Searching for Fair Joint Gains in Agent-based Negotiation

Minyi Li
Swinburne University of Technology
Faculty of Information and Communication Technologies
Hawthorn, 3122 Victoria, Australia
myli@swin.edu.au

Quoc Bao Vo
Swinburne University of Technology
Faculty of Information and Communication Technologies
Hawthorn, 3122 Victoria, Australia
bvo@swin.edu.au

Ryszard Kowalczyk
Swinburne University of Technology
Faculty of Information and Communication Technologies
Hawthorn, 3122 Victoria, Australia
rkowalczyk@swin.edu.au

ABSTRACT

In multi-issue negotiations, autonomous agents can act cooperatively to benefit from mutually preferred agreements. However, empirical evidence suggests that they often fail to search for joint gains and end up with inefficient results. To address this problem, this paper proposes a novel mediated negotiation procedure to support the negotiation agents in reaching an efficient and fair agreement in bilateral multi-issue negotiation. At each stage of negotiation, the mediator searches for the compromise direction based on a new E-DD (Equal Directional Derivative) approach and computes the new tentative agreement. The numerical analysis presented in this paper demonstrates that the proposed approach not only guarantees Pareto efficiency, but also produces fairer improvements for two negotiating agents compared with other existing methods.

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, joint gains, fairness, Pareto-improvement, mediator

1. INTRODUCTION

Negotiation is one of the most effective techniques for autonomous agents to resolve conflicts and reach mutually beneficial agreements [3, 6, 7, 8, 10, 13, 15]. It is being increasingly used in different domains including agent-based trading systems, resource allocation, service level agreements, etc. [6, 16]. When multiple issues are involved in negotiation simultaneously, like price, quality attributes, delivery time, etc., the agents with divergent preferences can benefit from making Pareto improvements: without hurting each other, the agents can gain by lowering their requirements on some negotiation issues while demanding more in other issues [5, 13, 15]. For instance, consider the problem of selling laptops. The seller hopes to increase sales without lowering the profit per item. She raises the price by 5%, while adding extended warranty, technical support and free delivery service into the contract which are of low cost for her. The buyer is also more willing to negotiate such a package because he values highly the extended warranty and prefers free delivery. Furthermore, the after - sales service is very important to him. Such situations where both parties are better off, are normally called “win-win” situations [13, 15, 17].

Cite as: Searching for Fair Joint Gains in Agent-based Negotiation, Minyi Li, Quoc Bao Vo, Ryszard Kowalczyk, *Proc. of 8th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2009)*, Decker, Sichman, Sierra and Castelfranchi (eds.), May, 10–15, 2009, Budapest, Hungary, pp. 1049–1056

Intuitively, agents should have a common interest to cooperate and search for possible joint gains. However, empirical evidence suggests that self-interested agents often fail to reach consensus or end up with inefficient results [4, 6, 14, 15, 17, 18]. Raiffa [15] provides analyses on the negotiators’ failure to achieve efficient agreements in practice. Lax and Sebenius [14] discuss the Negotiator’s Dilemma in deciding whether to pursue a cooperative or a competitive strategy at a particular time during negotiation. Fatima et al. [6, 8] point out that self-interested agents would like to reach an agreement that is as favorable to them as possible, whereas the final decision is jointly made and needs to be agreed to by both the agents. Consequently, the problems met by negotiation agents are not only to choose cooperative or competitive strategies, but also to consider how much they could gain individually if they cooperate and in which way of cooperation they could gain more, or at least receive a fair deal. Negotiation therefore, requires techniques that deal with rational agents fairly and lead them to mutually beneficial agreements. In this paper, we propose a cooperative framework for bilateral multi-issue negotiations. A novel negotiation procedure is presented, in which a non-biased mediator agent searches for the compromise directions based on a new E-DD (Equal Directional Derivative) approach and supports negotiation agents in reaching the efficient and fair agreement. The paper is organized as follows: Section 2 provides a background on negotiation theory and discusses some related work. Section 3 presents the proposed mediation procedure. In Section 4 the experiments and numerical analysis are provided, and finally Section 5 presents the concluding remarks and discusses the future work.

2. BACKGROUND

2.1 Bilateral Multi-issue Negotiation Model

Formally, the bilateral multi-issue negotiation model can be defined as follows [17]:

1. Negotiators: Let $N = \{N_1, N_2\}$ be a set of two agents who are involved in negotiation.

2. Attributes: Let $Attr = \{x_1, x_2 ... x_m\}$ be a set of attributes which are the issues that agents are negotiating over, such as price, length of warranty, storage capacity, delivery time, etc.. Each attribute, denoted by x_j can take a value such as '\$1500' or '3 years' from its corresponding interval $X_j, j \in \{1, 2...m\}$. We will assume in this paper that all attributes are continuous, i.e., $X_j \subseteq \Re$.

3. Outcomes: An outcome (or alternative) o is represented by an assignment of values to the corresponding attributes in $Attr$: $o = (x_1, x_2 ... x_m)$ and $o \in O$, where O indicates the set of all possible outcomes. We denote by O^{In} the set of interior points of the outcome space O.

4. Preference: Preference indicates the ranking (or order, precedence) of possible outcomes based on satisfaction or utility they could provide for a negotiation agent. In the negotiation theory, the standard way to model the negotiator's preference is with his preference relations, also called a binary relation [11]. Let o_1, o_2 be two possible outcomes $o_1, o_2 \in O$. Then '$\succeq$' is a preference relation on O such that $o_1 \succeq o_2$ if and only if $o_1$is at least as preferable as o_2 (or, o_1 is weakly preferred to o_2). And o_1 is strictly preferred to o_2 (notation $o_1 \succ o_2$) if and only if $o_1 \succeq o_2$ but not $o_2 \succeq o_1$. When $o_1 \succeq o_2$ and $o_2 \succeq o_1$, we say that the agent is indifferent between these two outcomes o_1 and o_2, denoted by $o_2 \prec\succ o_1$.

5. Utility: Utility is a measure of the agent's relative satisfaction with a particular outcome $o \in O$. An agent's utility is based on its value function and some negotiation-specific regulations or factors, such as participation fees, transaction costs etc.

   While the preference relation is usually a conventional foundation to examine the behavior of a negotiator, it is often convenient to represent preferences with a utility function and reason indirectly about the preferences using the utility function [11]. Without loss of generality, we will normalize the utility value of agents as a real number over the interval $[0, 1]$. Consider a set of alternative outcomes O, a utility function $u_i(o)$ of an agent N_i assigns a numerical value to o ($o \in O$), such that the rank ordering of these alternatives is preserved. Formally, $u_i(o) : O \to [0, 1]$ is a utility function representing agent N_i's preference relation '$\succeq$' if the following holds for all $o_1, o_2 \in O$: $u_i(o_1) \geq u_i(o_2) \Leftrightarrow o_1 \succeq o_2$ [11].

2.2 Pareto Efficiency

A goal of negotiations is to achieve an outcome as "Pareto Efficient" as possible [9, 14]. Pareto efficiency states that no one could gain more without making others worse off. By definition, an outcome o^* is Pareto efficient if there is no other Pareto dominant alternative[1] that could improve the utility of at least one agent without lowering the utility of others.

From another point of view, an outcome that is not Pareto efficient implies that certain changes in assignment of attribute values may result in some agents being made better off with no other agents being made worse off. It can be more efficient through a series of Pareto improvements, called the Pareto efficient enhancement process.

2.3 Fairness

Fair division has been extensively studied in economics, political science and mathematics for a long time [1]. As agents are autonomous, fairness is also an important criterion in the joint gain division scenario: the outcome needs to respect and balance agents' individual utility gains. Under incomplete information settings, it is difficult to find out an ideal outcome that everyone would be given an equal share of the joint gains after the Pareto efficient enhancement process. Nevertheless, in mediated negotiation the mediator should be able to assign joint gains as fair as possible.

2.4 Related Research

A large amount of earlier research work have been dealing with the joint gains seeking methods, but most of them have one of two key limitations: either they do not address fairness, or are based on the assumption of complete information. In [15] Raiffa develops an approach for making moves and finding joint gains. Each step of improvement increases one party's utility but keeps another unchanged. Ehtamo et al. [4] present a *bisecting approach* to choose the compromise direction over continuous issues in bilateral negotiation. It is essentially a gradient search method, involving an iterative procedure with which two negotiators can search for joint gains and finally reach an efficient outcome. The gradient of agent N_i's utility at the current point o, denoted by $\nabla u_i(o)$, points in the most preferred direction in which agent N_i would like to move as to make immediate gains. The authors examine how the compromise directions should be chosen and conclude that the simple 50-50 split between the two negotiators' gradient directions turns out to be a suitable and fair compromise direction (T^{BS} in Fig. 1). However, after conducting adequate experiments with different preference settings we argue that the bisecting approach is in fact, not as fair as intuitively expected. We analyze the main reason behind this problem from the payoff space of marginal utilities in Section 3.

Another mediation-based negotiation model with incomplete information is given by Lai et al. [13]. In their approach, the mediator conducts a Pareto efficient enhancement for a proposal in each negotiation period. It doesn't address the fairness issue between utility gains of two negotiating agents. The algorithm they develop (we call it ϵ-*Satisfying* approach in the following sections) is of high efficiency in the two-attribute cases, however, the computational complexity grows rapidly as the number of dimensions increases. Furthermore, the approach acts on the premise that the Pareto frontier search direction could be found by

[1] An alternative o Pareto dominates another agreement o' if for all agents involved in the negotiation, $u_i(o) \geq u_i(o')$ and the inequality is strict for at least one agent.

comparing the values of one attribute with the other attribute x_2 equal to $x_2 + \epsilon$ (or $x_2 - \epsilon$) on two agents' initial indifference curves. But sometimes there might be more such points (See Section 4.1) and sometimes such point may not exist (e.g. the initial agreement is the maximum or minimum point of the indifference curve). In either case, the Pareto frontier search direction cannot be determined and the procedure is incapable of proceeding.

In this paper we propose a novel mediator-based negotiation procedure, which guarantees Pareto optimality, while addressing the fairness issue efficiently.

3. THE PARETO EFFICIENT ENHANCEMENT NEGOTIATION PROCEDURE

In this section, we describe the procedure to help the negotiation agents search for joint gains and finally reach a Pareto efficient agreement with the improved fairness. In mediation, the initial tentative agreement is important: it defines the beginning for the enhancement process and affects the fairness of the final agreement to be reached through mediations. Ehtamo et al. [2] present several methods to choose the initial tentative agreement (called reference point in their paper). Vo et al. [17] present a procedure to find out the fair initial tentative agreement by maximizing the utilities of agents meanwhile reducing the difference between the agents' valuations of each attribute in the initial tentative agreement. In this paper, we focus on the enhancement process which iteratively improves the initial agreement and finally achieves Pareto Optimality. It is based on the following assumptions:

1. Two negotiation agents N_1 and N_2 are acting cooperatively to solve a problem involving multiple continuous attributes. They are currently at an initial tentative agreement point $o = (x_1, x_2 ... x_m)$ $(o \in O)$, which is a particular point in the m-dimensional space of real numbers.

2. Agents' utility functions $u_1(o)$ and $u_2(o)$ can be described by monotonic, continuously differentiable and concave functions $u_i(o) : \Re^m \rightarrow \Re, (i = 1, 2)$.

3. Agents are honest. This assumption excludes the possibility that the agents act strategically to obtain more gains from the enhancement process. For instance, if an agent knows that there is still a large space for Pareto improvement, it could give out a smaller gradient magnitude and mislead the mediator to determine a compromise direction which is much closer to its gradient direction.

The Pareto efficient enhancement is an iterative procedure that converges to a Pareto efficient agreement. In each iteration, the mediator firstly asks the agents for their utility gradients $\nabla u_1(o)$ and $\nabla u_2(o)$ respectively, then performs two major tasks: *i)*checks whether the current agreement is efficient: if yes, it terminates the enhancement process; otherwise, it chooses a fair compromise direction; *ii)*determines a new tentative agreement along the compromise direction which is more efficient than the current one. This iterative process continues until an efficient agreement is found. The following sections 3.1 and 3.2 describe the procedure and technical details for calculating the compromise direction and determining the new tentative agreement. For the simplicity of explanation, we consider the two attributes case $o = (x_1, x_2)$, however, the approach is feasible for the m-dimensional $(m \geq 2)$ cases based on the same principle.

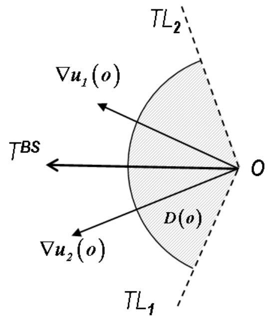

Figure 1: Bisecting approach

3.1 Searching for Fair Improvement Direction

The direction choosing problem will be analyzed in the payoff space of marginal utilities and the corresponding procedure will now be described in technical details. In this paper we concentrate on the set of interior points O^{In} except for the boundary points of the outcome space O. Let ϕ denote the angle between two agents' utility gradients $\nabla u_1(o)$ and $\nabla u_2(o)$ at the current point o:

$$\phi = arccos\left(\frac{\nabla u_1(o) \cdot \nabla u_2(o)}{\|\nabla u_1(o)\| \cdot \|\nabla u_2(o)\|}\right) \tag{1}$$

We call this angle an opposition angle, since it can be considered as a measurement to examine the opposition situation of the negotiation agents. The following definitions are adapted from [12]:

Definition 1. The agents are in local weak opposition at a point o iff their utility gradients $\nabla u_1(o)$, $\nabla u_2(o)$ form an acute or right opposition angle at point o: $0 \leq \phi \leq \frac{\pi}{2}$ (See Fig. 2(a)). When the opposition angle $\phi = 0$, the two gradients point in the same direction $\nabla u_1(o) = k\nabla u_2(o), k \in \Re^+$.

Definition 2. The agents are in local strong opposition at a point o iff their utility gradients $\nabla u_1(o)$, $\nabla u_2(o)$ form an obtuse opposition angle at point o: $\frac{\pi}{2} < \phi < \pi$ (See Fig. 2(b)).

Definition 3. The agents are in local strict opposition at a point o iff the opposition angle between their utility gradients is equal to π: $\phi = \pi$ (See Fig. 2(c)). In such case, the two gradients are in the opposite direction $\nabla u_1(o) = -k\nabla u_2(o), k \in \Re^+$.

Moreover, the following theorem has been proved in [12]:

THEOREM 1. *An outcome $o^* \in O^{In}$ is efficient iff the agents are in local strict opposition at o^*: $\phi = \pi$.*

Using the geometric interpretation of the gradients, these efficient (local strict opposition) points may also be called tangential points, i.e., points where the two indifference curves (or indifference surfaces) tangent each other[12].

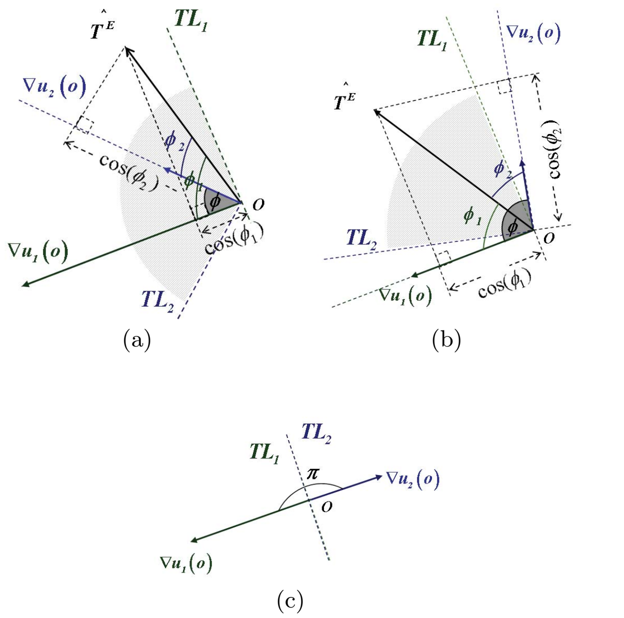

Figure 2: Compromise direction within the jointly improving directions set

The utility gradient $\nabla u_i(o)$ of agent N_i represents the most preferred direction in which agent N_i would like to move to make immediate gains from o. The gradient magnitude, also called length or value, represents the greatest rate of increase of the utility. Any direction making an acute angle with the gradient (i.e., any direction $\vec{d}$ such that $\vec{d} \cdot \nabla u_i(o) > 0$) is a direction in which agent N_i's utility can be improved [3, 12]. According to [3], for each point $o \in O^{In}$ these directions form two sets D_1 and D_2, which contain all the improving directions for the agents N_1 and N_2, respectively. Formally,

$$\begin{aligned} D_i(o) \;&=\; \Big\{ \vec{d} \mid \vec{d} \cdot \nabla u_i(o) > 0,\ \exists \delta > 0 \ \ s.t. \\ &\qquad o + \lambda \cdot \vec{d} \in O,\ 0 < \lambda < \delta \Big\},\ i = 1,\ 2 \end{aligned}$$

Then the set of jointly improving directions is: $D(o) = (D_1(o) \bigcap D_2(o)) \bigcup \{0\}$. The zero vector is added for notational convenience and a point o is Pareto optimal if and only if $D(o) = \{0\}$ [3]. In Fig. 1, $D(o)$ is the set of directions which go through point o and locate between the two intersecting tangent lines (TL_1 and TL_2) of the agents' indifference curves at point o.

Given a compromise direction $\vec{d}$ at point o, the *utility gains* of agent N_i moving from point o to point $o + \lambda \cdot \vec{d}$ is:

$$\Delta u_i = u_i\left(o + \lambda \cdot \vec{d}\right) - u_i(o) \tag{2}$$

where $\lambda\,(\lambda > 0)$ is the moving distance. Moreover, when the moving distance $\lambda \to 0$, the ratio of agent N_i's *marginal utility gains* in the compromise direction $\vec{d}$ is:

$$\lim_{\lambda \to 0} \frac{\Delta u_i}{\lambda} = \lim_{\lambda \to 0} \frac{u_i\left(o + \lambda \cdot \vec{d}\right) - u_i(o)}{\lambda} \tag{3}$$

The ratio of marginal utility gains, in essence, is equal to the directional derivative of function u_i, which represents the instantaneous rate of increase of the function u_i, moving through point o, in the direction of $\vec{d}$:

$$\nabla_{\vec{d}}\, u_i(o) = \nabla u_i(o) \cdot \vec{d} = \lim_{\lambda \to 0} \frac{u_i\left(o + \lambda \cdot \vec{d}\right) - u_i(o)}{\lambda} \tag{4}$$

If the two gradients $\nabla u_1(o)$ and $\nabla u_2(o)$ point in the same direction (i.e. $\phi = 0$) at point o, we would choose this direction as the compromise direction at that stage: $\vec{d} = \nabla \hat{u_i}(o)$, $i = 1, 2$ ($\nabla \hat{u_i}(o)$ is a unit vector representing the direction of the gradient vectors). When the opposition angle $\phi \neq 0$, let us firstly review the method presented by Ehtamo et al [3], in which the direction choosing problem was formulated in the payoff space of marginal utilities using axiomatic bargaining theory. For the case of two parties, the authors define a bisecting choosing function $T^{BS}(o) : o \in O^{In} \to \Re^2$

$$T^{BS}(o) = \frac{\nabla u_1(o)}{\|\nabla u_1(o)\|} + \frac{\nabla u_2(o)}{\|\nabla u_2(o)\|}$$

then T^{BS} chooses the direction that bisects $D(o)$ and it yields:

$$\frac{\nabla u_1(o) \cdot T^{BS}(o)}{\nabla u_2(o) \cdot T^{BS}(o)} = \frac{\|\nabla u_1(o)\|}{\|\nabla u_2(o)\|} \tag{5}$$

Hence, in bisecting approach the marginal gains of agents are always divided at the ratio of their gradient magnitudes. But more often than not, different agents would have different gradient magnitudes at a point and this should not be neglected. When one's magnitude far outweighs the other, the bisecting approach will give a high percentage of marginal gains to one agent but leave the other too little. For instance, in a two dimensional case, agent N_1's utility gradient $\nabla u_1(o) = \{1.44, 0.36\}$ and agent N_2's utility gradient $\nabla u_2(o) = \{-0.04, -0.16\}$. Then the bisecting direction $\vec{d} = \{0.727607, -0.727607\}$ divides marginal gains at the ratio of $\|\nabla u_1(o)\| : \|\nabla u_2(o)\| = 9 : 1$. 90% of marginal gains is given to agent N_1 but just 10% is left for agent N_2. Therefore, we argue that the bisecting approach is not as fair as intuitively expected. Based on a tentative agreement, the fair compromise direction should produce the equal shares of marginal gains between the negotiation agents, i.e. the compromise direction that has the equal directional derivatives on two agents' utility functions at the tentative agreement.

$$\nabla u_1(o) \cdot \hat{T^E} = \nabla u_2(o) \cdot \hat{T^E} \tag{6}$$

where $\hat{T^E}$ is a unit vector representing the compromise direction. In order to guarantee both negotiation agents can improve their utilities, the compromise direction must meet the requirement that $\hat{T^E} \in D(o)$.

Theorem 2. *At any interior point $o \in O^{In}$, there exists a compromise direction $\hat{T^E} \in D(o)$ such that:*

$$\nabla u_1(o) \cdot \hat{T^E} = \nabla u_2(o) \cdot \hat{T^E}$$

Proof. Let α be the plane (or hyperplane) defined by the two gradient vectors $\nabla u_1(o)$ and $\nabla u_2(o)$, TL_1 and TL_2 be the tangent lines[2] that intersect two agents' indifference

[2]For a m-dimensional case ($m > 2$), these would become tangent hyperplanes.

curves[3] at the current point o (See Fig. 2). According to [3], the set of jointly improving directions $D(o)$ is the set of directions between TL_1 and TL_2 (See the shaded area in Fig 1).

If $\phi = \pi$, then the tangent lines coincide with each other and $D_1(o) \bigcap D_2(o) = \emptyset \Leftrightarrow D(o) = \{0\}$ (See Fig. 2(c)). In such case, let $\vec{d} = 0 \Rightarrow \vec{d} \in D(o)$ and it satisfies $\nabla u_1(o) \cdot \vec{d} = \nabla u_2(o) \cdot \vec{d} = 0$.

Otherwise when $\phi \neq \pi$, for any jointly improvement direction $\vec{d} \in D(o)$, it forms two acute angles with the gradient vectors $\nabla u_1(o)$ and $\nabla u_2(o)$, denoted by ϕ_1 and ϕ_2 respectively; and two angles with the tangent lines TL_1 and TL_2, denoted by $\bar{\phi_1}$ and $\bar{\phi_2}$ respectively. By definition, $\nabla u_1(o)$ is perpendicular to $TL_1 \Rightarrow \phi_1 = \frac{\pi}{2} - \bar{\phi_1}$ and $\nabla u_2(o)$ is perpendicular to $TL_2 \Rightarrow \phi_2 = \frac{\pi}{2} - \bar{\phi_2}$. Let θ ($0 < \theta < \pi$) be the angle between these two tangent lines, then $\bar{\phi_2} = \theta - \bar{\phi_1}$. Let function $f(\bar{\phi_1}) = \frac{sin(\bar{\phi_1})}{sin(\theta - \bar{\phi_1})}$, then f is a continuous function over $\bar{\phi_1} \in (0, \theta)$ and the range of this function is $(0, +\infty)$. Based on the *Intermediate Value Theorem*, there must exist $\bar{\phi_1^*} \in (0, \theta)$, such that $0 < f(\bar{\phi_1^*}) = \frac{sin(\bar{\phi_1^*})}{sin(\theta - \bar{\phi_1^*})} = \frac{\|\nabla u_2(o)\|}{\|\nabla u_1(o)\|}$. Because $\frac{sin(\bar{\phi_1^*})}{sin(\theta - \bar{\phi_1^*})} = \frac{sin(\bar{\phi_1^*})}{sin(\bar{\phi_2^*})} = \frac{cos(\phi_1^*)}{cos(\phi_2^*)} \Rightarrow \frac{cos(\phi_1^*)}{cos(\phi_2^*)} = \frac{\|\nabla u_2(o)\|}{\|\nabla u_1(o)\|}$, and thus $\nabla u_1(o) \cdot cos(\phi_1^*) = \nabla u_2(o) \cdot cos(\phi_2^*)$. As $\nabla u_i(o) \cdot \hat{d^*} = \|\nabla u_i(o)\| \cdot cos(\phi_i^*)$ ($\hat{d^*}$ is the unit vector in plane α that forms acute angles ϕ_1^*, ϕ_2^* with two gradient vectors and $\bar{\phi_1^*}$, $\bar{\phi_2^*}$ with two tangent lines respectively), it meanwhile satisfies $\nabla u_1(o) \cdot \hat{d^*} = \nabla u_2(o) \cdot \hat{d^*}$ and the proof is completed. □

Now the problem is to define a choosing function for this compromise direction $\hat{T^E}$. As $\hat{T^E}$ is in the plane α defined by the two gradient vectors, based on *Coplanar Vector Theorem*, $\hat{T^E}$ could be considered as the sum of two vectors which are in the same or opposite direction of two gradients $\nabla u_1(o)$ and $\nabla u_2(o)$ respectively,

$$\hat{T^E} = k_1 \cdot \frac{\nabla u_1(o)}{\|\nabla u_1(o)\|} + k_2 \cdot \frac{\nabla u_2(o)}{\|\nabla u_2(o)\|} \quad (k_1, k_2 \in \Re) \quad (7)$$

Substituting $\hat{T^E}$ in Equation (6) with Equation (7), we obtain:

$$\frac{k_1}{k_2} = \frac{\|\nabla u_2(o)\| - \|\nabla u_1(o)\| \cdot cos(\phi)}{\|\nabla u_1(o)\| - \|\nabla u_2(o)\| \cdot cos(\phi)} \quad (8)$$

Let $k_1 = t \cdot [\|\nabla u_2(o)\| - \|\nabla u_1(o)\| \cdot cos(\phi)]$, $k_2 = t \cdot [\|\nabla u_1(o)\| - \|\nabla u_2(o)\| \cdot cos(\phi)]$ ($t \in \Re$ and $\left\|\hat{T^E}\right\| = 1$). Consequently, we can define a *collinear vector* $\bar{T^E}$ of $\hat{T^E}$, $\hat{T^E} = t \cdot \bar{T^E}$:

$$\begin{aligned} \bar{T^E} &= [\|\nabla u_2(o)\| - \|\nabla u_1(o)\| \cdot cos(\phi)] \cdot \frac{\nabla u_1(o)}{\|\nabla u_1(o)\|} \\ &+ [\|\nabla u_1(o)\| - \|\nabla u_2(o)\| \cdot cos(\phi)] \cdot \frac{\nabla u_2(o)}{\|\nabla u_2(o)\|} \quad (9) \end{aligned}$$

Then the direction choosing function $\hat{T^E}(o) : o \in O^{In} \rightarrow \Re^m$ can be obtained as follows:

If two gradients point in the same direction ($\phi = 0$):

$$\hat{T^E}(o) = \hat{\nabla u_i}(o) \quad i = 1, 2$$

Else:

If $\bar{T^E} \cdot \nabla u_i(o) > 0$, $i = 1, 2$:

$$\hat{T^E}(o) = \frac{\bar{T^E}}{\|\bar{T^E}\|}$$

Else:

$$\hat{T^E}(o) = -\frac{\bar{T^E}}{\|\bar{T^E}\|}$$

3.2 Choosing A New Tentative Agreement

Once a compromise direction $\vec{d} = \hat{T^E}(o)$ has been chosen, the mediator needs to choose a new tentative agreement along $\vec{d}$ so that both agents benefit. Ehtamo et al. [3] describes this procedure as the maximization of decision makers' utilities. The decision makers consider all the possible agreements $o' = o + \lambda \cdot \hat{T^{BS}}$ ($\hat{T^{BS}}$ is a unit vector representing the bisecting direction) and announce one that optimizes its own utility: $o_i^* = o + \lambda_i^* \cdot \hat{T^{BS}}$, where $\lambda_i^* = argmax_{\lambda \in (0, +\infty]} u_i\left(o + \lambda \cdot \hat{T^{BS}}\right)$. And then the mediator chooses the one which is closer to the current point o such that it guarantees both parties could improve their utilities: $\bar{\lambda} = min\{\lambda_1^*, \lambda_2^*\}$. It is also pointed out by Ehtamo et al. [3] that the complete optimization could be replaced by small improvements (a shorter distance is moved at each stage). Moreover, we argue that the enhancement by small improvements could produce a fairer outcome.

THEOREM 3. *At each improvement stage when the opposition angle $\phi \neq 0$, with compromise direction $\vec{d} = \hat{T^E}(o)$, if moving distance $\lambda \rightarrow 0$, the utility gains for two agents at that stage is equal to each other:*

$$\lim_{\lambda \rightarrow 0} \frac{\Delta u_1}{\Delta u_2} = 1$$

PROOF. Based on Equation (2), $\lim_{\lambda \rightarrow 0} \frac{\Delta u_1}{\Delta u_2} = \lim_{\lambda \rightarrow 0} \frac{u_1(o + \lambda \cdot \vec{d}) - u_1(o)}{u_2(o + \lambda \cdot \vec{d}) - u_2(o)}$ $= \lim_{\lambda \rightarrow 0} \frac{\frac{u_1(o + \lambda \cdot \vec{d}) - u_1(o)}{\lambda}}{\frac{u_2(o + \lambda \cdot \vec{d}) - u_2(o)}{\lambda}} = \frac{\lim_{\lambda \rightarrow 0} \frac{u_1(o + \lambda \cdot \vec{d}) - u_1(o)}{\lambda}}{\lim_{\lambda \rightarrow 0} \frac{u_2(o + \lambda \cdot \vec{d}) - u_2(o)}{\lambda}} = \frac{\nabla u_1(o) \cdot \vec{d}}{\nabla u_2(o) \cdot \vec{d}}$
According to Equation (6) $\frac{\nabla u_1(o) \cdot \vec{d}}{\nabla u_2(o) \cdot \vec{d}} = 1$, so $\lim_{\lambda \rightarrow 0} \frac{\Delta u_1}{\Delta u_2} = \frac{\nabla u_1(o) \cdot \vec{d}}{\nabla u_2(o) \cdot \vec{d}} = 1$ and the proof is completed. □

Consequently at every stage, the smaller λ is, the more balanced the joint gains distribution will be. Nevertheless, the computational complexity should also be considered, because more steps and data transfers between the mediator and negotiation agents may be needed as λ reduces. In our approach, the moving distance $\widetilde{\lambda}$ is predetermined based on an overall consideration of various system goals(e.g. computational complexity, the accuracy of fairness, etc.). On the other hand, when the tentative agreement is very close to Pareto frontier, moving $\widetilde{\lambda}$ may go beyond the Pareto frontier (particularly when $\widetilde{\lambda} > \bar{\lambda}$, $\bar{\lambda} = min\{\lambda_1^*, \lambda_2^*\}$) and does not guarantee that both agents could increase their utilities. To address this issue, we employ a binary search approach by iteratively decreasing $\widetilde{\lambda}$ with $\widetilde{\lambda} = \widetilde{\lambda}/2$.

Regarding the standard of Pareto efficiency, Lai et al. [13] introduce the definition of *ϵ-satisfying Pareto efficient solution* in a two dimensional case. According to this definition, how close the final outcome is to Pareto frontier is defined by

[3]For a m-dimensional case ($m > 2$), these would become indifference surfaces.

the accuracy parameter ϵ, particularly when $\epsilon \to 0$, the solution is completely Pareto efficient. Similarly, we proposed the following definitions and theorem for the stopping point of the binary search process:

Definition 4. Given a point o in the m-dimensional space, we call the set of possible alternatives

$$\sigma_\tau(o) = \{o' \mid |o' - o| \leq \tau,\ o, o' \in O\}$$

as the *τ-space*[4] of point o.

Definition 5. A point o is a *τ-satisfying Pareto efficient solution* if one of the following two properties is satisfied:

- there does not exist any point that is mutually better than o for both agents;
- all the points mutually better than o, if existing, are located in the *τ-space* of point o.

THEOREM 4. *When $\tau \to 0$, a τ-satisfying Pareto efficient solution is completely Pareto efficient.*

PROOF. If $\tau \to 0$, then $\sigma_\tau(o) \to \emptyset$. There will be no mutually preferred alternatives for both agents, then o is Pareto efficient and the proof is completed. □

Now, we describe an iteration to reach the next tentative agreement from the current point o. The mediator firstly asks the two agents for their gradients $\nabla u_1(o)$ and $\nabla u_2(o)$ respectively:

1. If the opposition angle $\phi = \pi$, then o is a Pareto-optimal outcome according to Theorem 1 and the process is ended by the final agreement point o.
2. Otherwise, the mediator chooses the compromise direction $\vec{d} = \hat{T^E}(o)$ and asks the two agents to consider the possible new tentative agreement $o + \widetilde{\lambda} \cdot \vec{d}$. If both of them are willing to move, then $o + \widetilde{\lambda} \cdot \vec{d}$ is chosen to be the new tentative agreement. Otherwise, the mediator begins a binary search process to find out the possible mutually preferred alternative along $\vec{d}$ by iteratively decreasing $\widetilde{\lambda}$ with $\widetilde{\lambda} = \widetilde{\lambda}/2$. This binary search process continues, until either it finds out a point that both agents are willing to go; or if $\widetilde{\lambda} \leq \tau$, then the current point o is a τ-satisfying Pareto efficient solution and the enhancement process ends.

4. NUMERICAL ANALYSIS

In this section, we compare the experimental results over different types of utility functions using bisecting, ϵ-Satisfying and the proposed E-DD (equal directional derivative) approach. The experiment of negotiation over fishing rights in [3] will be discussed in Section 4.1. We will compare the final result made by small improvements with that of complete optimization method presented in [3]. Section 4.2 will provide a clear view of overall performance of these three approaches over 5000 random preference combinations with the quadratic utility functions.

[4]Within a two (or three) dimensional space, the *τ-space* is a circle (or a sphere) centered at o and with radius τ.

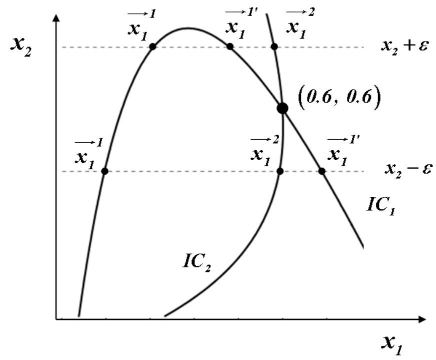

Figure 3: ϵ-Satisfying approach

4.1 Negotiation over Fishing Rights

A well known real world example of negotiation over fishing rights is given by [3], in which two countries affect each other at time t through the size of the remaining fish population Q_t ($x_1 + x_2 \leq Q_t$) and they wish to maximize their overall discounted utility. Suppose that, if uninterrupted, the fish population would grow according to $Q_{t+1} = {Q_t}^{\alpha}$ ($0 < \alpha < 1$). In accordance with [3], let $o = (x_1, x_2)$ be a possible alternative, $o \in O$ and $O = \{(x_1, x_2) \mid x_1 + x_2 \leq Q_t\}$, the utility functions of two countries are as follows:

$$u_i(o) = ln x_i + \beta_i ln(1/2)(Q - x_1 - x_2)^\alpha$$

where β_i is the discount factor of the country i and Q is the initial stock of fish. In [3] the optimization function along a compromise direction d is given by:

$$\lambda_i = \frac{d_i(Q - x_1(k) - x_2(k)) - \alpha\beta_i(d_1 + d_2)x_i(k)}{(1 + \alpha\beta_i)(d_1 + d_2)d_i}$$

Table 1 and Table 2 compare the final outcomes and the joint gains divisions among two agents using three approaches. With the same setting of [3] ($Q = 1.259$; $\alpha = 0.2852$; $\beta_1 = 0.9$; $\beta_2 = 0.4$; the initial tentative agreement point $o = (0.6, 0.6)$) and $\tau = 0.0001, \epsilon = 0.0001$, Table 1 displays the result of Pareto efficient enhancement by complete optimization, while Table 2 shows the enhancement result with a pre-defined $\lambda = 0.01$. Both Table 1 and 2 indicate that E-DD approach achieves Pareto efficiency with better fairness than bisecting approach. When the Pareto efficient enhancement is made by small improvements, E-DD approach produces an efficient and fair outcome only with 5.39735% difference ratio (*DiffRatio*) between two agents.

$$DiffRatio = \left| 100\% \cdot \frac{G_1}{G_1 + G_2} - 100\% \cdot \frac{G_2}{G_1 + G_2} \right|$$

where G_i is the utility gains of agent N_i from the Pareto efficient enhancement process. Obviously, the smaller the difference ratio is, the fairer the outcome is, or vice versa.

Please note that ϵ-Satisfying approach is not applicable in this case. Either with $x_2 = 0.6 + \epsilon$ or $x_2 = 0.6 - \epsilon$, there are more than one points ($\vec{x}_1^{\,1}$ and $\vec{x}_1^{\,1'}$) on agent N_1's indifference curve IC_1 that are indifferent to the initial point $(0.6, 0.6)$ (See Fig. 3). Hence it's unable to determine the search direction for Pareto frontier and the procedure is incapable of proceeding.

Table 1: Enhancement by complete optimization

Approach	Final Agreement	Opposition Angle	Joint Gain Division		Fairness
			Agent 1	Agent 2	(*DiffRatio*)
E-DD	$(0.475638, 0.593761)$	3.14	0.0673713	0.122723	29.1179%
Bisecting	$(0.514405, 0.549899)$	3.14	0.152533	0.0490046	51.3692%
ϵ-Satisfying	$(0.6, 0.6)$	2.89	NA	NA	NA

Table 2: Enhancement by $\lambda = 0.01$

Approach	Final Agreement	Opposition Angle	Joint Gain Division		Fairness
			Agent 1	Agent 2	(*DiffRatio*)
E-DD	$(0.486452, 0.58133)$	3.14	0.092032	0.102533	5.39735%
Bisecting	$(0.506961, 0.558305)$	3.14	0.136684	0.0636108	36.4829%
ϵ-Satisfying	$(0.6, 0.6)$	2.89	NA	NA	NA

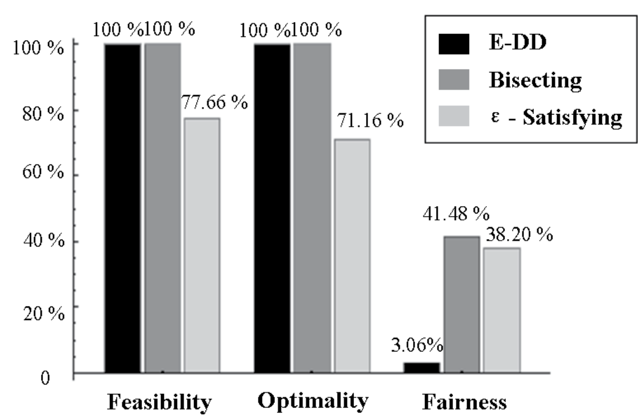

Figure 4: Overall performance of three approaches

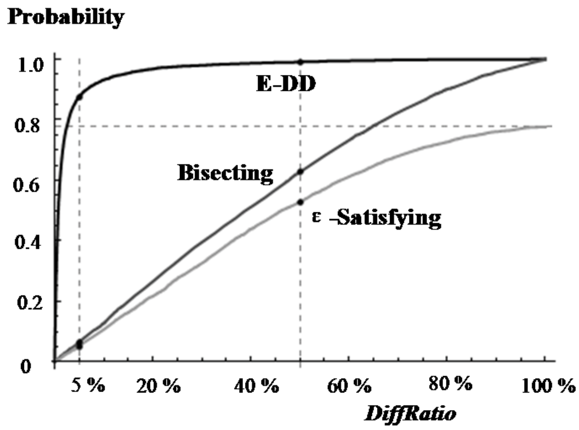

Figure 5: Cumulative Distribution of Difference Ratios

4.2 Quadratic Utility Function

This section describes the experiments with quadratic utility functions used in Lai et al. [13] and compares the overall performance of three approaches in the efficiency and fairness. In those experiments, the accuracy parameter ϵ and τ are still 0.0001 and moving distance $\lambda = 0.01$. The value of each attribute is normalized to $[0, 1]$ and the preferences of agents N_1 and N_2 is characterized by quadratic utility functions $U_1(x_1, x_2) = 1 - \sum_{j=1}^{2} w_{1j} x_j^2$ and $U_2(x_1, x_2) = 1 - \sum_{j=1}^{2} w_{2j}(1 - x_j)^2$ where w_{ij} is the weight that agent N_i puts on the attribute j, $\sum_{j=1}^{2} w_{ij} = 1,\ i \in \{1, 2\}$.

We run 5000 experiments with random preference combinations and random initial tentative agreements. The overall performance is evaluated by the following measures: *i*) the feasibility of the approach for the initial tentative agreements; *ii*) the optimality of the final outcomes *iii*) the difference ratio of the joint gains divisions.

On the one hand, feasibility is a fundamental characteristic of a possible solution: a solution is feasible if it can be applied to improve two agents' utility from the inefficient initial agreement. Moreover, since Pareto optimality is a central objective for most of negotiation systems, a good solution for the Pareto efficient enhancement should be able to help agents reach optimal final outcome. The first two series in bar chart Fig. 4 calculate the percentages of cases in which the approach is feasible, and in which the approach achieves optimal outcomes respectively. It could be seen that bisecting and the proposed E-DD approach are always feasible for all the initial tentative points meanwhile both of them guarantee Pareto optimality. However, in more than 20% of those experiments, ϵ-Satisfying could not be used and only in 71.16% cases it reaches optimal final outcomes.

On the other hand, joint gains are actually made by agents' cooperation behaviors, thus fairness is an important goal in the negotiation system. If the system is unfair, the weaker side would not be willing to cooperate. The series of Fairness in Fig 4 compares the average difference ratio among three approaches in those experiments. It displays the fact that the proposed E-DD approach produces the smallest average difference ratio (3.06%) between agents' utility gains. Bisecting approach is not as fair as intuitively expected. Its average difference ratio between two agents' utility gains is 41.48%, even higher than that of ϵ-Satisfying approach (38.20%). Additionally, Fig 5 represents the cumulative distributions of difference ratios. The most striking feature is that the difference ratios of E-DD approach are mostly controlled within 5.0% (the probability that the difference ratio of E-DD approach is smaller than 5% is nearly 0.9). While looking at the cases of bisecting and ϵ-Satisfying approaches, the difference ratios are relatively evenly distributed. The probability that the difference ratio is smaller than 5% of these two approaches are both less than 0.05. It should be noted that the maximum value of cumulative distribution function of ϵ-Satisfying approach is lower than 0.8. This is because, as we mentioned before, the approach is not applicable in more than 20% cases of those experiments.

5. CONCLUSION AND FUTURE WORK

In this paper, we have proposed a framework for Pareto efficient enhancement process, in which a non-biased mediator agent supports the negotiation agents in reaching an efficient and fairer agreement. The whole procedure of Pareto efficient enhancement is an iterative process of finding out the fair compromise direction and then computing a new tentative agreement at each stage of negotiation. We have presented an equal directional derivative approach, called E-DD, to search for the fair improvement direction which produces equal marginal gains for two agents at each tentative agreement point. It has been empirically evaluated that the proposed approach has three features: *i*) it is always feasible for the interior set of tentative agreements; *ii*) it guarantees optimal outcomes; *iii*) it produces fairer outcomes compared with other existing methods.

We are planning to extend this work to cope with multilateral negotiation and investigate the strategy of non-truthful agents.

6. REFERENCES

[1] Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet. Preference handling in combinatorial domains: From AI to social choice. *AI Magazine, Special Issue on Preferences*, 29(4):37–46, 2008.

[2] H. Ehtamo, R. Hamalainenen, P. Heiskanen, J. Teich, and M. V. S. Zionts. Generating Pareto solutions in a two party setting: Constraint proposal methods. *Management Science*, 45:1697–1709, 1999.

[3] H. Ehtamo, E. Kettunen, and R. Hamalainen. Searching for joint gains in multi-party negotiations. *European Journal of Operational Research*, 127(1):54–69, April 2001.

[4] H. Ehtamo, M. Verkamaand, and R. Hamalainen. How to select fair improving directions in a negotiation model over continuous issues. *IEEE Transactions on Systems, Man, and Cybernetics Part C: Applications and Reviews*, 29:26 – 33, 1999.

[5] P. Faratin, C. Sierra, and N. R. Jennings. Using similarity criteria to make negotiation trade-offs. In *ICMAS '00: Proceedings of the Fourth International Conference on MultiAgent Systems (ICMAS-2000)*, pages 119–126, Washington, DC, USA, 2000. IEEE Computer Society.

[6] S. Fatima, M. Wooldridge, and N. R. Jennings. An agenda-based framework for multi-issue negotiation. *Artif. Intell.*, 152(1):1–45, 2004.

[7] S. Fatima, M. Wooldridge, and N. R. Jennings. Optimal negotiation of multiple issues in incomplete information settings. In *AAMAS '04: Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 1080–1087, Washington, DC, USA, 2004. IEEE Computer Society.

[8] S. Fatima, M. Wooldridge, and N. R. Jennings. Approximate and online multi-issue negotiation. In *AAMAS '07: Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 1–8, New York, NY, USA, 2007. ACM.

[9] R. Fisher and W. Ury. *Getting to YES*. Houghton Mifflin Harcourt, 1991.

[10] R. J.Lin and S. cho T.Chou. Mediating a bilateral multi-issue negotiation. In *2003 IEEE International Conference on E-Commerce Technology (CEC'03)*, pages 76– 83, 2003.

[11] R. L. Keeney and H. Raiffa. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. Cambridge University Press, 1993.

[12] G. E. Kersten and S. J. Noronha. Rational agents, contract curves, and non-efficient compromises. *IEEE Transactions on Systems, Man, and Cybernetics*, 28(3):326 – 338, 1998.

[13] G. Lai, C. Li, and K. Sycara. Efficient multi-attribute negotiation with incomplete information. *Group Decision and Negotiation*, 15:511–528, 2006.

[14] D. Lax and J. Sebenius. The manager as negotiator: The negotiator's dilemma: Creating and claiming value. *in Goldberg, Stephen, Frank Sander and Nancy Rogers, eds. Dispute Resolution. 2nd ed. Boston, MA: Little, Brown and Co.*, pages 49–62, 1992.

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

[16] S. Saha and S. Sen. An effcient protocol for negotiation over multiple indivisible resources. In *IJCAI 2007: Proceedings of the Twentieth International Joint Conference on Artificial Intelligence*, pages 1494–1499, Hyderabad, India, 2007.

[17] Q. B. Vo and L. Padgham. Searching for joint gains in automated negotiations based on multi-criteria decision making theory. In *AAMAS '07: Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 1–8, New York, NY, USA, 2007. ACM.

[18] Q. B. Vo, L. Padgham, and L. Cavedon. Negotiating flexible agreements by combining distributive and integrative negotiation. *Intelligent Decision Technologies*, 1(1-2):33–47, 2007.