.939	200.057
17	210.404	192.392
18	201.195	184.935
19	196.657	177.713
20	187.383	170.584

Table 1: Order-up-to vs. Separating Plane Policies for $N(20, 10^2)$, $h = 5$, $p = 30$

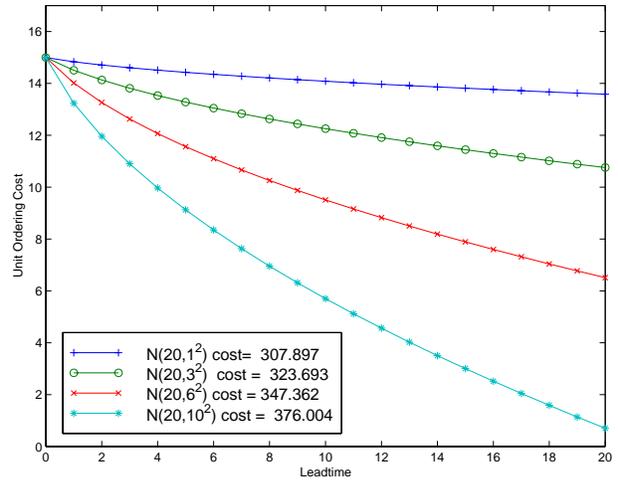


Figure 1: One-Mode Iso-Utility Curves for Normal Demands $N(20, \sigma^2)$, $h = 5$, $p = 30$

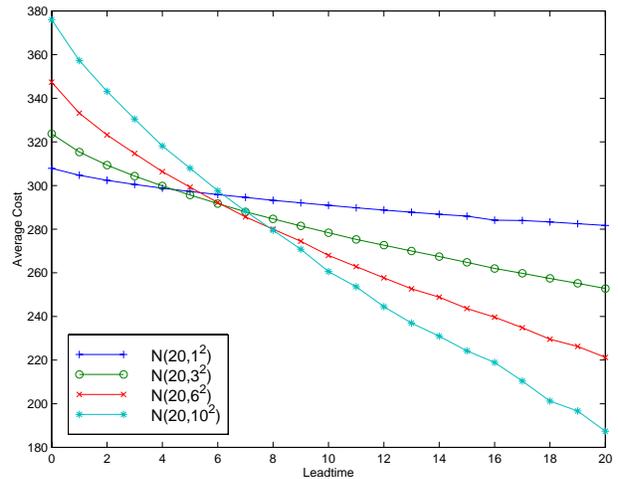


Figure 2: Savings of Using Two Modes for Normal Demands $N(20, \sigma^2)$, $h = 5$, $p = 30$

References

- [1] Mordecai Avriel. *Nonlinear Programming: Analysis and Methods*. Prentice-Hall, Inc., 1976.
- [2] M. Ben-Daya and Abdul Raouf. Inventory models involving lead times as a decision variable. *Journal of Operational Research Society*, 45(5):579–582, 1994.
- [3] J. D. Blackburn. *Time-based Competition: The Next Battleground in American Manufacturing*. Irwin, Homewood, IL., 1991.

- [4] J. A. Bockerstette and R. L. Shell. *Time Based Manufacturing*. Industrial Engineering and Management Press, Norcross, GA, 1993.
- [5] E. Bulinskaya. Some results concerning optimum inventory policies. *Theory of Probability and its Applications*, 9(3):389–403, 1964.
- [6] M. Cantamessa. Agent-based modeling and management of manufacturing systems. *Computers in Industry*, 34(2):173–186, 1997.
- [7] Soon-Yong Choi, Dale O. Stahl, and Andrew B. Whinston. *The Economics of Electronic Commerce*. Macmillan Technical Publishing, 1997.
- [8] A. Clark and H. Scarf. Optimal policies for a multi-echelon inventory problem. *Management Science*, 6(4):475–490, 1960.
- [9] Y. Fukuda. Optimal policies for the inventory problem with negotiable leadtime. *Management Science*, 10(4):690–708, 1964.
- [10] Paul Glasserman and Sridhar Tayur. Sensitivity analysis for base-stock levels in multiechelon production-inventory systems. *Management Science*, 41(2):263–281, February 1995.
- [11] D. Goldfarb. Extension of davidon’s variable metric method to maximization under linear inequality and equality constraints. *SIAM J. Appl. Math.*, 17:739–764, 1969.
- [12] Yu-Chi Ho and Xi-Ren Cao. *Perturbation Analysis of Discrete Event Dynamic Systems*. Kluwer Academic Publishers, 1991.
- [13] Arthur Hsu. Modeling the tradeoff between leadtime and cost in supply chain management. Technical report, Graduate School of Industrial Administration, Carnegie Mellon University, 1997.
- [14] Arthur Hsu and Dajun Zeng. Finding ordering policies given multiple leadtime options. Graduate School of Industrial Administration, working paper, Carnegie Mellon University, 1999.
- [15] Ravi Kalakota, Jan Stallaert, and Andrew B. Whinston. Implementing real-time supply chain optimization systems. Technical report, University of Texas at Austin, Dept. of MSIS, 1996.
- [16] Ravi Kalakota and Andrew B. Whinston. *Electronic Commerce: A Manager’s Guide*. Addison-Wesley, 1996.
- [17] Hon-Shiang Lau and Long-Geng Zhao. Optimal ordering policies with two suppliers when lead times and demands are all stochastic. *European Journal of Operational Research*, 68:120–133, 1993.
- [18] Hau L. Lee and Corey Billington. Managing supply chain inventory: Pitfalls and opportunities. *Sloan Management Review*, pages 65–73, Spring 1992.
- [19] J.S. Liu and K. Sycara. Coordination of multiple agents for production management. *Annals of Operations Research*, 75:235–289, 1997.
- [20] Michael E. Porter. *Competitive advantage: creating and sustaining superior performance*. The Free Press, 1985.
- [21] Carl Shapiro and Hal R. Varian. *Information Rules: A Strategic Guide to the Network Economy*. Harvard Business School Publishing, 1999.
- [22] Jing-Sheng Song. The effect of leadtime uncertainty in a simple stochastic inventory model. *Management Science*, 40(5):603–613, May 1994.
- [23] J. Swaminathan, S. Smith, and N. Sadeh. Modeling supply chain dynamics: A multiagent approach. *Decision Sciences*, 29(3), summer 1998.
- [24] Katia Sycara and Dajun Zeng. Coordination of multiple intelligent software agents. *International Journal of Cooperative Information Systems*, 5(2 & 3):181–211, 1996.
- [25] S. Wang. Analyzing agents for electronic commerce. *Information Systems Management*, 16(1):40–47, 1999.
- [26] M. Wooldridge and N. R. Jennings. Intelligent agents: Theory and practice. *The Knowledge Engineering Review*, 10(2):115–152, 1995.
- [27] P. R. Wurman, W. E. Walsh, and M. P. Wellman. Flexible double auctions for electronic commerce: Theory and implementation. *Decision Support Systems*, 24, 1998.
- [28] Dajun Zeng and Katia Sycara. How can an agent learn to negotiate? In Jörg Müller, Michael J. Wooldridge, and Nicholas R. Jennings, editors, *Intelligent Agents III: Theories, Architectures, and Languages*, number 1193 in Lecture Notes in Artificial Intelligence, subseries in Lecture Notes in Computer Science, pages 233–244. Springer, 1996.
- [29] Daniel Dajun Zeng. *Multi-Issue Decision Making in Supply Chain Management and Electronic Commerce*. PhD thesis, Carnegie Mellon University, GSIA, 1998.