An influence diagram model for multi-agent negotiation

Chhaya Mudgal and Julita Vassileva *
Department of Computer Science
University of Saskatchewan Saskatoon, Canada
{chm906, jiv}@cs.usask.ca

Abstract

Bilateral negotiation is an important aspect of e-commerce to bring satisfactory agreement in business transactions. Though auction-based models are more common currently, they are unable to accommodate tradeoffs between multiple parameters. This makes them inappropriate when trading with intangible goods, like information, expertise and help. 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 information resources (help) from other agents. Negotiation is modeled using influence diagram, a decision theoretic tool. Agents create models of their opponents during negotiation, which help them predict better the opponent actions. The experiments show that the proposed approach brings better deals to the agent who uses it. However, if both negotiating agents use this approach, the negotiation fails more often.

1. Introduction

Most business transactions in e-commerce involve negotiation to settle on the most suitable price for both parties. Bilateral agent negotiation has advantage over other negotiation mechanisms (e.g. auctions) in markets where the agents are expected to have repeated interactions, since it allows agents to learn about each other and develop relationships. Therefore a basic ingredient of bilateral negotiation is the correct anticipation of the other side's actions.

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 [2]. 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). It provides a ranked list of potential helpers and a standard market 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 negotiating with the agent of the first potential helper from the list about the price and when a deal is made both agents inform their users. If the knowledgeable user agrees to help, a chat window opens for both users. 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. In this way, similarly to Kasbah, a one-to-many or many-to-many negotiation problem is modeled as a series of disconnected 1-1 negotiation problems.

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.

2. Negotiation Mechanism

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

2.1 Modeling Decision

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*. 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 *preferred price* (this is, depending on the current role of the agent, the *maximum price* that the agent is willing to pay if it is in the role of a buyer /helpee or

the *minimum price* it is willing to accept, if it is in the role of a seller / helper),

- 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 varies according to their risk behavior.

The decision process is modeled with an influence diagram. 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. The uncertainty about the opponent can be represented with a conditional node (see Figure 1) with outcomes - the probabilities that an opponent can be in any of the states *accept, reject* and *counteroffe*r. If the agent has no information about the opponent these three states will be equally likely. In case the agent wants to model the opponent's behavior, the probability for the opponent's action can be predicted by the opponent's model. We have adopted the second approach.

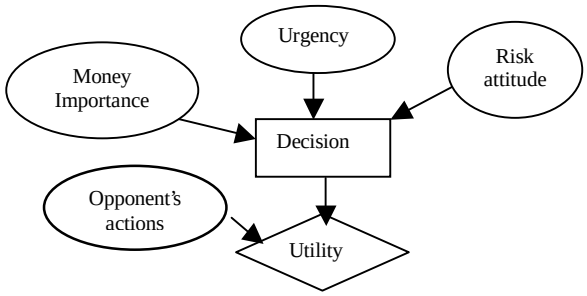

Figure 1: Decision Model: Influence Diagram.

2.2 Modeling the Opponent

In most negotiation approaches it is assumed that the negotiating parties have full knowledge about each other. This is not the case when agents are self-interested. It is unlikely that the user will be willing to share her preferences with other users (or their agents). In order to deal with this problem each agent creates and maintains a model of the opponent represented as a probabilistic influence diagram.

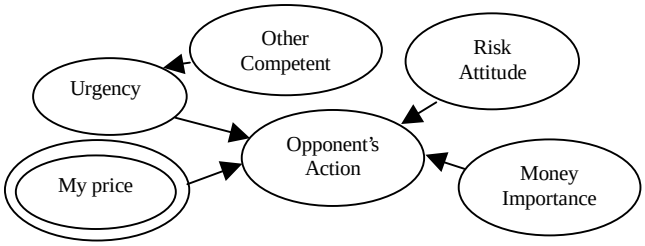

Figure 2: Opponent Model: Probabilistic Influence Diagram

In Figure 2 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. If the next move of the opponent does not match with the predicted action, Bayes' update rule is used to update the probabilities. The predicted action is fed into the decision model of the agent. For more information about evaluating probabilistic influence diagrams see [1].

3. Evaluation

We evaluated the proposed mechanism by comparing it with a simple increment / decrement approach for the offered prices and with a decision theoretic approach without modeling the opponent, i.e. the influence diagram containing a chance node for the opponent's action. The results showed that the decision theoretic approach brings benefit (better deals achieved in 80% of the cases) when the preferred price intervals of the negotiating agents overlap. Modeling the opponent is always advantageous when the agent negotiates with a "less smart" agent (agent which doesn't model the opponent, or who uses a simple increment/ decrement approach). When both agents model each other, the one that starts the negotiation is in a disadvantaged position (the other agent's model is more adequate since it uses one more observation in comparison with the opponent). The number of rejections (failed negotiations), however, increases drastically when both agents model each other (53.3% of all negotiations). In the case the seller (helper) rejects in 80% of the cases when there is small overlap between the preferred prices. The reason might be that the seller suspects the opponent in "buffing".

References:

[1] Shachter, R., Evaluating Influence Diagrams. *Operations Research.* Vol. 34, No. 36, 1988, 871-882.

[2] 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.