

Intelligent Decision Support for the e-Supply Chain

Richard Goodwin, Pinar Keskinocak, Sesh Murthy, Frederick Wu, Rama Akkiraju

IBM T. J. Watson Research Center

Yorktown heights, NY 10598

{rgoodwin, keskinocak, murthy, fwu, akkiraju}@watson.ibm.com

Abstract

Much of the attention in artificial intelligence (AI) for e-Business has focused on business to consumer transactions. Shopping bots, systems to recommend movies and books based on similar opinions by other users and news filtering agents, are just some examples. However, we feel that AI can have a larger impact on the supply chain that delivers goods and services to the end consumer. Reductions in costs and the pervasiveness of the Internet have encouraged companies to move towards using e-commerce for transactions with their business partners. Companies are willing to invest resource because of the reduced product cycle times and the lower transaction costs that they expect. A result of this movement is that companies can afford to interact with a larger number of trading partners and form project and customer specific partnerships that would have been too costly in the past. To manage a larger and more dynamic set of partnerships and to be able to take advantage of transient opportunities, business users will need decision-support systems to identify and analyze the opportunities in terms of their business objectives. In this paper, we describe our agent-based decision-support framework for creating systems to support trading partners in the e-supply chain. In particular, we will focus on the issues that need to be addressed in order to create a viable and useful decision-support system.

Introduction

As companies interact with their customers, suppliers, and people within their own enterprises, they are more and more taking advantage of easy and real-time access that the internet/intranets are providing. Forester Research in Cambridge, Mass., estimates that businesses bought and sold \$43 billion in goods over the Internet last year, vs. only \$8 billion bought by consumers. By 2003 more than 90% of the predicted \$1.4 trillion in Internet commerce will be conducted between businesses. Partnerships can be built on-the-fly as needed, and different suppliers may be added or removed from the chain in a relatively short amount of time. Using technology and information, companies can "connect" to a larger number of customers and suppliers, they can virtually integrate their supply chains and blur the traditional boundaries in the value chain among suppliers, manufacturers and end users (Magretta, 1998).

One reason why so many companies embrace e-business is because it makes locating and interacting with potential business

partners (both buyers and suppliers) much easier. The possibility of nearly instantaneous updates of suppliers' inventory status plus access to the latest requirements of customers can give the fully networked supply-chain manufacturer a significant advantage over businesses that rely on traditional fax/phone communications. However, the amount of increase in connectivity and access to numerous information sources also tremendously increases the number of alternatives one needs to consider before making business decisions. To take full advantage of the new opportunities offered by e-business, companies need to be able to respond quickly to opportunities as they arise, making smart decisions based on the voluminous data flows. Hence, we envision that companies will need decision support in many areas, including information gathering, locating potential suppliers and buyers, and deciding in what to buy and what to sell.

We are building an agent-based decision support system which can assist a decision-maker to sift through the data and respond to opportunities in a timely manner. In effect this system finds 'good' matches between demand and supply, taking into account the constraints of the production schedule. It is customizable to the objectives of the business, which generally include maximizing profit, maximizing customer satisfaction, and minimizing production bottlenecks.

Current Implementation

Our current implementation of the decision support system is part of an "electronic trading floor" we are building for the Colombier Group of Myllykoski Oy, a pan-European paper trading company. The trading floor Web site hosted by Colombier lets sellers list excess inventory and production capacity, and buyers list items and capacity they want - all anonymously if they like. The "decision support system" analyzes the consequences of various trades on the buyer's or seller's manufacturing schedule, inventory and finances (depending on who is using it), and ultimately recommends trades. Finally, the auction component offers sellers a variety of auction methods to unload surplus inventory or capacity. People or automated agents at the sellers and buyers could monitor recommendations and close deals. The system will significantly alter Colombier's current supply-chain model, in which suppliers fax offers to Colombier, which then negotiates a price, buys the paper, stores it and figures out how to sell it.

In the new model, Colombier earns a transaction fee, and reduces its financial risk because it does not hold inventory. The

seller does not have to let buyers know it has excess paper or capacity, and thus can't be bargained down on price. By the same token, sellers don't know buyers are in a jam, or are small, and thus won't artificially raise the price. Hence, all parties benefit.

The prototype trading system began running in March. The decision support software, for which IMB has filed a patent, is currently being developed. The actual Colombier trading site is scheduled to go online in September.

Decision Support Issues

Several issues need to be considered while building a decision-support system for managing the e-supply chain. We briefly discuss some of these issues, standards for information exchange, representations of products, services and offers of capacity in integrated catalogs and methods for matching supply with demand.

The electronic marketplace will be fully symmetric with respect to procurement and sales. A manufacturer will find offers to sell the materials they need as well as offers to purchase goods they produce. All this information will supplement the data already in the firm's ERP system, which includes inventory, customer orders, and the state of the manufacturing schedule. Taken together, these data can enable much better decisions as to what products to make, what raw materials to purchase, what orders to accept, etc., but only if the huge volumes of data are fully analyzed and the numerous possibilities considered.

Information exchange

A key enabler for this vision is the mundane task of information exchange between disparate systems. Standards must be established for exchanging industry-specific data, and they must gain wide acceptance; a promising enabling candidate is XML. We use XML to simplify data sharing between supply-chain partners over the Web and for integration between a supply-chain manufacturer's enterprise resource planning system and their Website, as well as their partners' Websites. These standards are emerging and will continue to evolve under pressure to standardize across industries for economies of scale and pressures to specialize within industries for supporting market specific functionality.

Building an integrated catalog

In supply-chains, buyers and suppliers need to be able to fully characterize products and services they wish to buy or sell. On-line catalog items need not be restricted to existing inventory, but could also include offers of capacity to produce a product or set of products. The representation of capacity requires a much richer, more flexible characterization of products and production capabilities. Initially, we are representing items as attribute-value pairs where the values can be simple values, ranges or set of values. This seems sufficient to capture most important aspects of the items and is flexible enough for supporting business-to-business trading activities such as posting a item to the catalog, browsing and searching the catalog, placing an order for an existing item, and placing an order for an item to be manufactured from capacity. As an example, let us consider the representation of inventory and capacity for paper products. An inventory item, say a roll of paper, can be represented by attributes such as grade, grammage, width, diameter, core

diameter, core type, finish, and coating, where the values for these attributes are M-Brite for grade, 65.00 for grammage, 838.00 for width, 1250.00 for diameter, 76.00 for core diameter, semi-matt for finish and light weight coated for coating. Similarly, these attributes can also be used to represent capacity to produce rolls of paper on a paper-making machine. For example, a capacity item for manufacturing paper could have values M-Brite for grade, [65, 75] for grammage, [400,600] or {650, 700, 750} for width, [500, 1250] for diameter, [60, 80] for core diameter, {wood, metal} for core type, {matt, semi-matt, glossy} for finish, and {none, light weight} for coating. In this representation, [400,600] is the range of values from 400 to 600 and {650, 700, 750} is a set of three values. Using this representation, this capacity item is an offer to make any roll of paper with attribute values within these sets or ranges.

In this simple representation of capacity, we are specifying capacity to manufacture goods in terms of the products we are willing to manufacture rather than the capabilities of the mill or machines. In a service oriented business, capacity might be offered as time on a machine, with the characteristics of the machine specified as the entry in the catalogue. In addition to specifying the capacity offered for sale, a manufacturer might indicate a preference for manufacturing certain types of goods, through discriminatory pricing for example.

Decision Support for Various Entities in the e-Supply Chain

The decision support system we are building can assist a decision-maker to sift through the data and respond to opportunities in a timely manner. In effect this system finds 'good' matches between demand and supply, taking into account the constraints of their own product schedule. The following are some of the features of the decision support system, discussed in the context of manufacturers:

1. *Decision Support for what to sell:* At any moment, a manufacturer may have some product inventory in stock and some unfulfilled orders or forecast orders. It may be advantageous to offer some of the inventory for sale, especially those items which are not committed to customer orders or forecasts, or which cannot be profitably reworked to meet a customer order or forecast. In addition, a manufacturer has the capacity to produce product in the future. It may be advantageous to offer some capacity for sale, especially capacity that is not already committed to customer orders or forecasts. The decision is a complex tradeoff between maintaining flexibility in being able to use capacity for the most profitable types of production and in making sure capacity does not go unused. This is similar to the airline yield problem where seats on plane are sold with various restrictions at various times for different prices. Manufacturers have the added complication that all or part of an order can be outsourced even after it is sold, if it becomes more cost efficient. A decision-support system for deciding what to sell must make recommendations about what types of inventory and capacity to offer, to whom, for what price and with what restrictions at what time.
2. *Decision Support for what to buy.* In addition to its own inventory, a manufacturer will have access to inventory being offered for sale on the Internet (e-inventory), as well as some unfulfilled orders or forecast orders. It may be advantageous to buy some of the inventory or capacity for sale, especially those items which can be profitably used to satisfy customer orders or forecasts, or which would be difficult to make because of production constraints. In addition, over the Internet a manufacturer will be able to purchase

other manufacturers' capacity to produce product in the future (e-capacity).

3. Decision Support for what to promote. In addition to offering raw capacity for sale, a manufacturer can 'push' sales of products which are both advantageous to make and desirable to customers. Since raw capacity can be used to make so many different items, each promotion should be limited to items of particular interest to an individual customer.

A decision-support system must be customizable to the objectives of the business, which generally include maximizing profit, maximizing customer satisfaction, and minimizing production bottlenecks. It must also be configurable for different entities in the supply chain such as manufacturers, dealers, brokers and consumers.


Decision Support Framework

In decision support for e-supply chain management, given "demand" and "supply" our goal is to find "matches" between demand and supply to satisfy the company's business objectives. Demand includes company's orders, forecasts, request-for-quotes (RFQs), and other requests posted on the Internet, possibly through various trading sites, whereas supply includes company's own inventory and capacity, combined with e-inventory and e-capacity. The resulting "matches" will suggest alternatives such as which orders should be satisfied from the companies own inventory and capacity, which orders should be outsourced by using e-inventory and e-capacity, which orders should be satisfied by combining inventory with e-capacity or by combining e-inventory with capacity, and which promotions should be sent to which customers. In short, the decision support system will suggest which supplies should be used to satisfy which demands.

Figure 1: Decision support framework

Figure 1 shows a high level architecture of the decision support system. In this agent-based architecture, given demand items, query-generating agents generate queries to search for supply items that could possibly be used to satisfy the demand items. Search agents take these queries and search various sites for items that (approximately) match these queries. The matching agents take the results of the search agents items and try to find matches between these supply items and the demand items. The search agents, as well as the decision support agents, collaborate within an agent architecture called Asynchronous Team (A-Team) [Talukdar, deSouza, and Murthy 1993]. In an A-Team, agents cooperate by sharing results in population(s) and by dynamically evolving a shared population of solutions towards the Pareto-frontier. Since agents can work on solutions created by other agents, the best solution alternatives are usually the result of contributions made by many agents embodying multiple problem-solving approaches. This architecture supports our aim of providing an intelligent assistant that can work cooperatively with schedulers to produce multiple solutions that illustrate tradeoffs. In addition, the A-Team architecture allows us to combine multiple problem-solving approaches, which increases the likelihood that good solutions can be found in a reasonable amount of time. After the matching agents create a population of "good" solutions by collaborating in an A-Team, each of these

solutions is evaluated in business terms, and the filtering agents select a small subset of the "better" solutions and present them to the decision maker. The decision-maker can then examine these



alternatives, modify and reevaluate them, and finally choose the "best" solution by considering tradeoffs.

The decision-support optimization problem in our framework can be viewed as a matching problem where we match supply with demand. Supplies can take the form of inventory or could be a configurable capacity. Demand could be an order for a fully specified product or a request for proposal (RFP). To perform the optimization, we use a portfolio of algorithms including bipartite-graph matching algorithms, minimum cost-flow algorithms, linear programming approximations and greedy algorithms. In addition to matching, we take into account conversion facilities that can transform one product into another. For example, a paper converter, such as Colombier can transform one large roll of paper into multiple smaller rolls. Our optimization algorithms take conversion costs and waste into account when creating matches with product transformations. The variety of the algorithms and the use of local iterative improvement algorithms helps to reliably generate near-optimal solutions.

To summarize, the components of our decision support approach include generating multiple alternative solutions, evaluating the alternatives in business terms, selecting a small subset of these alternatives and presenting them to the decision maker, collaborating with the decision maker to generate new or improved alternatives, and helping the decision maker to understand the tradeoffs and select the "best" alternative.

Conclusions and Future Research

The vast quantities of information and the dynamic nature of the Internet enables supply-chain will require better decision support systems to enable enterprises to take advantage of the new opportunities. Better product representations, information exchange protocols and better matching of demand and supply are but of few of the critical issues that need to be address.

In addition to exploring these catalogue-related issues, we are examining the use of auctions for both sales and procurement. Auctions are a time-honored means of determining a fair price for goods under certain conditions, and Web-based consumer auctions are rapidly growing in popularity. In the supply-chain environment, the eventual role of electronic auctions remains to

be seen. However, we expect that auctions will initially be used in the traditional situations, e.g. for non-commodity items that do not have established markets, or for fast liquidation of old or off-quality product. We are also working on building decision support for selecting items to auction, designing auctions, and for auction participation strategy.

References

Magretta J. 1998. The Power of Virtual Integration: An Interview with Dell Computer's Michael Dell. *Harvard Business Review*, March-April 1998, pp.73-84.

Murthy S., Rachlin, J., Akkiraju R., Wu F. 1997. Agent-Based Cooperative Scheduling. In *Constraints & Agents*. Technical Report WS 1997-97-05. Menlo Park:AAAI Press.

Talukdar, S.N.; deSouza, P.; and Murthy, S. 1993, "Organizations for computer-based agents," *Engineering Intelligent Systems*, Vol. 1, No.2.