Chapter 16

BILATERAL NEGOTIATION WITH INCOMPLETE AND UNCERTAIN INFORMATION

Trading Help in a Distributed Peer Help Environment

Julita Vassileva
Department of Computer Science, University of Saskatchewan, 57 Campus Drive, Saskatoon S7N 5A9 Canada.
jiv@cs.usask.ca

Chhaya Mudgal
Simplified TeleSys Inc., 11100 Metric Blvd. #770, Austin, TX,78758, USA
cmudgal@hotmail.com

Abstract The paper describes a bilateral negotiation approach with incomplete and uncertain information that has been implemented 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. The agents negotiate iteratively using an 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. The evaluation shows the advantages and disadvantages of the proposed negotiation mechanisms.

Keywords: Negotiation, bargaining, decision-theory, personal agents.

1. INTRODUCTION

Despite the convenience that e-commerce has brought, humans are still involved in most of the important processes 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 decision-making and considering tradeoffs between multiple factors. The benefit of dynamically negotiating a price is that the resource is allocated to a consumer who values it the most. Negotiation varies in duration and complexity and can extend over a long period of time, which is a disadvantage for time-bounded consumers. Agent technology has helped human users by supporting or even replacing them in negotiation (Maes et al., 1997). Therefore, it is very important that users understand and trust the mechanism by which agents negotiate on their behalf.

A basic ingredient of the negotiation process is the correct anticipation of the opponent's actions. In open multi agent systems (i.e. the systems in which the agents can come and go at any time) there is always an element of uncertainty about the participants.

We propose a negotiation mechanism that imitates human behaviour in the process of negotiation, is time-sensitive, anticipates the opposing party's actions and considers the agent's risk attitude towards money. This mechanism has been implemented in personal agents trading peer help within a multi-agent based web-based system called I-Help (Vassileva et al., 1999). The paper presents the negotiation mechanism, an evaluation showing the advantage of modelling the opponent and finally talks about the real-world test of negotiation within I-Help.

2. THE APPLICATION AREA

I-Help is a web-based multiagent system that provides a student of a university course with human and electronic help resources. Human help and time can be considered as non-tangible differentiated goods, which are traded for money. As a part of the system, there is a matchmaking service to find a competent peer-student online to help (Vassileva et al., 1999). When the students in the class request peer help their personal agents contact a centralized matchmaker who knows who is online and has the required knowledge resources; it 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 bilateral negotiation (bargaining) 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 (Maes et al., 1997) a one-to-many or many-to-many negotiation problem is modelled as a series of disconnected one-to-one negotiation problems.

The reason for choosing a bilateral negotiation model over an auction is that it seems to be more understandable and natural from a human user point of view. First, it allows for considering multiple factors in the negotiation, not only the price (even though there are multi-value auctions that allow this too). Second, it allows a more personal way of making deals, since the agent is not obliged to give the help-request to the lowest-price bidder (helper in our case). Users always have the final say and they may decide not to accept/give help to a given person or at a given moment. Finally, bilateral negotiation avoids the problem with building cliques and corrupting the bidding process.

Agent negotiation in I-Help reduces the burden on the user (in our case the students) by helping them concentrate on their work rather than thinking about how to get a better deal. Viewed on a macro-level, I-Help's personal agents participate in an economy, designed to motivate the knowledgeable students to help their peers. For this they receive payment in virtual currency, called ICU (I-Help Currency Unit), which can later be exchanged for some real world equivalent (gift certificates, marks, etc.). In I-Help the personal 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, comprising the users' preference model, such as the urgency of their current task, importance of money and the risk behaviour. 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 I-Help environment is dynamic (users log in and out all the time, and when they are on line they have with different tasks, priorities etc.) and since the personal agents represent real users, it is hard to predict the priorities of the opponent agent on the basis of its past behaviour (since the user's preferences can change from session to session). However, it is useful to try to model the opponent's behaviour during one session, since this can help predict better the opponent's reaction.

3. THE NEGOTIATION MECHANISM

We define bilateral negotiation as an iterative process in which the agents make offers and counteroffers based on the preferences of their users. The negotiation protocol is a straightforward iterative process of agents making offers and counteroffers. During negotiation the agent can be in *Offer* or *Counter-offer* state repeatedly. The final state will be *Accept* or *Reject*. Once the agent is in a final state, it cannot retreat back. In order to do so the whole negotiation process has to start again.

Thus, from a microscopic (i.e. from an individual agent's) point of view, we can consider negotiation as a decision problem that requires a decision-maker to weigh preferences and to choose an action (from the set of actions allowed by the negotiation protocol) that gives the maximum utility.

3.1. MODELING DECISION

Decision making in the negotiation process is modelled using an influence diagram (ID) (Shachter, 1988). ID is a Bayesian network, extended with utility functions and with variables representing decisions and especially suited for modelling decision-making problems. IDs can represent multiple objectives and allow tradeoffs in one area against costs in another. IDs allow accounting for uncertainty and are able to represent it in a quantitative way. Unlike decision trees, IDs do not grow exponentially; they suppress minute details and hence are suitable for getting an overview of a complex problem. IDs have been widely used in modelling decision-making processes (for example, in (Suryadi and Gmytrasiewicz, 1999)), however, not in the context of agent negotiation.

The agent decision making (the ID) takes into account the preferences of the user, which depend on the domain of the negotiation. The preferences in the domain of I-Help include:

- the maximum price of the buyer (i.e. the helpee),
- the urgency of demand of the resource for the buyer, or the urgency of the seller's current work (which has to be interrupted in order to help),
- the importance that either agent attaches to money (either to get more money or to save money), and
- the user's risk behaviour (risk-averse or a risk-seeking person).

We have incorporated utility in order to represent how the decision-maker values different outcomes and objectives. The utility of a decision depends on the role in which the agent is at the moment of decision making (each agent in I-Help can be in the role of a buyer or a seller of help). The utilities for the buyer (helpee) and the seller (helper) for the actions *accept*, *reject* and *counter-propose* vary according to their risk behaviour.

It is important to note that the agent's risk behaviour does not overlap with the money importance. In literature these factors have often been considered as tightly connected, but in our case this is not necessarily true. Money importance and risk-behaviour are two different variables in the user preference model and they are set by the user. Risk behaviour affects the increment and the decrement of the proposed/asking price.

Similarly to (Zheng and Sycara, 1997), we use "negotiation strategies" to denote the actions, which the agents take in every iteration depending on the preference model. 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 ($max\_price > std\_price$) then
    Offered price := $std\_price - \Delta$
Else
    Offered price := $max\_price - \Delta$

For Sellers
If ($min\_price > std\_price$) then
    Offered price := $min\_price + \Delta$
Else
    Offered price := $std\_price + \Delta$

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, which provides 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
        $\Delta := 1 - e^{\frac{-x}{R}}, x > R$
    If ($risk\_behavior = risk\ averse$) then
        $\Delta := 1 - e^{\frac{-x}{R}}, x < R$

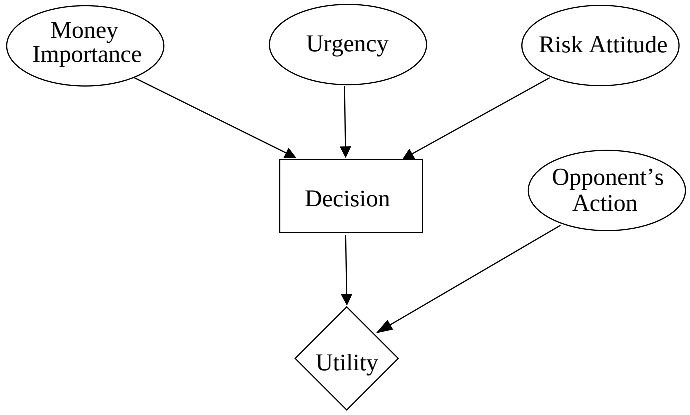

Figure 16.1. An influence diagram for the decision model.

For Sellers

If ($urgency$ = $very\ urgent$) then
    If ($risk\_behavior$ = $risk\ seeking$) then
        $\Delta := \sqrt{min\_price}$
    If ($risk\_behavior$ = $risk\ averse$) then
        $\Delta := \log(min\_price)$

We use an influence diagram that has a conditional node representing the uncertainty about the opponent's action (see Figure 16.1). The outcomes of this node are the probabilities that an opponent can decide to accept, reject and counteroffer. At every step the agents choose between these protocol actions by calculating the maximum expected utility for the actions, which are represented as the possible choices for the decision node in the influence diagram. 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. If the agent does not know anything about the opponent, the node corresponding to the opponent's action can be considered conditional and the states in which the opponent are equally likely. However, if the agent has some knowledge about the opponent (either from pervious encounters or from learning from the opponent's behaviour during the negotiation), the probability in the node representing the opponent's action will be calculated using the model of the opponent.

3.2. MODELLING THE OPPONENT

Ideally (as often is assumed in cooperative environments (Zlotkin and Rosenschein, 1991)) negotiating parties have full knowledge about the opponent. This, however, is not the case when agents are self-interested. In I-Help, it is unlikely that the user will be willing to share her private preferences with other users (or their agents). In order to negotiate more effectively, the personal agents model the preferences of the opponent using a probabilistic influence diagram. Unfortunately, in a dynamic market where new buyers and sellers can enter and leave the system at any time, it is very costly for agents to create and maintain models of all agents they have ever encountered. For this reason, the agents do not keep opponent models across sessions (when the user logs out). After the first round of offers in the negotiation each agent creates a model of its opponent and uses it to predict its reaction to the counteroffer that he is going to make. However, unlike (Gmytrasiewicz and Durfee, 1995), for computational efficiency reasons, there is no recursive agent modelling.

Figure 16.2 shows a probabilistic influence diagram (Shachter, 1988); 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 (Shachter, 1986). For instance, 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. 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 recalculate the probabilities. 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, negotiation might become computationally inefficient.

4. EVALUATION

The goals of the evaluation are to test the performance of the negotiation mechanism in terms of the number of successful negotiations and the quality of deals that agents make for their users. Another goal is to see if modelling the opponent brings benefits in negotiation. Finally,

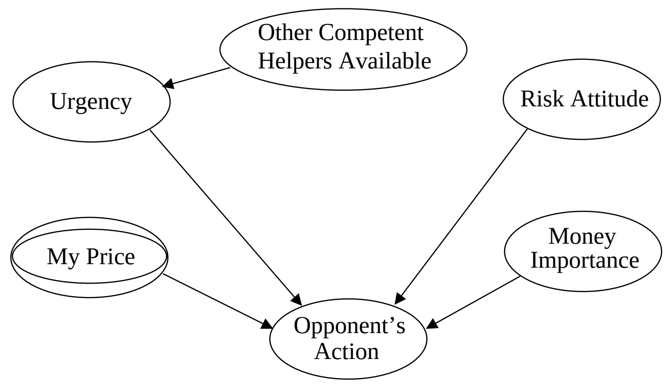

Figure 16.2. An influence diagram for the opponent model.

we want to see if the agent negotiation approach is acceptable for their human users. In order to answer the first two questions, we compare the deals obtained by agents using the proposed approach (decision theoretic with modelling the opponent) with deals obtained by the same agents under identical circumstances using other negotiation approaches.

4.1. EXPERIMENTAL SETUP

Since negotiation is done on behalf of the user, it is necessary to take into account her preferences. There are 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 (Session 1/S1), another one in which the preferred prices overlap partially (Session 2/S2) and one where there is no overlap in the preferred prices (Session 3/S3). The experiments were carried out between agents implemented as Java servlets using simple messaging to communicate. The experimental setup follows several assumptions related to the peer-help application. First, the issue of negotiation is the price per unit of help-session time. 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 and the availability of competent helpers/helpees requesting help on 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 price-offer decay and raise function are dependent on the money importance and the urgency of the current task/help-request. 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 do not use decision theoretic approach for negotiation. Both 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 price offer increment/decrement strategy. In this way no agent wins and no one looses. In Approach 2 the agents use an influence diagram in their decision-making processes, with a chance node corresponding to the opponent's action (thus, they are uninformed about the opponent; the probabilities of the opponent's actions are taken as equal). In Approach 3 the agents create and take into account the model of the opponent in their own decision-making, as explained in section 3.2.

4.2. PERCENTAGE OF REJECTIONS (FAILED NEGOTIATION)

Our experiments showed that increasing the "intelligence" of the negotiating parties leads to increase in the percentage of rejections in negotiation (as can be seen in Table 16.1). This could be interpreted as a negative result since the percentage of failed negotiations increases. However, it can be viewed also positively, since it means that instead of accepting, the agents are rejecting deals that are not profitable. As it was expected, most of the rejections happened in session 3, where there is no overlap between the preferred price intervals of the negotiators. That means that the agents are able to correctly infer that there is no big benefit to be gained in continuing negotiation. From the results one could generalize that the agent using the more advanced approach is usually the one that rejects, which is explainable, since the more "intelligent" negotiator decides to quit when it sees no sense to continue. However, there is an exception - in case (S:2, B:3), where the seller, though using less advanced approach, is responsible for 66.6% of the

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 16.1 . 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.

rejections (in session 3 when there is no overlap between the prices). It is hard to explain this result delving into speculations. The fact that in the case when both agents use the approach 3, the one that starts the negotiation (the buyer) is responsible for most of the rejections in session 3 is also hard to explain, since the buyer is one step behind the seller in modelling the opponent, i.e. the seller should be able to predict better that there is no point to continue and to reject.

4.3. QUALITY OF DEAL

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, 15 experiments were carried out for each negotiation approach, 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 16.2) show that using the decision theoretic approach and modelling the opponent (Approach 3) brings better quality deals in most of the cases when there is a large preferred price overlap (session 1 and session 2).

Approach 1 is only better when there is little or no scope for negotiation (session 3), which is easy to explain: the agents are in fact co-operative since they reveal their preferred prices. When both agents use Approach 3, the helper/seller gets always a better deal (the last three

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

Table 16.2 . Comparison of the percentage of better deals achieved by using a different strategy. The top part of the table shows the better deals achieved by the helpee/buyer when switching strategies, and the bottom part shows the percentage of better deals achieved by the helper/seller when switching strategies.

lines in the right side of the table), because of its advantageous position (one step ahead) in modelling its opponent. The agent that starts the negotiation is in a disadvantaged position, since it reveals its starting price first.

There are some anomalies in the cases when the seller switches to a more advanced approach while the buyer uses the same more advanced approach. For example, when the seller switches from Approach 1 to Approach 2 while the buyer uses Approach 2 in both cases, we see (in the 2nd line on the right side of Table 2) that the more advanced approach does not bring any benefits for the seller in sessions 1 and 3 and brings only a minor benefit (better deals in 1/3 of the cases) in session 2. In the case that the seller switches from Approach 2 to Approach 3 while the buyer uses Approach 3 all the time, we also see that there is not a big benefit in session 1 (only in 33.3% of the cases the seller achieves a better deals). However, in sessions 2 and 3, it achieves always better deals than with Approach 2.

4.4. USER ACCEPTANCE

This negotiation mechanism has been applied in two different deployments of the I-Help system at the University of Saskatchewan (available at `http://www.cs.usask.ca/i-help/`). Deployment 1 in fall 1999 was in a third year computer science class (operating systems) with about 60 students, and deployment 2 took place in fall 2000, in an introductory computer science class (Programming in C++) with 350 second year engineering students. The goal of both experiments was to see if the economy introduced by making personal agents trade help resources for currency helps to motivate human users to offer help. We also wanted to know if human users find negotiation of their behalf by their agents useful and if they trust their agents that they arrange the best deals for them. In both deployments trace data has been collected by I-Help itself as learners interact with the system. In addition we passed out questionnaires at the end to collect data about the students subjective feelings about the system. Unfortunately, the students rarely used the system and most of the usage concerned the public discussion forum provided by I-Help and the electronic help resources. Very few peer-help requests were made in both of the experiments, which did not allow us to test at all the user acceptance of the negotiation mechanism.

A number of social factors seem to affect I-Help usage. The choice of the group had a strong influence on the amount of use. Smaller or more cohesive groups do not need the system. Deployment 1 was with 3rd year students who knew each other very well, had established multiple ways of interacting with one another in the course and in the labs and hence they did not find any need to login to the system to get help. The reasons for this choice were purely pragmatic: time until the beginning of the term was short and implementation for this course required the least adaptation effort, for the domain representation and student modelling were already developed. A similar effect, however, appeared in deployment 2, event though we had a large group (3 parallel sections) of second year Engineering students. Due to the culture of the College of Engineering, involving a lot of group work and extra-curricular activities, the students knew each other well and had established knowledge networks, and they shared the same laboratory space with one another so there was ready access to face-to-face help. In addition, there were technical issues that prevented the students to use the peer-help component of the system. They usually accessed the system from a lab where the software required for their course assignments (Visual C++) is installed, rather than from their own computers. Unfortunately, the lab is using old and slow computers (Pentium I). Running the programming environment simultaneously with a browser consumes the operative memory entirely, which unacceptably slows down the performance in both I-Help and the programming environment.

Has the economy worked? There were a number of help-requests in deployment 1, but it isn't clear that the economy was the motivating factor. Respondents to the questionnaire administered after deployment 1 were evenly split as to whether they found the virtual currency motivating or not. Some students mentioned that it would be good to be able to exchange the accumulated help-currency in marks towards their final grade in the course. Two people were particularly negative. One problem may be that the currency exchange rate of I-Help credit units into things of value in the real world isn't very favourable (minimal prizes have been given for top helpers). Another problem may be that rewarding students solely on the level of their bank account does not take into account the quality of the help given. For example, in deployment 1, one student who was involved in many help sessions in the role of helper (17), abandoned 5 of his helpees during an ongoing discussion. Finally, perhaps the I-Help currency has to be converted into other things than material goods. Several students revealed that their main motivation for posting answers on the public discussion forum was "glory", that in this way they became recognised as "authorities" among their peers. Some other students mentioned that they hoped to attract the attention of the instructor, another form of recognition. Perhaps the I-Help currency has to be mapped onto fame and social status, not prizes. In fact, it seems to be generally recognised that social recognition is an efficient reward system also in many newsgroups and on the Internet for the developers of free software (Raymond, 1999). Though our I-Help data is inconclusive, we believe that some form of reward is necessary to stimulate student participation. The crucial question is the choice of the real world equivalent. The reward should be based on the social values of the group.

5. RELATED WORK ON NEGOTIATION

Earlier work on negotiation in DAI (Durfee and Lesser, 1987; Sasthi and Fox, 1987) is 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 are not centralized. The application areas for this type of negotiation include manufacturing, planning, scheduling, meeting scheduling, task and resource allocation in subcontracting networks. DAI approaches typically assume cooperative agents. However, in commerce applications, agents are usually self-interested.

Negotiation protocols and strategy optimality have been studied extensively (Zlotkin and Rosenschein, 1991), from a game theoretic perspective. Game-theoretic approaches are based on a sound mathematical model, but they have some important 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. they tend to focus on the "macro-level", thus modelling the general behaviour of the system, but not on the "micro-level", i.e. the reasoning and decision making process of an individual agent;

4. not much work in this area can be applied to build agents working in practical market based systems,

5. game theory assumes a win-lose situation, while in reality negotiation usually aims at achieving a win-win situation.

Micro-level negotiation mechanisms based on case-based reasoning, argumentation and persuasion techniques have been proposed by Sycara (Sycara, 1988). Recently, the focus of automated negotiation has become saving the users time and lifting the burden of information and decision overload, as well as the studying the impact of different negotiation mechanisms on the outcome of negotiation. For ensuring success in this type of negotiation it is essential that the parties are aware of each other's moves and motives. In the past researchers have used the history of past negotiations to learn about the opponent. However, using the history of negotiation to learn about the opponent is not efficient when the environment is dynamic, i.e. when new agents appear and leave, and agents change their strategies frequently. Various approaches for modelling the opponent, learning the opponent's strategies and modelling the environment of the system have been studied to see their effect on negotiation. Modelling the opponent in negotiation has also been proposed in game theory and DAI (Carmel and Markovitch, 1998; Vidal and Durfee, 1996). The agents of Carmel and Markovich (Carmel and Markovitch, 1998) use model-based learning and explicit models (based on deterministic finite automata) of their opponents' strategies to generate expectations about their behaviour. Durfee and his students (Gmytrasiewicz and Durfee, 1995; Vidal and Durfee, 1996) have studied the impact of agent modelling each other (recursive modelling) using reinforcement learning in an information economy. Combined approaches to model the opponent and the environment have been proposed, for example, using a Bayesian belief update mechanism (Zheng and Sycara, 1997) and stochastic modelling using Markov chains (Park et al., 1996).

Dealing with risk attitudes is an important feature of decision-making. A few researchers (Zlotkin and Rosenschein, 1991) have taken into account risk attitudes for negotiation, but their work has mainly utilized the Zeuthen's Principle (Harsanyi, 1977) 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.

Finally, we are not aware of much work aimed at evaluating agents negotiating on behalf of human users in a social context and in a real application with respect to user acceptance, trust and appropriateness of the economic interaction. Even though, we currently have no strong results in this area, we believe our work is an important first step in the right direction.

6. CONCLUSIONS

In open environments with self-interested agents, a decision-theoretic approach to negotiation with modelling the opponent has proven to be quite advantageous. In this paper we have presented a negotiation mechanism that utilizes an influence diagram taking into account the preferences and the risk behaviour 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 usually a better deal for the agent who uses it when there is space for negotiation. We have implemented this negotiation mechanism in personal agents representing human users in an Internet-based virtual market environment for peer help (I-Help). The negotiation mechanism proposed in this paper is not restricted to intangible goods. It can be generalized to other market domains.

Acknowledgements

The authors would like to thank Silvia Breban, a graduate student at the University of Saskatchewan, for re-implementing the negotiation mechanism described in this paper so that it could be incorporated in the I-Help agents. Thanks to Lori Kettel and Susan Bull for analysing the data from the two deployments of I-Help. This work has been partially supported by Telelearning Network of Centres of Excellence under Project .28 and by Canadian Foundation for Innovation under Project 3811.

References

Carmel, D. and Markovitch, S. (1998). Learning models of intelligent agents. In *Proceedings of the 3rd International Conference on Multiagent Systems*, pages 64–71.

Durfee, E. and Lesser, V. (1987). Negotiating task decomposition and allocation using partial global planning. In Huhns, M. and Gasser, L., editors, *Distributed Artificial Intelligence, Volume 2*, pages 229–243. Morgan Kaufmann, San Mateo, CA.

Gmytrasiewicz, P. and Durfee, E. (1995). A rigorous operational formalization of recursive modelling. In *Proceedings of the 1st International Conference on Multiagent Systems*, pages 125–132.

Harsanyi, J. (1977). *Rational Behaviour and Bargaining Equilibrium in Games and Soocial Situations*. Cambridge University Press, Cambridge, UK.

Maes, P., Guttman, R., and Moukas, G. (1997). Agents that buy and sell. *Communications of the ACM*, 42(3):81–83.

Park, S., Durfee, E., and Birmingham, W. (1996). Advantages of strategic thinking in multi-agent contracts. In *Proceedings of the 2nd International Conference on Multiagent Systems*, pages 259–266.

Raymond, E. (1999). Homesteading the Noosphere, the Cathedral and the Bazaar: Musings on Linux and Open Source by an Accidental Revolutionary. O'Reilly and Associates. available at `http://www.tuxedo.org/~esr/writings`.

Sasthi, A. and Fox, S. (1987). Constraint directed negotiation of resource allocation. In Huhns, M. and Gasser, L., editors, *Distributed Artificial Intelligence, Volume 2*, pages 163–194. Morgan Kaufmann, San Mateo, CA.

Shachter, R. (1986). Evaluating influence diagrams. *Operations Research*, 34(36):871–882.

Shachter, R. (1988). Probabilistic inference and influence diagrams. *Operations Research*, 36(4):589–604.

Suryadi, D. and Gmytrasiewicz, P. (1999). Learning models of other agents using influence diagrams. In *Proceedings of the 7th International Conference on User Modelling*, pages 223–232.

Sycara, K. (1988). Resolving goal conflicts via negotiation. In *Proceedings of the 7th National Conference on Artificial Intelligence.*

Vassileva, J., Greer, J., McCalla, G., Deters, R., Zapata, D., Mudgal, C., and Grant, S. (1999). A multi-agent approach to the design of peer-help environments. In *Proceedings of the International Conference on AI in Eductaion*, pages 38–45.

Vidal, J. and Durfee, E. (1996). The impact of nested agent models in an information economy. In *Proceedings of the 2nd International Conference on Multiagent Systems*, pages 377–384.

Zheng, D. and Sycara, K. (1997). Benefits of learning in negotiation. In *Proceedings of the 14th National Conference on Artificial Intelligence*, pages 36–41.

Zlotkin, G. and Rosenschein, J. (1991). Cooperation and conflict resolution via negotiation among autonomous agents in non cooperative domains. *IEEE Transactions on Systems, Man and Cybernetics*, 21(6):1317–1332.