Autonomous Agents and Multiagent Systems manuscript No.
(will be inserted by the editor)

Automated Multiagent Negotiation on Multiple Issues with Private Information

Ronghuo Zheng · Nilanjan Chakraborty · Tinglong Dai · Katia Sycara

the date of receipt and acceptance should be inserted later

Abstract In this paper, we design negotiation strategies for non-mediated (i.e., totally decentralized) negotiation among multiple agents on multiple issues where the agents have no knowledge about the preferences of other agents. Most of the existing literature on non-mediated (i.e., totally decentralized) negotiation considers agents with either full information or with probabilistic beliefs about other agents' preferences on the issues. However, in reality, it is usually not possible for agents to have complete information about other agents' preferences or accurate probability distributions. We design a reactive negotiation strategy for general *multiagent multi-attribute negotiation*, where the agents have *non-linear utility functions and no information* about the utility functions of other agents. We prove that agents generating offers using our negotiation strategy reaches to an agreement acceptable to all the agents. We also prove that rational agents do not have any incentive to deviate from the proposed

R. Zheng
Tepper School of Business
Carnegie Mellon University
Pittsburgh, PA 15232
E-mail: ronghuoz@andrew.cmu.edu

N. Chakraborty
Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15232
E-mail: nilanjan@cs.cmu.edu

T. Dai
Tepper School of Business
Carnegie Mellon University
Pittsburgh, PA 15232
E-mail: dai@cmu.edu

K. Sycara
Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15232
E-mail: katia@cs.cmu.edu

strategy. We further present simulation results to demonstrate that on randomly generated problem instances, the negotiation solution obtained by using our strategy is quite close to the Nash bargaining solution.

Keywords Multiagent Negotiation · Distributed Decision Making · Decision Making under Uncertainty

1 Introduction

In multi-attribute negotiation, two or more parties (or agents), with limited common knowledge about each others' preferences, want to arrive at an agreement on a set of issues over which they have possible conflicting preferences [45,17,35–37,26]. During a negotiation, there are both cooperative and competitive objectives that drive the behavior of an agent. Any agreement is better than the best alternative to negotiated agreement (or BATNA), which implies that an agent would prefer to reach an agreement rather than the negotiation breaking down. On the other hand, she would like to reach an agreement that is most beneficial to herself. Over the years, different models with diverse objectives have been developed for negotiation. Some models try to explain human behavior in negotiation. Other models are for designing autonomous software agents for application in negotiation support systems and/or distributed decision making in autonomous multiagent systems. We study the negotiation problem from the perspective of designing systems of autonomous intelligent agents.

The extant literature on the mathematical study of negotiation can be divided into two broad classes, namely, mediated negotiation and non-mediated negotiation. In a mediated negotiation, the presence of a non-biased mediator is presumed and agents interact with each other through the mediator. In a non-mediated negotiation agents interact with each other directly. In this paper, we consider non-mediated negotiation. Within the non-mediated negotiation literature, researchers make different assumptions about the number of negotiating agents, the number of issues they are negotiating on, and the knowledge that an agent has regarding other agent's preferences (modeled using utility functions). Most work to date has focused on two-agent single issue negotiation, although there has been some work on two-agent, multi-issue negotiation (e.g., [15]) or with multiagent, single issue negotiation (e.g., [4]). Furthermore, computational modeling of multi-attribute negotiation has either assumed (a) complete knowledge of the preference structure of the opponents, i.e., the utility of the agents are assumed to be known, (e.g., [34]) or (b) a probability distribution over the preferences of the agents is known (e.g., [22,6,30]). Much of the literature also assume linear additive utility functions for the agents.

When the utility function is assumed to be linear and information about the opponent's utility function is known, monotonic concession strategies and Zeuthen strategies [11] have been proposed for negotiation. Rational strategies that correspond to sequential equilibrium of a game have been proposed for probabilistic knowledge about opponent (e.g., [13]). However, these strategies cannot be used if there is no knowledge about the opponents' utility functions and the utility functions are nonlinear. In general, a negotiation may involve multiple agents and multiple issues, the utility functions of the agents may be nonlinear and the agents may not have *any*

knowledge about the utility of the other players. A fundamental open question in multiagent negotiation in such a general setting, that we study in this paper, is the following: *Is it possible to design negotiation strategies for agents so that they provably come to an agreement given that they have no prior knowledge about the utility functions of other agents*?

More formally, we consider m (≥ 2) agents negotiating on a set of N (≥ 1) issues. We assume that the agents propose sequentially in a pre-specified order (that may be decided before negotiation begins). When there are two agents, the agents propose alternately. Each agent has a (strictly) concave nonlinear private utility function (known only to herself). Therefore, she can compute the utility of any offer to herself but cannot compute the utility of an offer to any other agent. Each agent also has a private reservation utility. A reservation utility is the lowest utility of offers an agent can accept in a negotiated agreement. Consequently, any offer with a utility less than the reservation utility is not acceptable to that agent. For agents with general utility functions, the utility of an offer is not simply a sum of the utilities of the individual issues. Therefore, we allow the agents to negotiate with package offers where agents negotiate on multiple issues simultaneously (instead of issue by issue negotiations). Although issue by issue offers are more convenient mathematically, packaged offers have the advantage that they allow agents to make trade-offs over different issues, which is a realistic feature of many negotiations [14]. Our goal is to design strategies for generating packaged offers for agents with private information that provably lead to an acceptable agreement for all the agents.

To illustrate the difficulty of generating an acceptable solution in a multi-issue negotiation, it is instructive to consider the geometry of a negotiation problem. Figure 1 gives a geometric view of the *offer space* for three agents (called A, B, C) negotiating on two-issues. The 3 curves, R_A, R_B, and R_C denote the set of all offers with value equal to the reservation utility of the agents A, B, and C respectively (i.e., their reservation curves). For each afent, the convex set bounded by the reservation curve is the feasible offer set and any offer within this set is an acceptable offer to an agent since it's utility is not less than the reservation utility. For example, $O_A^{'}$ is a feasible offer for agent A. The *zone of agreement* (the hatched region in Figure 1) is the set of offers that is acceptable to all agents. Any point within the zone of agreement is called a *satisficing agreement*. The objective of the agents is to find a satisficing agreement. Note that the agents do not know the other agents' utility functions, and therefore, the zone of agreement is unknown to the agents. Thus, geometrically speaking, in negotiation, the goal of the agents is to find a point in the zone of agreement, under the restriction that none of the agents have any explicit knowledge of the zone of agreement. Note that if the zone of agreement is empty, no agreement can be achieved.

Let us consider two agents negotiating on a single issue (e.g., a buyer and a seller negotiating on the price of a house). Here, if the zone of agreement is non-empty (i.e., if the lowest price at which the seller is willing to sell is less than the highest price the buyer is willing to pay), there will always be an agreement reached in the negotiation. This is because an offer of an agent with utility equal to her reservation utility is acceptable to the other agent. Even for multi-issue negotiation where the agents (with linear additive utility functions) negotiate issue by issue and have a different reserva-

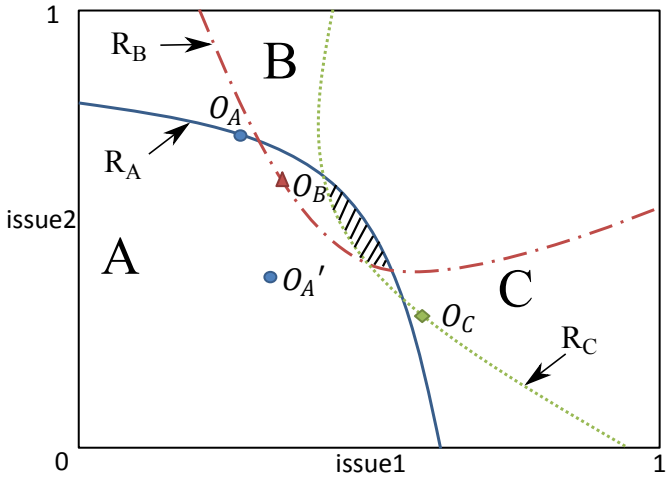

Fig. 1 Illustrative sketch of the offer space of 3 agents, A, B, and C, negotiating on 2 issues. The 3 curves, R_A, R_B, and R_C are the reservation curves of the agents A, B, and C respectively. All offers that lie on the reervation curve has utility equal to the reservation utility. The convex sets bounded by the three curves are the feasible offer sets (e.g., $O_A^{'}$ is a feasible offer for agent A). The zone of agreement is the common intersection of the three sets which is the hatched region.

tion price[1] for each issue (e.g., [13]), an agreement can be reached trivially by one of the parties proposing offers corresponding to her reservation price for each issue. However, for packaged multiagent multi-attribute negotiation with nonlinear utility functions, it is non-trivial for an agent to find an offer acceptable to other agents. Even if she makes an offer that is on her reservation utility, it may still not be acceptable to the other agent. Figure 1 shows that although the offers O_A, O_B, and O_C gives the agents A, B, C, their least possible utilities (i.e., they concede as much as they can), O_A, O_B, and O_C do not lie in the (unknown) zone of agreement and hence neither of them are acceptable to all the agents. In general, there are possibly infinite number of offers that lie on an agent's reservation utility that are unacceptable to the other agents. Thus, in contrast to single issue negotoiation, for multi-issue negotiation with private information and nonlinear utility function, it is non-trivial to generate offers acceptable to all agents.

Contributions: In this paper, we present a class of strategies called *sequential projection strategies* for generating offers and prove that the agents following such strategies will reach an agreement. The sequential projection strategies consists of two steps: (a) A concession step in which the agents reduce the utility of offers that are acceptable to them (unless they have reached their reservation utility) (b) An offer generation step in which the agents use the previous offers of their opponents to generate a new offer with utility equal to their current acceptable utility. Note that in step (a), although the agents will concede, the amount by which they concede (or

[1] Note that we use "price" here to be consistent with the literature [13].

the rule by which they decide on the amount to concede) is not specified. Thus, there is a degree of freedom in the choice of concession rule (or concession strategy). We show that the convergence holds for general concave utility functions irrespective of the specific concession strategy the agents adopt (as long as the agents concede up to their individual reservation utilities). This class of sequential projection strategies is a generalization of the alternate projection heuristic that was proposed in the literature for two agents [27,44].

The concession step is always present (either directly or indirectly) in non-mediated negotiation strategies where the agents take turns to make offers. It is assumed that the utility agents obtain from reaching an agreement decreases with time. There are two rationales given for this assumption. The first is that the value of an outcome may be time-sensitive and may decrease with time, The second is that in negotiations with deadlines, the agents must be willing to propose an offer with their reservation utility at the deadline, since any agreement is better than no agreement. Furthermore, the agents also want to maximize their own utility and thus they should start at their maximum utility and concede as time progresses. A common feature of the concession strategies in the literature [27,44] is that they are assumed to be properties of the agents themselves and not reactive to the concession strategies of the other agents. Hence it is not clear in those works whether the agents had any incentive to concede. Although in real life negotiations, negotiators do concede with time, the concession behavior is seen even in negotiations where the utility of issues do not decrease with time or there are no hard deadlines for negotiation. This is because negotiators do want to come to an agreement and know that if others realize that they are not conceding then there is a chance that the negotiation may stall and one party may walk out. Here, we design concession strategies for negotiation that conform to this intuition. In other words, we show that during negotiation, concession is rational even without hard deadlines or issues where utility does not decrease with time. We prove that if the agents use a *reactive* concession strategy, i.e., each agent concedes by an amount equal to her estimate of the amount of concession of her opponents, then the agents have no incentive to deviate from the concession strategy.

To the best of our knowledge, *for non-mediated negotiation, this is the first paper that provides negotiation strategies with guaranteed convergence to a satisficing solution for general multi-attribute, multilateral negotiation with agents having nonlinear utility functions and no knowledge about other agents' preferences*. We also demonstrate the performance of the reactive sequential projection strategy through simulations for two agent negotiation as well as for more general multiagent negotiation.

The remainder of this paper is organized as follows. In Section 2 we give an overview of the related literature. In Section 3 we outline the framework of automated negotiation and state the problem that we are studying more formally. In Section 4 we show our convergence proof of the the sequential projection strategy that we develop here. Thereafter, in Section 5 we prove that it is rational for the agents to use a reactive concession strategy in the absence of any knowledge about the utility functions of other agents. In Section 6, we present simulation results on randomly generated negotiation instances. Finally in Section 7, we summarize our contributions and outline avenues of future work.

2 Related Literature

The literature on negotiation can be divided into two broad categories, namely, mediated negotiation and non-mediated negotiation [43]. In mediated negotiation, it is assumed that there is an unbiased mediator who collects the agents preferences and proposes offers to the agents [27,5,31]. In non-mediated negotiation agents interact with each other and exchange their offers and preferences directly. In this paper, we are concerned with non-mediated negotiation and therefore we will restrict our literature review to papers that do not assume the presence of a mediator.

Earlier papers in negotiation or bargaining focused on two-agent single-issue negotiation. In a complete information setting, different axiomatic solution concepts were proposed, including, Nash bargaining solution [34], Kalai-Smorodinsky solution [25], egalitarian solution [24], pareto-optimal solution. However, the process by which agents should arrive at such a solution was not detailed. It is usually difficult for agents to arrive at the solutions mentioned above, especially in the absence of information about the other agents' preferences. In such settings the notion of a satisficing solution has been used (see [33] where the authors use the notion of a satisficing solution for solving a distributed planning problem). In our paper, we assume that agents have no information about other agents' preferences. Hence, we will be using a satisficing solution as our solution concept.

To model the iterative negotiation process, the alternating-offer game (or protocol) was proposed by Rubinstein [38]. The alternating offer protocol is one of the most popular negotiation protocols for bilateral single-issue setting. Work in economics using the framework of the alternating-offer game often focuses on the single issue problem. In the original alternating-offer game in [38], as well as subsequent literature (e.g., [32]), the two players (or agents) with complete information have incentive to concede because it is assumed that the utility of the negotiation outcome decreases with time. Transaction cost of bargaining is another reason for the players to concede (see [8]). Some studies (e.g., [10], [7]) consider outside options as an alternative incentive for the players (or agents) to concede over time. These studies also extend the alternating-offer game to the setting where the two players have incomplete and asymmetric information, i.e., they are uncertain about the opponent's type [18,10,9]. In [40], the authors prove that for single issue negotiation with deadline, the rational strategy is to wait for the deadline and make an offer corresponding to one's reservation utility. Note that in this case, there is always an agreement. In other words, in single issue negotiation continuous concession is not a rational strategy. In single issue negotiation with non-empty zone of agreement there is always an agreement and in the presence of hard deadline and private information waiting is the best strategy even for risk averse agents [40]. However, for the multi-attribute setting, one cannot guarantee that an agreement will be reached using the above strategy (as illustrated by Figure 1). In our setting, the agents have no information about the opponent's utility structure or type, and we do not assume that the utility of negotiated agreement decreases with time. In our setting, the players concede as a part of the search process to achieve a possible agreement in the absence of any information about the opponent's utility.

The literature using AI methods focuses on developing tractable heuristics for negotiating agents to generate offers. Although there is a large body of automated negotiation literature (see [43] for a comprehensive review), most prior work assumes either full information or commonly known random distributions. The alternating-offer game has also been extended to multi-agent or multi-issue negotiation (e.g., [14] , [15], [4]). They usually assume that there are two issues in the negotiation and that the agents utility functions are linear and additive on the values of the two issues (e.g. [2], [28]). In the presence of incomplete information, Bayesian learning has been proposed in agents' negotiation strategy [29]. A classification method for learning an opponent's preferences during a bilateral multi-issue negotiation using Bayesian techniques was developed in [39]. However, the Bayesian updating rule is only applicable when the agents are of certain set of types.

A common feature of the concession strategies in the literature is that they are assumed to be properties of the agents themselves and not reactive to the concession strategies of the other agents (notable exceptions being [1,41]). In [11], the necessity and difficulty of setting up a monotonic concession protocol is discussed. The author proposes several definitions of multilateral concession and analyzes their properties. However, all of these definitions are in the framework of complete information. In [41], the author proposes a reactive tit-for-tat negotiation strategy and [1] proposes an extension of contract net protocols to negotiations. Empirical work where reactive strategies for agents have been proposed include [3,16]. Experimental comparisons of human negotiation vs automated negotiation have been made in [16] and it was shown that reactive concession strategies performed much better than non-reactive strategies. Therefore, we present a reactive strategy for the agents to concede and prove that it is a rational strategy.

Mathematically, our work is most closely related to [23,27,44]. In [23] presents a constraint decentralized proposal method for multi-agent negotiations under the assumption of quasi-linear utility function with a neutral coordinator. Our paper, by contrast, *theoretically* addresses incomplete information and does not need the presence of a mediator. In [44], the authors consider more than two agents negotiating on multiple issues assuming neither prior information of agents nor linearly additive utility functions. However, they restrict their setting to be with three agents and two issues, and their analysis about the convergence is solely via simulations. In our work, we provide a general theoretical proof of convergence of the sequential projection protocol that is valid for any number of agents and issues. This paper extends our previous work [46] on bilateral multi-issue negotiation. A sucinct version of this work without the proofs and details appeared as an extended abstract in [47].

3 The Negotiation Framework

In this section, we will define the terminology that will be used throughout the paper and present a formal statement of the negotiation problem that we are studying. For designing negotiating agents, we need the following key elements: (a) Notion of solution to the negotiation problem (b) Negotiation protocol and (c) Negotiation strategy. The solution to a negotiation problem helps in defining the termination criterion of

the negotiation. The negotiation protocol defines when an agent should provide her offers to the other agents. The negotiation strategy gives a method for agents to decide on the offers that they want to propose.

We consider m self-interested agents $i \in \{1, 2, ..., m\}$, negotiating on a set of issues $j \in \{1, 2, ..., N\}$. Let $[0, 1]$ denote the unit interval in $\mathbb{R}$ and $[0, 1]^N$ be the unit hypercube in $\mathbb{R}^N$. We assume that the issues take on continuous values and the negotiation domain for each issue is $\Omega_j = [0, 1]$ with 0 and 1 corresponding to the extreme values of the issues. Any point within the unit hypercube is a *package offer* or simply an *offer*. We assume that the utility function of agent i, $u_i(x)$, $i = 1, 2, ..., m$, is continuous and concave $\forall x \in [0, 1]^N$. Without loss of generality, we can normalize the range of the agent i's utility function to $[0, 1]$. Each agent, i, has a *reservation utility*, ru_i. Any offer with utility less than its reservation utility is not acceptable to that agent. The set of all feasible offers that an agent i can accept is $A^i = \{x \in [0, 1]^N | u_i(x) \geq ru_i\}$. The set A^i is strictly convex for each i. The *zone of agreement*, $\mathcal{Z}$, is defined as the common intersection of the feasible offer sets of all agents, i.e., $\mathcal{Z} = \cap_{i=1}^{m} A^i$. Since the zone of agreement is the intersection of a finite number of convex sets, it is a convex set. Any offer (i.e., point) within the zone of agreement is acceptable to all the agents. From the above definitions, it follows that for a solution to exist to any negotiation problem, the zone of agreement has to be non-empty. Any point within the zone of agreement is acceptable to every agent and we call such a solution a *satisficing solution* to the negotiation. Note that the zone of agreement is fixed by the utility functions and reservation utilities and cannot change during a negotiation.

There has been different definitions proposed for a *proper* negotiation solution. Axiomatic solution concepts has been proposed for bargaining games (e.g., Nash bargaining solution [34], Kalai-Smorodinsky solution [25], egalitarian solution [24], pareto-optimal solution). The set of points that satisfy these different solution requirements are all subsets of the zone of agreement. However, computing them requires that all the agents know each others utility functions. Since an agent does not know the utility function of her opponent, we use a satisficing solution as our solution concept. A satisficing solution is any agreement that gives the negotiators a utility greater than or equal to their reservation utility. The use of a satisficing solution in this very general setting where the agents have no information about their opponents is in the spirit of Herbert Simon [42].

A key issue in designing negotiating software agents is to choose a protocol for negotiation. For two agent negotiation, we will assume that the agents use an *alternating-offer protocol* [38], where an agent proposes its offer and the other agent responds to the offer by accepting it or proposing a new offer. For general multi-agent negotiation, we will use a generalization of the alternating offer protocol, namely, a *sequential-offer protocol*. In a sequential-offer protocol, each agent proposes an offer in a fixed sequence. We use a sequential protocol in the multi-agent setting and assume that the agents propose their offers in a given order. An agent computes her own offer using the latest offers of all the other agents (including herself) and either proposes a new offer or accepts the current offer, if the offers made by the previous agent is within her acceptable offer set. When all agents accept the current offer the negotiation ends. Given the negotiation protocol, the problem in designing a negotiat-

ing agent is to compute a strategy for generating offers, which is stated more formally below.

Problem Statement: *Given m agents negotiating on n issues where (a) each agent, i, has a strictly concave private utility function, u_i, and (b) the zone of agreement has a nonempty interior, find a method for computing the offer an agent should propose such that it is guaranteed that the agents will follow the offer generation method and eventually reach an agreement.*

3.1 Overview of Solution Approach

Informally speaking, a negotiating agent not only wants to reach an agreement with the other agent but also may want to obtain as much utility as possible. Thus, when agents start out in a negotiation, they want to propose offers that have the highest utility for them and gradually move towards offers with lower utility. However, they will neither propose nor accept any offer with utility lower than their reservation utility. This intuition implies that, during negotiation, agents gradually reduce the utility of offers acceptable to them (which is very often seen in practice). Thus agents use a *concession strategy* to determine their current utility at time t (denoted by $s_i(t)$). This concession continues until an agent reaches her reservation utility. In other words, $s_i(t)$ is a *monotonically decreasing* function of t and $s_i(t) \geq ru_i, \forall t$.

In the literature the concession strategy is usually assumed to be non-reactive. It is assumed that there is a decay parameter for the utility and this decay parameter is just a property of the agent and does not depend on the other agents concession behavior (e.g., [12]). In contrast, we assume a reactive strategy of concession, where an agent concedes according to her perception of how much the other agents present in the negotiation concede. For agent i, let A_t^i be the set of all offers that have utilities higher than $s_i(t)$ at time t. The set, $A_t^i = \{x \in [0,1]^n \,|u_i(x) \geq s_i(t)\}$, is called the *current feasible offer set* of agent i. For all t, A_t^i is a convex set and $A_1^i \subseteq A_2^i \subseteq \cdots \subset A^i$. The boundary of the set A_t^i is called the *indifference surface (or curve)* of agent i at time t.

Agent Strategy: The negotiation strategy that we will use consists of two steps [27]. When it is the turn of agent i to make an offer, she accepts the current offer if it is satisficing. Otherwise, agent i uses the last two offers of every other agent j to compute the difference in the utilities of the offers to her (we denote this by Δu_{ij}. The amount by which agent i reduces her utility to compute her current utility is equal to the minimum Δu_{ij} computed over all the other agents j. She then generates an offer on the indifference surface corresponding to her current utility by projecting the mean of all the other agents latest offers to her current indifference surface. Note that this method generates an offer that is satisficing to the agent and closest (in terms of Euclidean distance) to an *average offer* made by the other agents.

In the next section, we will first present the sequential projection method for computing offers for an agent and then give a convergence proof for the method. For presentation purposes, we will first assume that agents use some strategy for concession without making any assumptions of what the strategy is. We will show that under this very general assumption the sequential projection method ensures that the

agents converge to an agreement. We will then restrict our attention to the class of concessions strategies that ensures that the agents reach their reservation utility in finite time. We will show that in these conditions, the agents can reach an agreement in finite time. Finally, we will show that agents using a reactive concession strategy is a rational strategy since it allows the agents to determine if any of the other agents stop conceding. Although the results in this paper on multiagent negotiation is applicable to two agent negotiation, we will discuss the two agent negotiation separately since this problem is important by itself. For the presentation of the proofs we will present the proofs for the general multiagent case.

4 Offer Generation Method

We now present the sequential projection method for an agent to generate her offer, given the latest offers of all the other agents. As stated before, the agents make their proposals in a fixed pre-determined sequence. We assume at period $t = 0$, each of the agents propose an offer maximizing her own utility. After initialization, the agents proposes sequentially, such that at time $t = 1$, agent 1 proposes, at time $t = 2$, agent 2 proposes and so on. Let x_t^j be the *standing offer* (i.e., the last offer the agent j made) of agent j in period t. let $P_A[x]$ be the projection of point x on the set A with P being the projection operator. If a proposal by an agent i is not the same as the standing offers made by all the other agents in period t, agent $i + 1$ proposes her own offer at time $t + 1$. *The agent determines her offer by projecting the convex combination of all of the agents' standing offers to her current indifference surface in period* $t + 1$. More specifically, agent $i + 1$ at period $t + 1$ proposes

$$x_{t+1}^{i+1} = P_{A_{t+1}^{i+1}}\Big[\sum_{j=1}^{m} a_t^{i,j} x_t^j\Big] \tag{1}$$

where A_{t+1}^{i+1} is the set of acceptable offers for agent $i + 1$ at time $t + 1$, the weight that agent i puts on the standing offer of agent j is $a_t^{i,j}$ and $\sum_{j=1}^{m} a_t^{i,j} = 1$.

In order to understand the working of the offer generation method, let us first look at the two-agent case. In this case, as stated below, we can choose the weights $a_t^{i,j}$ such that the agent uses only the last offer of her opponent to determine her offer. This method was presented in [27] and we call it the alternating projection method. Figure 2 presents an example of the alternating projection proposing method for two agents negotiating on two issues. In this example, the solid indifference curves belong to agent 1 and the dashed indifference curves belong to agent 2. In period $t-4$, agent 1 proposes an offer x_{t-4}^1. Agent 2 rejects this offer and both agents update their indifference curves. In period $t - 3$, agent 2 identifies x_{t-3}^2 on her indifference curve such that x_{t-3}^2 is the projection of x_{t-4}^1 to her indifference curve. Agent 2 offers x_{t-3}^2, agent 1 rejects this, and both agents update her indifference curve. In period $t - 2$, agent 1 identifies x_{t-2}^1 by projection of x_{t-3}^2 to her current indifference curve and proposes it to agent 2. The process continues until an offer is accepted or the deadline is reached.

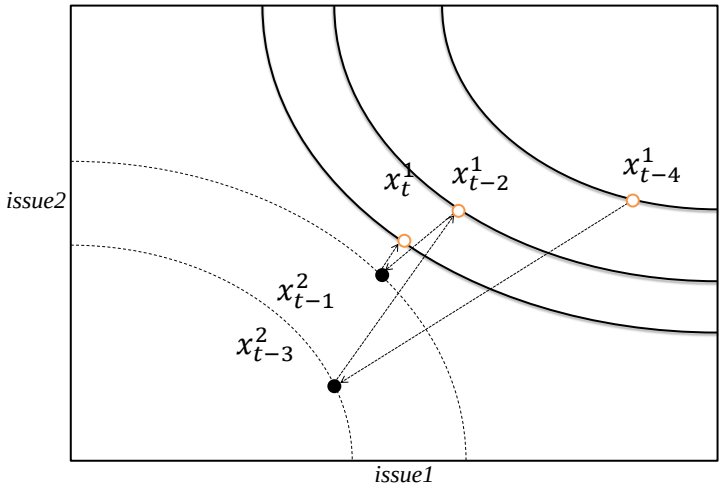

Fig. 2 The alternating projection protocol for two issues and two agents

For general multiagent negotiation, to simplify the notations and presentation, we assume $a_t^{i,j} = 1/m$ for $i, j \in \{1, 2, \cdots, m\}, t \geq 0$ (although the discussion and results below holds for general values of $a_t^{i,j}$ satisfying $\sum a_t^{i,j} = 1$). Therefore, agent $i{+}1$ at period $t{+}1$ proposes an offer according to $x_{t+1}^{i+1} = P_{A_{t+1}^{i+1}}\left[\frac{1}{m}\sum_{j=1}^m x_t^j\right]$, where x_t^j is the newest offer proposed by agent j until period t. For notational convenience, we define $w_t := \frac{1}{m}\sum_{j=1}^m x_t^j$. Figure 3 illustrates the sequential projection method for three agents negotiating on two issues. At time $t-1$, it is agent 1's turn to make an offer. The standing offers at period $t-1$ from agent 2 is $x_{t-1}^2 = x_{t-3}^2$ since agent 2's previous proposal was at time $t-3$. Similarly, the standing offer of agent 3 is $x_{t-1}^3 = x_{t-2}^3$. Agent 1 combines the most recent offers of every agent, i.e., $x_{t-4}^1, x_{t-3}^2, x_{t-2}^3$ to compute the point w_{t-2} and project onto her indifference surface at $t-1$ (the solid curve which is obtained using her concession strategy). Similarly, at time t, it is agent 2's turn to offer and she computes her own offer by projecting the point w_{t-1} (computed by averaging $x_{t-1}^1, x_{t-3}^2, x_{t-2}^3$) and the negotiation proceeds.

4.1 Convergence of Offer Generation Method

In this section, we prove that if the agents follow a sequential projection strategy along with a concession strategy for offer generation the agents can converge to an agreement if the zone of agreement is non-empty. For multi-issue negotiation with private utility function, the agents don't know the non-empty zone of agreement, even if one exists. Therefore, the existence of non-empty zone of agreement cannot guarantee that agents will reach an agreement, even if they are given enough time. Thus, we need to examine whether the negotiation strategy is convergent or not. The

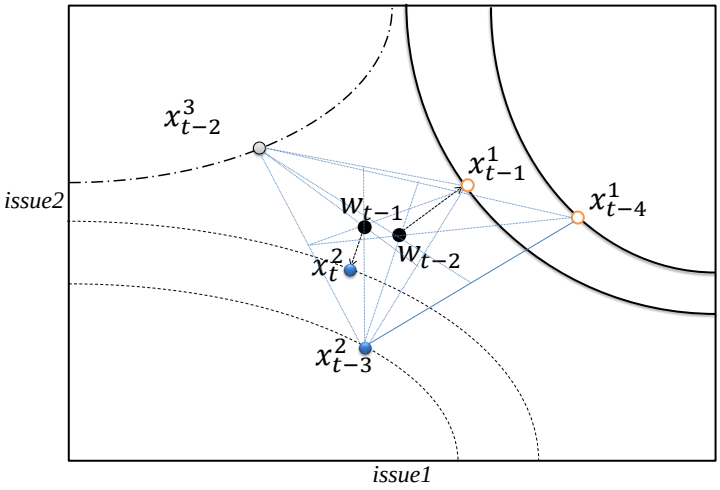

Fig. 3 The sequential projection method for three agents negotiating on two issues.

convergence results are valid for any concession strategy and we will not make any assumption about the rate of concession. This implies that our results hold even if each agent in the negotiation can have a different concession strategy. Since the alternating projection strategy for two agents is a special case of the multiagent sequential projection strategy, it also implies that the alternating projection strategy for two agent negotiation is a provably convergent strategy.

We will now state the main convergence theorems and lemmas required to prove the results. The proofs of all the theorems and the lemmas are presented in Appendix A. We will first need a classical result from convex geometry [].

Lemma 1 *Let A be a nonempty closed convex set in $[0,1]^N$. Then for any $x \in [0,1]^N$, $y \in A$, we have the following:*

$$(P_A[x]-y)'(y-x) \leq -\|P_A[x]-y\|^2, \tag{2}$$

$$\|P_A[x]-x\|^2 \leq \|x-y\|^2 - \|P_A[x]-y\|^2. \tag{3}$$

Proof See Appendix A.

Lemma 1 is a classical result from the convexity of the set that will be used for proving our main results. Theorem 1 states that the offer generation method ensures that the distance between the new offer generated by an agent and the mean of all the previous offers never increases. This fact is used in proving our main claim in Theorem 2

Theorem 1 *Let x^i_t be the offer of agent i at time t and w_t be mean of the standing offers at time t from all agents. Then the sequence $\{\sum_{i=1}^m \left\|x^i_t - w_t\right\|^2\}$ is monotonically non-increasing with t.*

Proof See Appendix A.

Theorem 2 *If the zone of agreement has a non-empty interior, the sequential projection proposing protocol will always converge to an agreement.*

Proof See Appendix A.

4.2 Finite Time Convergence of Offer Generation Method

In the previous subsection, we did not make any assumptions about the time an agent takes to concede to the reservation utility. In theory, an agent can adopt a concession strategy such that the reservation utility is reached asymptotically (or in infinite time). However, for practical purposes, we can assume that agents will use a concession strategy such that they reach their reservation utility in finite time. We now study the finite time convergence properties of the sequential projection method for offer generation.

To understand finite time convergence properties of a method for computing offers, we study the following question: Given that the concession strategies of the agents are such that all the agents reach their reservation utilities in finite time, say T_0, do the agents converge to an agreement in finite time (provided the zone of agreement has a non-empty interior)? Note that here also, we do not make any assumptions about the specific concession strategy used by the agents. The answer to this question is yes in general and we prove this in Theorem 3 below.

Theorem 3 *For m agents negotiating on N issues, if the agents use concession strategies such that they reach their reservation utility in finite time, then they can reach an agreement in finite time (assuming that the zone of agreement has a non-empty interior).*

Proof See Appendix A.

Theorem 3 is an immediate result from Theorem 1 and Theorem 2. Intuitively speaking, as we can guarantee the convergence of the sequence $\{\sum_{i=1}^{m} \left\|x_t^i - w_t\right\|^2\}$ using the fact that the sequence is non-increasing (by Theorem 1) and that the sequence has finite sum (shown in Theorem 2) , we have $\forall \varepsilon > 0$, $\exists T > 0$, s.t., $\forall t > T$, $\sum_{i=1}^{m} \left\|x_t^i - w_t\right\|^2 < \varepsilon$. In words, the distance between the offers generated by an agent and the mean of the current offers of all the agents will decrease to zero (within a numerical error tolerance ε) in finite time.

5 Incentive of agents to concede

In the previous sections, we described a class of convergent negotiation strategies for automated agents. We proved that if agents use any strategy from the class of conceding sequential projection strategies then it is guaranteed that the agents will reach an agreement provided the zone of agreement is non-empty. As stated earlier, any

concession strategy in the class consists of determining the utility of the offer that an agent should propose and choosing one offer from all the alternatives with the same utility. We first note that if an agent i concedes, there is no incentive for her to propose an offer on the indifference surface that is not the projection of the convex combination of other agents' offers. This is because all points on the indifference surface of agent i have the same utility and by proposing another point she may decrease the chance of reaching an agreement (since the convergence proof holds only for projections). However, it is not clear what should be the choice of the concession rate, i.e., a choice of concession rate that doesn't leave the agent vulnerable to be exploited by other agents who may be using different concession rates. Here we show that if the agents use a reactive concession strategy, then the agents are not vulnerable to exploitation.

We prove that there is a *reactive concession strategy*, namely, conceding by an amount equal to the minimum of the perceived change in utility of the other agents' offers that is rational. More precisely, we prove that if any of the agents do not concede, it is possible for the other agents to determine this within a finite number of rounds and hence stop conceding. This combined with the fact that an agent does not know other agents' utility provides the threat of the negotiation coming to a stall, even if the zone of agreement is non-empty. Since the utility of a negotiated agreement is not worse than the utility for breakdown, it is rational for an agent to concede.

We now prove that if agent i stops conceding, and all other agents use a reactive strategy, the negotiation can stall. As shown in Figure 4, let agent i propose x_t^i at

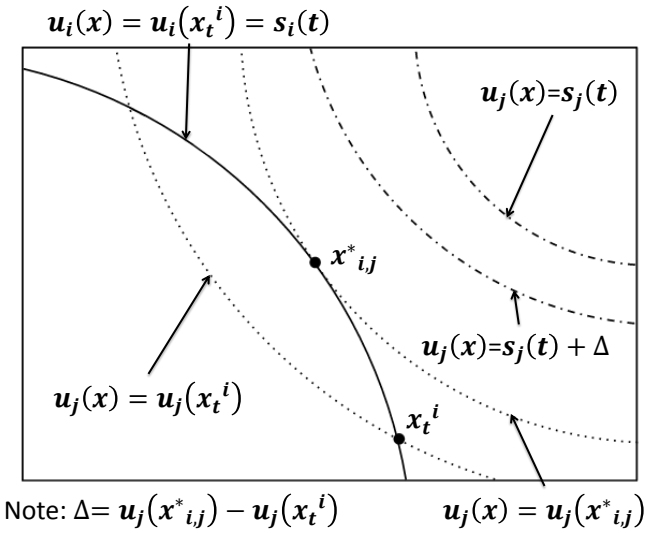

Fig. 4 Figure for proving that there is an incentive to concede

time t. If agent i stops to concede from time t, all offers proposed by agent i after time t are on the indifference curve $u_i(x) = u_i(x_t^i) = s_i(t)$. Let $x_{i,j}^*$ be the point

on the indifference surface $u_i(x) = s_i(t)$ such that $u_j(x) = u_j\left(x^*_{i,j}\right)$ is the highest possible perceived utility by agent j. Therefore, without loss of generality, the minimum of the perceived utility improvement from other agents' offers (including agent i) for agent j would be smaller or equal to $\Delta_j = u_j\left(x^*_{i,j}\right) - u_j\left(x^i_t\right)$. Thus,if agent i stops conceding, using the reactive strategy agent j would concede by no more than Δ_j. If $\Delta_j < s_j(t) - u_j\left(x^*_{i,j}\right)$, where $s_j(t)$ is the current utility level of agent j at time t, the negotiation will stall (see Figure 4 for the two-issue case), as there would be no agreement between agent i and agent j. Since agent i has no knowledge about the utility function of other agents, she is uncertain about whether $\forall j,\ j \in \{1, 2, \cdots, m\},\ j \neq i$, the largest possible perceived utility improvement from agent i's offers, $\Delta_j = u_j\left(x^*_{i,j}\right) - u_j\left(x^i_t\right)$, is larger than $s_j(t) - u_j\left(x^*_{i,j}\right)$. Thus, agent i is not sure about whether there will be an agreement or not if she stops to concede at any time t. Since an agreement would provide higher utility than her reserved utility for no agreement, agent i would not stop conceding. Therefore, all of the agents would keep conceding through the negotiation process. The above discussion shows that concession is rational for agents even in negotiations without any hard deadlines or when the utility of negotiated outcome does not decrease with time.

6 Simulation Results

In this section we present simulation results depicting the practical performance of our reactive negotiation strategy. We evaluate our solution with respect to the Nash bargaining solution [34]. An alternative is to evaluate the solution by measuring its distance from the Pareto optimal solution set. However, in higher dimensional settings that appear in general multilateral multi-attribute negotiation, there are no known efficient algorithms to compute the Pareto optimal set. Hence, we use the Nash solution, which is also Pareto optimal. The Nash solution maximizes the joint utility (i.e., the product of the utilities) of the agents. For the class of (strictly) concave utility functions that we consider, the Nash solution can be obtained by solving a convex optimization problem and hence we can easily find this solution irrespective of the number of negotiating agents or number of negotiation issues.

For the utility function of the agents we have assumed a very general *hyperquadric* function [21]

$$u_k(x) = 1 - \sum_{i=1}^{Q} |H_i(x)|^{n_i},$$

where x is the n-dimensional proposal vector, $H_i(x) = \sum_{j=1}^{N} a_{ij}x_j, n_i = l_i/m_i, l_i, m_i \in \mathbb{Z}^+$; $f(x)$ is strictly concave if $1 < n_i < \infty$. Hyperquadrics are a very general class of functions used in computer graphics [21] and can model a wide range of convex functions. The feasible set of offers for an agent k at time t is the intersection of the unit n-dimensional hypercube $[0, 1]^N$ with $u_k(x) \geq s_k(t)$. Popular convex functions for modeling utilities in economics like the Cobb-Douglas functions can be shown to be special cases of the hyperquadric function. The sole reason for using this function

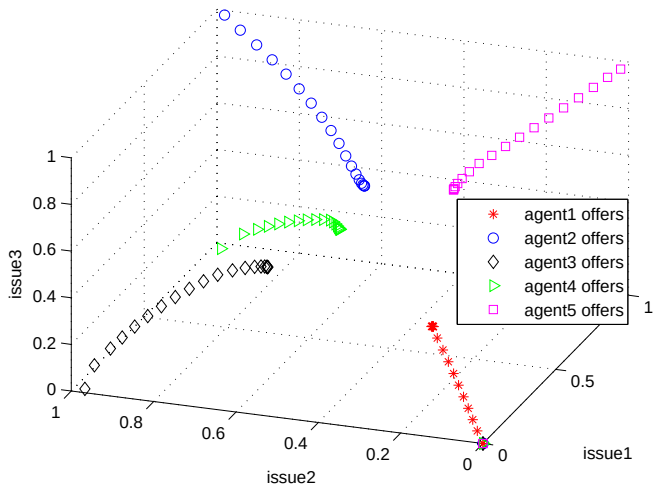

Fig. 5 Sequence of offers made by 5 agents in a three-issue negotiation scenario when agent 1 stops conceding. The agents use sequential projection along with reactive concession for generating offers. The agents come to a deadlock and no agreement is obtained although in this case the zone of agreement in non-empty (to avoid clutter, the zone of agreement is not shown here). This shows that an agent may not reach an agreement at all if she does not concede.

is that it is possible to generate a wide variety of preference structures for the agents with these functions.

The perceived utility of agent i for agent$j$'s offer, x_j is $u_i(x_j)$. Thus the minimum perceived utility change from all other agents' offers by agent i at round t is

$$\Delta u_i(t) = \min_j \left| u_i(x_j^{[-1]}) - u_i(x_j^{[-2]}) \right| \tag{4}$$

where $x_j^{[-1]}$ and $x_j^{[-2]}$ denote the last two offers proposed by agent j. The current utility of agent i in round t is given by $s_i(t) = s_i(t-N) - \Delta u_i(t)$, where $s_i(t)$ is the acceptable utility of agent i at time t. Figure 5 shows a simulation where the agent 1 stops conceding after reaching half of its reservation utility. Since all other agents are reactive, they realize within a few steps that agent 1 is not conceding and they also stop conceding. Hence the agents do not reach an agreement, as the concession of the agents stop, although their zone of agreement is nonempty. Thus, in the results that follow, we do not allow any agents to stop conceding, since that may result in no agreement.

Table 1 shows the performance of the algorithm when the adjusted time-dependent concession strategy is used. The reservation utility of the agents are assumed to be 0.2 here. The number of agents are varied between 2 and 9. The results are averaged over 100 random runs for each row of the table. The numerical tolerance used for convergence is 0.001. As can be seen from Table 1 (second and third columns), the number of rounds required for convergence are fairly constant. The fourth column gives the ratio of our solution to the Nash solution. For randomly generated instances,

Table 1 Performance of the sequential projection algorithm with reactive concession strategy.

Number of agents	Number of rounds		Ratio of Joint Utility	
	Mean	SD	Mean	SD
2	72.07	10.30	0.9278	0.0874
3	86.39	15.57	0.8825	0.0865
5	94.68	7.49	0.8819	0.0851
7	95.46	6.74	0.8846	0.0740
9	96.72	5.17	0.9023	0.0691

Table 2 Performance of the sequential projection algorithm with 9 agents negotiating on different number of issues.

Number of issues	Number of rounds		Ratio of Joint Utility	
	Mean	SD	Mean	SD
2	79.00	4.67	0.9162	0.0653
3	89.00	4.06	0.9267	0.0247
4	91.56	5.46	0.9377	0.0348
5	94.60	4.16	0.9014	0.0539

Table 3 Performance of the sequential projection algorithm with 5 agents negotiating on 5 issues with different reservation utilities.

Reservation utilities	Number of rounds		Ratio of Joint Utility	
	Mean	SD	Mean	SD
0.1	77.70	6.60	0.8546	0.0717
0.2	79.60	5.02	0.8938	0.0570
0.3	84.20	5.94	0.8993	0.0474
0.4	80.30	5.62	0.8785	0.0551

the solution obtained is quite close to the Nash bargaining solution (fourth and fifth columns).

In order to check the robustness of the algorithm, we check the sensitivity of the algorithm to the number of issues and the reservation utility of the agents. Table 2 shows the performance of the algorithm for 9 agents negotiating on different number of issues. The results in Table 2 shows the number of rounds is quite stable if we increase the number of issues. Table 3 shows the performance of the algorithm for 5 agents negotiating on 5 issues with different reservation utilities. The results in Table 3 shows the number of rounds is quite stable if we increase the reservation utilities of the agents as long as the zone of agreement is non-empty.

It should be noted that during each round, each agent simply computes the closest point on a convex set to a given convex set. This computation is usually very fast and can be done at a rate of the order of 1000 computations per second for less than 10 issues on current processors. For 9 agents it takes less than a second to reach the agreement. Thus, the negotiation strategy is computationally simple and outcomes can be obtained without much of a delay.

7 Concluding Remarks

In this paper, we propose a class of sequential projection strategies for general multilateral multi-attribute negotiation where agents have no knowledge about the other players' utility functions. We prove that if all agents use any strategy from this class they are guaranteed to arrive at an agreement. Further, a sequential projection strategy with a reactive concession function is a rational strategy for the agents. The convergence guarantees hold for any nonlinear concave utility function. We also performed computational experiments to demonstrate that, in practice, the quality of solution obtained by our algorithm is quite close to the Nash bargaining solution (that maximizes the joint utility of the agents). The negotiation converges in a reasonable number of iterations and scales well as the number of agents or number of issues are increased.

This work can be extended in several directions. One direction is to design rational strategies for agents to negotiate in the presence of hard deadlines. Another possibility is to extend this sequential projection method to negotiation between multiple negotiation teams. At present our agents are myopic in nature and do not try to learn the other agents utility function from the sequence of offers. It would also be interesting to investigate whether the agents can be incorporated with some learning capability so that they converge faster to an agreement or to a more efficient solution.

Acknowledgment

This research was made possible by generous funding from The Army Research Institute ARO grant W911NF-08-1-0301.

References

1. Aknine, S., Pinson, S., Shakun, M.F.: An extended multi-agent negotiation protocol. Autonomous Agents and Multi-Agent Systems **8**(1), 5–45 (2004)
2. Bac, M., Raff, H.: Issue-by-issue negotiations: the role of information and time preference. Games Econ Behav **13**(1), 125–134 (1996)
3. Bartos, O.J.: Simple model of negotiation: A sociological point of view. Journal of Conflict Resolution **21**(4), 565–579 (1977)
4. Binmore, K.: Bargaining and coalitions. Game-Theoretic Models of Bargaining pp. 269–304 (1985)
5. Chalamish, M., Kraus, S.: Automed - an automated mediator for multi-issue bilateral negotiations. Journal of Autonomous Agents and Multi-Agent System (2011)
6. Chatterjee, K., Samuelson, W.: Bargaining under incomplete information. Operations Reasearch **31**(5), 835–851 (1983)
7. Compte, O., Jehiel, P.: Gradualism in bargaining and contribution games. Review of Economic Studies **71**, 975–1000 (2004)
8. Cramton, P.: Dynamic bargaining with transaction costs. Management Science **37**, 1221–1233 (1991)
9. Cramton, P.C.: Strategic delay in bargaining with two-sided uncertainty. Review of Economic Studies **59**(1), 205–225 (1992)
10. D. Fudenberg, D.L., Tirole, J.: Incomplete information bargaining with outside opportunities. Quarterly Journal of Economics **102**, 37–50 (1987)
11. Endriss, U.: Monotonic concession protocols for multilateral negotiation. In: Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems, AAMAS '06, pp. 392–399 (2006)
12. Faratin, P., Sierra, C., Jennings, N.R.: Negotiation decision functions for autonomous agents. Int. Journal of Robotics and Autonomous Systems **24**(3-4), 159–182 (1998)

13. Fatima, S., Wooldridge, M., Jennings, N.: An agenda-based framework for multi-issue negotiation. Artif Intell **152**, 1–45 (2004)
14. Fatima, S., Wooldridge, M., Jennings, N.: Optimal negotiation of multiple issues in incomplete information settings. Proceedings of the 3rd International Conference on Autonomous Agents and Multi-agent Systems pp. 1080–1087 (2004)
15. Fatima, S., Wooldridge, M., Jennings, N.: Multi-issue negotiation with deadlines. Journal of Artificial Intelligence Research **27**, 381–417 (2006)
16. Filzmoser, M.: Automated vs. human negotiation. International Journal of Artificial Intelligence **4**, 64–77 (2010)
17. Fisher, R., Ury, W.: Getting to yes: Negotiating agreement without giving in. Harvard University Press, Boston, MA (1981)
18. Fudenberg, D., Tirole, J.: Sequential bargaining with incomplete information. Review of Economic Studies **50**, 221–247 (1983)
19. Grant, M., Boyd, S.: Graph implementations for nonsmooth convex programs. In: V. Blondel, S. Boyd, H. Kimura (eds.) Recent Advances in Learning and Control, Lecture Notes in Control and Information Sciences, pp. 95–110. Springer-Verlag Limited (2008)
20. Grant, M., Boyd, S.: CVX: Matlab software for disciplined convex programming, version 1.21 (2011)
21. Hanson, A.J.: Hyperquadrics: smoothly deformable shapes with convex polyhedral bounds. Computer Vision, Graphics, and Image Processing **44**(2), 191–210 (1988)
22. Harsanyi, J., Selten, R.: Two-person bargaining games with incomplete information. Manage Sci **18**(5), 80–106 (1972)
23. Heiskanen, P., Ehtamo, H., Hamalainen, R.: Constraint proposal method for computing pareto solutions in multi-party negotiations. Eur J Opera Res **133**, 44–61 (2001)
24. Kalai, E.: Proportional solutions to bargaining situations: Intertemporal utility comparisons. Econometrica **45**(7), 1623–1630 (1977)
25. Kalai, E., Smorodinsky, M.: Other solutions to nash's bargaining problem. Econometrica **43**(3), 513–518 (1975)
26. Kraus, S.: Strategic negotiation in multi-agent environments. The MIT Press, Cambridge, MA (1982)
27. Lai, G., Sycara, K.: A generic framework for automated multi-attribute negotiation. Group Decis Negot **18**, 169–187 (2009)
28. Lang, K., Rosenthal, R.: Bargaining piecemeal or all at once. Econ J **111**, 526–540 (2001)
29. Li, C., Tesauro, G.: A strategic decision model for multi-attribute bilateral negotiation with alternating offers. In: Proceedings of the ACM Conference on Electronic Commerce pp. 208–209 (2003)
30. Lin, R., Kraus, S., Wilkenfeld, J., Barry, J.: Negotiating with bounded rational agents in environments with incomplete information using an automated agent. Artificial Intelligence **172**(6-7), 823–851 (2008)
31. Lopez-Carmona, M.A., Marsa-Maestre, I., Klein, M., Ito, T.: Addressing stability issues in mediated complex contract negotiations for constraint-based, non-monotonic utility spaces. Autonomous Agents and Multi-Agent Systems **24**(3), 485–535 (2012)
32. Merlo, A., Wilson, C.: A stochastic model of sequential bargaining with complete information. Econometrica: Journal of the Econometric Society pp. 371–399 (1995)
33. Moehlman, T.A., Lesser, V.R., Buteau, B.L.: Decentralized negotiation: An approach to the distributed planning problem. Group Decision and Negotiation **1**, 161–191 (1992)
34. Nash, J.: The bargaining problem. Econometrica **18**, 155–162 (1950)
35. Pruitt, D.G.: Negotiation Behavior. Academic Press (1981)
36. Raiffa, H.: The art and science of negotiation. Harvard University Press, Cambridge, MA (1982)
37. Rosenschein, J.S., Zlotkin, G.: Rules of Encounter - Designing Conventions for Automated Negotiation among Computers. MIT Press (1994)
38. Rubinstein, A.: Perfect equilibrium in a bargaining model. Econometrica **50**, 97–109 (1982)
39. S, S.B., Spencer, B.: Learning opponents' preferences in multi-object automated negotiation. Published at the Seventh International Conference on Electronic Commerce (ICEC'05) (2005)
40. Sandholm, T., Vulkan, N.: Bargaining with deadlines. In: AAAI/IAAI, pp. 44–51 (1999)
41. Shakun, M.F.: Multi-bilateral multi-issue e-negotiation in e-commerce with a tit-for-tat computer agent. Group Decision and Negotiation **14**, 383–392 (2005)
42. Simon, H.A.: Rational choice and the structure of the environment. Psychological Review **63**(2), 129–138 (1956)
43. Sycara, K., Dai, T.: Agent reasoning in negotiation. In: D. Kilgour, C. Eden (eds.) Handbook of Group Decision and Negotiation, pp. 437–451. Springer Netherlands (2010)

44. Wu, M., de Weerdt, M.M., Poutre, H.L.: Efficient methods for multi-agent multi-issue negotiation: Allocating resources. In: J.J. Yang, M. Yokoo, T. Ito, Z. Jin, P. Scerri (eds.) Proceedings of the PRIMA, *Lecture Notes in Computer Science*, vol. 5925, pp. 97–112. Springer (2009)
45. Young, O.R.: Bargaining: Formal Theories of Negotiation. Urbana: University of Illinois Press (1975)
46. Zheng, R., Chakraborty, N., Dai, T., Sycara, K., Lewis, M.: Automated bilateral multi-issue negotiation with no information about opponent. In: Proceedings of Hawaii International Conference on Systems Sciences (2013)
47. Zheng, R., Chakraborty, N., Dai, T., Sycara, K., Lewis, M.: Multiagent negotiation on multiple issues with incomplete information. In: Proceedings of the Twelfth International Conference on Autonomous Agents and Multiagent Systems (2013)

8 Appendix A

Proof (Lemma 1) :
For any $x \in [0,1]^N$, $y \in A$, $P_A[x]$ is the projection of point x on set A. We have

$$\begin{aligned}(P_A[x]-y)'(y-x) = (P_A[x]-y)'(y-P_A[x]) \\ + (P_A[x]-y)'(P_A[x]-x).\end{aligned} \tag{5}$$

Since A is a convex set, by the property of the projection operator, we have

$$(P_A[x]-y)'(P_A[x]-x) \leq 0, \tag{6}$$

Using (6) in (5) we obtain inequality (2), namely,

$$(P_A[x]-y)'(y-x) \leq -\|P_A[x]-y\|^2.$$

We also have

$$\begin{aligned}\|P_A[x]-x\|^2 &= \|x-y\|^2 + \|P_A[x]-y\|^2 \\ &\quad + 2(P_A[x]-y)'(y-x) \\ &\leq \|x-y\|^2 - \|P_A[x]-y\|^2.\end{aligned} \tag{7}$$

where the last inequality was obtained using (2).

Proof (Theorem 1) :
Let agent $i+1$ be the agent proposing an offer x_{t+1}^{i+1} in period $t+1$. Then we have

$$x_{t+1}^{i+1} = P_{A_{t+1}^{i+1}}[w_t], x_{t+1}^j = x_t^j, j \neq i, \text{ and} \tag{8}$$

$$w_{t+1} = w_t + \frac{1}{m}\left(x_{t+1}^{i+1} - x_t^{i+1}\right). \tag{9}$$

Using the results above, we can get

$$\begin{aligned}
&\sum_{j=1}^{m} \|x_{t+1}^j - w_{t+1}\|^2 \\
&= \sum_{\substack{j=1 \\ j\neq i+1}}^{m} \|x_{t+1}^j - w_{t+1}\|^2 + \|x_{t+1}^{i+1} - w_{t+1}\|^2 \\
&= \sum_{j=1}^{m} \|x_t^j - w_t\|^2 + \frac{m-1}{m}\|x_{t+1}^{i+1} - x_t^{i+1}\|^2 \\
&+ 2\left(x_{t+1}^{i+1} - x_t^{i+1}\right)'\left(x_t^{i+1} - w_t\right). \qquad (10)
\end{aligned}$$

By Lemma 1, we have

$$\left(x_{t+1}^{i+1} - x_t^{i+1}\right)'\left(x_t^{i+1} - w_t\right) \leq -\left\|x_{t+1}^{i+1} - x_t^{i+1}\right\|^2,$$

which in turn gives

$$\begin{aligned}
&\sum_{j=1}^{m} \|x_{t+1}^j - w_{t+1}\|^2 \\
&\leq \sum_{j=1}^{m} \|x_t^j - w_t\|^2 - \frac{m+1}{m}\|x_{t+1}^{i+1} - x_t^{i+1}\|^2 \\
&\leq \sum_{j=1}^{m} \|x_t^j - w_t\|^2.
\end{aligned}$$

Proof (Theorem 2) :
Suppose $\exists\ s$, s.t. $A_s = \cap_{i=1}^m A_s^i \neq \varnothing$, (otherwise $\forall t, \cap_{i=1}^m A_t^i = \varnothing$, which implies there is no interior point contained in the set $\lim_{t\to\infty} \cap_{i=1}^m A_t^i$) then as $A_t^i \subset A_{t+1}^i$ for $i \in \{1, 2, ..., m\}$, we have $\forall t \geq s, \cap_{i=1}^m A_t^i \neq \varnothing$. Let

$$e_t^i = P_{A_t^i}\left[w_{t-1}\right] - w_{t-1}. \qquad (11)$$

Without loss of generality, we assume agent 1 proposes in period $s+1$. Then by Lemma 1, $\forall i \in \{1, 2, \cdots, m\}$,

$$\begin{aligned}
\left\|e_{s+i}^i\right\|^2 &\leq \|w_{s+i-1} - x\|^2 - \left\|x_{s+i}^i - x\right\|^2 \\
&\leq \frac{1}{m}\sum_{j=1}^{m}\left\|x_{s+i-1}^j - x\right\|^2 - \left\|x_{s+i}^i - x\right\|^2. \qquad (12)
\end{aligned}$$

Thus, by summing (12) over all agents, we get

$$\sum_{i=1}^{m}\left\|e_{s+i}^i\right\|^2 \leq \frac{1}{m}\sum_{i=1}^{m}\sum_{j=1}^{m}\left\|x_{s+i-1}^j - x\right\|^2 - \sum_{i=1}^{m}\left\|x_{s+i}^i - x\right\|^2. \qquad (13)$$

Moreover, $x_s^i = x_{s+1}^i = \cdots = x_{s+i-1}^i$ and $x_{s+i}^i = x_{s+i+1}^i = \cdots = x_{s+m}^i$. Therefore,

$$\begin{aligned}
&\sum_{i=1}^{m}\sum_{j=1}^{m}\left\|x_{s+i-1}^j - x\right\|^2 \\
&= \sum_{i=1}^{m}\Big[\sum_{j=1}^{i-1}\|x_{s+j}^j - x\|^2 + \sum_{j=i}^{m}\|x_{s+j-1}^j - x\|^2\Big] \\
&= \sum_{j=1}^{m}\sum_{i=j+1}^{m}\|x_{s+j}^j - x\|^2 + \sum_{j=1}^{m}\sum_{i=1}^{j}\|x_{s+j-1}^j - x\|^2 \\
&= \sum_{j=1}^{m}(m-j)\,\|x_{s+j}^j - x\|^2 + \sum_{j=1}^{m} j\|x_{s+j-1}^j - x\|^2
\end{aligned} \tag{14}$$

By substituting (14) into (13), we obtain

$$\sum_{i=1}^{m}\|e_{s+i}^i\|^2 \le \sum_{i=1}^{m}\frac{i}{m}\left[\|x_{s+i-1}^i - x\|^2 - \|x_{s+i}^i - x\|^2\right] \tag{15}$$

The inequality (15) holds for $s = s + km$ and $x_{s+(k-1)m+i}^i = x_{s+km+i-1}^i$, $\forall k \in N$, so

$$\begin{aligned}
&\sum_{k=1}^{R}\sum_{i=1}^{m}\left\|e_{s+km+i}^i\right\|^2 \\
&\le \sum_{k=1}^{R}\sum_{i=1}^{m}\frac{i}{m}\left[\left\|x_{s+i-1}^i - x\right\|^2 - \left\|x_{s+i}^i - x\right\|^2\right] \\
&= \sum_{i=1}^{m}\frac{i}{m}\left[\left\|x_{s+i-1}^i - x\right\|^2 - \left\|x_{s+Rm+i}^i - x\right\|^2\right]
\end{aligned} \tag{16}$$

When $R \to \infty$, the inequality (16) implies

$$\lim_{k\to\infty}\sum_{i=1}^{m}\left\|e_{s+km+i}^i\right\|^2 = 0.$$

Hence, $\lim_{t\to\infty}\left\|e_t^i\right\| = 0$ for all i.

Proof (Theorem 3) :
From inequality (16) in Theorem 2, we have

$$\begin{aligned}
&\sum_{k=1}^{R}\sum_{i=1}^{m}\left\|e_{s+km+i}^i\right\|^2 \\
&\le \sum_{i=1}^{m}\frac{i}{m}\left[\left\|x_{s+i-1}^i - x\right\|^2 - \left\|x_{s+Rm+i}^i - x\right\|^2\right]
\end{aligned}$$

Moreover, from the definition of e_t^i, we can obtain

$$\sum_{i=1}^{m} \left\| e_{s+(k+1)m+i}^i \right\|^2 \leq \sum_{i=1}^{m} \left\| e_{s+km+i}^i \right\|^2, \ \ \forall k \in N.$$

Thus,

$$\begin{aligned}
&\sum_{i=1}^{m} \left\| e_{s+Rm+i}^i \right\|^2 \\
&\leq \frac{1}{R} \sum_{i=1}^{m} \frac{i}{m} \left[\left\| x_{s+i-1}^i - x \right\|^2 - \left\| x_{s+Rm+i}^i - x \right\|^2 \right] \\
&\leq \frac{1}{R} \sum_{i=1}^{m} \frac{i}{m} \left\| x_{s+i-1}^i - x \right\|^2
\end{aligned}$$

which implies, $\forall \varepsilon > 0$, there exists $T > 0$, where

$$T = s + R \left\lceil \frac{\sum_{i=1}^{m} \frac{i}{m} \left\| x_{s+i-1}^i - x \right\|^2}{\varepsilon} \right\rceil + i$$

such that $\forall t > T$, $\sum_{i=1}^{m} \left\| e_t^i \right\|^2 < \varepsilon$.

9 Appendix B

For completeness purpose, we present the convex optimization formulation for finding the Nash bargaining solution. In [34], Nash gave an axiomatic approach to define reasonable outcomes in a negotiation. This discussion is available in the original paper and many subsequent works. In this paper, we are using a convex optimization approach for computing the Nash bargaining solution. Hence, we will restrict our discussion to the formulation of the optimization problem. Let there be m agents negotiating on n issues with the issues taking on continuous values between 0 and 1 . Let $u_i(x)$ be the utility function of agent i, which is assumed to be concave. Without loss of generality, we assume that no-agreement results in a utility of 0. The objective function to be maximized is the joint utility, namely,

$$f(x) = \prod_{i=1}^{m} u_i(x) \tag{17}$$

Since $u_i(x)$ is concave and non-negative, $f(x)$ is non-negative, and hence maximizing $f(x)$ is equivalent to maximizing $\log(f(x))$. Let x_j denote the jth component of x. The convex optimization problem to be solved for computing the Nash equilibrium is

$$\begin{aligned}
\text{maximize} \ \ & \sum_{i=1}^{m} \log(u_i(x)) \\
\text{s.t.} \ \ & u_i(x) \geq r u_i \quad i = 1, \ldots, m, \\
& 0 \leq x_j \leq 1, \quad \forall j = 1, \ldots, n.
\end{aligned} \tag{18}$$

where ru_i is the ultimate reservation utility of agent i. Since each u_i is a concave function of x, the set of constraints in (18) forms a convex set. The objective function to be maximized is a sum of log-concave functions and hence the problem is a convex optimization problem. In the paper, we have used the solver CVX [20, 19] implemented in MATLAB for obtaining the Nash bargaining solutions.