A Multi-issue Negotiation Framework for Non-monotonic Preference Spaces

(Extended Abstract)

Miguel A. Lopez-Carmona, Ivan Marsa-Maestre, Juan R. Velasco and Enrique de la Hoz
Computer Engineering Department, Universidad de Alcala
Edificio Politecnico, Ctra. N-II, Km 36.500
28871 Alcala de Henares (Madrid), Spain
{miguelangel.lopez, ivan.marsa, juanramon.velasco, enrique.delahoz@uah.es}

ABSTRACT

We present a framework for non-mediated bilateral multi-issue negotiation under non-monotonic preference spaces. The framework is based on a region-based recursive bargaining mechanism. Preliminary experimental evaluation shows that our approach may obtain approximate Pareto-optimal results in acceptable negotiation time with a low failure rate.

Categories and Subject Descriptors

I.2.8 [**Artificial Intelligence**]: Problem Solving, Control Methods, and Search—*heuristic methods*; I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*multiagent systems*

General Terms

Algorithms, Performance, Experimentation

Keywords

automated negotiation, multi-issue, non-monotonic

1. INTRODUCTION

We seek to address the challenges of issue interdependencies in negotiation, which yield intractably large contract spaces and utility functions with multiple local optima. In the existing research, nearly all the models which assume issue interdependencies rely on monotonic utility spaces, binary valued issues or low-order dependencies [1, 2]. We propose a novel generic framework for non-mediated two-agent automated negotiations, which is able to operate in complex non-monotonic utility spaces. It is based on the exchange of offers, defined as regions of the negotiation space. The joint exploration of the solution space is recursive, which means that, when agents agree on a given region, a new negotiation on lower-sized regions is performed within the agreed region. If a new agreement cannot be found, agents return to upper level regions to perform a new search. Preliminary experimental evaluation shows that the proposed negotiation framework achieves in approximate Pareto-optimal outcomes obtained in reasonable times.

Cite as: A Multi-issue Negotiation Framework for Non-monotonic Preference Spaces (Extended Abstract), Miguel A. Lopez-Carmona et al., *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. 1611-1612

2. NEGOTIATION FRAMEWORK

2.1 Preference Structure

Let $X = \{x_i | i = 1, \ldots, n\}$ be the *issues* under negotiation, where each issue x_i can be normalized to a continuous or discrete range d_i, and $s = \{x_i^s | i = 1, \ldots, n\}$ a *contract* defined by the issues' values. A *region* will be formed by the set of contracts lying within the hypersphere defined as a 2-tuple $R =< c, r >$, where $c \in D$ and $r \in \mathbb{R}$ define the center and the radius. Each agent $A_{i\in\{b,s\}}$ embeds a *utility function* $U_i : D \to \mathbb{R}$, which can be non-monotonic and non-differentiable. We define the *overall satisfaction degree* (OSD) of R as an estimate of its overall utility. Let $\{s_k \in D | k = 1, \ldots, nsc\}$ be a set of nsc uniformly distributed sample contracts in R, $u_{th} \in [0, 1]$ the reservation value for any contract, and l the number of acceptable contracts which satisfy $U_i(s_k) \geq u_{th}$, then $OSD(R, u_{th}) = \frac{l}{nsc}$.

2.2 Negotiation Protocol

Our negotiation protocol is formalized as a negotiation dialogue composed of a sequence of *bargaining threads* (BTHs): $N_d = \{b_{r_{i1}}^{t_0} \to b_{r_{i2}}^{t_1} \to \ldots\}$. Each BTH

$$b_{r_{im}}^{t_n} = \{(R_b, R_s)_{r_{im}}^{t_n} \to (res_b, res_s)_{r_{im}}^{t_n+1} \to \ldots \\ \to (R_b, R_s)_{r_{im}}^{t_{n+1}-2} \to (res_b, res_s)_{r_{im}}^{t_{n+1}-1}\}$$

is a sequential exchange of offers (regions) of size r_{im} and responses to the offers. $(R_b, R_s)_{r_{im}}^{t_n+a}$ represents the simultaneous exchange of offers of size r_{im}, and $(res_b, res_s)_{r_{im}}^{t_n+a+1}$ the responses to these offers. The dialogue admits three types of responses: *Accept*, *Reject*, and *Request*. Before the beginning of a negotiation dialogue agents agree on an ordered finite set of region sizes $RS = \{r_i | i = 1, \ldots, m; \forall l < k, r_l > r_k\}$, and on the number of possible BTHs of a given size. The r_m parameter represents the lowest region size and r_1 the highest region size. A negotiation starts with a BTH of size r_1, the goal of the agents being to reach an agreement on a region of size r_m. It is worth noting that $r_m \to 0$ represents a contract. Every time a region (offer) is accepted by the opponent, the current BTH ends, and negotiation moves towards a new thread of lower size. The search in the new thread is restricted by the domain of the reached agreement in the previous thread. However, if agents abort the dialogue in the current thread because of the impossibility of reaching an agreement, they return to negotiate regions of a higher size. This process may be seen as a search tree, which structures the joint exploration of the negotiation space. The acceptability of regions and offers and the unfeasibility of agreements in a bargaining thread govern the transitions between the different threads.

2.3 Responding and Proposing Mechanisms

The *responding mechanism* depends on *acceptance* $(ath_{r_{im}})$ and *quality* $(qth_{r_{im}})$ thresholds. Each agent individually and privately defines the thresholds for each region size: $ATH^i = \{ath^i_{r_1}, ath^i_{r_2}, \ldots, ath^i_{r_m}\}$, $QTH^i = \{ qth^i_{r_1}, qth^i_{r_2}, \ldots, qth^i_{r_m}\}$. The responsive strategy for agent A_b will be (for A_s the strategy is similar):

$$\begin{cases} \text{Accept} & OSD((R_s)^{t_n+a}_{r_{im}}, u_{th}) \geq ath^b_{r_{im}} \\ \text{MovementRequest} & OSD((R_s)^{t_n+a}_{r_{im}}, u_{th}) < ath^b_{r_{im}} \text{ AND} \\ & OSD((R_s)^{t_n+a}_{qr_{im}} - (R_s)^{t_n+a}_{r_{im}}) \geq qth^b_{r_{im}} \\ \text{Reject} & \text{Otherwise} \end{cases}$$

If the offer is accepted, the current BTH ends. If the opponent's offer is not accepted, then the agent evaluates its quality by computing the OSD of the surroundings of the offer. If the OSD is higher than $qth_{r_{im}}$ then the agent responds with a message containing a request for offer movement; otherwise, the agent rejects the opponent's offer. The *request for offer movement* is defined as a vector $\bar{v}q_{(R_s)^{t_n+a}_{r_{im}}}$ which indicates the preferred direction expressed for an agent for the movement of $(R_s)^{t_n+a}_{r_{im}}$. In order to obtain $\bar{v}q$, we use the center of mass of the filtered samples (those above the utility threshold u_{th}) taken in the computation of the OSD for the surroundings of the opponent's offer.

The *proposing mechanism* is based on three basic mechanisms by which an agent generates regions: 1) *Root Region Mechanism*, A_b applies simulated annealing to her utility function to find a local maximum s^{t_n+a}, and generates the region $(R_b)^{t_n+a}_{r_{im}} =< s^{t_n+a}, r_{im} >$; 2) *Directed Child Region Mechanism*, A_b generates a child region $(R_b)^{t_n+a+2}_{r_{im}} =< s^{t_n+a+2}, r_{im} >$, where $s^{t_n+a+2} = s^{t_n+a} + 2 * r_{im} * \bar{v}q_{(R_b)^{t_n+a}_{r_{im}}}$; 3) *Random Child Region Mechanism*, A_b generates a child region $(R_b)^{t_n+a+2}_{r_{im}} =< s^{t_n+a+2}, r_{im} >$, where $s^{t_n+a+2} = s^{t_n+a} + 2 * r_{im} * \bar{v}q_{random}$. The child region is generated on a random direction from the center of the parent region.

To prepare any offer, an agent generates a region $(R_i)_{r_{im}}$ by means of any of the three mechanisms described above, and then evaluates if its OSD is above the acceptance threshold $ath_{r_{im}}$. The rules which govern the generation of offers within a BTH are: 1) The *first region* in a BTH is always a root region; 2) Any *unacceptable region* is discarded and then a new search is performed for finding a new region. If the discarded region is a root, the agent searches for a root region; otherwise, the agent generates a random child region; 3) The *rejection* by the opponent of an offer implies that the agent moves to the rejected offer's parent, and then searches for a new random child region; 4) An agent tries to generate a directed child region upon the reception of a *movement request*; 5) The number of negotiation rounds within a BTH is bounded. If this limit is reached, the BTH is considered *unfeasible*; 6) The opponent's *acceptance* of a previous offer implies the end of the current BTH, which in turns implies a final agreement or the beginning of a new BTH of lower size.

3. CONCLUSIONS

We have followed a sequential evaluation of parameters. For each evaluation we used the best configuration obtained in the previous parameter analysis. In the setting we study, we considered two agents bargaining on three issues, and simulated the non-monotonic preference scenario with an aggregation of *Bell* functions. This type of functions are usually used in the construction of landscapes for the performance evaluation of evolutionary optimizers.

The experimental results show that the generation of root regions plays an important role in the quality of the solutions. It is expected that an agent can obtain better agreements if she searches for better local optima when generating root offers. However, if the agent is too precise in the search, the probability of getting stuck in a restricted set of local optima or in a global optimum which comprises a zone of no agreement may increase. We analyzed this issue by performing experiments with two different values for the maximum number of iterations in the optimization process which generates the root region centers. From the results we conclude that there are two main advantages of using very few iterations: the negotiation times are considerably reduced at a minimum cost in terms of optimality, and the probability of negotiation failure decreases.

To investigate the influence of the highest-sized region, we varied the configuration of the region r_1 from 2% to 60% (with respect to the issues' domain length), for a fixed search depth $m = 15$. The best results in terms of negotiation time and optimality were obtained for $r_1 = 40\%$. Regarding the acceptance threshold, the experiments show that the results improve when the agent starts searching with low acceptance thresholds, and progressively increases the thresholds for lower sized regions. It was also expected that the distribution of region sizes had an influence on the outcomes. We tested several distributions of region sizes, obtaining the best results for an equispaced distribution.

Finally, we conducted several experiments in which the distribution of quality thresholds varied. The aim was to evaluate the usefulness and performance of the quality measure mechanism. From the results we conclude that the quality evaluation mechanism may contribute to a significant improvement in the negotiation results when using a constant quality threshold and this threshold is not excessively low.

4. ACKNOWLEDGMENTS

This work has been supported by the Spanish Ministry of Education grant TIN2008-06739-C04-04.

5. REFERENCES

[1] M. Klein, P. Faratin, H. Sayama, and Y. Bar-Yam. Protocols for negotiating complex contracts. *IEEE Intelligent Systems*, 18(6):32–38, 2003.

[2] I. Marsa-Maestre, M. A. Lopez-Carmona, J. R. Velasco, T. Ito, M. Klein, and K. Fujita. Balancing utility and deal probability for auction-based negotiations in highly nonlinear utility spaces. In *21st International Joint Conference on Artificial Intelligence (IJCAI 2009)*, pages 214–219, Pasadena, California, USA, July 2009.