

Bilateral Negotiation with Incomplete and Uncertain Information: a Decision-Theoretic Approach using a Model of the Opponent

Chhaya Mudgal, Julita Vassileva

Department of Computer Science, University of Saskatchewan, Canada
{chm906, jiv}@cs.usask.ca

Abstract. The application of software agents to e-commerce has made a radical change in the way businesses and consumer to consumer transactions take place. Agent negotiation is an important aspect of e-commerce to bring satisfactory agreement in business transactions. We approach e-commerce and negotiation in the context of a distributed multiagent peer help system, I-Help, supporting students in a university course. Personal agents keep models of student preferences and negotiate on their behalf to acquire resources (help) from other agents. We model negotiation among personal agents by means of influence diagram, a decision theoretic tool. To cope with the uncertainty inherent in a dynamic market with self-interested participants, the agents create models of their opponents during negotiation, which help them predict better their opponents' actions. We carried out experiments comparing the proposed negotiation mechanism with influence diagram, one using in addition a model of the opponent and one using a simple heuristic approach (as a base for comparison). The results show some of the advantages and disadvantages of the proposed negotiation mechanisms.

1. INTRODUCTION

With the advent of e-commerce online businesses have become much more popular than before. Placing online orders, making payments electronically and finding information about the products and the vendors have become easier. Despite the comfort that e-commerce has brought with itself, humans are still involved in most of the important process of business, for example, in making decisions in all phases of buying and selling. Negotiation is one of the key factors in commerce systems, which involves a lot of decision making and tradeoffs between various factors. Agent technology has helped consumers by supporting their negotiation strategies. Some of the practical applications that assist users in negotiation are Auctionbot, Kasbah, Tete-a-tete, e-Bay [4].

Most business transactions in e-commerce involve negotiation to settle on the most suitable price for both parties. During negotiation individuals or organizations have to make decisions of varied nature to attain their objectives. The benefit of dynamically negotiating a price is that the resource is allocated to a consumer who values it the most instead of fixing the price in advance. Negotiation varies in duration and

complexity depending on the domain of application [4]. Negotiation protocols can extend over a long period of time, which is disadvantageous for time-bounded consumers.

A basic ingredient of the negotiation process is the correct anticipation of the other side's actions. In open multi agent systems (i.e. the systems in which the agents can be added dynamically) it is hard to know or predict about the state of the environment. In such systems there is always an element of uncertainty about the participants. We propose a negotiation mechanism that utilizes the elements of human behavior in the process of negotiation and which provides for anticipating the opposing party's actions and considering one's risk attitude towards money. The negotiation mechanism described in this paper has been developed as a part of a peer help system called I-Help. The purpose of negotiation is to find the best deal for the user independently on whether she requires help on a certain topic or is playing the role of a helper or another topic. I-Help is an online multiagent system that provides a student of a university course with a matchmaking service to find a peer-student online who can help [12]. Human help and time can be considered as non-tangible differentiated goods, which are traded for money. When the students in the class need help their agents contact a centralized matchmaker who knows which other agents are online and have the required resources (knowledge) and provides a ranked list of potential helpers. It also provides a standard marked price for the help request, which is calculated based on the difficulty of the topic and the number of knowledgeable users on this topic who are on line at the moment. The agent of the person requesting help starts a negotiation with the agent of the first potential helper from the list about the price (in our case this is the payment rate per unit of help time) and when a deal is made both agents inform their users. If the knowledgeable user agrees to help, a chat window opens for both sides and the help session is started. If the agents fail to achieve a deal, the agent of the person seeking help starts a negotiation with the second agent on the list etc... In this way, similarly to Kasbah [4] a one-to-many or many-many negotiation problem is modelled as a series of disconnected 1-1 negotiation problems.

I-Help's personal agents form an economic society designed to motivate the students who are knowledgeable to help their fellow students by receiving payment in cyber currency, which can later be evaluated or traded in terms of gift certificates or other means. In I-Help the agents make decisions on behalf of their users about the price to offer and how to increase or decrease the price to strike a better deal depending on user specified constraints, such as the urgency of the current work, importance of money and the risk behavior. Agent negotiation in I-help reduces the burden on the user (in our case the students) by allowing them to concentrate on their work rather than making them think about how to get a better deal.

In order to be able to make reasonable decisions for their users the agents consult their users' preference model. This model is initialized by the user and can be updated by the user, after receiving feedback from the agent about the success rate in negotiation.

The negotiation mechanism proposed in this paper is not restricted to intangible goods. It can be generalized to other market domains. Our approach is based on decision theory, which allows the agent to make rational choices. We will show how the agent benefits by using the decision theoretic approach. In the I-Help system the

environment is dynamic and since the agents represent real users, it is hard to predict the actions of the opponent¹ agent on the basis of its past behavior (since the user's preferences can change in the meantime). However, it is useful to try to model the opponent's behavior during one session, since this can help predict better the opponent's reaction.

2. RELATED WORK

Earlier work on negotiation in DAI was concerned with bringing cooperation and coordination among distributed nodes to improve the global efficiency of the system where the goals and the information of the system were not centralized. The application areas for negotiation included manufacturing, planning, scheduling, meeting scheduling, task and resource allocation in subcontracting networks. Various search techniques have been proposed to improve the efficiency of negotiation [2], [8]. On the other hand, some researchers studied negotiation from a theoretical perspective to find how agents should react to each other during their interaction using a game theoretic approach [15]. Negotiation mechanisms based on human negotiating techniques, using case based reasoning, argumentation and persuasion techniques have been proposed too [11]. Much work has been done on using agents to negotiate on behalf of their users, either in the market place [7] or for consumer to consumer negotiation. However, the focus of negotiation has been either to study the market performance or to search for appropriate negotiation strategies.

Recently, saving the users time and lifting the burden of information and decision overload, as well as the studies of the impact of different negotiation mechanisms on the outcome of negotiation has become the focus of automated negotiation. However, for a successful negotiation it is essential that the negotiating parties are aware of each other's moves. In the past researchers have used the history of negotiation to learn about the opponent. This approach works if the agents are interacting in a static environment, where no new participants appear. However, using the history of negotiation to learn about the opponent is not efficient when the environment is dynamic. Various approaches for modeling the opponent, learning the opponent's strategies and modeling the environment of the system have been studied to see their effect on negotiation. A Bayesian belief update mechanism has been applied in Bazaar [14] to update the beliefs of each agent about the environment and the opponent agent in the negotiation. Deterministic finite automata have been used to model the opponent's strategy in a game theoretic approach [1], Stochastic modeling using Markov chain has been proposed to capture the environment factors which influence the expected utility of a negotiating agent [6]. A recursive modeling approach based on reinforcement learning has been used to model other agents in order to adapt to each other and to the market system [13].

When the agents represent real users and negotiate on their behalf one of the requirements is that the negotiation should be transparent for the user and should take

¹ We will use the word "opponent" to denote the other agent in negotiation, though we don't imply necessarily an adversary or strongly competitive negotiation.

into account the user's preferences. Game-theoretic approaches have been applied for negotiation, but they have limitations: 1) the payoff matrix size grows exponentially with the number of agents; 2) it is very hard to find equilibrium points for strategies to be favourable, 3) though game theoretic approaches have been applied in theoretical research on strategy optimality, we have not seen much work in this area applied to practical market based systems. In addition, game theory assumes a win-lose situation, while our aim is towards win-win situation.

Therefore, negotiation mechanism we propose is based on sequential decision making in which the agent utilizes a preference model of the user. The user preference model is built by assessing the utility function that incorporates the user's risk attitude. Dealing with risk attitudes is an important feature of decision making. A few researchers [16] have taken into account risk attitudes for negotiation, but their work has mainly utilized the Zeuthen's Principle [3] in a game theoretic situation to determine who is the person more willing to make concessions. We utilize the risk attitudes in a different way, as discussed in section 3.

In the proposed work decision making in a negotiation process is modeled using an influence diagram. An influence diagram is a Bayesian network extended with utility functions and with variables representing decisions. An influence diagram is solved by computing the action yielding the highest expected utility. Influence diagrams have been applied in modelling decision making processes (for example, in [5]), however, not in the context of agent negotiation. We chose to use influence diagram for negotiation because even though a Bayesian network is a flexible tool for constructing models, it mainly considers causal impact between the events and hence is well suited for forecasting and diagnosing. Influence diagrams have been developed especially for making decisions. We view negotiation as a decision problem that requires a decision-maker to weigh his preferences and to perform an action that gives him** the maximum utility. Therefore influence diagrams serve as a good tool for representation of the problem as well as for solving it. It is also intuitively better understandable for a user than a Bayesian network, and since our agents represent human user, understandability is an important factor. Unlike decision trees, influence diagrams don't grow exponentially; they suppress minute details and hence are ideal for getting an overview of a complex problem.

We model the opponent's actions using a probabilistic influence diagram. Modelling the opponent in negotiation has been proposed in game theory and DAI [1, 13]. The agents of Carmel and Markovich [1] use model-based learning and explicit models of their opponent's strategies to generate expectations about their behaviour. Vidal and Durfee [13] have studied the impact of agent modeling each other on an information economy. Agents use reinforcement learning to model each other. Our agents use probabilistic influence diagram to infer the preferences of their opponents.

3. NEGOTIATION MECHANISM

Negotiation is an iterative process in which the agents make offers and counteroffers based on the preferences of their users.

3.1. Modeling Decision

We believe that a negotiation model based on sequential decision making should be based on following characteristics:

1. It should provide effective methods for organizing the domain problem into a structure. Structuring tools like influence diagrams [9] are an effective and efficient way to represent a negotiation context. A decision model also provides means to capture the nature of the problem, identify important objectives and to generate alternative courses of action.
2. The model should account for uncertainty and be able to represent it in a quantitative way, because there are inherent uncertainties in any practical environment.
3. The model should be able to support the dynamics of the situation.
4. The model should be able to deal with multiple objectives and allow tradeoff in one area against costs in another.
5. Finally, the model should be such that it allows the decision-maker to change their beliefs about the likelihood of uncertainties and change their preferences.

The negotiation protocol is a straightforward iterative process of making offers and counteroffers. So, during negotiation the agent can be in Offer or Counter-offer state repeatedly. The final state will be Accept or Reject. Similarly to [14], we use "negotiation strategies" to denote the actions, which the agents take in every iteration depending on the preference model. In our model once the agent is in a final state, he cannot retreat back from it. In order to do so the whole negotiation process has to start again. The negotiation mechanism takes into account the preferences of the user, which usually depend in the domain of the negotiation context. The preferences include:

- the maximum price of the buyer (i.e. how much the helpee is willing to pay),
- the urgency of demand of the resource for the buyer, or the urgency of the seller's current work (which she has to interrupt in order to help),
- the importance that either agent attaches to money, and
- the user's risk behavior (risk-averse or a risk-seeking person).

We have incorporated utility in order to model the way in which the decision-maker values different outcomes and objectives. Each agent in our system can be in any role; he can be a buyer or the seller of help. The utility for the buyer (helpee) and the seller (helper) for the actions of accept, reject and counter-propose vary according to their risk behavior.

It is important to note that the agent's risk behavior considered in this paper does not overlap with the money importance. In literature these entities have often been considered as tightly connected, but in our case this is not necessarily true. Money importance and risk-behavior are two different entities and they are set independently by the user (in the user preference model). The risk behavior of the user instructs the personal agent² about the increase or decrease in the price offers to be made. A risk-seeking person will try to counter-propose an offer rather than accepting any price offered. A risk-averse person will accept whatever minimum price he/she is offered

² Throughout the paper we will refer to the personal agent (buyer or seller) as "he" and to the user (student in our case) as "she".

and will refrain from counter proposing in fear of losing. We are calculating the utility values of the action alternatives that an agent has at any time during negotiation. Utility of actions is dependent on the money that the seller gets and the buyer has to pay. The utility of action also varies with the specified risk behavior of the user. For instance, as shown in the Figure 1 the utility of accepting an offer for a risk-averse buyer increases much slower as the difference between the offered price and the preferred price decreases. That means that as long as the price of the opponent comes closer to the preferred price of the buyer, he will be more willing to accept it, since there is not significant growth in utility if he continues to counter-propose. For a risk-seeking agent, the utility of continues to grow fast in this case, since he is willing to take the risk of contra-proposing, hoping to get a price even lower than his preferred price.




Fig. 1. Variation of U_{accept} for a **buyer**

Risk behavior also affects the increment and the decrement of the buyer and the seller. For a risk-averse buyer, if the urgency of current task is very high and the importance of money is also high, he will start by offering a price which is not too low compared to the maximum price his user is willing to pay. A risk-seeking buyer will start from a very low price and will try to get the lowest price he can get. For a risk-seeking seller the utility of accepting an offer increases, if he gets more money than what he has specified as his minimum price.

The functions that the agents use to increase or decrease their offers and counteroffers as a buyer and as a seller are defined as follows:

For Buyers

If $\text{max_price} > \text{std_price}$ then
 Offered price := $\text{std_price} - \Delta$
 Else
 Offered price := $\text{max_price} - \Delta$

For Sellers

If $\text{min_price} > \text{std_price}$ then
 Offered price := $\text{min_price} + \Delta$
 Else
 Offered price := $\text{std_price} + \Delta$

where std_price is the market price provided by the matchmaker. It is calculated based on the current situation of the market of help on this topic and on the difficulty of the topic, thus providing some measure for the actual worth of the resource. For both the buyer and the seller the values of Δ should not exceed their preferred prices, R . Δ is determined as follows (x is the offered price):

For Buyers

If urgency = very urgent then
If risk_behavior = risk seeking then
 $\Delta := 1 - e^{-x/R} \quad x > R$
If risk_behavior = risk averse then
 $\Delta := 1 - e^{-x/R} \quad x < R$

For Sellers

If urgency = very urgent then
If risk_behavior = risk seeking then
 $\Delta := \sqrt{\text{min_price}}$
If risk_behavior = risk averse then
 $\Delta := \log(\text{min_price})$

We use an influence diagram that has a conditional node representing the uncertainty about the other party (see Figure 2). The outcomes of this node are the probabilities that an opponent can be in any of the states accept, reject and counteroffer. Since the agent does not know anything about the environment or about the other agent's user, we consider that all the states in which the opponent can be are equally likely. At every step the agents have to choose between three actions: accept, reject and counter-offer. They do so by calculating the maximum expected utility for the actions, which are represented as the possible choices for the decision node in the influence diagram. In any practical application of negotiation there are often multiple objectives involved and there has to be tradeoff between one over the other. Before the decision is made the factors that are already known and affect the decision (deterministic nodes) are taken into account as they affect the actions to be made. The node corresponding to the opponent's action can be considered conditional since nothing is known about him. For the first experiment we treat the outcomes of the opponent node as equally likely. In the second experiment we replace the equal likelihood of the opponent's actions with the outcome of a model of the opponent using a probabilistic influence diagram.




Fig. 2. Influence Diagram for the decision model

4.2 Modeling the Opponent

Ideally (and as often is assumed in cooperative environments [15]) negotiating parties have full knowledge about the opponent. This, however, it is not the case when agents are self-interested. In trade and commerce there can be hidden intentional assumptions. It is unlikely that the user will be willing to share his preferences with other users (or their agents). In order to deal with such situation and still be able to know as much as possible about the opponent we deploy a model of the opponent by using a probabilistic influence diagram to model the opponent. Since in a dynamic environment like a market place where the situation is changing all the time and new buyers and sellers keep on entering and leaving the system, it is very costly for agents

to create and maintain models of the other the participants in the environment. Our agents have no prior knowledge about each other. After the first round of offers made the agent starts using his opponent's move to predict his reaction to the counteroffer that he is going to make. It is also important to note that we are not doing recursive agent modeling.

Figure 3 shows a probabilistic influence diagram; the oval nodes are conditional and the double-circled node is deterministic. Conditional probability distribution of the conditional nodes over the outcomes is assessed on the basis of the first offer. Probability distribution of the Opponent's action node can be calculated by performing reductions over the nodes. For instance, by performing the arc reversal from the Money Importance node to the Opponent's Action node, makes Money Importance a barren node. Hence, it can be removed from the diagram and a new conditional probability distribution is calculated. Conditional predecessors of the nodes (if any) are inherited. In a similar way the diagram can be simplified by using arc reversal operation and barren node removal, which finally gives the probability distribution for the Opponent's Action node. If the next move of the opponent does not match with the predicted action, Bayes' update rule is used to update the information. For more information about probabilistic influence diagrams refer to [9].




Fig. 3. Probabilistic Influence Diagram for the opponent's actions

Evaluating both influence diagrams and probabilistic influence diagrams is NP-hard. In our case the number of nodes is relatively small and there are no complexity problems. However, if the number of nodes increases drastically, negotiation might become computationally inefficient.

4. EXPERIMENT

The goals of the evaluation are to test the performance of the negotiation mechanism in terms of quality of deals that agents make for their users and to see if modeling the opponent brings benefits in negotiation. In order to answer these questions, we compare the deals obtained by agents using the proposed approach (decision theoretic with modeling the opponent) with deals obtained by the same agents under identical circumstances using other negotiation approaches. Since negotiation is done on behalf of the user, it is necessary to take into account her preferences. There is an immense number of possible preference combinations. For this evaluation, 5 different combinations of user preferences for helper and helpee are tested with 3 different combinations of preferred prices: one, in which the preferred price intervals of the negotiators overlap widely (called session 1 and denoted as S1), another one in which

the preferred prices overlap but not widely (session 2 / S2) and one where there is no overlap in the preferred prices (session 3 / S3). The experiments were carried out on a simulation.. The experimental setup follows several assumptions related to the peer-help application:

- The issue of negotiation is the price per unit of time at which the help session will occur.
- The matchmaker provides to both parties the standard price for help on each topic, which is calculated by a central component based on the difficulty of the topic.
- The agent whose user needs help (helpee / buyer) begins the negotiation with the agent of the user who is on the top of the list of possible helpers.
- The agents do not know anything about the other agents in the system at the start of the negotiation process.
- Each agent makes an offer depending on the user's risk attitude. The decay function and the raise function are dependent on the money importance and the urgency of the work / help.
- The agents are allowed to offer the same price more than once -- in this way the agents stay rigid on a price, if they do not want to increment or decrement it. However, a seller is not allowed to ask for price greater than his previous offered value and a buyer is not allowed to ask for price less than his previous value within one negotiation session.
- Each agent ensures that it does not exceed the preferred price limit set by the user (the maximum price for the helpee / buyer and minimum price for helper/ seller).
We compared three negotiation approaches.

In Approach 1 the agents don't use decision theoretic approach for negotiation. Both the parties make offers and counteroffers by making an increment or a decrement by a fixed amount based on the difference between their preferred prices. This corresponds to an "ideal" case of fair trade between agents that are not self-interested, reveal their preferred prices and use the same strategy. In fact, there is no need of negotiation: the agents can immediately agree on a price which is the average between the preferred prices of the agents. In this way no agent wins and no one loses.

In Approach 2 the agents use an influence diagram to model the decision-making process, with a chance node corresponding to the opponent's action. The probabilities of the opponent's actions are taken as equal (i.e. no model of the opponent is created).

In Approach 3 the agents take into account the model of the opponent in their own decision-making, as explained in section 3.2.

Our experiments showed that increasing the "intelligence" of the negotiating parties increases the percentage of rejections in negotiation (as can be seen in Table 1). This could be interpreted as negative result, since the percentage of failed negotiations increases. However, it can be viewed also positively, since it means that the agents are rejecting deals that are not profitable, instead of accepting them. As it can be expected most of the rejections happened in session 3, where there is no overlap between the preferred price intervals of the negotiators.

A series of experiments were made to investigate the influence of the negotiation mechanism proposed (Approach 3) on the quality of deals achieved by using it. The preferences of the helpee and the helper were kept constant for all configurations and across the three sessions only the negotiation approach of one of the agents was

changed. Again, for each negotiation approach 15 experiments were carried out for each of the 5 different settings of user preferences and in 3 sessions depending on the preferred prices. The deals were analysed from the point of view of the helpee /buyer and of the helper / seller (i.e. depending on which agent changes its negotiation approach). The results (see Table 2) show that using the decision theoretic approach with modeling the opponent (Approach 3) brings better quality deals in most of the cases when there is a large preferred price overlap. Approach 1 is only better when there is little or no scope for negotiation. When both agents use Approach 3, the helper/seller gets a better deal, since he is in an advantageous position (one step ahead) in modeling its opponent. The agent that starts the negotiation is in a disadvantaged position.

Table 1: Rejections in the various configurations. The first column describes the negotiation approaches used by the agents in the configuration, e.g. S:1, B:1 means that both the seller and the buyer use approach 1.

Configuration	Total Rejections (in %)	% Rejected by Helpee (Buyer)			% Rejected by Helper (Seller)		
		S1	S2	S3	S1	S2	S3
S:1, B:1	0	-	-	-	-	-	-
S:1, B:2	0	-	-	-	-	-	-
S:1, B:3	20	-	-	100	-	-	0
S:2, B:1	0	-	-	-	-	-	-
S:2, B:2	6.67	-	-	100	-	-	0
S:2, B:3	20	-	-	33.3	-	-	66.6
S:3, B:1	26.6	-	-	0	-	-	100
S:3, B:2	26.6	-	0	0	-	100	100
S:3, B:3	53.3	50	0	66.6	50	100	33.3

Table 2: Comparison of the percentage of better deals achieved by using a different strategy. The left part of the table shows the better deals achieved by the helpee /buyer when switching strategies, and the right part shows the percentage of better deals achieved by the helper / seller when switching strategies.

Case comparisons from helpee's viewpoint	Best deals for S1 (in %)	Best deals for S2 (in %)	Best deals for S3 (in %)	Case comparisons from helper's viewpoint	Best deals for S1 (in %)	Best deals for S2 (in %)	Best deals for S3 (in %)
B: 1→2, S: 1	60	25	20	S: 1→2, B:1	60	40	25
B: 1→2, S: 2	100	80	0	S: 1→2, B:2	0	33.3	0
B: 1→2, S: 3	60	0	0	S: 1→2, B:3	80	60	0
B: 2→3, S: 1	100	80	100	S: 2→3, B:1	80	100	100
B: 2→3, S: 2	60	40	50	S: 2→3, B:2	100	100	100
B: 2→3, S: 3	100	50	0	S: 2→3, B:3	33.3	100	100
B: 1→3, S: 1	80	60	50	S: 1→3, B:1	100	100	100
B: 1→3, S: 2	80	80	50	S: 1→3, B:2	100	100	100
B: 1→3, S: 3	66.6	0	0	S: 1→3, B:3	66.6	100	100

5. CONCLUSIONS

In open environments with self-interested agents, a decision-theoretic approach to negotiation with modeling the opponent has proven to yield good deals. In this paper we have presented a negotiation mechanism that utilizes decision model (influence diagram) taking into account the preferences and the risk behavior of the user. Such a decision model allows to take into account and to handle tradeoffs among the factors that affect decision-making in negotiation and can be applied to any domain.

We have extended this negotiation mechanism to create and use a model of the opponent, represented with a probabilistic influence diagram. Our experimental results show that this mechanism finds a better deal for the agent who uses it when there is space for negotiation. We are currently implementing the proposed negotiation mechanisms in personal agents representing human users in an Internet-based virtual market environment for peer help (I-Help).

Acknowledgement: This work has been partially supported by NSERC under TL-NCE Project 6.28.

6. REFERENCES

1. Carmel, D., and Markovitch, S. Learning Models of Intelligent Agents in Proceedings of Third International Conference on Multiagent Systems.(1998), 64-71.
2. Durfee, E., and Lesser, V. Negotiating Task Decomposition and Allocation Using Partial Global Planning. Distributed Artificial Intelligence. Volume 2 Huhns, L., and Gasser, M.(eds) Morgan Kaufmann: San Mateo, California (1987), 229-243.
3. Harsanyi, J., Rational Behaviour and Bargaining Equilibrium in Games and Social Situations. Cambridge University Press, Cambridge, (1977).
4. Maes, P., Guttman, R., Moukas, G., Agents that Buy and Sell. Communications of the ACM. 42, 3, (1997) 81-83.
5. Suryadi, D., Gmytasiewicz, P. Learning Models of Other Agents using Influence Diagrams, in Proceedings of the Seventh International Conference on User Modeling, (1999), 223-232.
6. Park, S., Durfee, E., Birmingham, W. Advantages of Strategic Thinking in Multiagent Contracts. in Proceedings of Second International Conference on Multi-Agent Systems. (1996), 259-266.
7. Preist, Chris., Commodity Trading Using An Agent-Based Iterated Double Auction in Proceedings of the Third Annual Conference on Autonomous Agents, (1999), 131-138.
8. Sathi, A., and Fox S. Constraint Directed Negotiation of Resource Allocation. Distributed Artificial Intelligence, Volume 2, Morgan Kaufmann: San Mateo, CA (1987), 163-194.
9. Shachter, R., Probabilistic inference and influence diagrams. Operations Research. 36,4 (1988), 589-604.

10. Shachter, R., Evaluating Influence Diagrams. *Operations Research*. 34, 36, (1986), 871-882.
11. Sycara, K. Resolving goal conflicts via Negotiation in Proceedings Seventh National Conference on Artificial Intelligence. (1988).
12. Vassileva J., Greer J, McCalla G., Deters R., Zapata D., Mudgal C., Grant S. A Multi-Agent Approach to the Design of Peer-Help Environments, in Proceedings of AIED'99, Le Mans, France, July, (1999), 38-45. also available on line at: <http://julita.usask.ca/homepage/Agents.html>
13. Vidal, J., and Durfee, E., The Impact of Nested Agent Models in an Information Economy in Proceedings of Second International Conference on Multi-Agent Systems. (1996), 377-384.
14. Zheng, D., and Sycara, K. Benefits of Learning in Negotiation in Proceedings of Fifteenth National Conference on Artificial Intelligence, (1997). 36-41.
15. Zlotkin, G., and Rosenschein, J. Cooperation and Conflict Resolution via Negotiation among Autonomous Agents in Non Cooperative Domains. *IEEE Transactions on Systems, Man and Cybernetics*, 21, 6, (1991). 1317-1332.
16. Zlotkin, G., and Rosenschein, J. Negotiation and Task Sharing among Autonomous Agents in Cooperative Domains. In Proceedings of Eleventh International Joint Conference in Artificial Intelligence. (1989). 912-917.