e-Everything: e-Commerce, e-Government, e-Household, e-Democracy

14th Bled Electronic Commerce Conference

Bled, Slovenia, June 25 - 26, 2001

Automated Negotiation on Agent-Based e-Marketplaces: An Overview

Susanne Klaue, Karl Kurbel, Iouri Loutchko

Europe University Viadrina Frankfurt (Oder), Business Informatics,
P.O. BOX 1786, D-15207, Frankfurt (Oder), Germany
SKlaue@euv-frankfurt-o.de, Kurbel@euv-frankfurt-o.de, Loutchko@euv-frankfurt-o.de

Abstract

This paper first describes a few existing models for product and merchant brokering and automated negotiation in the framework of agent-based e-marketplaces. Our analysis shows that only simplest types of automated negotiation protocols are currently presented online, whereas some powerful systems have been already implemented in the product and merchant brokering field. Then we give an overview of theoretical approaches to automated negotiation including game theory based negotiation, multi-attribute utility theory based negotiation and auction based negotiation. We also discuss some further challenges for researchers and practitioners working in the field of automated negotiation in the framework of agent-based electronic marketplaces.

1. Introduction

An e-marketplace can be shortly described as online parallel to the physical marketplace. It can exist in different forms including auctions, product exchanges, online shopping markets and e-catalogs. There are also some specific characteristics of e-marketplaces compared to the traditional marketplaces. In an online marketplace, potential buyers and sellers exchange information about goods or services, reaching agreements through information alone; physical goods and services may be delivered to the customer later on, outside of marketplace. That allows the buyers to search for the best from the huge number of products available on the e-marketplace; merchants can promptly react on the situation on the e-marketplace to satisfy potential customers in a more efficient manner. These new opportunities and very rapidly changing environment of the e-marketplaces lead to the need of ever more sophisticated services to assist both buyers and sellers. In particular, the buyers have to be provided with the mechanisms that first help them to specify their preferences and then to search for the best possible choice based on preferences specified by the buyers. In some cases, when the preferences both on the buying and selling sides are vague, some automation of the negotiation between buyers and sellers is useful. Automating negotiations gives rise to new challenges for developers of architecture and software technologies underlying e-marketplaces.

One of the most popular and promising answers to the new challenges posed by arising e-marketplaces is the software agents' paradigm. The term 'software agent' as used in this paper refers to a program that acts independently in the framework of an e-marketplace in the interests of its user. It should be noted that there are contexts in e-commerce in which both software agents and automated negotiations are not adequate. If negotiations get too complex then automation is not the right choice. Instead there should be support of human users who then have to conduct the negotiation themselves.

According to [Nwana, 1996], the most important characteristics of software agents are Autonomy (agents process their work independently and proactively without the need for human management), Cooperation (agents are able to communicate with one another, negotiating on certain issues), and Learning (agents are able to learn as they react or interact with their environment and other agents). We point out that agents acting on e-marketplaces are usually autonomous. Their capabilities to cooperate differ from marketplace to marketplace. The focus of this paper will be on the cooperation aspect. Namely, we consider a few existing models for automated negotiation in the framework of agent-based systems and the most important theoretical approaches to automated negotiation including game theory based negotiation, multi-attribute utility theory based negotiation, and auction based negotiation.

The term 'negotiation' as used in this paper refers to automated negotiation where all parties involved are software agents. However, we do not assume that negotiation is always a complex process including reasoning, proposals and counter-proposals. For convenience a simple exchange of information between two agents is also considered a negotiation.

Negotiation is an important part of commerce today, and automated negotiation will be a vital part of electronic commerce in the future [Beem and Segev, 1997]. Despite this fact current e-commerce technology hardly supports automated negotiation. The main reason is that negotiation is difficult, and automated negotiation is even more so. To give only two of many examples highlighting this difficulty, we refer to the following issues: the need for an ontology to describe the negotiation's objectives, and the need for the agents' strategies for negotiation.

The remainder of this paper is organized as follows. In section 2 we review some practical implementations of agent-based e-marketplaces designed for product and merchant brokering and for automated negotiation. In section 3, we analyze some of existing theoretical approaches to automated negotiation. Section 4 concludes the paper.

2. Practical Implementations

According to [Guttman et al., 1998] agents in the field of e-commerce can be used and are especially important for the three stages in the framework of their Consumer Buying Behavior (CBB) model: Product Brokering, Merchant Brokering, and Negotiation corresponding to *what* to buy, *who* to buy it from, and *how* to determine the terms of the transaction respectively.

Ideally, all three stages of the CBB model should be implemented equally well in the framework of an e-marketplace. Though in the paper we are mainly concerned with the negotiation aspects a short analysis of those e-marketplaces that only provide tools for product and merchant brokering is given too.

2.1 e-Marketplaces for Product and Merchant Brokering

As the first example of an e-marketplace designed for product and merchant brokering we consider the *Frictionless Commerce* online shop [http://retail.frictionless.com/].

In the framework of the *Frictionless Commerce* system, the consumer is given an opportunity to choose what to buy and who to buy a certain product from. The technology used by the Frictionless Commerce enables users to search for merchants selling the product based on value, not just price. Online consumers use the Frictionless' Engine to initialize their agents by selecting individualized criteria, such as product features and merchant services. Consumers' preferences can be vague. In addition, consumers are asked to rate their preferences ranging from 'must have' to 'not important'. On the merchant's side all preferences are crisp. The main problem to solve is to rank the crisp proposals coming from the merchants according to the consumer's vague preferences. The Frictionless' Engine solves this problem using multi-attribute utility theory [Raiffa, 1982] thus helping the consumer to determine what to buy and who to buy the product from.

Another system in the Product and Merchant Brokering field is the MATE (Multi-Agent Trading Environment) agent-based e-marketplace presented in [Owen et al., 1999]. The main difference of the MATE e-marketplace from the system described above is that a purchasing assistant and the merchants’ agents are added to the e-marketplace architecture.

The purchasing assistant supports a buyer both at the specification and the evaluation stages of the product and merchant brokering. At the specification stage the user inputs his preferences for the product features and the relative weights of these features using a user interface. It is important to note that the user interface specification is directly connected with the product ontology and thus dynamically changes according to the product under specification. The user’s preferences for some features may be vague, say, he or she specifies the price as an interval between the reservation price and the maximum price. At the evaluation stage the purchasing agent produce a ranking of merchants’ offers by using an evaluation algorithm that takes as input the buyer’s preferences for product features, the relative weights for these futures and the merchant's specifications. The evaluation algorithm used by the purchasing assistant on the MATE e-marketplace is derived from the multi-attribute utility theory [Raiffa, 1982].

One more important characteristic of the MATE e-marketplace is that every merchant is also represented by an agent acting on his behalf on the e-marketplace. The merchant's agent provides a user interface for creation and modification of product data in accordance with the product ontology as defined by the market. Another, probably even more important, functionality of the merchant's agent is the automation of the complex process of selection the products from the merchant's database which satisfy both the buyers' criteria and the business goals of the merchant. On the MATE e-marketplace, fuzzy-logic matching [Zadeh, 1992] is employed by the merchant's agent to select the best match against a buyer's requirement specification, even when no exact match is possible.

Because the offers of the merchants in the framework of the MATE e-marketplace all have crisp preferences there is no matter to negotiate; the main problems to handle are to select the best offer and merchant on the base of buyer's preferences and relative weights for the buyer's agent and to find the most suitable product in the product database on the base of buyers' criteria and merchant's business goals for the merchant's agent. The methods and technologies used to solve these problems (multi-attribute utility theory, fuzzy-logic matching) can be useful in the process of negotiation too. Moreover the whole process of product and merchant brokering can be considered to be a searching for the suitable partners to negotiate, i.e. a 'pre-negotiation' step.

As we see several powerful agent-based e-marketplaces in the fields of Product Brokering and Merchant Brokering have been already implemented. As to the Negotiation step of the CBB model the situation is not so good here. In the rest of the section we review some of the agent-based e-marketplaces with the negotiation between agents.

2.2 Auctions

Online auctions are no doubt the largest class of electronic marketplaces [Beam and Segev, 1997]. Because many online auctions can be regarded as multi-agent e-marketplaces we discuss some of their common features.

The first example we consider is an English auction in the framework of *eBay's Auction Web* [http://pages. ebay.com]. To sell something one has to provide a description of an item together with some preferences including payment method, where to ship, who will pay for shipment, minimum bid and reserve price (real amount the seller is willing to accept for this item). In fact, by providing this information the seller initializes an agent to negotiate about one issue (price) with vague preferences (the price should not be less than the reserve price). For a bidder, the auction system offers an optional "phantom" bidding service, called proxy bidding, which is equivalent to initialization of a buying agent. In proxy bidding, a bidder can confidentially enter the value of the maximum bid – his only vague preference (the price should be not higher than the maximum bid).

The *eBay's Auction Web's* English auction with "phantom" bidding service can be considered a multi-agent e-marketplace with negotiation between agents about one issue (price), although of a very simple type.

One more example is the *Fishmarket* electronic auction house [Rodriquez et al., 1997]. This academic prototype of an online auction has the age-old institution of a fish market as the underlying model. Both software agents and human beings can trade there. According to the trading conventions of Spanish fish markets the downward bidding protocol or Dutch protocol is used. What makes Fishmarket different from most other agent-based auction sites is that not only market-owned agents with simple architecture and strategies are employed. In addition, user-encoded buying and selling agents of arbitrary complexity may compete in the auction through standardized Java agent interface applets.

Another interesting academic prototype is a free-to-use online auction called *eAuctionHouse* that is one of the services of the *eMediator* system [http://ecommerce.cs.wustl.edu/emediator/]. This auction supports several important features including combinatorial auctions and bidding via software agents. In a combinatorial auction bidders may place bids on combinations of items. This feature is very useful if a bidder has preferences for bundles of items, as it may be the case in electricity markets, equities trading and bandwidth auctions.

	Kind of negotiation	**Negotiation issues**
eBay's Auction Web	English auction	Price
Fishmarket auction house	Dutch auction	Price
eAuction House	Combinatorial auction	Prices for combination of items

Table 1: *Agents' negotiation in auctions*

When more relevant issues than just the price exist, it is perhaps not the right choice in such complex circumstances to use auctions at all. Especially in the business-to-business area additional issues like terms and conditions, payment method, who will pay for shipment, timings, penalties, etc. may play an essential role; sometimes they are even more important than the price. In these cases the multi-issue negotiation with vague preferences on both is more suitable, than an auction.

2.3 Multi-Agents e-Marketplaces

As an example of an agent-based system for negotiation we consider first the *Kasbah* marketplace implemented at the MIT Media Laboratory [Chavez and Maes, 1996]. Users wishing to buy and users wishing to sell certain products in that marketplace can initialize agents by specifying what they want to buy or sell, the desired price, the highest acceptable price (for buying agents) or the lowest acceptable price (for selling agents), the date they want the transaction to be completed and a negotiation strategy. The negotiation strategy for a selling agent is specified by the 'decay' function the agent uses to lower the price over the given time frame. Similarly, a buying agent is given a 'raise' function for the bids.

Negotiation in Kasbah is straightforward. After buying and selling agents have been matched, buying agents offer to selling agents their current bids. Selling agents compare the bids with their current prices and reply with either "yes" (if one of bids is accepted) or "no" (if all bids are rejected). In the last situation, in a specified time interval, selling agents lower the prices and buying agents raise the bids according to their negotiation strategies. The procedure of offering the bids and comparing the bids with current prices is repeated until an agreement is reached or the time is up.

The Kasbah e-marketplace can be considered a multi-agent e-marketplace with multilateral and one-issue (price) negotiation and vague preferences on both sides. Selling agents have a span from the desired price to the lowest acceptable price, and buying agents have a span from the desired price to the highest acceptable price.

	Kind of negotiation	**Negotiation issues**
Kasbah	Negotiation through proposals; no critique and counter-proposals	Price
Tete-a-Tete (not finished)	Negotiation through proposals, critique and counter-proposals	Price and other value-added services (warranties, delivery times, return policies)
ZEUS-based e-marketplace	Negotiation through proposals; no critique and counter-proposals	Price

Table 2: *Agents' negotiation on multi-agent e-marketplaces*

Many real-world commercial negotiations have multi-issue character. As a response to this challenge and as continuation of the Kasbah project, a new project, *Tete-a-Tete*, was started at the MIT Media Laboratory [Guttman and Maes, 1998]. The aim of this project was to extend Kasbah towards a multi-agent system in retail sales with multi-issue negotiation, including warranties, delivery times, service contracts, return policies and other merchant's value-added services.

Initially it was intended to allow vague preferences on both sides (consumer and merchant) and negotiation with proposals, critique and counter-proposals. Due to several reasons the Tete-a-Tete project was later transformed into the Frictionless Commerce project discussed above.

One reason for the transition from the Tete-a-Tete marketplace to the Frictionless Commerce Engine was that each of the two negotiating sides in business-to-consumer e-commerce has a different status and non-degenerate multi-issue negotiation with vague preferences on both sides is neither practical nor really necessary. Secondly, the complexity of the negotiation problem in the Tete-a-Tete marketplace is substantially higher than in the Kasbah marketplace so that the solutions used in the Kasbah project were not directly applicable to the initial Tete-a-Tete project.

Finally we mention a recent project described in [Collis and Lee, 1998]. This project has been resulted in a prototype of a multi-agent e-marketplace with one-issue (price) negotiation and vague preferences on both sides. The new and interesting thing about this project is that the authors used the *ZEUS Agent Building Toolkit* for implementation of the multi-agent e-marketplace. The description of the ZEUS Agent Building Toolkit can be found in [Collis and Lee, 1998].

3. Theoretical Approaches

In this section, we give an overview and a short analysis of some theoretical approaches to automated negotiation including game theory based negotiation, auction based negotiation, and multi-attribute utility theory based negotiation. These approaches are of great importance both for the existing implementations of the agent-based e-marketplaces and for the future projects in this field.

3.1 Game Theory Based Bargaining

The sub-field of microeconomics known as game theory is widely acknowledged to provide a useful set of tools for the design of multi-agent architectures (Binmore and Vulkan, 1999).

In the framework of game theory traditionally two main branches are distinguished: cooperative and non-cooperative game theory. A topic like coalition formation is typically analyzed by cooperative game theory. Non-cooperative game theory deals largely with how intelligent individuals interact with one another in an effort to achieve their own goals. That is the branch of game theory we discuss here.

One way to describe a game is by listing the players (or individuals) participating in the game, and, for each player, listing the alternative choices (called actions or strategies) available to that player. In the case of a two-player game, the actions of the first player form the rows, and the actions of the second player the columns, of a matrix. The entries in the matrix are two numbers representing the utility or payoff to the first and second player respectively. A play consists of choosing certain strategies by the players; an outcome of the play is a pair of numbers representing the utilities of the players. To determine “rational” outcomes non-cooperative game theory defines the notion of an equilibrium strategy. Among most widely used concepts of equilibrium strategies are the Nash equilibrium and ‘dominant’ strategies. A dominant strategy is optimal for all players independent of what the strategies of the other players are. The strategies of players are in Nash equilibrium if no player can benefit by unilaterally changing his strategy.

Unfortunately, not every game possesses a dominant strategy and the strategies being in Nash equilibrium cannot always be interpreted as the most suitable ones. A very famous example is the Prisoner's Dilemma game. In this game, there exists a pair of strategies (non-dominant and not in Nash equilibrium) that are more suitable for both players than another pair of strategies that are in Nash equilibrium. The

Prisoner's Dilemma demonstrates that a mathematical definition of a “rational” outcome of a game and thus results of mathematical analysis of negotiation strategies supplied by the game theory cannot always be applied to practical problems.

Another limitation appearing by applications of game theory is that, at least in classical game theory, frequently simplifying assumptions are made for analysis of a game like the assumption about the full rationality of the players and the assumption that the rules of the play and the preferences and beliefs of the players are their ‘common knowledge’. Such assumptions limit the practical applicability of game-theoretic results. In particular, private information such as reservation prices, preferences for different features of the products and the relative importance of these futures as well as time preferences and limitations are usually hidden from the opponent in real-life negotiations.

In the past decades a variety of bargaining models has been developed in the framework of game theory. In particular, a lot of research was made in the fields of cooperative and non-cooperative bargaining with incomplete and vague information, and bargaining over multiple issues. One of the approaches is to specify a limited number of player “types”, each with a set of preferences and beliefs about the behavior of the other agents. Giving the distribution of beliefs the probability theory is applied to represent the uncertainty of outcomes, and the Bayesian law is frequently used to model the way in which new information revises beliefs. For the overview of the results and literature in other fields of game theory we refer to [Gerding et al., 2000].

Advantages	Critical Points
Provides a classification of negotiation	The notion of equilibrium strategy is not unique
In many situations, complete mathematical analysis is possible	Mathematically optimal solution is not always the most suitable one for practice
Bargaining with uncertain and vague information can be modeled	Simplifying assumptions are frequently made

Table 3: *Game theory based bargaining*

3.2 Auction Based Negotiation

The common view among economists is that an auction is the most effective way of resolving the “one to many” bargaining problem, both in theory and in practical applications. In this section we give a short overview of the auction types that are most widely used in online auctions and analyze them from the point of view of agents’ negotiation.

According to the results of a field survey of 100 auctions presented in [Beam and Segev, 1997] four basic types of single auctions (auctions with a single seller) can be found online: the ascending bid auction (English auction), the descending bid auction (Dutch auction), the first price sealed bid auction, and the second price sealed bid auction (the Vickrey auction).

In the English auction, the buyers successively raise the bids up to their reservation prices until only one buyer remains. In a standard terminology, the buyer's reservation price is the maximum price he is willing to pay for the item and the seller's reservation price is the lowest acceptable sale price for the item. Then, if a buyer's maximum bid meets or exceeds the seller's reservation price, the transaction between the seller and the buyer is completed. All buyers can see the high bid and, in some auctions, other buyers' bids. The rules of the English auction are easy to understand and to implement in the framework of an online auction. As a result, this auction type is the most popular among online auctions both without and with agent support.

In the descending bid auction (Dutch auction) the seller starts at a very high price, and then lowers the price continuously. All buyers can see the current price and then decide if the price is still too high or if they wish to purchase at that price. The first bidder at the current price wins the auction. As shown in [Rodriquez et al., 1997], the Dutch auction can be implemented as an agent-based online auction, too.

In the first price sealed bid auction each buyer independently submits in secret a single bid. Then bids are opened simultaneously and the item is sold to the buyer who makes the highest bid. Nobody is allowed to update a bid once submitted, and the winner pays the highest price bid.

The Vickrey auction is similar to the first sealed bid auction. Each buyer independently submits a single bid, without seeing bids of other buyers, and the winner is the buyer with the highest bid. However, the price paid by the winner is the second-highest bid.

Both the first price sealed bid auctions and the Vickrey auctions don't allow buyers to update dynamically their bids. That makes agent technology in general and agents' negotiation in particular hardly advantageous for these auction types.

Advantages	**Critical Points**
One-to-many negotiation	One-issue (price) negotiation
Mathematically optimal strategies can be found both for sellers and for buyers	Theoretical results can not always be applied in practice
Bargaining with uncertain and vague information can be modeled	Simplifying assumptions are frequently made

Table 4: *Auction based negotiation*

Auction theory provides us with an understanding of the conditions under which the various auction types are optimal and which strategies are optimal for participants in different auctions. For example, if all buyers participating in a first price sealed bid auction are perfectly rational then, under some additional conditions, the buyer's optimal bid is (n-1)r/n, where n is the number of buyers participating in the auction and r is the buyer's reservation price. For more details about optimal strategies of buyers and sellers participated in different auctions we refer to [Beam, Segev and Shanthikumar, 1996].

However, in the real word, all buyers are seldom equally well-informed, skilled and perfectly rational, so that the results obtained from the mathematical analysis of auctions not always lead to the best bidding strategies in practice.

3.3 Multi-Attribute Utility Theory

As we have seen in Section 2.1, some e-marketplaces for product and merchant brokering have been designed for helping the consumer to determine what to buy and who to buy the product from according to consumer's preferences. Some of consumer's preferences can be vague. In this case, the multi-attribute utility theory is normally used to rank the crisp proposals coming from the merchants according to the consumer's vague preferences. In the more general situation of multi-lateral negotiation with many issues and with vague preferences both by buyers and sellers, the same problem should be treated, namely, to rate the offers coming from the negotiation partners according to own vague preferences.

In this section, following [Raiffa, 1982], we present some elements of the multi-attribute utility theory. We suppose that m participants take part in the negotiation and the negotiation subject can be characterized by n issues, all of them of numerical nature. Let x_j^i denote the value for issue j (j = 1,...,n) offered to the negotiation participant i (i = 1,m) by another participant at some moment. In general, an interval of values is acceptable by each participant, i.e., $a_j^i \le x_j^i \le b_j^i$, j =1,...,n, i = 1,...,m. Different values from this interval can be of different worth for every negotiation participant. The worth of values of negotiation issues is modeled by scoring functions:

$$S_j^i : [a_j^i, b_j^i] \to [0,1], \text{ j = 1,...,n, i = 1,...,m.}$$

The bigger the value of a scoring function for a certain value of an issue is, the more suitable is this value for the negotiation participant.

In a real negotiation, different negotiation issues are of different importance for every negotiation participant. To model this, we introduce the notion of relative importance that a participant assigns to each issue under negotiation. Let ω_j^i be the relative importance of issue j, j = 1,...,n, for the participant i, i = 1,...,m. For convenience, we assume that the normalization relation is valid:

$$\sum_{j=1}^{n} \omega_j^i = 1, \text{ i = 1,...,m.}$$

Now we suppose that negotiation participant i, i = 1,...,m, is given an offer. Because the negotiation is characterized by n issues, the offer can be represented by a vector

$x = (x_1,...,x_n)$. Using the scoring functions and relative importance of issues under negotiation we can introduce the notion of a general scoring function:

$$S^i(x) = F^i(x, S_1^i(x),..., S_n^i(x), \omega_1^i,..., \omega_n^i), F^i : R^{3n} \to R, \text{ i = 1,...,m.}$$

The exact form of scoring functions depends on a concrete situation. In many cases, linear function can be used to model the utility:

$$S^i(x) = \sum_{j=1}^{n} \omega_j^i S_j^i(x), \text{ i = 1,...,m.}$$

This situation is the simplest one from the mathematical point of view. If all negotiators use the linear scoring functions, it is possible to compute the optimum value of x (see [Raiffa, 1982]) giving theoretical value for the 'best deal'.

In a real negotiation, however, the final result achieved in the process of negotiation will depend on negotiation strategies, even in the case of linear scoring functions.

Advantages	**Critical Points**
Many-to-many multilateral negotiation	Scoring functions are problem- and user-dependent
For linear scoring functions, optimal value of the 'best deal' can be found	For non-linear scoring functions, the mathematical analysis is very difficult
Bargaining with uncertain and vague information can be modeled	No general recommendations how to construct the negotiation strategies

Table 5: *Multi-attribute utility based negotiation*

4. Summary

Most of today's multi-agent systems on the Internet use simple and static negotiation rules and are designed to negotiate about the price, not about the warranties, delivery times, service contracts, return policies and other merchant's value-added services. Methods and tools to process vague preferences and uncertain information are rarely used. This is definitely a shortcoming since the ability to negotiate and make decisions in a vague environment can be regarded as one of the main characteristics of an agent acting on an electronic marketplace.

The conventional approaches often used to model and analyze negotiation, like game theory or auction theory, are hardly useful for modeling of multi-lateral negotiation about many issues and with vague preferences and uncertain information on either negotiating side. New ideas in this field are required. Using multi-attribute utility theory in the framework of the multi-agent negotiation seems to be very promising.

References

Beam, C., Segev, A. (1997): Automated negotiations: A survey of the state of the art, *Wirtschafts-informatik* 3/97, pp. 263-267.

Beam, C., Segev, A. and Shanthikumar, J.G. (1996): Automated negotiations through Internet-based Auctions, CITM working paper 96-WP-1019, available at http://haas.berkeley.edu/~citm/nego.

Binmore, K. and Vulkan, N.(1999): Applying game theory to automated negotiation, *Netnomics* No. 1, pp. 1-10.

Chavez, A., Maes, P. (1996): Kasbah: An agent marketplace for buying and selling goods, Proceedings of the First International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology (PAAM-96), London, 22-24 April 1996, pp. 75-90.

Collis, J. C. and Lee, L. C. (1998): Building electronic marketplaces with the ZEUS toolkit, *Proceedings of the Agent Mediated Electronic Trading Workshop* (*AMET'98*), pp. 17-32.

Gerding, E.H., van Bragt, D.D.B. and La Poutre, J.A. (2000): Scientific approaches and techniques for negotiation. A game theoretic and artificial intelligence perspective. *CWI report on negotiation*, available at http://www.cwi.nl/projects/TA/.

Guttman, R.H., Maes, P. (1998): Cooperative vs. Competitive multi-agent negotiations in retail electronic commerce, *Proceedings of the Second International Workshop on Cooperative Information Agents (CIA'98),* Paris, France, July 3-8.

Guttman, R.H., Moukas, A.G., Maes, P. (1998): Agent-mediated electronic commerce: a survey, *The Knowledge Engineering Review* **13**, pp. 147-159.

Nwana, H. (1996): Software agents: An overview; The Knowledge Engineering Review **11,** pp. 205-244.

Owen, M.J., Lee., L., Sewell, G., Steward, S. and Thomas, D. (1999): Multi-agent trading environment, *BT Technology Journal* **17**, pp. 33-43.

Raiffa, H. (1982): *The Art and Science of Negotiation,* Cambrigde: Harvard University Press.

Rodriquez, J. A., Noriega, P., Sierra, C., Padget, J. (1997): FM96.5 – A Java-based electronic Auction house, Proceedings of the Second International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology (PAAM-97), London, UK, pp. 207-224.

Sandholm, T. (1999): e-Mediator: a next generation electronic commerce server, *AAAI Workshop on AI in Electronic Commerce,* pp. 46-55.

Zadeh, L.A. (Ed.) (1992): *Fuzzy Logic for the Management of Uncertainty*, John Wiley & Sons.