Efficient Methods for Multi-agent Multi-issue Negotiation: Allocating Resources[1]

Mengxiao Wu [a] Mathijs de Weerdt [b] Han La Poutré [a]

[a] *Center for Mathematics and Computer Science (CWI), Amsterdam*
[b] *Delft University of Technology, Delft*

1 Introduction

In this paper, we propose an automated *multi-agent multi-issue* negotiation solution to solve a resource allocation problem. We present a *multilateral* negotiation model, by which agents bid *sequentially* in consecutive rounds. Issues are bundled and negotiated concurrently, so win-win opportunities can be generated as trade-offs exist between issues. We develop heuristics of negotiation strategies for three-agent two-issue cases where the agents have *non-linear* utility functions and *incomplete information* about his opponents' preferences, deadlines, etc. The strategies are composed of a Pareto-optimal-search method and concession strategies. An important technical contribution of this work lies in the development of the Pareto-optimal-search method for three-agent multilateral negotiation. Moreover, we present the identification of agreements and Pareto-optimal outcomes achieved by our methods in the mathematical way. Compared to game-theoretic solutions, our heuristic methods are practical and tractable; the whole solution is very efficient such that (near) Pareto-optimal outcomes can be achieved.

2 The Negotiation Model

Suppose three agents $N = \{1, 2, 3\}$ partition two issues (resources) $M = \{1, 2\}$ through negotiation. The range of each issue is normalized to a *continuous* range $[0, 1]$. Each agent $i \in N$ requires a combination of a part of every issue $j \in M$, and only a unanimous agreement can be accepted. The negotiation takes place round by round $n \in \mathbb{N}$ until an agreement is reached or some agent quits. In each round, three agents bid their own desired parts of two issues, $\mathbf{x}_i = (x_{i,1}, x_{i,2})$, sequentially in some (pre-specified) order. If the bid profile $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3)$ forms an agreement, in which the sum of all bids of each issue is no more than the total value 1, agent i will get a utility $u_i(\mathbf{x}, n) = v_i(\mathbf{x}_i)$. We assume the valuation function v_i to be *continuous* and *strictly monotonically increasing* in each of the issues and the utility function u_i is *strictly convex*. All agents get zero utility without agreements. In this model, every agent's preference and negotiation deadline are private information.

3 The Negotiation Strategies

When it is an agent's turn to bid, given his opponents' latest bids, he needs to determine i) a desired utility c_i and ii) one bid $\mathbf{x}_i$ of the utility. Given agent i's utility function, a utility can be represented by an indifference curve, which is a graph showing different combinations of issues, between which the agent is indifferent. Given set C_i of all points on the curve, i.e., $C_i = \{\mathbf{x}_i \mid v_i(\mathbf{x}_i) = c_i\}$, agent i can choose a most satisfying bid which benefits his opponents most and gives himself the same maximum utility. This is the semi-cooperative part of a competitive game to generate win-win opportunities and reach Pareto-optimal outcomes possibly. For the first bid, the agent just chooses one point on the curve of his initial desired utility randomly.

[1]The full version of this paper appeared in: Proceedings of the 12th International Conference on Principles of Practice in Multi-Agent Systems (PRIMA'09), Nagoya, Japan, December 2009.

First, we present our Pareto-optimal-search method, the *orthogonal bidding strategy*, which lets agent i find the most satisfying bid on his current indifference curve. The idea is to make a *reference point* r_i based on the other two agents' latest bids and bid the point in C_i which is closest (measured in the Euclidean distance) to r_i. We propose the notion of reference point $r_i = (r_{i,1}, r_{i,2}) = (1 - \sum_{k \in N-\{i\}} x_{k,1},\ 1 - \sum_{k \in N-\{i\}} x_{k,2})$, because the rest of the issues (represented by r_i) left by his opponents imply the two agents' joint expectation of agent i's partition and bid $\mathbf{x}_i$ closest to r_i can make them most satisfied. Figure 1 left illustrates how agents use the orthogonal bidding strategy to make bids sequentially. In this figure, the bidding order is agent 2 (green), agent 3 (red) and agent 1 (blue); the odd steps determine reference points and the even steps determine bids based on the reference points.

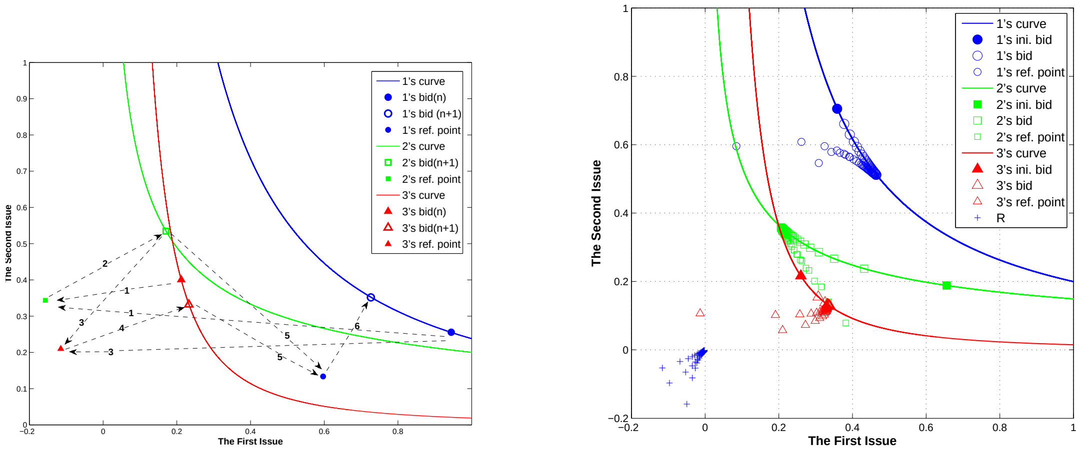

Figure 1: Examples of the Orthogonal Bidding Strategies

Given relatively low utilities, suppose three agents keep bidding (with the orthogonal strategy) on the indifference curves of the utilities without concession and no agent quits. Figure 1 right illustrates the process, in which the bids and reference points move closer round by round. Finally, each agent's bids and reference points are converged into one point on his indifference curve. That means, each agent's final bid completely satisfies his opponents' desires and an agreement is reached. We give the following lemma.

Lemma 1. *A profile of bids $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3)$ is an agreement $\iff$ each reference point r_i Pareto dominates bid $\mathbf{x}_i$, i.e., $r_{i,j} \geq x_{i,j}$, where $i \in N$ and $j \in M$.*

However, an agreement that all issues are exactly partitioned does not necessarily indicate a Pareto-optimal solution in multi-issue negotiation. Agents may still have chances to get Pareto improvements by making trade-offs between issues. We analyse that every combination of two points (bids) on the other two agents' curves introduces a reference point to agent i. If agent i chooses any one of those reference points as his bid where $r_{i,j} > 0$ $(j \in M)$, his opponents can get their desired utilities and agent i can get a utility larger than zero. We call the area composed of all such reference points the reference area of agent i and let X_i denote the set of points in it. For each agent i, once his reference area and his indifference curve have intersections, there are agreements. We give the following lemma.

Lemma 2. *$|X_i \cap C_i| > 1$ where $i \in N$ $\iff$ there is an agreement $\mathbf{x}$ such that $u_i(\mathbf{x}, n) > c_i$.*

Only if the reference area and the indifference curve of every agent i has a unique intersection, the profile of three intersections (bids) is a Pareto-optimal solution. We give the following theorem with a condition that the indifference curves are strictly convex.

Theorem 1. *$|X_i \cap C_i| = 1$ where $i \in N$ $\iff$ there is a Pareto-optimal solution $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3)$ where $u_i(\mathbf{x}, n) = c_i$.*

Second, we study how agents make concession by lowering their desires of utilities, if no agreement can be reached on their current indifference curves. In this work, we develop several concession strategies. Agent i can concede a fixed amount utility, or concede a fixed fraction of current utility c_i, or concede a fixed fraction of the difference between his current utility c_i and the utility introduced by the latest reference point r_i, or concede a fixed fraction of the remaining issues, given all agents' latest bids.