Accepted for Presentation at the International Conference on Internet Computing, Las Vegas, Nevada, June 25-29, 2001

Intelligent Trading Networks: Peer-to-Peer Trading in a Distribution Marketplace

Bruce M. McLaren, Philip J. Hayes, and Aidan J. McKenna
OpenWebs Corporation
2403 Sidney Street
Pittsburgh, Pennsylvania 15203-2116

Abstract. We have developed an intelligent, peer-to-peer trading technology that facilitates trading in a distribution-intensive marketplace and improves upon existing business-to-business exchanges. This technology, tradingNetwork™, enables independent businesses to collaborate with one another to improve service to their customers, strengthen relationships between partners, minimize excess inventory, and increase sales opportunities. Each trading partner buys and sells autonomously and automatically, using a private set of rules, with full access to all of its enterprise data. We have created a generic, extensible, and object-oriented infrastructure that may be customized by each partner to intelligently manage and control decision-making. Objects in the infrastructure include decisions (e.g., Buy-Side Partner Selection), criteria (e.g., price, preferred partners, promise dates), and criteria parameters (e.g., preferred price point, importance of individual criterion). Messages and data are passed between partners and between the intelligent trading component of a partner and that partner's Enterprise Resource Planning (ERP) system using a customized protocol of XML documents.

Keywords: E-commerce; Business-to-Business; Peer-to-Peer; Intelligent Decision Support

1. Introduction

Business-to-business (B2B) e-commerce is fast becoming the ubiquitous engine driving the commerce of our times. Market analysts such as the Gartner Group and Forrester Research predict huge increases in the use of and investment in B2B technology.

- Gartner predicts that over 50 percent of large enterprises will have implemented a large-scale application to support electronic trading with external partners by 2002.
- Forrester envisions the volume of trade to reach $1.3 trillion by 2003 [6].

Clearly, the need for new technologies to support B2B e-commerce is great and will only increase in the future.

The first wave of B2B technology has produced the *B2B exchange* [9]. In a B2B exchange, many buyers and sellers are brought together in a virtual sense in a central market space on the Internet. This central market allows the participants to buy and sell from one another at a dynamic price that is controlled by the rules of the exchange. Examples of companies that provide B2B exchanges are Ariba and CommerceOne.

B2B exchanges differ from other B2B e-commerce by the participation of *multiple* buyers and sellers, instead of one-to-one Internet connections between buyer and seller or auctions (or reverse auctions) that support one-to-many relationships (i.e., one buyer, many sellers or one seller, many buyers). B2B exchanges work on the principle that "the more competing buyers and sellers that can be brought together in one place, the more liquid a market becomes and the more efficient the price-setting mechanism is" [9, p. 9]. Because of the availability of the Internet and the advent of B2B exchanges, trading no longer requires people to come together physically.

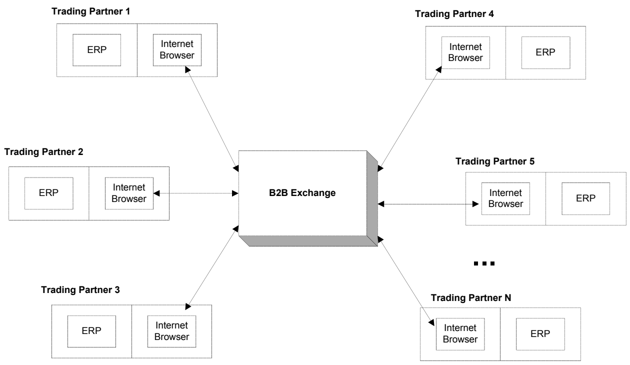

Figure 1: The B2B Exchange Model

Participants in the simplest and most common kind of B2B exchange interact with the exchange site directly via an Internet browser, as depicted in Figure 1.

The exchange site is responsible for linking buyers to sellers and vice versa, according to criteria set by the exchange. Thus, exchanges are centralized and controlled by a single organization. Although a trading partner in an exchange may use many databases and an ERP system to support its buying and selling decisions, the B2B exchange process is typically decoupled from the rest of the enterprise.

Thus, while simple B2B exchanges have strengths and represent an important architectural step in the evolution of e-commerce, they also suffer from several shortcomings, including:

- Trading rules are centralized so each partner receives only those responses that are facilitated by the protocol and rules defined by the exchange coordinator.
- Decision-making consistency and strategic business objectives are not enforced by a published set of rules that govern the buying/selling behavior of each partner.
- Privacy and confidentiality are compromised because the exchange coordinator has complete access to all transactions.
- There is no integration between the Internet browsers of the buyers and sellers and their ERP systems, leading to re-keying and errors.

All of these shortcomings can and have been addressed to some extent within the exchange model. Some exchanges provide ways for partners to submit or receive orders electronically from or to their ERP systems. Many exchanges have strong confidentiality policies. Nevertheless, the exchange model remains an imperfect e-commerce solution.

In order to address the shortcomings of the B2B exchange model, we have developed a radically different trading model. In our model, trading partners interact in peer-to-peer (P2P) fashion, rather than through an intermediary. Partners choose whom they wish to trade with and under what conditions. Transaction information is available only to the participating parties. Further, the trading component at each partner site is seamlessly integrated with that partner's enterprise systems and data. Each partner can establish confidential rules that control and automate all aspects of buying and selling and allow the partner to meet its private business objectives.

The issues and hurdles involved in applying P2P to business-to-business trading are varied and challenging. In addition to addressing ERP integration problems, it is also necessary to solve the problems of decentralizing data, securing the privacy of each partner's enterprise information, and gaining sufficient participation to make the concept viable.

Our strategy is to develop and install an initial deployment in a large but private network of partners in a single vertical market. Such an approach allows us to:

a) validate the concept in a controlled manner,

b) penetrate and achieve success in a single market as a starting point,

c) integrate with a smaller number of ERP systems (initially a single one),

d) ramp up a critical mass of partners necessary to support a trading network,

e) enlist a small number of major players (e.g., manufacturers) who can create incentives for smaller-size partners to participate.

Our initial deployment of tradingNetwork™ is in the tire industry, which operates as a three-tier distribution channel: manufacturer, wholesale distributor, and retail store. tradingNetwork™ can operate across, up, and down this channel, giving the wholesale and retail participants greater freedom to source their products than in traditional hierarchical distribution.

2. Our Approach

The architecture of tradingNetwork™ is shown in Figure 2. Each partner is composed of an ERP/OM system (i.e., an ERP system "wrapped" with order management software), an Intelligent Trading Module (ITM), and a message broker. The ERP/OM system is responsible for maintaining all enterprise data (e.g., inventory, order processing) and provides the GUI used at the point of sale. The ERP/OM is a full three-tier architecture (i.e., presentation, business logic, and data). The ITM contains and applies the business rules that guide all external buying and selling. The message broker is responsible for facilitating intra-node communication between the ERP/OM and ITM and for inter-node communication between partners. All communication occurs via a customized protocol of XML messages.

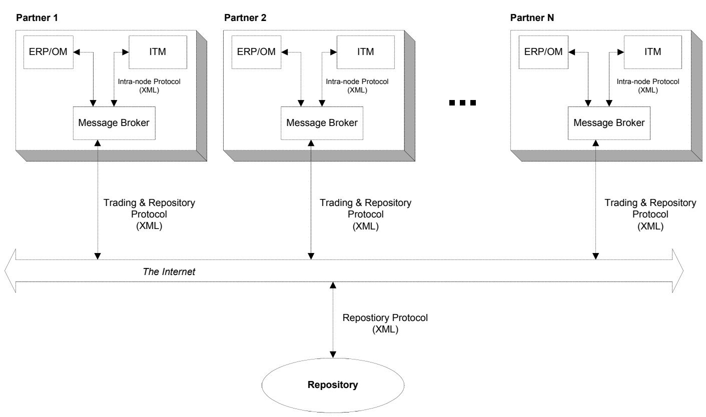

Figure 2: The tradingNetwork™ Architecture

In addition to the partners, the network includes a repository that contains a registry of the network participants. To participate in the trading network, an organization's URL must be recorded in this repository, which also contains certain public information (e.g., the physical location of each partner's business, brands a dealer specializes in). All partners in the network have access to the repository via a set of XML messages.

Our architecture is Microsoft-centric. The OM software, GUI, and ITM are implemented in Visual Basic and, in cases of efficiency-critical functionality, C++. All of the components are object-oriented and implemented as COM and DCOM objects. Microsoft's BizTalk Server is the message broker (although other messaging software could easily be used) and Solomon IV is the ERP system.

Such an architecture is loosely coupled or *cohesive*, rather than coupled [6, pg. 36-37]. By necessity, the independent trading partners must be loosely coupled, as their systems and data are fully independent and confidential. The loose coupling within each partner's node, while not necessary, confers the greatest degree of flexibility. Through its independence from the ERP/OM, the Intelligent Trading Module may be integrated with different ERP and OM systems that adhere to the intra-node protocol.

In contrast to B2B exchanges, decision-making is decentralized in our model. Each partner directly contacts other participants in the network without the aid of an intermediary. The repository is the only centralized aspect of the model, and it acts as a passive data store, not as an active agent during trading. Because no intermediary is involved, trading and information exchange are private and confidential between involved partners.

A further advantage of our model, and the primary emphasis of the remainder of this paper, is the automated and intelligent decision-making capability of the ITM. The ITM at each partner site applies pre-defined buying and selling rules, ensuring that an organization's business objectives are uniformly addressed. The knowledge employed by the ITM may be customized and extended at each partner site.

This automated decision-making capability is key to the practicality of our P2P approach to B2B. With automated decisions, a buyer can query multiple potential sellers and quickly receive a response tailored to their relationship. In the absence of intelligent decision-making, sellers cannot differentiate their response appropriately without human intervention, resulting in slow and/or variable response time.

Currently, our system is deployed at a single large tire dealer ($60M in annual sales). During this deployment we will refine the functionality and performance of the software within a single node of the trading network, focusing on the ERP/OM capabilities. In late 2001, we expect to begin wider deployment and to use the trading capabilities of the system in practice.

3. Intelligent Trading

In building an intelligent trading capability we had a number of goals, including:

- to define a generic decision process and framework that could be easily extended.
- to establish a library of decision-making objects that could be easily (i.e., without programming) added to a partner's ITM. Not all partners apply the same types of information to their decisions, and we wanted to provide a capability to tailor each partner's ITM to account for these differences.
- to enable users to easily define and modify their business rules. We initially considered the use of traditional expert system rules [2, 5] but decided on a simpler approach.
- to employ a simple scoring algorithm that would be both intuitive and comprehensible to a non-technical user.

3.1. The Representation

In devising a general decision process for the ITM, we implemented the following object-oriented decision objects:

A *Decision* is a point in the trading process in which a partner must make a buying or selling choice given a list of alternatives, be it partners, inventory items, or offers. For the initial version of the system, we identified three decisions: *Buy-Side Partner Selection, Sell-Side Offer Selection*, and *Buy-Side Offer Evaluation*.

A decision is made by evaluating *Criteria* that relate to that decision. For instance, the selection of partners to contact with a purchase request (i.e., Buy-Side Partner Selection) may be made by considering preferred partners. The ranking of offers received from a seller (i.e., Buy-Side Offer Evaluation) may be impacted, in part, by the offer price. Basically, each decision relies on the evaluation and combination of a set of such underlying criteria.

Criteria Parameters are qualifiers of the criteria. These range from data required by the criterion, such as a single value (e.g., a preferred price) or a list of values (e.g. preferred partners), to parameters that specify the importance of a criterion.

Global Parameters affect the overall decision-making process. For instance, a score threshold may be used to indicate the cut-off for consideration of scored alternatives.

The three decisions currently defined in the model follow a typical B2B purchase scenario. A customer contacts a dealer or retailer (Partner A) to purchase a product. If Partner A is unable to service the request from its own inventory, its ITM would send purchase requests for the desired product to a set of candidate partners. Deciding which partners to contact with a purchase request (Buy-Side Partner Selection) might require the evaluation of criteria such as preferred partners, preferred brands, and proximity. Each of the partners contacted by Partner A would then decide whether they can and want to respond to Partner A's request with one or more offers (Sell-Side Offer Selection). This decision may involve criteria such as partner preferences, inventory preferences (i.e., a minimum stocking level), and margin preferences. Finally, after the partners reply with offers, Partner A must decide which offer or offers to accept (Buy-Side Offer Evaluation). Criteria that might impact this decision include brand preferences, promise time, and price.

3.2. The Decision Process

Each of the decisions applies a generic *decision process* and produces a *decision result*, i.e., a scored and sorted set of partners, inventory items, or offers. The specific actions taken at each step of the decision process vary based on the specific criteria and parameters associated with that decision. The canonical steps of the decision process are:

1. *Select candidates*. Candidates are either (1) retrieved via a decision-specific process (provided as a specialized method) or (2) provided to the decision as input. Both filter undesirable candidates from the list. If this step filters out all candidates, the remaining steps of the decision process are skipped with a decision result of "Empty."
2. *Rank candidates by weighted-average score*. Using the criteria associated with the

decision and the parameters associated with each criterion, a score between 0 and 100 is assigned to each of the candidates. "Importance" is a special parameter that is used to apply a weight to each criterion. More formally, if a decision has n criteria, and the i^{th} criterion is represented by a function C_i that yields values in the range 0 to 100, then the score for each candidate is:

$$W_1 * C_1 + W_2 * C_2 + \ldots + W_n * C_n$$

where each W_i is the weight (i.e., "importance") of the corresponding C_i. W_i is between 0 and 1 and $W_1 + W_2 + \ldots + W_n = 1$.

3. *Select finalists*. The scored and sorted list of candidates is then evaluated for finalists. Finalists can be selected by taking the top N candidates or by taking all candidates that have a weighted average above a given threshold. The selected finalists (zero or more) are returned as the decision result.

3.3. Editing and Configuration

A program called the *Decision Editor* allows a partner to modify the decision objects and customize the decision process. Each partner in the network has its own Decision Editor, an off-line program used to maintain the business rules for that site. This enforces consistent application of priorities and rules for all buy/sell decisions made by that partner.

A Decision Editor window depicting Buy-Side Offer Evaluation is shown in Figure 3. This window contains the criteria and parameters that are applied to a decision at run-time and allows the user to change this information. The other decisions are similarly represented in the Decision Editor.

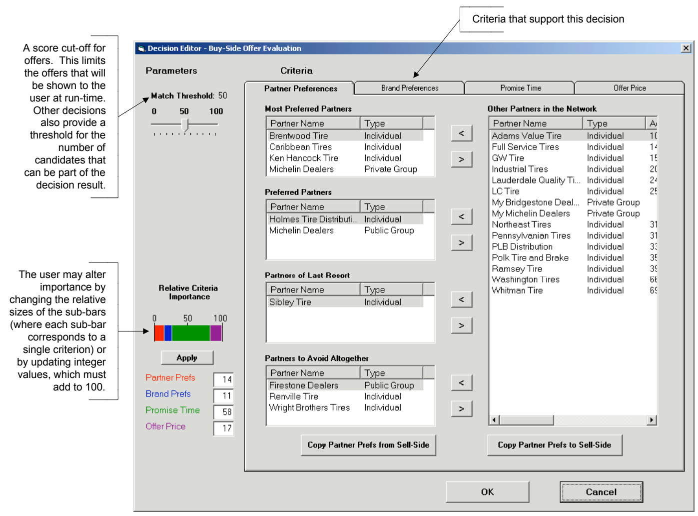

Figure 3: A Screen Shot of the Decision Editor

Criteria are implemented as either a set of lists or a numeric range (i.e., minimum, maximum). For instance, as shown in Figure 3, partner preferences are represented as five lists: most preferred, preferred, other partners in the network, partners of last resort, and partners to avoid altogether. The user can define the "rules" representing partner preferences by simply moving partners between lists using the move buttons. The scores assigned by the decision process for these categories are 100, 75, 50, 25, and 0, respectively. An example of a numeric range criterion is offer price. For instance, a preferred minimum tire price might be $30 and a maximum acceptable price might be $200. A numeric range is scored using a scaling function between the best value (score of 100) and the worst value (score of 0).

A companion tool, the *Decision Configurator*, has been designed but not yet implemented. The Decision Configurator will allow each partner to choose criteria from a library and associate the criteria with each of their private decisions. For instance, Partner X may choose the criteria preferred partners, preferred brands, promise time, and offer price as criteria of the decision shown in Figure 3. Partner Y may choose preferred partners, offer price, and fulfillment history for the same decision. This results in a configuration with only three tabs and three importance elements (lower left of window), corresponding to each of the chosen criteria. Both run-time criteria evaluation and the

Decision Editor will be dynamically configured by this tool. Over time, a large library of criteria will be compiled.

4. Intelligent Trading in Action

An example of intelligent trading in action illustrates the overall architecture of tradingNetwork™ and, in particular, the decision-making capability of the model.

Figure 4 depicts the typical purchasing scenario discussed in section 3.1. The overall network includes many partners, but only two are shown: Jackson Brothers, the buyer in the scenario, and Sibley Tire, a potential seller. The arrows between system components and partners depict the XML message protocol that occurs between and within these partners[1]. The decisions that have been implemented in tradingNetwork™ are also indicated by the dashed ovals.

Suppose that a customer enters Jackson Brothers and requests four tires for a 1996 Ford Contour. The retail salesperson uses the ERP/OM system to determine the tire size (in this case 205/60R15) and checks whether the desired tires are in local inventory.

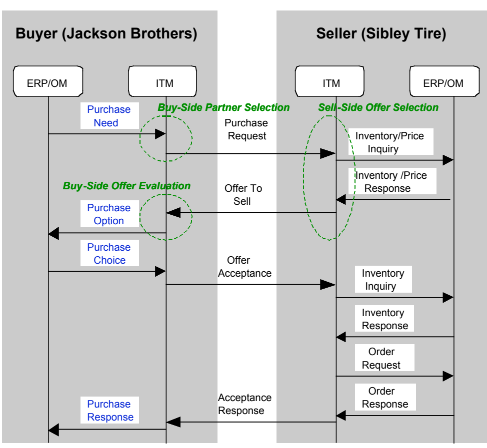

Figure 4: A Typical Purchasing Scenario

The tires are not available in Jackson Brothers' inventory, so a "Purchase Need" message is sent by Jackson Brothers' ERP/OM to its ITM (see the upper left of Figure 4). This triggers Jackson Brothers' Buy-Side Partner Selection, which uses its criteria and parameters to score and rank its partners as follows:

Score	Partner	Address	City	State	County
75	GW Tire	1532 Lime Hollow Rd.	New York	NY	New York
75	Sibley Tire	5544 Bryant Street	Chester	PA	Delaware
63	Full Service Tires	1411 Browning Road	Rochester	NY	Monroe
50	Adams Value Tire	1000 Wade Lane	Uxbridge	MA	Worcester
50	Brentwood Tire	1146 Duffield St.	Sturbridge	MA	Worcester

Notice that Sibley Tire is tied for top-ranked partner in the list. Sibley is assigned a score of 75 because it is most preferred (criterion score of 100) and has a moderate overlap of preferred brands with Jackson Brothers (criterion score of 50). The relative weights of the two criteria are 0.5, thus the weighted calculation for Sibley is: (0.5 * 100) + (0.5 * 50) = 75.

The results of the Buy-Side Partner Selection are restricted to the partners shown above because others are filtered by either (a) the match threshold of 50 or (b) Jackson Brothers' preference for partners located in preferred (probably nearby) counties. Physical location is used as a preliminary filter in the first step of the decision process to avoid needlessly scoring large numbers of partners.

Next, Jackson Brothers' ITM sends out a "Purchase Request" to all of the selected partners in the Buy-Side Partner Selection decision result, one of which is Sibley Tire. Processing this request invokes the second decision: Sell-Side Offer Selection. Sibley's ITM immediately sends an "Inventory/Price Inquiry" to its ERP/OM system to check whether it has the requested tires in its inventory. Sibley's ERP/OM replies with an "Inventory/Price Response," containing several possible offers, and the ITM produces the following decision result:

Score	Brand	Part #	Qty Avail.	Qty Req.	Unit Cost	Total Cost	Margin
76	Michelin	47503	57	4	$ 107.91	$ 431.64	52.34%
70	BF Goodrich	33582	45	4	$ 75.74	$ 302.96	59.80%
68	Michelin	64401	5	4	$ 125.41	$ 501.64	35.25%
67	BF Goodrich	71072	109	4	$ 94.58	$ 378.32	53.61%

Four types of tires in Sibley's inventory match the request and score above the threshold of 65. Sibley's criteria for this decision and their respective weights are: (brand preferences 0.4) (margin preferences 0.5) (partner preferences 0.1). Michelin part #47503 scores highest because Michelin is a most-preferred brand (criterion score of 100), its margin is reasonably high (criterion score of 52.34, based on a range of 0 to 100), and Jackson Brothers is a most-preferred buyer (criterion score of 100). The score is calculated as: (0.4 * 100) + (0.5 * 52.34) + (0.1 * 100) = 76 (rounded).

[1] All XML messages would actually go through each partner's Message Broker, but the intermediate steps are not shown.

Notice that the second offer in the list provides a slightly higher margin than the first, but BF Goodrich is not a most-preferred brand, resulting in a lower overall score. Notice also that the third offer, part #64401, virtually wipes out Sibley's inventory of this tire (i.e. quantity available is 5, quantity requested is 4). If an inventory preference would have been applied, this offer may not have been so highly rated.

Sibley's ITM next sends an "Offer to Sell" message, containing all of the above offers, to Jackson Brothers. Jackson Brothers' ITM receives the message, along with offers from other partners in the network, and executes the final decision, Buy-Side Offer Evaluation. This produces the following results:

Score	Partner	Brand	Part #	Qty	Unit Cost	Total Cost	Promise Date
90	Sibley	BF Goodrich	33582	4	$ 75.74	$ 302.96	Feb. 2, 2001
84	Full Service	Bridgestone	36064	4	$ 94.95	$ 379.80	Feb. 3, 2001
84	Sibley	BF Goodrich	71072	4	$ 94.58	$ 302.96	Feb. 5, 2001
83	GW Tire	Michelin	41402	4	$ 99.23	$ 396.92	Feb. 5, 2001
82	Sibley	Michelin	47503	4	$ 107.91	$ 431.64	Feb. 5, 2001

Notice that not all of Sibley's offers appear in the Jackson Brothers' decision result because a threshold of 82 is applied. Also notice that Jackson Brothers' ITM scores Sibley's offers different than Sibley's ITM does because it applies different criteria to this decision (i.e., (promise time 0.45) (cost 0.45) (partner preferences 0.1)). In particular, the best offer from Sibley's perspective, Michelin part #47503, is rated lower than two other Sibley offers by Jackson Brothers. This is because both the promise time and cost of this part are less favorable from the perspective of the buyer.

At this point, intelligent decision processing is complete, but the message protocol completes the purchasing scenario. All of the scored offers are sent as "Purchase Options" from Jackson Brothers' ITM to its ERP/OM. It is only at this point that the offers are displayed to the salesperson in the ERP/OM user interface. The salesperson discusses the options with the customer, perhaps emphasizing offers with the highest scores. When the customer chooses one of the options (in this scenario, the top-rated offer from Sibley Tires) the salesperson selects that choice in the ERP/OM user interface. Next, Jackson Brothers' ITM sends an "Offer Acceptance" message to Sibley. Sibley's ITM checks whether the tires are still in inventory (to ensure they were not sold during the intervening time) by sending an "Inventory Inquiry" to the ERP/OM. The ERP/OM replies that the tires are still in stock, and the ITM sends an "Order Request," confirmed by the ERP/OM with an "Order Response." Finally, the sale is consummated when Sibley's ITM sends an "Acceptance Response" to Jackson Brothers, which allows Jackson Brothers' ITM to enter this order into its enterprise via a "Purchase Response."

Note that a transmission failure of any of these messages results, ultimately, in the abortion of the transaction between Jackson Brothers and Sibley. In that event, the Jackson Brothers' salesperson would have the option of accepting one of the other offers.

5. Discussion of the Model

In this paper we have emphasized the intelligent decision-making aspect of our architecture. As mentioned earlier, we decided to develop a "simple" intelligence. We believe that a number of advantages accrue from this approach. First, it is readily understood and easy to use. This is substantiated by numerous meetings we have had with tire dealers, who have easily comprehended the model. Second, it is straightforward to customize and extend the model. The Decision Editor and Decision Configurator allow a user to customize existing criteria and add/subtract criteria, respectively. Finally, it is much easier to deploy and maintain than a complicated rule-based system or neural network would have been.

We considered using an industry-standard message protocol, such as proposed in [3]. Unfortunately, there are many competing standards (e.g., xCBL from Commerce One, cXML from Ariba). Many standards are tantamount to no standard. In the absence of a standard, we designed our XML documents to be generic, with as much domain independence as possible. For instance, instead of including product attributes that are specific to our initial deployment (e.g., tire size, speed ratio), we have defined generic attribute-value pairs that can be used in different domains.

Perhaps the most fundamental design decision was building a peer-to-peer trading model. While this aspect of the architecture has not been the primary emphasis of this paper, we have at least made it clear why we chose P2P and how it improves on the centralized approach of the B2B exchange. In particular, the P2P model directly supports our goals of autonomy, confidentiality, and decentralization.

Several advanced features have been considered for the system; however, we decided to first complete an initial version of the software that would allow us to establish a presence in the market and experiment with the software. For instance, we plan to extend the message protocol and decision-making infrastructure to support "negotiation." Currently, an offer from a seller to a buyer is final. In practice, however, negotiation could occur and eventually we will support this. We also plan to define messages and functionality to support "unsolicited sales offers," i.e., a seller actively trying to sell its inventory rather than waiting for purchase requests.

Another future development is integration with different ERP systems. While this task is nontrivial, we have designed our message protocol and software so that integration with new ERP software, such as SAP or PeopleSoft, will be as straightforward as possible. A new integration project is planned for late 2001.

Eventually, we want our software to improve its performance over time. For example, a buyer's ITM may be able to recognize that a particular seller always offers the best price, even though the seller is not highly rated by Buy-Side Partner Selection. In this situation, the buyer may want to automatically improve the seller's rating. Case-based reasoning [4, 7, 11] or machine learning [8] may enable such functionality. The trading protocol and results of a trade are similar to the steps and result of a plan, so a method used by a case-based planner [10] may be applicable to our problem. To achieve this capability, the system stores all transaction information, including the content of XML messages, the scores assigned by the decision process, and the selections made by retail clerks. Once sufficient data is collected, we will be better positioned to analyze and implement such a capability.

6. Conclusions

Although we are in the midst of our first deployment, we believe that the architecture described in this paper provides potential benefits over the B2B exchange model. The "proof in the pudding" will be in how well it benefits the business process it was designed to support. Our anticipation is that tradingNetwork™ will:

- *Improve customer service*. Our model provides each partner in a network with access to a large virtual inventory, thus expanding the customer's choices at the point-of-sale and potentially improving customer satisfaction.
- *Strengthen existing relationships*. AMR Research notes "exchanges won't be successful until they can support the existing relationships between buyer and seller." [1, p. 20] With tradingNetwork™, each partner can give preference to certain partners and streamline business interactions with those partners.
- *Minimize excess inventory*: The automated sell-side feature of the software is a means for reducing inventory, thus allowing trading partners to minimize stock.
- *Increase sales*: The automated sell-side feature reduces the sales cycle for all non-manual sales, thus leading to more total sales.
- *Promote usability*: Since the technology is integrated with enterprise software, it is available to retail salespeople. The integration is transparent to the salesperson, thus making it more usable.

Finally, the intelligent decision-making capability, the feature of the software most emphasized in this paper, allows a business to standardize its buying and selling process, in accordance with its business objectives. Through the autonomy and confidentiality provided by the P2P model, the buying and selling preferences of each partner remain the private province of that partner.

7. References

[1] Fontanella, J. and O'Brien, D., "The AMR Research Survey on Trading Exchanges: B2B Means Back to Basics." AMR Research, October, 2000.

[2] Giarratano, J. C. and Riley, G., *Expert Systems: Principles and Programming*. PWS Publishing Company, 1998.

[3] Glushko R. J., Tenenbaum, J. M., and Meltzer, B., "An XML Framework for Agent-Based E-commerce." *Comm. of the ACM*, Vol. 42, No. 3, March, 1999.

[4] Kolodner, J., *Case-Based Reasoning*. Morgan Kaufmann Publishers, 1993.

[5] Kowalski, T. J. and Levy, L. S., *Rule-Based Programming*. Kluwer Academic Publishers, 1996.

[6] Linthicum, D. S., *B2B Application Integration: e-Business-Enable Your Enterprise*. Addison-Wesley, 2000.

[7] McLaren, B. M., *Assessing the Relevance of Cases and Principles Using Operationalization Techniques*. Ph.D. thesis, University of Pittsburgh, 1999.

[8] Mitchell, T., *Machine Learning*. WCB/McGraw-Hill, 1997.

[9] Sculley, A. B. and Woods, W. W. A., *B2B Exchanges: The Killer Application in the Business-to-Business Internet Revolution*. ISI Publications, 1999.

[10] Veloso, M. V., *Learning by Analogical Reasoning in General Problem Solving*. Ph.D. thesis, Carnegie Mellon University, 1992.

[11] Watson, I., Applying Case-Based Reasoning. Morgan Kaufmann, 1997.