A Pareto Optimal Model for Automated Multi-attribute Negotiations

Guoming Lai
Tepper School of Business
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213, USA
guomingl@andrew.cmu.edu

Katia Sycara
Robotics Institute
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213, USA
katia@cs.cmu.edu

Cuihong Li
School of Business
University of Connecticut
2100 Hillside Road
Storrs, CT 06269, USA
Cuihong.Li@business.uconn.edu

ABSTRACT

This paper presents an applicable model for complex multi-attribute negotiations between autonomous agents. The model adopts a novel protocol which decomposes the original n-dimensional negotiation space into a series of negotiation base lines and in each period agents negotiate locally based on a given base line. A belief based negotiation strategy and an offer enhancement process are proposed for agents to make base offer on the negotiation base line and search for Pareto optimal enhancements of the base offer. The model achieves asymptotic Pareto optimality.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent systems

General Terms

Algorithm, Experimentation, Performance

Keywords

Multi-attribute negotiation, Pareto optimal

1. INTRODUCTION

Multi-attribute negotiation is important in practice. There exist situations where negotiators have to negotiate multiple issues together. For instance, an employer and a union usually need to simultaneously negotiate wage level, health care and vacations, etc. because those issues together determine the utility of the final contract for each other. Moreover, in some situations negotiators may also be willing to introduce additional issues into their negotiation because they may benefit from trading off the multiple issues when they have different preferences such that they can reach "win-win" outcomes. However, to negotiate multiple issues simultaneously is difficult.

AAMAS'07, May 14–18 2007, Honolulu, Hawai'i, USA.

This paper presents a formal model for automated multi-attribute negotiations which considers Pareto optimality and tractability at together. The paper introduces a protocol that breaks down the original space into a series of lines which are called *negotiation base lines* and the autonomous agents can negotiate based on those lines, which makes it tractable for agents to negotiate in the n-dimensional space without intensive searching and reasoning. The paper also proposes a belief-based negotiation strategy and an offer enhancement process which can help agents negotiate under incomplete information and achieve asymptotic Pareto optimal outcomes.

2. THE NEGOTIATION SETUP

There are two self-interested agents $i \in \{b, s\}$ negotiating a set of issues $j \in \{1, 2, ..., n\}$ with the range normalized to $[0, 1]$ for each issue. $\mathbf{0}^n/\mathbf{1}^n$ $(\mathbf{1}^n/\mathbf{0}^n)$ is the best/worst solution for agent $b(s)$. Each agent has a normalized utility function $U_i(\vec{x})$ on this space, with $U_b(\mathbf{0}^n) = 1$ and $U_b(\mathbf{1}^n) = 0$ for agent b and $U_s(\mathbf{0}^n) = 0$ and $U_s(\mathbf{1}^n) = 1$ for agent s. The preference of each agent is monotonic and strictly quasi-concave, i.e., for any solution $\vec{x}$, the set of solutions that an agent prefers to $\vec{x}$ is strictly convex. This implies that each Pareto optimal solution of a multi-attribute negotiation is on a joint tangent hyperplane of a pair of indifference curves (or surfaces) of the two agents, for tractability. Each agent sets a deadline (T_b/T_s) and an absolute reservation utility (r_b/r_s) for the negotiation–the least utility at which an agent can accept an offer. Each agent has a time preference parameter (δ_b/δ_s) for the negotiation, based on which she derives the relative reservation utility (r_b^t/r_s^t) of negotiation period t which represents the least utility she desires to get in period t. The ranges of the negotiation space and the best/worst point of each agent on the negotiation space are common knowledge, however, the preferences, the reservation utilities and the negotiation strategies are all private information and only known to an agent herself.

3. NEGOTIATION PROTOCOL

We propose a negotiation protocol which breaks down the n-dimensional space into a series of negotiation base lines and in each period agents negotiate locally based on a given negotiation base line. The protocol can be described:

Step 0: At the beginning, the two best points in the negotiation space for the two agents are connected by a line as the initial negotiation base line. Then one agent is chosen (randomly or by some rule) to be the first mover in this

978-81-904262-7-5 (RPS) ©2007 IFAAMAS

negotiation.

Step 1: The proposer in the current period first makes a base offer on the present negotiation base line by her negotiation strategy. The responder makes her decision whether to accept it or not. If it is accepted, the negotiation ends; otherwise, it goes to *Step 2*.

Step 2: The proposer finds new offers based on the offer enhancement process. The responder responds whether her utility has been improved and whether she decides to accept an offer. If an offer is accepted, the negotiation ends; else if the threshold of the enhancement process is reached, the negotiation procedure goes to *Step 3*; otherwise, *Step 2* is repeated.

Step 3: If one of the deadlines is reached, the negotiation ends; otherwise, the negotiation will proceed to the next period, where the agents first update their beliefs and exchange their roles; the negotiation base line is updated by connecting the current best offer (made by the proposer to the responder) with the one in the previous negotiation period; and then the procedure goes back to *Step 1*.

4. NEGOTIATION STRATEGIES

Based on the protocol, the negotiation within one period consists of two parts: (1) to propose and respond the base offer on the negotiation base line at the beginning, and (2) to propose and respond the enhancement offers. Since this paper focuses on the negotiations under incomplete information where the agents can have nonlinear preferences over the multiple issues, it is intractable for the agents to reason and get the opponent's preference and reservation utility precisely, even if they negotiate on the negotiation base lines. However, in practice, the negotiators can have some subjective beliefs over the opponent's characteristics which usually play an important role in a negotiation. With such a perspective, this paper proposes a computational negotiation strategy. This strategy is based on the agents' beliefs over the opponent's reservation price and the beliefs are updated in each period. The goal of the strategy is to maximize the agents' own expected utilities. In the following, we first present the responding strategy, and based on it we detail the proposing strategy.

The responding strategy adopted is (based on the perspective of agent s): if the offer $\vec{x}$ made by agent b in period t is no worse than $\vec{x}_{r^t_s}$, agent s will accept it; otherwise, it will be rejected. Thus,

$$a^t_s = \begin{cases} accept, & \text{if } \vec{x} \succeq_s \vec{x}_{r^t_s}; \\ reject, & o/w, \end{cases}$$

where a^t_s represents the reaction of agent s in period t.

The proposing strategy is (based on the perspective of agent b): to maximize her expected utility, agent b would take the offer $\vec{x}^*$ that follows:

$$\vec{x}^* = \arg \max_{\vec{x} \in \mathcal{S}_t \cap \vec{x} \succeq \vec{x}_{r^t_b}}$$

$$\left[\begin{array}{c} U_b(\vec{x}) \int_{[\vec{a}_{t-2}, \vec{x}]} f^t_b(\vec{m}) d\vec{m} + \delta_b \times \int_{[\vec{x}, \vec{b}_{t-1}]} f^t_b(\vec{m}) d\vec{m} \times \\ \left(\int_{[\vec{x}, \vec{y}_{r^{t+1}_b}]} U_b(\vec{z})\, g^t_b(\vec{z}) d\vec{z} + U_b(\vec{y}_{r^{t+1}_b}) \int_{[\vec{y}_{r^{t+1}_b}, \vec{b}_{t-1}]} g^t_b(\vec{z}) d\vec{z} \right) \end{array} \right],$$

where (a) $\mathcal{S}_t = [\vec{a}_{t-2}, \vec{b}_{t-1}]$ is the current negotiation base line with the endpoints $\vec{a}_{t-2}$ and $\vec{b}_{t-1}$; (b) $\vec{x}_{r^t_b}$ is agent b's current reservation offer (corresponding to her relative reservation utility in current period) and $\vec{y}_{r^{t+1}_b}$ will be agent b's next period reservation offer if agent s rejects agent b's current offer; (c) $f^t_b(.)$ is agent b's current belief (probability density function) on agent s's reservation offer and $g^t_b(.)$ is agent b's updated belief function on the offer that agent s will make if agent s rejects agent b's current offer contingent on the information revealed. Therefore, the first term in the formula is the expected utility of agent b if the offer $\vec{x}$ she chooses is acceptable for agent s and they reach agreement in period t, and the second term is the expected utility of agent b if agent s rejects the offer and the negotiation proceeds to period $t+1$.

5. THE OFFER ENHANCEMENT PROCESS

The above proposing strategy provides an approximating method for agents to choose the base offer at the beginning of each negotiation period. To improve the offer without further conceding, we divide a negotiation period into multiple sessions and allow the agents to have additional chances. The proposing agent in one negotiation period still can make further offers if the base offer is rejected by the responding agent–to look for enhancements of the base offer–until the session limit is reached. However, since the information is incomplete and both agents may have complex preferences over the issues, to find enhancement offers effectively is not a trivial task. We present a binary search approach which can assist agents to look for enhancement offers effectively and has the property of asymptotic Pareto optimality.

For clarity of presentation, we describe the enhancement process for a two-dimensional negotiation case (the procedure holds for $n (> 2)$ issues), in which we again assume agent b is the proposing agent and agent s is the responding agent. For a point $\vec{x}$ in the two-dimensional space, we use $\vec{x}(1)$ to represent the value of the first issue and $\vec{x}(2)$ to represent the value of the second issue, and we define the following concept.

Definition 1. Given a point $\vec{x}$, we call the range that contains the enhancement offers of $\vec{x}$ the **enhancement range** ER of $\vec{x}$.

We here use the value of the second issue[1] to characterize the enhancement range. The offer enhancement process then can be described as follows:

Step 0: Agent b–the proposing agent–sets the initial enhancement offer $\vec{x}_0$ to the base offer $\vec{x}^*$, picks her indifference curve $\mathcal{C}_b$ which passes through the base offer $\vec{x}^*$, and sets the original enhancement range ER_0 with the lower bound equal to ERL and upper bound equal to ERU. Note: ERL can be 0 or some positive value, which is determined by the intersection of the current indifference curve with the boundary of the negotiation space. It is similar for ERU, which can be 1 or some value smaller than 1.

Step 1: Given the latest enhancement offer $\vec{x}_n = (\vec{x}_n(1), \vec{x}_n(2))$ and the enhancement range ER_n in the process history[2],

[1]The range also can be characterized by the first issue.

[2]The process history stores the series of points $(\vec{x}_0, \vec{x}_1, ..., \vec{x}_n)$ that agent b has found in the current negotiation period, where $\vec{x}_0$ is the base offer and $\vec{x}_0 \preceq \vec{x}_1 \preceq \vec{x}_2 ... \preceq \vec{x}_n$ for both agents.

agent b chooses two points $\vec{x}_n^{eL}$ and $\vec{x}_n^{eU}$ on the indifference curve $\mathcal{C}_b$ which satisfy $\vec{x}_n^{eL}(2) = \frac{\vec{x}_n(2)+ERL}{2}$ and $\vec{x}_n^{eU}(2) = \frac{\vec{x}_n(2)+ERU}{2}$. Agent b proposes these two points as the potential enhancements of $\vec{x}_n$ to agent s.

Step 2: Agent s responds to these two proposals. Agent s first determines whether one of them is acceptable. If one is acceptable, the two agents reach agreement on it, and the process and the negotiation end; otherwise, agent s further determines whether one of them is better than $\vec{x}_n$ and responds to agent b.

Step 3: After receiving the response from agent s, agent b first checks whether the session limit is reached. If agent b will not propose any more, the two agents change their roles and the negotiation proceeds to the next period; otherwise, agent b updates the history of the enhancement process and the process then goes back to *Step 1*. To update the history, there are three scenarios depending on agent s's response. If the two potential enhancements are both worse than $\vec{x}_n$, then agent b sets $\vec{x}_{n+1} = \vec{x}_n$ and updates the enhancement range ER_n to ER_{n+1} with $ERL = \vec{x}_n^{eL}(2)$ and $ERU = \vec{x}_n^{eU}(2)$; else if $\vec{x}_n^{eL}$ is better than $\vec{x}_n$, then agent b sets $\vec{x}_{n+1} = \vec{x}_n^{eL}$ and updates the enhancement range ER_n to ER_{n+1} with ERL keeping the same and $ERU = \vec{x}_n(2)$; finally if $\vec{x}_n^{eU}$ is better than $\vec{x}_n$, then agent b sets $\vec{x}_{n+1} = \vec{x}_n^{eU}$ and updates the enhancement range ER_n to ER_{n+1} with ERU keeping the same and $ERL = \vec{x}_n(2)$. See Figure 1 for a depiction.

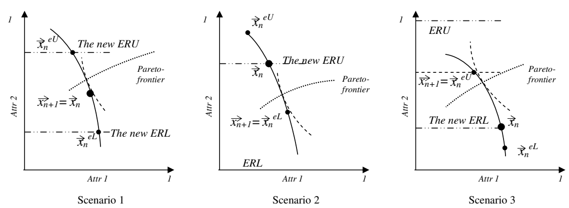

Figure 1: The scenarios of offer enhancement

The binary search approach adopted in this offer enhancement process can guide agents to look for enhancement offers and thus trade off the issues without sacrificing their utilities instantly. Moreover, this approach can provide asymptotic Pareto optimality property. To show this property, we first define the following concepts.

6. EXPERIMENTAL ANALYSIS

We provide an experimental analysis to evaluate the performance of the model. In Experiment 1, the two agents negotiate two issues who have constant elasticity of substitution (CES) utility functions of $u_b(x_1, x_2) = 1 - [.8x_1^3 + .2x_2^3]^{1/3}$ and $u_s(x_1, x_2) = 1 - [.3(1-x_1)^2 + .7(1-x_2)^2]^{1/2}$. The absolute reservation utilities and the deadlines of both agents are set to 0.2 and 12. Figure 2 shows the negotiation procedure. The dashed lines are the negotiation base lines which are updated in each period. Agents make the base offers on those lines and use the offer enhancement process to find Pareto optimal enhancements of those base offers. The dashed curve in the figure is the Pareto frontier. The two agents reach agreement in the 6th period with the agreement very close to the Pareto frontier.

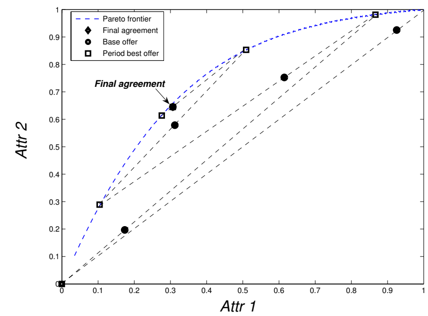

Figure 2: The negotiation process of experiment 1

7. CONCLUSION

Multi-attribute negotiation is important for agents to reach agreements on multiple issues, but it is much more complicated than a single-attribute negotiation. This paper provides a formal model that autonomous agents can use to negotiate multiple issues in general negotiation contexts with considering Pareto optimality. We show that this model is applicable in the situations where agents have *nonlinear* utility functions and the information is *incomplete*. The model also simplifies a multi-attribute negotiation by transforming it into negotiations on a series of base lines, thus avoiding searching the whole negotiation space and considers Pareto optimality at the same time. The numerical analysis shows the model achieves near Pareto optimality. In future work, we plan to take more sophisticated strategic behaviors of the agents into account.

8. ACKNOWLEDGEMENT

This research has been sponsored in part by AFOSR Grant F49620-01-1-0542 and in part by NSF grant IIS-0205526.

9. REFERENCES

[1] S. Fatima, M. J. Wooldridge and N.R. Jennings. An agenda-based framework for multi-issue negotiation. *Artificial Intelligence*, 152:1-45, 2004.

[2] K. Sycara. Problem restructuring in negotiation. *Management Science*, 37:1248-1268, 1991.

[3] P. Faratin, C. Sierra and N. R. Jennings. Using similarity criteria to make issue trade-offs in automated negotiations. *Artificial Intelligence*, 142(2):205–237, 2002.

[4] H. Ehtamo, R. Hamalainen, P. Heiskanen, J. Teich, M. Verkama and S. Zionts. Generating pareto solutions in a two-party setting: constraint proposal methods. *Management Science*, 45:1697–1709, 1999.

[5] M. Klein, P. Faratin, H. Sayama and Y. Bar-Yam. Negotiating complex contracts. MIT Sloan Working Paper No. 4196-01, Boston, USA, 2001.