Negotiating Efficient Outcomes over Multiple Issues

Sabyasachi Saha
Math & CS Department
University of Tulsa
sabyasachi-saha@utulsa.edu

Sandip Sen
Math & CS Department
University of Tulsa
sandip-sen@utulsa.edu

ABSTRACT

It is difficult to reach optimal outcomes in bilateral negotiations with multiple issues when the agents' preferences and priorities of the issues are not common knowledge. Self-interested agents often end up negotiating inefficient agreements in such situations. Some existing multiagent negotiation frameworks involve agents revealing their preferences to a trusted mediator. But in real-life situations, such a third party, trusted by both the agents may not be found. We design a protocol for bilateral multi-issue negotiation. This protocol guarantees envy-free and Pareto optimal agreement with minimum revelation of their preferences when both negotiators are rational and have the same ordinal preferences. This protocol also leads rational agents to envy-free and near-optimal solution in all cases.

Categories and Subject Descriptors

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

General Terms

Performance

Keywords

Negotiation, multi-issue, efficiency

1. INTRODUCTION

Negotiation is the most well-known approach for resolving conflicts in human and agent societies. With the rapid growth of E-commerce, research in automated negotiation has become increasingly important [4]. In automated negotiation, agents negotiate contracts on behalf of real-life negotiating parties they represent. An agent can negotiate resources and services with other agents or humans. Agents may bargain over a single issue of contention, but negotiations often involve bargaining over multiple issues [5], which may or may not be correlated. The issues are considered to be correlated if the utility of one issue depends on the consumption of any other issue(s). In this paper, we consider one-time negotiations with two agents negotiating over multiple, unrelated issues. We assume that the participating agents have no *a priori* knowledge about the preferences of the other agent. This situation is common in E-commerce negotiations, where the number and diversity of agents is so large that an agent may not have any estimate about the preferences of any particular agent it is negotiating with.

A negotiation process can be *distributive* or *integrative* in nature. In a *distributive* negotiation scenario, a gain for one participant implies a loss for the other. In *integrative* negotiation, agents can reach outcomes that are mutually beneficial. A single issue negotiation is usually a *distributive* negotiation. Multi-issue negotiations, on the other hand, where the ordinal and cardinal preferences are different, are integrative negotiations. In multi-issue negotiation, agents with known, divergent preferences can cooperate to reach agreements beneficial for both agents. But when the preferences are not known, using existing protocols, self-interested agents often fail to explore win-win possibilities and end up with inefficient agreements. By efficient or optimal solution, we refer to a solution which is *envy-free* and *Pareto optimal* [1].

Existing research shows that, if the agents are willing to reveal their complete preferences, it is possible to reach an efficient solution [3, 5]. If the agents, however, are unwilling to reveal their preferences at all, rational agents will end up with an equal split solution[1], which is not Pareto optimal unless both agents have identical preferences. So, there exists an efficiency-revelation trade-off. Our goal in this paper is to find out effective protocols for bilateral, multi-issue negotiation that can lead the participating agents to optimal or near-optimal outcomes while revealing only minimal information. Our work has some similarity with the work on contract negotiation by Sandholm [6] and Endriss *et al* [2] in the context of task allocation, where agents negotiate to reallocate tasks. But they want to identify, under complete information, necessary deals or exchanges required among agents to reach efficient allocation from an arbitrary initial allocation. We, on the contrary, want to find mechanisms that lead agents, without any prior information about opponent's preferences, to efficient outcomes without any third party intervention.

AAMAS'06 May 8–12 2006, Hakodate, Hokkaido, Japan.

[1]Each agent gets exactly half of each issue.

2. NEGOTIATION FRAMEWORK

We formally describe a representative negotiation scenario for allocation of available resources as a 3-tuple $\langle \mathcal{A}, R, \mathcal{U} \rangle$, where $\mathcal{A} = \{A_1, A_2\}$ is the set of agents, R is the resource whose allocation is being negotiated, and $\mathcal{U} = \{U_1, U_2\}$ is the set of utility functions, where U_i is the utility function of agent i. The resource $R = \{r_1, r_2, \ldots, r_H\}$ has $H \geq 2$ components and we assume that one continuously divisible unit of each component is available for allocation. Each component can be considered as a negotiation issue. Henceforth we will use the terms "component" and "issue" interchangeably. The negotiating agents must agree on the allocation of each component of the resource. An agent wants to obtain the entire unit of each components and some components are more important to it than others. Each agent has a complete preference order over the different components, which is not known to the other agent. Each agent can only assume that other agent is rational and maximizes its utility.

We now define some terms formally:

Outcome: An outcome O is an allocation of the resource components to different agents. Formally, $O : R \mapsto \mathbb{R}^{2\times H}_{[0,1]}$, where $\mathbb{R}^{2\times H}_{[0,1]}$ is a $2 \times H$ matrix, whose (i,j) cell, $x^O_{ij} \in [0,1]$, presents agent i's allocation of the j^{th} resource or component such that $x^O_{1j} + x^O_{2j} \leq 1$. If the agents cannot reach an agreement, they receive no share of the component, *i.e.*, $x^O_{ij} = 0$ for all i and j. Otherwise, $\sum_{i=1}^{2} x^O_{ij} = 1$, for $j \in \{1, \ldots, H\}$, *i.e.*, each component is exhaustively partitioned between the agents.

Total preference information: This is the relative importance of all the issues to a particular agent. For any agent i, total preference information is given by $\{w_{i1}, \ldots, w_{iH}\}$, where $w_{ij} \geq 0$ and $\sum_{j=1}^{H} w_{ij} = 1$. This is an agent's private information.

Utility: For agent i, utility of an outcome O, is defined as:

$$U^O_i = \sum_{j=1}^{H} w_{ij} * x^O_{ij}. \quad (1)$$

Pareto optimality: An outcome O is Pareto optimal, if there exists no other outcome O' such that all agent's utility in O' is at least as good as their utility in O and at least one agent's utility is more in O'. If outcome O is Pareto optimal, then there cannot exist any outcome O' such that, $U^{O'}_i \geq U^O_i$ for both $i = 1, 2$ and $U^{O'}_i > U^O_i$ for at least one i.

Envy-free: An outcome is envy-free, if and only if each agent's total utility for its own share is at least as much as its utility for the share assigned to the other agent [1].

3. PROPOSED NEGOTIATION PROTOCOL

In this section, we present the protocol that guarantees envy-free and Pareto optimal outcomes to two rational agents negotiating over multiple issues if their ordinal preferences are same. We refer to this protocol as the protocol leading to Pareto Optimal Solution through Equivalence (POSE).

3.1 POSE Protocol

In this protocol, two agents interact with a commonly observable random device. Two agents are negotiating over $H \geq 2$ components where each component has one continuously divisible unit. The protocol is as follows:

Step 1: Number the components from 1 to H and for each component the unit is initially split in equal halves between the two agents. Hence, each agent possesses 0.5 unit of all the components. Let, D^i_t is Agent A_i's current possession of component t. So, initially, $D^i_t = 0.5$, for all $t = 1 \ldots H$, and $i = 1, 2$.

Step 2: The random device chooses one of the two agents, S, to start. Let the set of components yet to be negotiated be G. Initially, $G = R$.

Step 3: Now, S will choose one of the remaining components from G, say C, and propose an exchange offer $E(C, Q_C)$. The offer means that S is ready to trade Q_C units of each of the $G - \{C\}$ components for one unit of C.

Step 4: The other agent, say S', can choose whichever of the options it prefers: either one unit of component C or Q_C units of each of the components in the set $G - \{C\}$.

For S', one unit of C produces $w_{S'C}$ utility and Q_C unit of each of the remaining $G - \{C\}$ components together produces $\sum_{j \in G-\{C\}} Q_C * w_{S'j}$ utility. S' will choose whichever one is more profitable to it. $D^{S'}_C$ and $D^S_C = 1 - D^{S'}_C$ are the amount or unit of component C currently possessed by agent S and S' respectively. If S' prefers C, then it can earn up to $(1 - D^{S'}_C)$ unit of C in exchange of its own $\{(1 - D^{S'}_C) * Q_C\}$ unit of each component in $G - \{C\}$. But if $\{(1 - D^{S'}_C) * Q_C\} > D^{S'}_{a \in G-\{C\}}$, then the entire exchange is not possible[2]. In this situation S' can give $D^{S'}_{a\in G-\{C\}}$ unit of each component in $G - \{C\}$ and earn $\frac{D^{S'}_{a\in G-\{C\}}}{Q_C}$ unit of C in return. On the other hand if it chooses the other option, it will receive $\{(1-D^S_C)*Q_C\}$ unit of each component in $G - \{C\}$ from S and S will receive $(1 - D^S_C)$ unit of C from S' and thereby own the entire unit of C. Similar to the earlier case, if $\{(1 - D^S_C) * Q_C\} > D^S_{a\in G-\{C\}}$, S' will receive $D^S_{a\in G-\{C\}}$ unit of each component in $G - \{C\}$ and give $\frac{D^S_{a\in G-\{C\}}}{Q_C}$ unit of C to S. D^i_a values are updated accordingly to present agent i's current possession of component a.

Step 5: The roles of agents S and S' are swapped and G is changed to $G - \{C\}$. If $|G| > 1$ return to Step 3, Otherwise stop. The final allocation is given by $x_{ij} = D^i_j$, where x_{ij} is the agent i's possession of component j.

In this protocol, agents only alternatively reveal their partial preference information.

[2]Since all the remaining issues are taken together for every exchange, for each agent, its possession of each of the issues yet to be negotiated will be same, i.e., for any agent S, D^S_a is same for all $a \in G - \{C\}$.

Proposition 1: *In POSE protocol rational agents produce an envy-free, Pareto-optimal solution if they have identical ordinal preferences*[3].

3.1.1 Rational strategy of agents in POSE

Now we analyze an agent's rational behavior when using the POSE protocol for negotiation. The negotiation strategy consists of two decisions. First, when an agent is proposing an exchange offer, denoted by E, which component, C, will it choose and what should be corresponding exchange ratio, Q_C. The other decision by an agent involves selecting between two options proposed by the other agent.

Statement 1: *When an agent has to select between two options, it will choose the one with higher utility.*

Proposition 2: *While proposing, an agent* i, *should choose the component* C *such that,* $C = \arg\min_{j\in G}(w_{ij})$ *and should propose a ratio* Q_C, *such that*

$$w_{iC} = \sum_{j\in G-\{C\}} w_{ij} * Q_C.$$

An agent proposes, according to its true preference, its least preferred option from the remaining issues.

3.1.2 An example

Let us illustrate the protocol with an example: two rational agents A_1 and A_2 are negotiating for a resource with three components. One unit of each of the three components are available for negotiation. Suppose ordinal preferences of both the agents are same and their total preference is given by $\{0.1, 0.4, 0.5\}$ and $\{0.2, 0.3, 0.5\}$ respectively. So, $w_{i1} \le w_{i2} \le w_{i3}$ for both the agents $i = 1, 2$. Initially each of the components is divided equally between the two agents, *i.e.,* each agent receives 0.5 unit of each component. Now, suppose, the random device chooses A_2 as the first agent (denoted by S in the POSE protocol). Then A_2 will choose the first issue to propose as this issue has the least relative weight to A_2 and it proposes an offer $E(1, \frac{0.2}{0.3+0.5} = 0.25)$ to A_1. So, A_2 is willing to trade 0.25 units of each of components 2 and 3 for one unit of component 1. For A_1, 0.25 unit of components 2 and 3 together is of higher utility than one unit of component 1. So, it will exchange its possession of 0.5 unit of component 1 for $0.25 * 0.5 = 0.125$ units of components 2 and 3 from A_2. So, after this exchange, the possession of components is $\{0, 0.625, 0.625\}$ for A_1 and $\{1, 0.375, 0.375\}$ for A_2 and the set of issues which are yet to be negotiated is $\{2, 3\}$. Now it is A_1's turn to propose and it chooses component 2, its least preferred issue among the remaining issues and offers $E(2, \frac{0.4}{0.5} = 0.8)$. A_2 prefers 0.8 units of component 3 than one unit of component 2. A_2 has 0.375 unit of component 2. So, it will relinquish 0.375 unit of component 2 and receive $0.375 * 0.8 = 0.3$ unit of component 3 from A_1. So, after this exchange the possessions of A_1 and A_2 are $\{0, 1, 0.325\}$ and $\{1, 0, 0.675\}$, respectively. The utility of A_1 and A_2 are 0.5625 and 0.5375, respectively. Now, the number of issues left to be negotiated is one. So, we stop this process and we claim that this allocation of the units is Pareto optimal as one agent's utility cannot be improved without decreasing the other agent's utility. In Figure 1, the corresponding utility pair is shown by point $P_1 = (0.5625, 0.5375)$.

[3]We exclude the proof of the propositions due to space limitation.

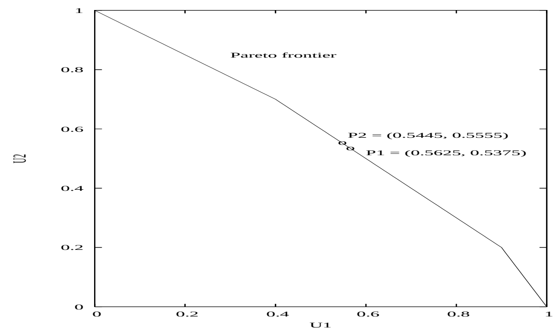

Figure 1: Bargaining solutions when weight of agent 1 is $\{0.1, 0.4, 0.5\}$ and weight of agent 2 is $\{0.2, 0.3, 0.5\}$.

If the random device chooses A_1 as the first agent instead of A_2 as in the above example, a similar process would lead the agents to the outcome where the possessions of A_1 and A_2 are $\{0, 1, 0.289\}$ and $\{1, 0, 0.711\}$, respectively, and the resultant utility pair is $P_2 = (0.5445, 0.5555)$, shown in Figure 1. Both the outcomes, P_1 and P_2, lie on the Pareto-frontier.

4. CONCLUSIONS AND FUTURE WORK

In multi-issue negotiation between two rational agents, there is a trade off between the quality of the agreement reached and the amount of information revealed. We proposed one protocol, where rational agents, without prior knowledge of the opponent preference, are guaranteed to reach a stable, envy-free, Pareto optimal solution if their ordinal preference rankings are identical. For arbitrary agent preferences, this protocol will lead rational agents to Pareto optimal solutions in most cases or otherwise very close to the Pareto optimal solutions while revealing minimal information.

Acknowledgments: This work has been supported in part by an NSF award IIS-0209208 and an NASA EPSCoR RIG.

5. REFERENCES

[1] S. J. Brams and A. D. Taylor. *Fair Division.* Cambridge University Press, 1996.

[2] U. Endriss, N. Maudet, F. Sadri, and F. Toni. On optimal outcomes of negotiations over resources. In *AAMAS-03*, pages 177–184, New York, NY, USA, 2003.

[3] Y. In and R. Serrano. Agenda restrictions in multi-issue bargaining. *Games and Economic Behavior*, 53:385–399, 2004.

[4] N. Jennings, P. Faratin, A. R. L. S. Parsons, C. Sierra, and M. Wooldridge. Automated negotiation: prospects, methods and challenges. *International Journal of Group Decision and Negotiation*, 10(2):199–215, 2001.

[5] H. Raiffa. *The Art and Science of Negotiation.* Harvard University Press, Cambridge, MA, USA, 1982.

[6] T. W. Sandholm. Contract types for satisficing task allocation: I theoretical results, *AAAI Sprint Symposium*, 1998.