Searching for joint gains in automated negotiations based on multi-criteria decision making theory

Vo, Bao; Padgham, Lin

https://researchrepository.rmit.edu.au/esploro/outputs/9921862917101341/filesAndLinks?institution=61RMIT_INST&index=null

Vo, B., & Padgham, L. (2007). Searching for joint gains in automated negotiations based on multi-criteria decision making theory. *Autonomous Agents And Multiagent Systems*, 508–515.

<https://doi.org/10.1145/1329125.1329321>

Document Version: Published Version

Published Version: <https://doi.org/10.1145/1329125.1329321>

Repository homepage: <https://researchrepository.rmit.edu.au>

© 2007 IFAAMAS

Downloaded On 2023/04/26 22:35:32 +1000

Searching for Joint Gains in Automated Negotiations Based on Multi-criteria Decision Making Theory

Quoc Bao Vo
School of Computer Science and IT
RMIT University, Australia
vqbao@cs.rmit.edu.au

Lin Padgham
School of Computer Science and IT
RMIT University, Australia
linpa@cs.rmit.edu.au

ABSTRACT

It is well established by conflict theorists and others that successful negotiation should incorporate "creating value" as well as "claiming value." Joint improvements that bring benefits to all parties can be realised by (i) identifying attributes that are not of direct conflict between the parties, (ii) tradeoffs on attributes that are valued differently by different parties, and (iii) searching for values within attributes that could bring more gains to one party while not incurring too much loss on the other party. In this paper we propose an approach for maximising joint gains in automated negotiations by formulating the negotiation problem as a multi-criteria decision making problem and taking advantage of several optimisation techniques introduced by operations researchers and conflict theorists. We use a mediator to protect the negotiating parties from unnecessary disclosure of information to their opponent, while also allowing an objective calculation of maximum joint gains. We separate out attributes that take a finite set of values (*simple attributes*) from those with continuous values, and we show that for simple attributes, the mediator can determine the Pareto-optimal values. In addition we show that if none of the simple attributes strongly dominates the other simple attributes, then truth telling is an equilibrium strategy for negotiators during the optimisation of simple attributes. We also describe an approach for improving joint gains on non-simple attributes, by moving the parties in a series of steps, towards the Pareto-optimal frontier.

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Multiagent Systems;
K.4.4 [Computers and Society]: Electronic Commerce

General Terms

Algorithms, Design

Keywords

automated negotiation, integrative negotiation, multi-criteria decision making

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
AAMAS'07 May 14–18 2007, Honolulu, Hawai'i, USA.
Copyright 2007 IFAAMAS .

1. INTRODUCTION

Given that negotiation is perhaps one of the oldest activities in the history of human communication, it's perhaps surprising that conducted experiments on negotiations have shown that negotiators more often than not reach inefficient compromises [1, 21]. Raiffa [17] and Sebenius [20] provide analyses on the negotiators' failure to achieve efficient agreements in practice and their unwillingness to disclose private information due to strategic reasons. According to conflict theorists Lax and Sebenius [13], most negotiation actually involves both integrative and distributive bargaining which they refer to as "creating value" and "claiming value." They argue that negotiation necessarily includes both cooperative and competitive elements, and that these elements exist in tension. Negotiators face a dilemma in deciding whether to pursue a cooperative or a competitive strategy at a particular time during a negotiation. They refer to this problem as *the Negotiator's Dilemma*.

We argue that the Negotiator's Dilemma is essentially information-based, due to the private information held by the agents. Such private information contains both the information that implies the agent's bottom lines (or, her walk-away positions) and the information that enforces her bargaining strength. For instance, when bargaining to sell a house to a potential buyer, the seller would try to hide her actual reserve price as much as possible for she hopes to reach an agreement at a much higher price than her reserve price. On the other hand, the outside options available to her (e.g. other buyers who have expressed genuine interest with fairly good offers) consist in the information that improves her bargaining strength about which she would like to convey to her opponent. But at the same time, her opponent is well aware of the fact that it is her incentive to boost her bargaining strength and thus will not accept every information she sends out unless it is substantiated by evidence.

Coming back to the Negotiator's Dilemma, it's not always possible to separate the integrative bargaining process from the distributive bargaining process. In fact, more often than not, the two processes interplay with each other making information manipulation become part of the integrative bargaining process. This is because a negotiator could use the information about his opponent's interests against her during the distributive negotiation process. That is, a negotiator may refuse to concede on an important conflicting issue by claiming that he has made a major concession (on another issue) to meet his opponent's interests even though the concession he made could be insignificant to him. For instance, few buyers would start a bargaining with a dealer over a deal for a notebook computer by declaring that he is most interested in an extended warranty for the item and therefore prepared to pay a high price to get such an extended warranty.

Negotiation Support Systems (NSSs) and negotiating software

agents (NSAs) have been introduced either to assist humans in making decisions or to enable automated negotiation to allow computer processes to engage in meaningful negotiation to reach agreements (see, for instance, [14, 15, 19, 6, 5]). However, because of the Negotiator's Dilemma and given even bargaining power and incomplete information, the following two undesirable situations often arise: (i) negotiators reach inefficient compromises, or (ii) negotiators engage in a deadlock situation in which both negotiators refuse to act upon with incomplete information and at the same time do not want to disclose more information.

In this paper, we argue for the role of a mediator to resolve the above two issues. The mediator thus plays two roles in a negotiation: (i) to encourage cooperative behaviour among the negotiators, and (ii) to absorb the information disclosure by the negotiators to prevent negotiators from using uncertainty and private information as a strategic device. To take advantage of existing results in negotiation analysis and operations research (OR) literatures [18], we employ multi-criteria decision making (MCDM) theory to allow the negotiation problem to be represented and analysed. Section 2 provides background on MCDM theory and the negotiation framework. Section 3 formulates the problem. In Section 4, we discuss our approach to integrative negotiation. Section 5 discusses the future work with some concluding remarks.

2. BACKGROUND

2.1 Multi-criteria decision making theory

Let A denote the set of feasible alternatives available to a decision maker M . As an act, or decision, a in A may involve multiple aspects, we usually describe the alternatives a with a set of attributes j ; ($j = 1, \dots, m$). (Attributes are also referred to as *issues*, or *decision variables*.) A typical decision maker also has several *objectives* $X_1, \dots, X_k$. We assume that X_i , ($i = 1, \dots, k$), maps the alternatives to real numbers. Thus, a tuple $(x_1, \dots, x_k) = (X_1(a), \dots, X_k(a))$ denotes the *consequence* of the act a to the decision maker M . By definition, objectives are statements that delineate the desires of a decision maker. Thus, M wishes to maximise his objectives. However, as discussed thoroughly by Keeney and Raiffa [8], it is quite likely that a decision maker's objectives will conflict with each other in that the improved achievement with one objective can only be accomplished at the expense of another. For instance, most businesses and public services have objectives like "minimise cost" and "maximise the quality of services." Since better services can often only be attained for a price, these objectives conflict.

Due to the conflicting nature of a decision maker's objectives, M usually has to settle at a compromise solution. That is, he may have to choose an act $a \in A$ that does not optimise every objective. This is the topic of the multi-criteria decision making theory. Part of the solution to this problem is that M has to try to identify the Pareto frontier in the consequence space $\{(X_1(a), \dots, X_k(a))\}_{a \in A}$.

DEFINITION 1. (Dominant)

Let $\mathbf{x} = (x_1, \dots, x_k)$ and $\mathbf{x}' = (x'_1, \dots, x'_k)$ be two consequences. $\mathbf{x}$ dominates $\mathbf{x}'$ iff $x_i > x'_i$ for all i , and the inequality is strict for at least one i .

The Pareto frontier in a consequence space then consists of all consequences that are not dominated by any other consequence. This is illustrated in Fig. 1 in which an alternative consists of two attributes d_1 and d_2 and the decision maker tries to maximise the two objectives X_1 and X_2 . A decision $a \in A$ whose consequence does not lie on the Pareto frontier is inefficient. While the Pareto

Figure 1: The Pareto frontier

frontier allows M to avoid taking inefficient decisions, M still has to decide which of the efficient consequences on the Pareto frontier is most preferred by him.

MCDM theorists introduce a mechanism to allow the objective components of consequences to be normalised to the payoff valuations for the objectives. Consequences can then be ordered: if the gains in satisfaction brought about by C_1 (in comparison to C_2) equals to the losses in satisfaction brought about by C_1 (in comparison to C_2), then the two consequences C_1 and C_2 are considered indifferent. M can now construct the set of *indifference curves*¹ in the consequence space (the dashed curves in Fig. 1). The most preferred indifference curve that intersects with the Pareto frontier is in focus: its intersection with the Pareto frontier is the sought after consequence (i.e., the optimal consequence in Fig. 1).

2.2 A negotiation framework

A multi-agent negotiation framework consists of:

1. A set of two *negotiating agents* $N = \{1, 2\}$.
2. A set of *attributes* $Att = \{\alpha_1, \dots, \alpha_m\}$ characterising the issues the agents are negotiating over. Each attribute α can take a value from the set Val_α ;
3. A set of alternative outcomes $\mathcal{O}$. An outcome $o \in \mathcal{O}$ is represented by an assignment of values to the corresponding attributes in Att .
4. *Agents' utility*: Based on the theory of multiple-criteria decision making [8], we define the agents' utility as follows:
 - *Objectives*: Agent i has a set of n_i objectives, or interests; denoted by j ($j = 1, \dots, n_i$). To measure how much an outcome o fulfills an objective j to an agent i , we use objective functions: for each agent i , we define i 's interests using the objective vector function $\mathbf{f}_i = [f_{ij}] : \mathcal{O} \rightarrow \mathbb{R}^{n_i}$.
 - *Value functions*: Instead of directly evaluating an outcome o , agent i looks at how much his objectives are fulfilled and will make a valuation based on these more basic criteria. Thus, for each agent i , there is a value function $\sigma_i : \mathbb{R}^{n_i} \rightarrow \mathbb{R}$. In particular, Raiffa [17] shows how to systematically construct an additive value function to each party involved in a negotiation.
 - *Utility*: Now, given an outcome $o \in \mathcal{O}$, an agent i is able to determine its value, i.e., $\sigma_i(\mathbf{f}_i(o))$. However, a negotiation infrastructure is usually required to facilitate negotiation. This might involve other mechanisms and factors/parties, e.g., a mediator, a legal institution, participation fees, etc. The standard way to implement such a thing is to allow money

¹In fact, given the k -dimensional space, these should be called *indifference surfaces*. However, we will not bog down to that level of details.

and side-payments. In this paper, we ignore those side-effects and assume that agent i 's utility function u_i is normalised so that $u_i : \mathcal{O} \rightarrow [0, 1]$.

EXAMPLE 1. There are two agents, A and B . Agent A has a task T that needs to be done and also 100 units of a resource R . Agent B has the capacity to perform task T and would like to obtain at least 10 and at most 20 units of the resource R . Agent B is indifferent on any amount between 10 and 20 units of the resource R . The objective functions for both agents A and B are *cost* and *revenue*. And they both aim at minimising costs while maximising revenues. Having T done generates for A a revenue $r_{A,T}$ while doing T incurs a cost $c_{B,T}$ to B . Agent B obtains a revenue $r_{B,R}$ for each unit of the resource R while providing each unit of the resource R costs agent A $c_{A,R}$.

Assuming that money transfer between agents is possible, the set Att then contains three attributes:

- T , taking values from the set $\{0, 1\}$, indicates whether the task T is assigned to agent B ;
- R , taking values from the set of non-negative integer, indicates the amount of resource R being allocated to agent B ; and
- MT , taking values from $\mathbb{R}$, indicates the payment p to be transferred from A to B .

Consider the outcome $o = [T = 1, R = k, MT = p]$, i.e., the task T is assigned to B , and A allocates to B with k units of the resource R , and A transfers p dollars to B . Then, $cost_A(o) = k \cdot c_{A,R} + p$ and $rev_A(o) = r_{A,T}$; and $cost_B(o) = c_{B,T}$ and $rev_B(o) = \begin{cases} k \cdot r_{B,R} + p & \text{if } 10 \leq k \leq 20 \\ p & \text{otherwise.} \end{cases}$

And, $\sigma_i(cost_i(o), rev_i(o)) = rev_i(o) - cost_i(o)$, ($i = A, B$).

3. PROBLEM FORMALISATION

Consider Example 1, assume that $r_{A,T} = \$150$ and $c_{B,T} = \$100$ and $r_{B,R} = \$10$ and $c_{A,R} = \$7$. That is, the revenues generated for A exceeds the costs incurred to B to do task T , and B values resource R more highly than the cost for A to provide it. The optimal solution to this problem scenario is to assign task T to agent B and to allocate 20 units of resource R (i.e., the maximal amount of resource R required by agent B) from agent A to agent B . This outcome regarding the resource and task allocation problems leaves payoffs of \$10 to agent A and \$100 to agent B .² Any other outcome would leave at least one of the agents worse off. In other words, the presented outcome is Pareto-efficient and should be part of the solution outcome for this problem scenario.

However, as the agents still have to bargain over the amount of money transfer p , neither agent would be willing to disclose their respective costs and revenues regarding the task T and the resource R . As a consequence, agents often do not achieve the optimal outcome presented above in practice. To address this issue, we introduce a mediator to help the agents discover better agreements than the ones they might try to settle on. Note that this problem is essentially the problem of searching for joint gains in a multilateral negotiation in which the involved parties hold strategic information, i.e., the integrative part in a negotiation. In order to help facilitate this process, we introduce the role of a neutral mediator. Before formalising the decision problems faced by the mediator and the

²Certainly, without money transfer to compensate agent A , this outcome is not a fair one.

negotiating agents, we discuss the properties of the solution outcomes to be achieved by the mediator. In a negotiation setting, the two typical design goals would be:

- *Efficiency*: Avoid the agents from settling on an outcome that is not Pareto-optimal; and
- *Fairness*: Avoid agreements that give the most of the gains to a subset of agents while leaving the rest with too little.

The above goals are axiomatised in Nash's seminal work [16] on cooperative negotiation games. Essentially, Nash advocates for the following properties to be satisfied by solution to the bilateral negotiation problem: (i) it produces only Pareto-optimal outcomes; (ii) it is invariant to affine transformation (to the consequence space); (iii) it is symmetric; and (iv) it is independent from irrelevant alternatives. A solution satisfying Nash's axioms is called a *Nash bargaining solution*.

It then turns out that, by taking the negotiators' utilities as its objectives the mediator itself faces a multi-criteria decision making problem. The issues faced by the mediator are: (i) the mediator requires access to the negotiators' utility functions, and (ii) making (fair) tradeoffs between different agents' utilities. Our methods allow the agents to repeatedly interact with the mediator so that a Nash solution outcome could be found by the parties.

Informally, the problem faced by both the mediator and the negotiators is construction of the indifference curves. Why are the indifference curves so important?

- To the negotiators, knowing the options available along indifference curves opens up opportunities to reach more efficient outcomes. For instance, consider an agent A who is presenting his opponent with an offer θ_A which she refuses to accept. Rather than having to concede, A could look at his indifference curve going through θ_A and choose another proposal θ'_A . To him, θ_A and θ'_A are indifferent but θ'_A could give some gains to B and thus will be more acceptable to B . In other words, the outcome θ'_A is more efficient than θ_A to these two negotiators.
- To the mediator, constructing indifference curves requires a measure of fairness between the negotiators. The mediator needs to determine how much utility it needs to take away from the other negotiators to give a particular negotiator a specific gain G (in utility).

In order to search for integrative solutions within the outcome space $\mathcal{O}$, we characterise the relationship between the agents over the set of attributes Att . As the agents hold different objectives and have different capacities, it may be the case that changing between two values of a specific attribute implies different shifts in utility of the agents. However, the problem of finding the exact Pareto-optimal set³ is NP-hard [2].

Our approach is thus to solve this optimisation problem in two steps. In the first steps, the more manageable attributes will be solved. These are attributes that take a finite set of values. The result of this step would be a subset of outcomes that contains the Pareto-optimal set. In the second step, we employ an iterative procedure that allows the mediator to interact with the negotiators to find joint improvements that move towards a Pareto-optimal outcome. This approach will not work unless the attributes from Att

³The Pareto-optimal set is the set of outcomes whose consequences (in the consequence space) correspond to the Pareto frontier.

are independent. Most works on multi-attribute, or multi-issue, negotiation (e.g. [17]) assume that the attributes or the issues are independent, resulting in an additive value function for each agent.⁴

ASSUMPTION 1. Let $i \in N$ and $S \subseteq \text{Att}$. Denote by $\bar{S}$ the set $\text{Att} \setminus S$. Assume that v_S and v'_S are two assignments of values to the attributes of S and $v_{\bar{S}}^1, v_{\bar{S}}^2$ are two arbitrary value assignments to the attributes of $\bar{S}$, then $(u_i([v_S, v_{\bar{S}}^1]) - u_i([v'_S, v_{\bar{S}}^1])) = (u_i([v_S, v_{\bar{S}}^2]) - u_i([v'_S, v_{\bar{S}}^2]))$. That is, the utility function of agent i will be defined on the attributes from S independently of any value assignment to other attributes.

4. MEDIATOR-BASED BILATERAL NEGOTIATIONS

As discussed by Lax and Sebenius [13], under incomplete information the tension between creating and claiming values is the primary cause of inefficient outcomes. This can be seen most easily in negotiations involving two negotiators; during the distributive phase of the negotiation, the two negotiators's objectives are directly opposing each other. We will now formally characterise this relationship between negotiators by defining the *opposition* between two negotiating parties. The following exposition will be mainly reproduced from [9].

Assuming for the moment that all attributes from Att take values from the set of real numbers $\mathbb{R}$, i.e., $\text{Val}_j \subseteq \mathbb{R}$ for all $j \in \text{Att}$. We further assume that the set $\mathcal{O} = \times_{j \in \text{Att}} \text{Val}_j$ of feasible outcomes is defined by constraints that all parties must obey and $\mathcal{O}$ is convex. Now, an outcome $o \in \mathcal{O}$ is just a point in the m -dimensional space of real numbers. Then, the questions are: (i) from the point of view of an agent i , is o already the best outcome for i ? (ii) if o is not the best outcome for i then is there another outcome o' such that o' gives i a better utility than o and o' does not cause a utility loss to the other agent j in comparison to o ?

The above questions can be answered by looking at the directions of improvement of the negotiating parties at o , i.e., the directions in the outcome space $\mathcal{O}$ into which their utilities increase at point o . Under the assumption that the parties' utility functions u_i are *differentiable concave*, the set of all directions of improvement for a party at a point o can be defined in terms of his *most preferred*, or *gradient*, *direction* at that point. When the gradient direction $\nabla u_i(o)$ of agent i at point o is outright opposing to the gradient direction $\nabla u_j(o)$ of agent j at point o then the two parties strongly disagree at o and no joint improvements can be achieved for i and j in the locality surrounding o .

Since opposition between the two parties can vary considerably over the outcome space (with one pair of outcomes considered highly antagonistic and another pair being highly cooperative), we need to describe the local properties of the relationship. We begin with the opposition at any point of the outcome space $\mathbb{R}^m$. The following definition is reproduced from [9]:

- DEFINITION 2.**
1. The parties are in *local strict opposition* at a point $\mathbf{x} \in \mathbb{R}^m$ iff for all points $\mathbf{x}' \in \mathbb{R}^m$ that are sufficiently close to $\mathbf{x}$ (i.e., for some $\epsilon > 0$ such that $\forall \mathbf{x}' \parallel \mathbf{x}' - \mathbf{x} \parallel < \epsilon$), an increase of one utility can be achieved only at the expense of a decrease of the other utility.
 2. The parties are in *local non-strict opposition* at a point $\mathbf{x} \in \mathbb{R}^m$ iff they are not in *local strict opposition* at $\mathbf{x}$, i.e., iff it is possible for both parties to raise their utilities by moving an infinitesimal distance from $\mathbf{x}$.

⁴Klein et al. [10] explore several implications of complex contracts in which attributes are possibly inter-dependent.

3. The parties are in *local weak opposition* at a point $\mathbf{x} \in \mathbb{R}^m$ iff $\nabla u_1(\mathbf{x}) \cdot \nabla u_2(\mathbf{x}) \geq 0$, i.e., iff the gradients at $\mathbf{x}$ of the two utility functions form an acute or right angle.
4. The parties are in *local strong opposition* at a point $\mathbf{x} \in \mathbb{R}^m$ iff $\nabla u_1(\mathbf{x}) \cdot \nabla u_2(\mathbf{x}) < 0$, i.e., iff the gradients at $\mathbf{x}$ form an obtuse angle.
5. The parties are in *global strict* (nonstrict, weak, strong) opposition iff for every $\mathbf{x} \in \mathbb{R}^m$ they are in *local strict* (nonstrict, weak, strong) opposition.

Global strict and nonstrict oppositions are complementary cases. Essentially, under global strict opposition the whole outcome space $\mathcal{O}$ becomes the Pareto-optimal set as at no point in $\mathcal{O}$ can the negotiating parties make a joint improvement, i.e., every point in $\mathcal{O}$ is a Pareto-efficient outcome. In other words, under global strict opposition the outcome space $\mathcal{O}$ can be flattened out into a single line such that for each pair of outcomes $\mathbf{x}, \mathbf{y} \in \mathcal{O}$, $u_1(\mathbf{x}) < u_1(\mathbf{y})$ iff $u_2(\mathbf{x}) > u_2(\mathbf{y})$, i.e., at every point in $\mathcal{O}$, the gradient of the two utility functions point to two different ends of the line.

Intuitively, global strict opposition implies that there is no way to obtain joint improvements for both agents. As a consequence, the negotiation degenerates to a distributive negotiation, i.e., the negotiating parties should try to claim as much shares from the negotiation issues as possible while the mediator should aim for the fairness of the division. On the other hand, global nonstrict opposition allows room for joint improvements and all parties might be better off trying to realise the potential gains by reaching Pareto-efficient agreements. Weak and strong oppositions indicate different levels of opposition. The weaker the opposition, the more potential gains can be realised making cooperation the better strategy to employ during negotiation. On the other hand, stronger opposition suggests that the negotiating parties tend to behave strategically leading to misrepresentation of their respective objectives and utility functions and making joint gains more difficult to realise.

We have been temporarily making the assumption that the outcome space $\mathcal{O}$ is the subset of $\mathbb{R}^m$. In many real-world negotiations, this assumption would be too restrictive. We will continue our exposition by lifting this restriction and allowing discrete attributes. However, as most negotiations involve only discrete issues with a bounded number of options, we will assume that each attribute takes values either from a finite set or from the set of real numbers $\mathbb{R}$. In the rest of the paper, we will refer to attributes whose values are from finite sets as *simple attributes* and attributes whose values are from $\mathbb{R}$ as *continuous attributes*. The notions of local oppositions, i.e., strict, nonstrict, weak and strong, are not applicable to outcome spaces that contain simple attributes and nor are the notions of global weak and strong oppositions. However, the notions of global strict and nonstrict oppositions can be generalised for outcome spaces that contain simple attributes.

DEFINITION 3. Given an outcome space $\mathcal{O}$, the parties are in *global strict opposition* iff $\forall \mathbf{x}, \mathbf{y} \in \mathcal{O}$, $u_1(\mathbf{x}) < u_1(\mathbf{y})$ iff $u_2(\mathbf{x}) > u_2(\mathbf{y})$.

The parties are in *global nonstrict opposition* if they are not in *global strict opposition*.

4.1 Optimisation on simple attributes

In order to extract the optimal values for a subset of attributes, in the first step of this optimisation process the mediator requests the negotiators to submit their respective utility functions over the set of simple attributes. Let $\text{Simp} \subseteq \text{Att}$ denote the set of all simple attributes from Att . Note that, due to Assumption 1, agent i 's

utility function can be characterised as follows:

$$u_i([\mathbf{v}_{Simp}, \mathbf{v}_{\overline{Simp}}]) = w_1^i * u_{i,1}([\mathbf{v}_{Simp}]) + w_2^i * u_{i,2}([\mathbf{v}_{\overline{Simp}}]),$$

where $\overline{Simp} = Att \setminus Simp$, and $u_{i,1}$ and $u_{i,2}$ are the utility components of u_i over the sets of attributes $Simp$ and $\overline{Simp}$, respectively, and $0 < w_1^i, w_2^i < 1$ and $w_1^i + w_2^i = 1$.

As attributes are independent of each other regarding the agents' utility functions, the optimisation problem over the attributes from $Simp$ can be carried out by fixing $u_i([\mathbf{v}_{\overline{Simp}}])$ to a constant C , and then search for the optimal values within the set of attributes $Simp$. Now, how does the mediator determine the optimal values for the attributes in $Simp$? Several well-known optimisation strategies could be applicable here:

- The *utilitarian solution*: The sum of the agents' utilities are maximised. Thus, the optimal values are the solution of the following optimisation problem:

$$\arg \max_{\mathbf{v} \in Val_{Simp}} \sum_{i \in N} u_i(\mathbf{v})$$

- The *Nash solution*: The product of the agents' utilities are maximised. Thus, the optimal values are the solution of the following optimisation problem:

$$\arg \max_{\mathbf{v} \in Val_{Simp}} \prod_{i \in N} u_i(\mathbf{v})$$

- The *egalitarian solution* (aka. the maximin solution): The utility of the agent with minimum utility is maximised. Thus, the optimal values are the solution of the following optimisation problem:

$$\arg \max_{\mathbf{v} \in Val_{Simp}} \min_{i \in N} u_i(\mathbf{v})$$

The question now is of course whether a negotiator has the incentive to misrepresent his utility function. First of all, recall that the agents' utility functions are bounded, i.e., $\forall o \in \mathcal{O}. 0 \leq u_i(o) \leq 1$. Thus, the agents have no incentive to overstate their utility regarding an outcome o : If o is the most preferred outcome to an agent i then he already assigns the maximal utility to o . On the other hand, if o is not the most preferred outcome to i then by overstating the utility he assigns to o , the agent i runs the risk of having to settle on an agreement which would give him less payoffs than he is supposed to receive. However, agents do have an incentive to understate their utility if the final settlement will be based on the above solutions alone. Essentially, the mechanism to avoid an agent to understate his utility regarding particular outcomes is to guarantee a certain measure of fairness for the final settlement. That is, the agents lose the incentive to be dishonest to obtain gains from taking advantage of the known solutions to determine the settlement outcome for they would be offset by the fairness maintenance mechanism. Firsts, we state an easy lemma.

LEMMA 1. *When $Simp$ contains one single attributes, the agents have the incentive to understate their utility functions regarding outcomes that are not attractive to them.*

By way of illustration, consider the set $Simp$ containing only one attribute that could take values from the finite set $\{A, B, C, D\}$. Assume that negotiator 1 assigns utilities of 0.4, 0.7, 0.9, and 1 to A, B, C , and D , respectively. Assume also that negotiator 2 assigns utilities of 1, 0.9, 0.7, and 0.4 to A, B, C , and D , respectively. If agent 1 misrepresents his utility function to the mediator by reporting utility 0 for all values A, B and C and utility 1 for

value D then the agent 2 who plays honestly in his report to the mediator will obtain the worst outcome D given any of the above solutions. Note that agent 1 doesn't need to know agent 2's utility function, nor does he need to know the strategy employed by agent 2. As long as he knows that the mediator is going to employ one of the above three solutions, then the above misrepresentation is the dominant strategy for this game.

However, when the set $Simp$ contains more than one attribute and none of the attributes strongly dominate the other attributes then the above problem diminishes by itself thanks to the integrative solution. We of course have to define clearly what it means for an attribute to strongly dominate other attributes. Intuitively, if most of an agent's utility concentrates on one of the attributes then this attribute strongly dominates other attributes. We again appeal to the Assumption 1 on additivity of utility functions to achieve a measure of fairness within this negotiation setting. Due to Assumption 1, we can characterise agent i 's utility component over the set of attributes $Simp$ by the following equation:

$$u_{i,1}([\mathbf{v}_{Simp}]) = \sum_{j \in Simp} w_j^i * u_{i,j}([\mathbf{v}_j]) \quad (1)$$

where $\sum_{j \in Simp} w_j = 1$.

Then, an attribute $\ell \in Simp$ *strongly dominates* the rest of the attributes in $Simp$ (for agent i) iff $w_\ell^i > \sum_{j \in (Simp - \ell)} w_j^i$. Attribute ℓ is said to be *strongly dominant* (for agent i) wrt. the set of simple attributes $Simp$.

The following theorem shows that if the set of attributes $Simp$ does not contain a strongly dominant attribute then the negotiators have no incentive to be dishonest.

THEOREM 1. *Given a negotiation framework, if for every agent the set of simple attributes doesn't contain a strongly dominant attribute, then truth-telling is an equilibrium strategy for the negotiators during the optimisation of simple attributes.*

So far, we have been concentrating on the efficiency issue while leaving the fairness issue aside. A fair framework does not only support a more satisfactory distribution of utility among the agents, but also often a good measure to prevent misrepresentation of private information by the agents. Of the three solutions presented above, the utilitarian solution does not support fairness. On the other hand, Nash [16] proves that the Nash solution satisfies the above four axioms for the cooperative bargaining games and is considered a fair solution. The egalitarian solution is another mechanism to achieve fairness by essentially helping the worst off. The problem with these solutions, as discussed earlier, is that they are vulnerable to strategic behaviours when one of the attributes strongly dominates the rest of attributes.

However, there is yet another solution that aims to guarantee fairness, the *minimax solution*. That is, the utility of the agent with maximum utility is minimised. It's obvious that the minimax solution produces inefficient outcomes. However, to get around this problem (given that the Pareto-optimal set can be tractably computed), we can apply this solution over the Pareto-optimal set only. Let $POSet \subseteq Val_{Simp}$ be the Pareto-optimal subset of the simple outcomes, the *minimax solution* is defined to be the solution of the following optimisation problem.

$$\arg \min_{\mathbf{v} \in POSet} \max_{i \in N} u_i(\mathbf{v})$$

While overall efficiency often suffers under a minimax solution, i.e., the sum of all agents' utilities are often lower than under other solutions, it can be shown that the minimax solution is less vulnerable to manipulation.

THEOREM 2. *Given a negotiation framework, under the min-max solution, if the negotiators are uncertain about their opponents' preferences then truth-telling is an equilibrium strategy for the negotiators during the optimisation of simple attributes.*

That is, even when there is only one single simple attribute, if an agent is uncertain whether the other agent's most preferred resolution is also his own most preferred resolution then he should opt for truth-telling as the optimal strategy.

4.2 Optimisation on continuous attributes

When the attributes take values from infinite sets, we assume that they are continuous. This is similar to the common practice in operations research in which linear programming solutions/techniques are applied to integer programming problems.

We denote the number of continuous attributes by k , i.e., $Att = Simp \cup \overline{Simp}$ and $|\overline{Simp}| = k$. Then, the outcome space $\mathcal{O}$ can be represented as follows: $\mathcal{O} = (\prod_{j \in Simp} Val_j) \times (\prod_{l \in \overline{Simp}} Val_l)$, where $\prod_{l \in \overline{Simp}} Val_l \subseteq \mathbb{R}^k$ is the continuous component of $\mathcal{O}$. Let $\mathcal{O}^c$ denote the set $\prod_{l \in \overline{Simp}} Val_l$. We'll refer to $\mathcal{O}^c$ as the *feasible set* and assume that $\mathcal{O}^c$ is *closed* and *convex*. After carrying out the optimisation over the set of simple attributes, we are able to assign the optimal values to the simple attributes from $Simp$. Thus, we reduce the original problem to the problem of searching for optimal (and fair) outcomes within the feasible set $\mathcal{O}^c$. Recall that, by Assumption 1, we can characterise agent i 's utility function as follows:

$$u_i([\mathbf{v}_{Simp}^*, \mathbf{v}_{\overline{Simp}}]) = C + w_2^i * u_{i,2}([\mathbf{v}_{\overline{Simp}}]),$$

where C is the constant $w_1^i * u_{i,1}([\mathbf{v}_{Simp}^*])$ and $\mathbf{v}_{Simp}^*$ denotes the optimal values of the simple attributes in $Simp$. Hence, without loss of generality (albeit with a blatant abuse of notation), we can take the agent i 's utility function as $u_i : \mathbb{R}^k \rightarrow [0, 1]$. Accordingly we will also take the set of outcomes under consideration by the agents to be the feasible set $\mathcal{O}^c$. We now state another assumption to be used in this section:

ASSUMPTION 2. *The negotiators' utility functions can be described by continuously differentiable and concave functions $u_i : \mathbb{R}^k \rightarrow [0, 1]$, ($i = 1, 2$).*

It should be emphasised that we do not assume that agents explicitly know their utility functions. For the method to be described in the following to work, we only assume that the agents know the relevant information, e.g. at certain point within the feasible set $\mathcal{O}^c$, the gradient direction of their own utility functions and some section of their respective indifference curves. Assume that a tentative agreement (which is a point $\mathbf{x} \in \mathbb{R}^k$) is currently on the table, the process for the agents to jointly improve this agreement in order to reach a Pareto-optimal agreement can be described as follows. The mediator asks the negotiators to discretely submit their respective gradient directions at $\mathbf{x}$, i.e., $\nabla u_1(\mathbf{x})$ and $\nabla u_2(\mathbf{x})$.

Note that the goal of the process to be described here is to search for agreements that are more efficient than the tentative agreement currently on the table. That is, we are searching for points $\mathbf{x}'$ within the feasible set $\mathcal{O}^c$ such that moving to $\mathbf{x}'$ from the current tentative agreement $\mathbf{x}$ brings more gains to at least one of the agents while not hurting any of the agents. Due to the assumption made above, i.e. the feasible set $\mathcal{O}^c$ is bounded, the conditions for an alternative $\mathbf{x} \in \mathcal{O}^c$ to be efficient vary depending on the position of $\mathbf{x}$. The following results are proved in [9]:

Let $B(\mathbf{x}) = 0$ denote the equation of the boundary of $\mathcal{O}^c$, defining $\mathbf{x} \in \mathcal{O}^c$ iff $B(\mathbf{x}) \geq 0$. An alternative $\mathbf{x}^* \in \mathcal{O}^c$ is efficient iff,

either

A. $\mathbf{x}^*$ is in the interior of $\mathcal{O}^c$ and the parties are in local strict opposition at $\mathbf{x}^*$, i.e.,

$$\nabla u_1(\mathbf{x}^*) = -\gamma \nabla u_2(\mathbf{x}^*) \quad (2)$$

where $\gamma > 0$; or

B. $\mathbf{x}^*$ is on the boundary of $\mathcal{O}^c$, and for some $\alpha, \beta \geq 0$:

$$\alpha \nabla u_1(\mathbf{x}^*) + \beta \nabla u_2(\mathbf{x}^*) = \nabla B(\mathbf{x}^*) \quad (3)$$

We are now interested in answering the following questions:

- (i) What is the initial tentative agreement $\mathbf{x}_0$?
- (ii) How to find the more efficient agreement $\mathbf{x}_{h+1}$, given the current tentative agreement $\mathbf{x}_h$?

4.2.1 Determining a fair initial tentative agreement

It should be emphasised that the choice of the initial tentative agreement affects the fairness of the final agreement to be reached by the presented method. For instance, if the initial tentative agreement $\mathbf{x}_0$ is chosen to be the most preferred alternative to one of the agents then it is also a Pareto-optimal outcome, making it impossible to find any joint improvement from $\mathbf{x}_0$. However, if $\mathbf{x}_0$ will then be chosen to be the final settlement and if $\mathbf{x}_0$ turns out to be the worst alternative to the other agent then this outcome is a very unfair one. Thus, it's important that the choice of the initial tentative agreement be sensibly made.

Ehtamo et al [3] present several methods to choose the initial tentative agreement (called *reference point* in their paper). However, their goal is to approximate the Pareto-optimal set by systematically choosing a set of reference points. Once an (approximate) Pareto-optimal set is generated, it is left to the negotiators to decide which of the generated Pareto-optimal outcomes to be chosen as the final settlement. That is, distributive negotiation will then be required to settle the issue.

We, on the other hand, are interested in a fair initial tentative agreement which is not necessarily efficient. Improving a given tentative agreement to yield a Pareto-optimal agreement is considered in the next section. For each attribute $j \in \overline{Simp}$, an agent i will be asked to discretely submit three values (from the set Val_j): the most preferred value, denoted by $pv_{i,j}$, the least preferred value, denoted by $wv_{i,j}$, and a value that gives i an approximately average payoff, denoted by $av_{i,j}$. (Note that this is possible because the set Val_j is bounded.) If $pv_{1,j}$ and $pv_{2,j}$ are sufficiently close, i.e., $|pv_{1,j} - pv_{2,j}| < \Delta$ for some pre-defined $\Delta > 0$, then $pv_{1,j}$ and $pv_{2,j}$ are chosen to be the two "core" values, denoted by cv_1 and cv_2 . Otherwise, between the two values $pv_{1,j}$ and $av_{1,j}$, we eliminate the one that is closer to $wv_{2,j}$, the remaining value is denoted by cv_1 . Similarly, we obtain cv_2 from the two values $pv_{2,j}$ and $av_{2,j}$. If $cv_1 = cv_2$ then cv_1 is selected as the initial value for the attribute j as part of the initial tentative agreement. Otherwise, without loss of generality, we assume that $cv_1 < cv_2$. The mediator selects randomly p values $mv_1, \dots, mv_p$ from the open interval (cv_1, cv_2) , where $p \geq 1$. The mediator then asks the agents to submit their valuations over the set of values $\{cv_1, cv_2, mv_1, \dots, mv_p\}$. The value whose the two valuations of two agents are closest is selected as the initial value for the attribute j as part of the initial tentative agreement.

The above procedure guarantees that the agents do not gain by behaving strategically. By performing the above procedure on every attribute $j \in \overline{Simp}$, we are able to identify the initial tentative agreement $\mathbf{x}_0$ such that $\mathbf{x}_0 \in \mathcal{O}^c$. The next step is to compute a new tentative agreement from an existing tentative agreement so that the new one would be more efficient than the existing one.

4.2.2 Computing new tentative agreement

Our procedure is a combination of the method of jointly improving direction introduced by Ehtamo et al [4] and a method we propose in the coming section. Basically, the idea is to see how strong the opposition the parties are in. If the two parties are in (local) weak opposition at the current tentative agreement $\mathbf{x}_h$, i.e., their improving directions at $\mathbf{x}_h$ are close to each other, then the compromise direction proposed by Ehtamo et al [4] is likely to point to a better agreement for both agents. However, if the two parties are in local strong opposition at the current point $\mathbf{x}_h$ then it's unclear whether the compromise direction would really not hurt one of the agents whilst bringing some benefit to the other.

We will first review the method proposed by Ehtamo et al [4] to compute the compromise direction for a group of negotiators at a given point $\mathbf{x} \in \mathcal{O}^c$. Ehtamo et al define a function $T(\mathbf{x})$ that describes the mediator's choice for a compromise direction at $\mathbf{x}$. For the case of two-party negotiations, the following *bisecting function*, denoted by T^{BS} , can be defined over the interior set of $\mathcal{O}^c$. Note that the closed set $\mathcal{O}^c$ contains two disjoint subsets: $\mathcal{O}^c = \mathcal{O}_0^c \cup \mathcal{O}_B^c$, where $\mathcal{O}_0^c$ denotes the set of interior points of $\mathcal{O}^c$ and $\mathcal{O}_B^c$ denotes the boundary of $\mathcal{O}^c$. The *bisecting compromise* is defined by a function $T^{\text{BS}} : \mathcal{O}_0^c \rightarrow \mathbb{R}^2$,

$$T^{\text{BS}}(\mathbf{x}) = \frac{\nabla u_1(\mathbf{x})}{\|\nabla u_1(\mathbf{x})\|} + \frac{\nabla u_2(\mathbf{x})}{\|\nabla u_2(\mathbf{x})\|}, \quad \mathbf{x} \in \mathcal{O}_0^c. \quad (4)$$

Given the current tentative agreement $\mathbf{x}_h$ ($h \geq 0$), the mediator has to choose a point $\mathbf{x}_{h+1}$ along $d = T(\mathbf{x}_h)$ so that all parties gain. Ehtamo et al then define a mechanism to generate a sequence of points and prove that when the generated sequence is bounded and when all generated points (from the sequence) belong to the interior set $\mathcal{O}_0^c$ then the sequence converges to a weakly Pareto-optimal agreement [4, pp. 59–60].⁵

As the above mechanism does not work at the boundary points of $\mathcal{O}^c$, we will introduce a procedure that works everywhere in an alternative space $\mathcal{O}^c$. Let $\mathbf{x} \in \mathcal{O}^c$ and let $\theta(\mathbf{x})$ denote the angle between the gradients $\nabla u_1(\mathbf{x})$ and $\nabla u_2(\mathbf{x})$ at $\mathbf{x}$. That is,

$$\theta(\mathbf{x}) = \arccos\left(\frac{\nabla u_1(\mathbf{x}) \cdot \nabla u_2(\mathbf{x})}{\|\nabla u_1(\mathbf{x})\| \cdot \|\nabla u_2(\mathbf{x})\|}\right)$$

From Definition 2, it is obvious that the two parties are in local strict opposition (at $\mathbf{x}$) iff $\theta(\mathbf{x}) = \pi$, and they are in local weak opposition iff $\pi/2 \geq \theta(\mathbf{x}) > \pi/2$, and they are in local strong opposition iff $\pi/2 \geq \theta(\mathbf{x}) \geq 0$. Note also that the two vectors $\nabla u_1(\mathbf{x})$ and $\nabla u_2(\mathbf{x})$ define a hyperplane, denoted by $h\nabla(\mathbf{x})$, in the k -dimensional space $\mathbb{R}^k$. Furthermore, there are two indifference curves of agents 1 and 2 going through point $\mathbf{x}$, denoted by $IC_1(\mathbf{x})$ and $IC_2(\mathbf{x})$, respectively. Let $hT_1(\mathbf{x})$ and $hT_2(\mathbf{x})$ denote the tangent hyperplanes to the indifference curves $IC_1(\mathbf{x})$ and $IC_2(\mathbf{x})$, respectively, at point $\mathbf{x}$. The planes $hT_1(\mathbf{x})$ and $hT_2(\mathbf{x})$ intersect $h\nabla(\mathbf{x})$ in the lines $IS_1(\mathbf{x})$ and $IS_2(\mathbf{x})$, respectively. Note that given a line $L(\mathbf{x})$ going through the point $\mathbf{x}$, there are two (unit) vectors from $\mathbf{x}$ along $L(\mathbf{x})$ pointing to two opposite directions, denoted by $L^+(\mathbf{x})$ and $L^-(\mathbf{x})$.

We can now informally explain our solution to the problem of searching for joint gains. When it isn't possible to obtain a compromise direction for joint improvements at a point $\mathbf{x} \in \mathcal{O}^c$ either because the compromise vector points to the space outside of the feasible set $\mathcal{O}^c$ or because the two parties are in local strong opposition at $\mathbf{x}$, we will consider to move along the indifference curve of one party while trying to improve the utility of the other party. As

⁵Let S be the set of alternatives, $\mathbf{x}^*$ is weakly Pareto optimal if there is no $\mathbf{x} \in S$ such that $u_i(\mathbf{x}) > u_i(\mathbf{x}^*)$ for all agents i .

the mediator does not know the indifference curves of the parties, he has to use the tangent hyperplanes to the indifference curves of the parties at point $\mathbf{x}$. Note that the tangent hyperplane to a curve is a useful approximation of the curve in the immediate vicinity of the point of tangency, $\mathbf{x}$.

We are now describing an iteration step to reach the next tentative agreement $\mathbf{x}_{h+1}$ from the current tentative agreement $\mathbf{x}_h \in \mathcal{O}^c$. A vector $\mathbf{v}$ whose tail is $\mathbf{x}_h$ is said to be *bounded* in $\mathcal{O}^c$ if $\exists \lambda > 0$ such that $\mathbf{x}_h + \lambda \mathbf{v} \in \mathcal{O}^c$. To start, the mediator asks the negotiators for their gradients $\nabla u_1(\mathbf{x}_h)$ and $\nabla u_2(\mathbf{x}_h)$, respectively, at $\mathbf{x}_h$.

1. If $\mathbf{x}_h$ is a Pareto-optimal outcome according to equation 2 or equation 3, then the process is terminated.
2. If $1 \geq \nabla u_1(\mathbf{x}_h) \cdot \nabla u_2(\mathbf{x}_h) > 0$ and the vector $T^{\text{BS}}(\mathbf{x}_h)$ is bounded in $\mathcal{O}^c$ then the mediator chooses the compromise improving direction $d = T^{\text{BS}}(\mathbf{x}_h)$ and apply the method described by Ehtamo et al [4] to generate the next tentative agreement $\mathbf{x}_{h+1}$.
3. Otherwise, among the four vectors $IS_i^\sigma(\mathbf{x}_h)$, $i = 1, 2$ and $\sigma = +/ -$, the mediator chooses the vector that (i) is bounded in $\mathcal{O}^c$, and (ii) is closest to the gradient of the other agent, $\nabla u_j(\mathbf{x}_h)$ ($j \neq i$). Denote this vector by $TG(\mathbf{x}_h)$. That is, we will be searching for a point on the indifference curve of agent i , $IC_i(\mathbf{x}_h)$, while trying to improve the utility of agent j . Note that when $\mathbf{x}_h$ is an interior point of $\mathcal{O}^c$ then the situation is symmetric for the two agents 1 and 2, and the mediator has the choice of either finding a point on $IC_1(\mathbf{x}_h)$ to improve the utility of agent 2, or finding a point on $IC_2(\mathbf{x}_h)$ to improve the utility of agent 1. To decide on which choice to make, the mediator has to compute the distribution of gains throughout the whole process to avoid giving more gains to one agent than to the other. Now, the point $\mathbf{x}_{h+1}$ to be generated lies somewhere on the intersection of $IC_i(\mathbf{x}_h)$ and the hyperplane defined by $\nabla u_i(\mathbf{x}_h)$ and $TG(\mathbf{x}_h)$. This intersection is approximated by $TG(\mathbf{x}_h)$. Thus, the sought after point $\mathbf{x}_{h+1}$ can be generated by first finding a point $\mathbf{y}_h$ along the direction of $TG(\mathbf{x}_h)$ and then move from $\mathbf{y}_h$ to the same direction of $\nabla u_i(\mathbf{x}_h)$ until we intersect with $IC_i(\mathbf{x}_h)$. Mathematically, let ζ and ξ denote the vectors $TG(\mathbf{x}_h)$ and $\nabla u_i(\mathbf{x}_h)$, respectively, $\mathbf{x}_{h+1}$ is the solution to the following optimisation problem:

$$\begin{aligned} & \max_{\lambda_1, \lambda_2 \in L} u_j(\mathbf{x}_h + \lambda_1 \zeta + \lambda_2 \xi) \\ & \text{s.t. } \mathbf{x}_h + \lambda_1 \zeta + \lambda_2 \xi \in \mathcal{O}^c, \text{ and } u_i(\mathbf{x}_h + \lambda_1 \zeta + \lambda_2 \xi) = u_i(\mathbf{x}_h), \\ & \text{where } L \text{ is a suitable interval of positive real numbers; e.g.,} \\ & L = \{\lambda | \lambda > 0\}, \text{ or } L = \{\lambda | a < \lambda \leq b\}, 0 \leq a < b. \end{aligned}$$

Given an initial tentative agreement $\mathbf{x}_0$, the method described above allows a sequence of tentative agreements $\mathbf{x}_1, \mathbf{x}_2, \dots$ to be iteratively generated. The iteration stops whenever a weakly Pareto optimal agreement is reached.

THEOREM 3. *If the sequence of agreements generated by the above method is bounded then the method converges to a point $\mathbf{x}^* \in \mathcal{O}^c$ that is weakly Pareto optimal.*

5. CONCLUSION AND FUTURE WORK

In this paper we have established a framework for negotiation that is based on MCDM theory for representing the agents' objectives and utilities. The focus of the paper is on integrative negotiation in which agents aim to maximise joint gains, or "create value."

We have introduced a mediator into the negotiation in order to allow negotiators to disclose information about their utilities, without providing this information to their opponents. Furthermore, the mediator also works toward the goal of achieving fairness of the negotiation outcome.

That is, the approach that we describe aims for both efficiency, in the sense that it produces Pareto optimal outcomes (i.e. no aspect can be improved for one of the parties without worsening the outcome for another party), and also for fairness, which chooses optimal solutions which distribute gains amongst the agents in some appropriate manner. We have developed a two step process for addressing the NP-hard problem of finding a solution for a set of integrative attributes, which is within the Pareto-optimal set for those attributes. For *simple attributes* (i.e. those which have a finite set of values) we use known optimisation techniques to find a Pareto-optimal solution. In order to discourage agents from misrepresenting their utilities to gain an advantage, we look for solutions that are least vulnerable to manipulation. We have shown that as long as one of the simple attributes does not strongly dominate the others, then truth telling is an equilibrium strategy for the negotiators during the stage of optimising simple attributes. For non-simple attributes we propose a mechanism that provides stepwise improvements to move the proposed solution in the direction of a Pareto-optimal solution.

The approach presented in this paper is similar to the ideas behind *negotiation analysis* [18]. Ehtamo et al [4] presents an approach to searching for joint gains in multi-party negotiations. The relation of their approach to our approach is discussed in the preceding section. Lai et al [12] provide an alternative approach to integrative negotiation. While their approach was clearly described for the case of two-issue negotiations, the generalisation to negotiations with more than two issues is not entirely clear.

Zhang et al [22] discuss the use of integrative negotiation in agent organisations. They assume that agents are honest. Their main result is an experiment showing that in some situations, agents' cooperativeness may not bring the most benefits to the organisation as a whole, while giving no explanation. Jonker et al [7] consider an approach to multi-attribute negotiation without the use of a mediator. Thus, their approach can be considered a complement of ours. Their experimental results show that agents can reach Pareto-optimal outcomes using their approach.

The details of the approach have currently been shown only for bilateral negotiation, and while we believe they are generalisable to multiple negotiators, this work remains to be done. There is also future work to be done in more fully characterising the outcomes of the determination of values for the non-simple attributes. In order to provide a complete framework we are also working on the distributive phase using the mediator.

Acknowledgement

The authors acknowledge financial support by ARC Discovery Grant (2006-2009, grant DP0663147) and DEST IAP grant (2004-2006, grant CG040014). The authors would like to thank Lawrence Cave-don and the RMIT Agents research group for their helpful comments and suggestions.

6. REFERENCES

- [1] F. Alemi, P. Fos, and W. Lacorte. A demonstration of methods for studying negotiations between physicians and health care managers. *Decision Science*, 21:633–641, 1990.
- [2] M. Ehrgott. *Multicriteria Optimization*. Springer-Verlag, Berlin, 2000.
- [3] H. Ehtamo, R. P. Hamalainen, P. Heiskanen, J. Teich, M. Verkama, and S. Zionts. Generating pareto solutions in a two-party setting: Constraint proposal methods. *Management Science*, 45(12):1697–1709, 1999.
- [4] H. Ehtamo, E. Kettunen, and R. P. Hmlinen. Searching for joint gains in multi-party negotiations. *European Journal of Operational Research*, 130:54–69, 2001.
- [5] P. Faratin. *Automated Service Negotiation Between Autonomous Computational Agents*. PhD thesis, University of London, 2000.
- [6] A. Foroughi. Minimizing negotiation process losses with computerized negotiation support systems. *The Journal of Applied Business Research*, 14(4):15–26, 1998.
- [7] C. M. Jonker, V. Robu, and J. Treur. An agent architecture for multi-attribute negotiation using incomplete preference information. *J. Autonomous Agents and Multi-Agent Systems*, (to appear).
- [8] R. L. Keeney and H. Raiffa. *Decisions with Multiple Objectives: Preferences and Value Trade-Offs*. John Wiley and Sons, Inc., New York, 1976.
- [9] G. Kersten and S. Noronha. Rational agents, contract curves, and non-efficient compromises. *IEEE Systems, Man, and Cybernetics*, 28(3):326–338, 1998.
- [10] M. Klein, P. Faratin, H. Sayama, and Y. Bar-Yam. Protocols for negotiating complex contracts. *IEEE Intelligent Systems*, 18(6):32–38, 2003.
- [11] S. Kraus, J. Wilkenfeld, and G. Zlotkin. Multiagent negotiation under time constraints. *Artificial Intelligence Journal*, 75(2):297–345, 1995.
- [12] G. Lai, C. Li, and K. Sycara. Efficient multi-attribute negotiation with incomplete information. *Group Decision and Negotiation*, 15:511–528, 2006.
- [13] D. Lax and J. Sebenius. The manager as negotiator: The negotiator's dilemma: Creating and claiming value, 2nd ed. In S. Goldberg, F. Sander & N. Rogers, editors, *Dispute Resolution*, 2nd ed., pages 49–62. Little Brown & Co., 1992.
- [14] M. Lomuscio and N. Jennings. A classification scheme for negotiation in electronic commerce. In *Agent-Mediated Electronic Commerce: A European Agentlink Perspective*. Springer-Verlag, 2001.
- [15] R. Maes and A. Moukas. Agents that buy and sell. *Communications of the ACM*, 42(3):81–91, 1999.
- [16] J. Nash. Two-person cooperative games. *Econometrica*, 21(1):128–140, April 1953.
- [17] H. Raiffa. *The Art and Science of Negotiation*. Harvard University Press, Cambridge, USA, 1982.
- [18] H. Raiffa, J. Richardson, and D. Metcalfe. *Negotiation Analysis: The Science and Art of Collaborative Decision Making*. Belknap Press, Cambridge, MA, 2002.
- [19] T. Sandholm. Agents in electronic commerce: Component technologies for automated negotiation and coalition formation. *JAAMAS*, 3(1):73–96, 2000.
- [20] J. Sebenius. Negotiation analysis: A characterization and review. *Management Science*, 38(1):18–38, 1992.
- [21] L. Weingart, E. Hyder, and M. Pietrula. Knowledge matters: The effect of tactical descriptions on negotiation behavior and outcome. Tech. Report, CMU, 1995.
- [22] X. Zhang, V. R. Lesser, and T. Wagner. Integrative negotiation among agents situated in organizations. *IEEE Trans. on Systems, Man, and Cybernetics, Part C*, 36(1):19–30, 2006.