Analyzing and exploiting the competitiveness of scenarios for negotiating convoy formation under time constraints

Yi Luo and Ladislau Bölöni
School of Electrical Engineering and Computer Science
University of Central Florida
Orlando, Florida
yiluo@mail.ucf.edu, lboloni@eecs.ucf.edu

November 15, 2009

Abstract

In the convoy formation problem, two embodied agents are negotiating the synchronization of their movement for a portion of the path from their respective sources to destinations. As equilibrium strategies are not practically possible, we are interested in strategies with bounded rationality, which achieve good performance in a wide range of practical negotiation scenarios. Naturally, the performance of a strategy is dependent on the strategy of the opponent and the characteristics of the scenario. The goal of this paper is to develop a *collaborativeness metric* of the negotiation scenario which formalizes our intuition of collaborative scenarios (where the agents' interests are closely aligned) versus competitive scenarios (where the gain of the utility for one agent is paid off with a loss of utility for the other agent).

We are using the Children in the Rectangular Forest (CRF) game as a canonical model of convoy formation, assume zero initial knowledge and a negotiation protocol requiring mandatory, but non-binding evaluations of the opponents offer. We also assume that the negotiation happens in physical time. We describe two negotiation strategies: the comparatively simple Internal Negotiation Deadline (IND) strategy and the computationally more expensive Uniform Concession (UC) strategy. Then, we describe how these strategies can be augmented by collaborativeness analysis: we approximate the collaborativeness metric in the first several negotiation rounds, and use the result to cut short the negotiation when the estimated collaborativeness is lower than a threshold. Through an experimental study, we show that augmenting the strategies with collaborativeness analysis significantly improves their performance for low collaborativeness scenarios, with only a minimal penalty in high collaborativeness scenarios.

1 Introduction

Collaboration between embodied agents often requires the temporal and spatial collocation of the agents. Agents need to coordinate their movements, agree on meeting points, time, common path and speed, as well as locations where they split and start moving on independent trajectories. Such convoy formation problems appear as sub-problems in many practical applications such as transportation and disaster rescue.

The evaluation (and creation) of offers in convoy formation problems is computationally expensive, as it often involves path planning. As the negotiation happens in real physical time, agents can not afford to evaluate a large number of offers for feasibility and utility. As equilibrium strategies are not practically possible, we are interested in developing strategies with bounded rationality, which achieve good performance in a wide range of practical negotiation scenarios. Naturally, the performance of a strategy is dependent on the strategy of the opponent and the characteristics of the scenario. The utility of a deal alone for a particular agent is not a good measure of the quality of the negotiation strategy; we also need to consider whether better deals were overlooked or whether the agent had "outsmarted" the opponent, by convincing it to accept a lesser deal. We also have an intuition of collaborative scenarios (where the agents interests are closely aligned) versus competitive scenarios (where the gain of the utility for one agent is paid of with a loss of utility for the other agent). Empirical observations of negotiation traces show that certain negotiation strategies perform better in collaborative scenarios while others in competitive scenarios. Thus, if we know the collaborativeness of a scenario, we can predict the performance of a negotiation strategy, and choose strategies accordingly. To accomplish this, we need a metric of the collaborativeness of a scenario.

To show the intuition behind collaborativeness in a negotiation scenario, let us first consider the split the pie game, frequently used to model worth oriented negotiations. Here two agents are negotiating over how to partition a pie into two disjoint pieces by making partitioning offers in each turn which can be accepted or declined by the opponent. As the parts allocated to one agent are lost for the other agent, the single pie game is fully competitive. A fully collaborative game would be one in which there would be a solution where both agents get the full pie.

In a multiple-pie game the agents are partitioning multiple pies over which they have different valuations. As long as the valuations are all positive, the game remains fully competitive. Note that although all zero-sum games are fully competitive, not all fully competitive games are zero sum[1]. For instance, a "split multiple pies" game where the agents value the different pies with different, positive values, is still fully competitive, but not zero sum. On the other hand, a fully cooperative game is one where there is a possible agreement which is individually optimal for both agents. An example of a fully collaborative game is a split the multiple pie game with two pies P_1 and P_2, where agent A values P_1 positively and P_2 negatively, while agent B values them the other way around. In this case the agents can easily agree on a partitioning where agent A gets the pie P_1, while agent B gets the pie P_2.

Similar considerations apply to the convoy negotiation, but the expressions of collaborativeness are more complex because the utility of an offer is a non-linear function of the issues and not all offers are feasible. The first goal of this paper is to develop a collaborativeness metric which matches well with our intuition of collaborativeness as outlined above.

To illustrate the ways in which the proposed metric can improve negotiation performance, we consider the practical negotiation setting of the Children in the Rectangular Forest (CRF) model, a

[1]Some game theory texts, such as [11] equate fully competitive with zero sum, by making the assumption that the utility function is just a convenient expression of the preference ordering. In convoy formation, however, the utility has the dimensionality of time, and it can not be arbitrarily scaled. There is a difference between a scenario where a 1 second utility decrease from one agent gives 1 second utility gain to the other agent, and the scenario where 1 second utility decrease gives 100 seconds utility gain for the opponent. In our terminology the first scenario is fully competitive and zero sum, while the second scenario is fully competitive but *not* zero sum.

convoy formation setting in which the convoy advantage is the ability to traverse a region inaccessible to individual agents. We consider that the negotiation happens in physical time, each negotiation round taking a fixed amount of time t_r. If the negotiation breaks down after a number of rounds, the time to reach the destination will be larger than the conflict deal; thus negative utility values are possible. We assume a zero knowledge starting point: all the information the agents have about each other needs to be acquired during the negotiation. We assume the negotiation protocol to be exchange of binding offers with mandatory, non-binding evaluations. The negotiation strategies for the convoy formation problem need to simultaneously solve the problem of managing their concession rate and search the offer space for solutions beneficial to both agents. We describe two negotiation strategies: the Internal Negotiation Deadline (IND) adapts its spatial concession pace to a preestablished deadline, while the Uniform Concession (UC) strategy pre-calculates pools of potential offers for various concession levels, then selects its own offer at various concession levels based on the similarity to the opponents' offer. As we will see in the experimental study, UC in general outperforms IND, at the cost of a much higher computational and memory requirements. Plotting the negotiation results function of the collaborativeness of the scenario, we find that both strategies perform badly for low collaborativeness scenarios, the utility dipping significantly below zero.

The next step is to identify how the collaborativeness metric can be used to improve negotiation performance. The collaborativeness depends on both agents, thus, it can be accurately calculated only by a full-knowledge supervisor. We describe an approach for approximating the collaborativeness in a zero-knowledge setting through the information acquired in the first several negotiation rounds. We augment the two proposed strategies with collaborativeness analysis based on this approximation. The augmented strategies IND+CA and UC+CA approximate the collaborativeness of the current scenario early in the negotiation, and, if the collaborativeness is lower than a threshold, bring the negotiation to the quick positive or negative conclusion. The experimental studies show that the augmented strategies significantly outperform the original strategies for difficult scenarios with low collaborativeness, while performing only minimally worse in favorable, high collaborativeness scenarios.

The remainder of this paper is organized as follows. Related work is discussed in Section 2. Section 3 describes a formal model of the convoy formation problem and introduces the proposed collaborativeness metric. Section 4 introduces the IND and UC negotiation strategies and describes their versions augmented with collaborativeness analysis (IND+CA and UC+CA). An experimental study comparing the performance of various negotiation strategies for scenarios of various levels of competitiveness is described in Section 5. We conclude in Section 6.

2 Related work

While automated negotiation [8] generated a lot of interest in recent years, negotiation about spatio-temporal issues in embodied agents has received relatively little attention. Nevertheless, many research results in multi-issue negotiation or collaborative robotics have relevance to our work.

Sandholm and Vulkan [13] analyze the problem of negotiating with internal deadlines where the deadlines are private information of the agents. The negotation problem is a "split a single pie", zero-sum negotation. They find that for rational agents, the sequential equilibrium is a strategy which requires agents to wait until their deadline, and at that moment, the agent with the earliest deadline concedes the whole cake.

Fatima, Wooldridge and Jennings [3, 4] extensively study the problem of multi-issue negotation under deadlines. The problem considered is the split multiple pie problem where the pie is assumed to shrink after every negotiation round, under both complete information and incomplete information assumptions. The authors compare three negotiation procedures: the package deal procedure where all the issues are discussed together, the simultaneous procedure where issues are discussed independently but simultaneously, and the sequential procedure where issues are discussed one after another. The authors show that the package deal is the optimal procedure for both agents.

Ito, Klein and Hattori [7] consider negotiations in real world settings where the utility values are non-linear. For instance, the value of the tires and the value of the engines can not be simply added up when designing a car, as the issues constrain each other. The authors propose an auction-based multi-issue negotiation protocol for negotiating among agents with a non-linear utility settings. The protocol also includes a mediator, which is responsible to choose the deal with the largest social utility from the deals made possible by the bids of the agents.

Golfarelli et al. [5] considers the case of robotic agents which are assigned a set of tasks which are attached to physical locations. The tasks carry precedence constraints (execute one specific task earlier than the other) and object constraints (fetch the object in order to execute the task). Agents need to determine, on a network of places and routes, a sequence of places to be visited in order to carry out a set of tasks. Through swapping tasks based on announcement-bid-award mechanism, the agents can decrease their tasks execution costs in the map. An extended version of this work [6], allows the agents to exchange clusters of tasks to avoid being stuck in local minima. To cluster similar tasks, the authors calculate spatial distance and temporal distance of tasks, and apply thresholds to differentiate between near and far tasks.

Saha and Sen [12] discuss the problem of negotiating efficient outcomes in a multi-issue negotiation where some of the parameters of the agent are not common knowledge. The "distributive" and "integrative" scenarios proposed by them are the equivalents of the "competitive" and "collaborative" scenarios we define for the spatio-temporal negotiation problem.

Crawford and Veloso [2] applied the "experts" algorithm to solve the multi-agent scheduling problem. In this algorithm the agent is helped by a number of "experts", but it needs to decide which experts' advice it should follow. The learning agent can dynamically change its strategy according to its opponents' behavior. The performance of each algorithm is measured in terms of total utility achieved over each of the trials.

3 Collaborativeness in the convoy formation problem

3.1 The general convoy formation problem

Let us start by defining the convoy formation problem for embodied agents. Two agents A and B move from their source positions S_A and S_B to their destinations D_A and D_B. We assume that the agents move along the paths given by the function $P_a(t) \rightarrow L$, which we read by saying that agent a is at the location L at time t.

At the initial timepoint t_0 we have $P_A(t_o) = S_A$ and we define the *arrival time* of A as the smallest time t_{arr} for which $P_A(t_{arr}) = D_A$. For every path we define the *unit cost* $c_P(t)$, and the cost of a time segment $C(t_1, t_2) = \int_{t_1}^{t_2} c_P(t)dt$. Most of the time, we are interested in the cost of the path $C_P(t_0, t_{arr})$. In the simplest case we are only interested in the time to reach the destination. This corresponds to a unit cost $c_P(t) = 1$, and the cost of the path $C_P(t_0, t_{arr}) = t_{arr} - t_0$. Many environmental factors can be modeled by the appropriate setting of the unit costs. For instance, the unit cost might be dependent on the location $c_P(t) = f(P_A(t))$ or on the speed of the agent $c_P(t) = f(P'_A(t))$. Locations or speeds which are unfeasible to the agent can be set to have an infinite unit cost.

Two agents form a *convoy* if they are following the same path $P_{A+B}(t)$ over the period of time $[t_{join}, t_{split}]$. An agent is motivated to join a convoy because of the *convoy advantage*: the unit cost for the convoy is smaller than for the individual agent over the same path. One example is the case when convoys can traverse areas which are not accessible to individual agents: $\exists t \in [t_{join}, t_{split}]\ \exists l\ P_{A+B} = l$ with $c_{P,A}(t) = \infty$ and $c_{P,A+B}(t) = c \in \mathbb{R}$. Naturally, convoy and non-convoy segments of the path need to be continuous in space: $P_A(t_{join}) = P_B(t_{join}) = P_{A+B}(t_{join}) = L_{join}$ and $P_A(t_{split}) = P_B(t_{split}) = P_{A+B}(t_{split}) = L_{join}$. We call L_{join} and t_{join} the join locations and time, and L_{split} and t_{split} the split locations and time, respectively.

We are considering self-interested agents which are searching for the path with the smallest cost from source to destination. This path might or might not include segments traversed as a convoy. In the following we assume that the agents are using negotiation to agree on the segment traversed as a convoy. The negotiation succeeds if an agreement is reached over a quadruplet $(L_{join}, t_{join}, L_{split}, t_{split})$. Convoy negotiation is thus a multi-issue negotiation, with two temporal and two spatial issues. It can be seen as a six-issue negotiation if we consider the spatial location $L = (x, y)$ as two issues.

3.2 Defining a collaborativeness metric

Each of us has an intuitive feel for negotiation scenarios which are "easy" because the negotiation partners have a strong incentive to form a deal and for scenarios which are "hard" because a rational agreement is difficult to find (or it might not exist). Also, we have an intuition of certain negotiation scenarios where one of the participants has "more to gain" from an agreement.

Our objective is to develop metrics which match well with these intuitions, while abstract away the other parameters of the scenario (such as the location and destination of the agents).

In the remainder of this paper we will assume that the unit cost is either unity ($c_p(t) = 1$) or infinity ($c_p(t) = \infty$). Under these conditions, the cost of every feasible path is equal to the time to destination, but not all paths are feasible. We will also assume that the agents negotiate in physical time, with each negotiation round taking time t_r, and the agent being immobile during negotiation.

We call the cost of an offer $C^{(A)}(O)$ of agent A for a particular offer $O = \{L_{join}, t_{join}, L_{split}, t_{split}\}$ the time it takes for the agent to reach its destination if it accepts the offer and follows the trajectory. The lower the time to destination, the more desirable is the offer for the agent. The time to destination is composed of three components: the time it takes for both agents to reach the meeting location, the time for traveling together in the forest, and time from the split location to the agent's destination. We assume $C^{(A)}(O) = \infty$ if the offer is unfeasible for the agent. The cost of the conflict deal $C^{(A)}_{conflict}$ is the time for the agent to reach its destination if it does not make any deal.

As time is passing during the negotiation, the actual cost of an offer made at negotiation round n will be $C^{(A)}_{r=n}(O) = n \cdot t_r + C^{(A)}(O)$. This also applies to the cost of the conflict deal at round n: $C^{(A)}_{conflict,r=n} = n \cdot t_r + C^{(A)}_{conflict}$.

Considering agents whose negotiation time is the physical time requires us to refine our definition of rationality of a deal. At the beginning of the negotiation, at time t_0, the agent has a conflict deal path with cost $C^{(A)}_{conflict}$. According to the *baseline rationality* definition, any offer which has a higher cost than $C^{(A)}_{\mathit{conflict}}$ is not rational and it will not be accepted by the agent. Any negotiation the agent might enter implies a risk of conflict. Thus, at negotiation round n the agent might find itself in the position that it has already incurred costs $C_x = n \cdot t_r$. If at this moment an offer with cost $C^{(A)}(O)$ is received, it will be called *pragmatically rational* if $C^{(A)}(O) < C^{(A)}_{\mathit{conflict}}$ and *baseline rational* if $C_{\mathit{offer}} + n \cdot t_r < C_{\mathit{conflict}}$. A rational agent will need to act based on the pragmatic rationality, as the original conflict deal alternative is not available any more at this moment in time. Occasionally, the agent might find it necessary to accept deals which are not baseline rational. In the rest of the paper, unless explicitly mentioned differently, the term rationality will mean pragmatic rationality.

However, when we are measuring the overall performance of the negotiation strategy / action strategy pairs, the term of comparison should be the original conflict deal. In order for a strategy pair to be acceptable, it needs to be baseline rational at least in the statistical average.

Definition 1 *The* **pragmatic utility of an offer O** *for agent A, denoted with $P_A(O)$, is the time the agent saves accepting the offer compared to the conflict deal, considering no time spent on the negotiation.*

$$U^{(A)}(O) = C^{(A)}_{\mathit{conflict}} - C^{(A)}(O) \tag{1}$$

The **baseline utility of the offer** *which has been made at the n-th negotiation round is:*

$$U^{(A)}(O) = C^{(A)}_{\mathit{conflict}} - C^{(A)}(O) - n \cdot t_r \tag{2}$$

For instance, let us consider an agent whose time to destination is 45 minutes proceeding alone. Let us assume that the agent spent 15 minutes negotiating a deal which takes it to destination in 40 minutes. The pragmatic utility of this deal is +5 while the baseline utility is -10. At time 15, the negotiation time being already spent, the agent is better off taking the deal (which makes it arrive at time 55) than taking the conflict deal (which makes it arrive at time 60). Thus, the deal is pragmatically rational (at time 15). The deal however, is not baseline rational, because the original conflict deal was 45, thus the agent would have been better off if it does not negotiate at all.

Definition 2 *We define the* **absolute best time to destination** $C^{(A)}_{ab}$ *for agent A the time it would take it to reach the destination assuming an ideally performant and ideally collaborative negotiation partner.*

For the CRF problem, the trajectory associated to the absolute best time to destination is a straight line from the source to destination traversed by the agent with its maximum velocity.

$$C^{(A)}_{ab} = \frac{|S_A, D_A|}{v_A} \tag{3}$$

This assumes that there is an ideal negotiation partner, who is (a) willing to accept any geometric location for meeting and splitting points proposed by the agent, (b) its velocity is greater than or equal of the current agent and (c) its current position is such that it can reach the meeting point at a time earlier or equal with the time it takes agent A to reach it. Note that for a practical scenario, the absolute best time to destination may not be feasible, even for an ideally cooperative negotiation partner.

Definition 3 *We define the* **ability constrained best time to destination** $C^{(A),\{B\}}_{acb}$*, of an agent A negotiating with an agent B, the time A can reach the destination assuming an ideally collaborative agent B.*

The ability constrained best time takes into account the physical limits of the negotiation partner and the scenario. The meeting and split point of the offer associated with the ability constrained best time might not be the one situated on the intersection of the straight line to destination with the forest. The offer(s) associated with $C_{acb}^{(A),\{B\}}$ might not be rational for agent B.

Let us consider an agent for which the absolute best deal involves meeting at point L_1 at time $t_1 = 20$, with the agent reaching its destination at time $t_{dest} = 100$. However, the opponent can not physically make it to the point L_1 in at time t_1, because it is too far away. We need to search for a deal which is feasible, for instance by extending the join time to $t_1' = 30$. This would also extend the time to destination to $t'_{dest} = 110$. Alternatively, we can also modify the location of the join point.

Definition 4 *The* **rationality constrained best time to destination** $U_{rcb}^{(A),\{B\}}$ *for agent A negotiating with agent B is the time to destination of agent A which can be obtained assuming that agent B will accept any offer, as long as it is rational for B.*

For instance, let us consider a case the ability constrained best deal for the agent A would have a time to destination 100, with meeting at point L_1 and splitting at L_2. Let us assume that this trajectory is also feasible for agent B. It is still possible, however, that this trajectory would result for the B in a deal which is worse than going around the forest alone. One reason for this might be that the split point L_2 is too far from the B's destination D_B. B will not accept such a deal. A different deal would need to be negotiated, which, however, would normally be less advantageous for agent A.

As $C_{acb}^{(A),\{B\}}$ and $U_{rcb}^{(A),\{B\}}$ introduce successive restrictions over $C_{ab}^{(A)}$, we have:

$$C_{conflict}^{(A)} \geq C_{rcb}^{(A)\{B\}} \geq C_{acb}^{(A)\{B\}} \geq C_{ab}^{(A)} \tag{4}$$

Each of these time to destination values define a set of one or more concrete offers which actually achieve them. Thus we define a rationality constrained best offer of A to be an offer $O_{rcb}^{(A),\{B\}}$ such that

$$C^{(A)}\left(O_{rcb}^{(A),\{B\}}\right) = C_{rcb}^{(A)\{B\}} \tag{5}$$

The metrics introduced until now characterize the scenario from the point of view of one of the agents. Let us now develop a metric which quantifies the desirability of a certain offer O from the point of view of the social good.

Definition 5 *We call the* **social cost of the offer O** *any function* $C_{social}(O) = C_{social}\left(C^{(A)}(O), C^{(B)}(O)\right)$ *which is monotonically increasing both with* $C^{(A)}$ *and with* $C^{(B)}$*:*

$$\begin{aligned} &\forall C^{(B)}, C_1^{(A)} \geq C_2^{(A)} \Rightarrow C_{social}(C_1^{(A)}, C^{(B)}) \geq C_{social}(C_2^{(A)}, C^{(B)}) \\ &\forall C^{(A)}, C_1^{(B)} \geq C_2^{(B)} \Rightarrow C_{social}(C^{(A)}, C_1^{(B)}) \geq C_{social}(C^{(A)}, C_2^{(B)}) \end{aligned} \tag{6}$$

We call denote with O_{social} the set of offers which minimize the social cost:

$$O_{social} = \underset{O}{\operatorname{argmin}} \left(C_{social}(O)\right) \tag{7}$$

Within the constraints of this definition, there are many possible functions which can serve as the social cost function. The choice of a specific function depends on the policy of the supervisor. One simple choice is to define the social cost as the sum of the individual costs.

$$C_{social}(O) = C^{A+B}(O) = C^{(A)}(O) + C^{(B)}(O) \tag{8}$$

Note however, that a social best offer might not be rational for both agents. We can define a rationality constrained social cost, which assumes a cost of plus infinity for the offers which are not rational for one of the agents:

$$C_{rcsoc}(O) =$$

$$\begin{cases} +\infty & \left(C^{(A)}(O) > C^{(A)}_{conflict}\right) \vee \left(C^{(B)}(O) > C^{(B)}_{conflict}\right) \\ C_{social}(O) & \text{otherwise} \end{cases} \tag{9}$$

Based on this definition, we can define the set of rationality constrained social best offers O_{rcsoc} as:

$$O_{rcsoc} = \underset{O}{\operatorname{argmin}} \left(C_{rcsoc}(O)\right) \tag{10}$$

Definition 6 *We define as the* **collaborativeness of the scenario** *from the point of view of agent A, negotiating with agent B, the ratio of the utility of the rationality constrained social best deal to the maximum rationally obtainable utility:*

$$\Xi^{(A),\{B\}} = \frac{C^{(A)}_{conflict} - C^{(A),\{B\}}_{rcsoc}}{C^{(A)}_{conflict} - C^{(A),\{B\}}_{rcb}} \tag{11}$$

Let us verify that this definition satisfies our intuition about the collaborativeness of a scenario. In a fully competitive scenario, there is no rational deal possible, thus the cost of the rational deal will be the conflict deal, thus we have $\Xi^{(A),\{B\}} = 0$. On the other hand, we say that a scenario is fully cooperative from the point of view of agent A if the rationality constrained social best offer is also the rationality constrained best offer for agent A. In this case $\Xi^{(A),\{B\}} = 1$.

Definition 7 *We define the* **relative utility of an offer** *for agent A as the ratio of the utility of the offer to the maximum rationally obtainable utility:*

$$U^{(A),\{B\}}_{rel}(O) = \frac{C^{(A)}_{conflict} - C^{(A)}(O)}{C^{(A)}_{conflict} - C^{(A),\{B\}}_{rcb}} \tag{12}$$

The relative utility of the agent can range from 0 to 1. Notice that the relative utility of a deal does not tell us whether the agent has negotiated "better" than the negotiation partner. There are situations when both agents can reach the maximum relative utility.

4 Negotiation strategies

To illustrate the ways in which the proposed collaborativeness characterization metrics can be exploited in a negotiation strategy, we will consider some negotiation strategies and augment them to take into consideration the collaborativeness of the scenario.

As the negotiation strategies are strongly dependent on the shape of the surface, we shall consider a simple version of the convoy formation problem, the Children in the Rectangular Forest (CRF) game, where the "convoy advantage" is the ability of the convoy to traverse a rectangular region inaccessible to the individual agents.

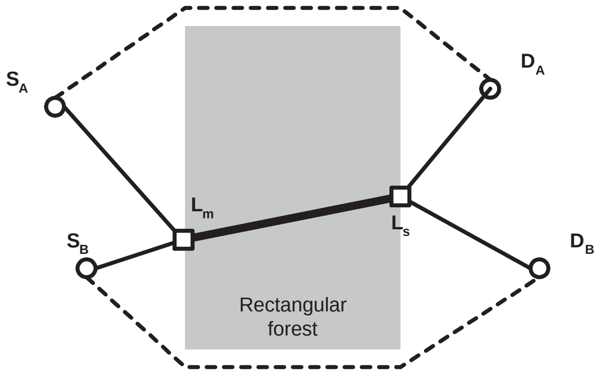

Figure 1: The Children in the Rectangular Forest problem. The trajectories associated with the conflict deal are shown with an interrupted line, while the trajectories corresponding to a possible agreement are shown with a continuous line.

4.1 The CRF problem

In [9, 10] we have considered a simplified convoy formation problem called Children in the Rectangular Forest (CRF), where the convoy advantage is represented by the convoys ability to traverse a rectangular obstacle which is not accessible to the individual agents (see Figure 1). The CRF game presents many challenges of the general problem such as the difficulty of establishing whether an offer is feasible to the opponent, whether it represents a concession or not, and the difficulty of simultaneously negotiating temporal and spatial issues. At the same time, the CRF problem simplifies away the path planning problem, as all the Pareto-optimal deals correspond to paths formed of at most three linear segments.

The four negotiation issues are not completely independent. For instance, if we know the maximum velocity of both agents, the split time t_{split} can be calculated from L_{join}, L_{split}, and t_{join}. Similarly, if all information is known about the current location and speed of the agents, the Pareto optimal value of t_{join} can be calculated, knowing L_{join}.

We call a *fully specified offer* a quadruple $O = \{L_{join}, t_{join}, L_{split}, t_{split}\}$ which specifies both the spatial and temporal components of an offer. A *spatially specified offer* specifies only the spatial components of the offer: $O = \{L_{join}, ?, L_{split}, ?\}$. An agent A can complete a spatially specified offer by calculating the timepoints $t^{(A)}_{join}$ and $t^{(A)}_{split}$ which are the earliest feasible ones for the agent. The resulting offer is the *best time completion* for A of the spatially specified offer O:

$$BTC^{(A)}(O) = BTC^{(A)}(\{L_{join}, ?, L_{split}, ?\}) = \{L_{join}, t^{(A)}_{join}, L_{split}, t^{(A)}_{split}\} \tag{13}$$

A CRF scenario is defined by the map of the CRF game (the size of the forest), the source points of the two agents S_A and S_B, the destination points of the two agents D_A and D_B, and the maximum velocities of the agents v_A and v_B. The path of the agents are series of segments together with the velocities of the vehicle on the different segments.

We call *time to destination*

$$C^{(A)}(O) = \max\left(\frac{|S_A, L_{join}|}{v_A}, \frac{|S_B, L_{join}|}{v_B}\right) + \frac{|L_{join}, L_{split}|}{\min(v_A, v_B)} + \frac{|L_{split}, D_A|}{v_A} \tag{14}$$

4.2 Negotiation protocol: exchange of binding offers with mandatory, non-binding evaluations

The negotiation setting we consider is a zero-knowledge setting: the only information the agents have about their opponent is acquired during the negotiation itself, no external sources of information exist. The negotiation protocol we assume is exchange of binding offers with mandatory, non-binding evaluations (EBOMNE).

The offers in convoy formation problems are quadruplets $O = \{L_{join}, t_{join}, L_{split}, t_{split}\}$, having two spatial and two temporal components. The offers are binding to the offering agent: if accepted by the opponent, they will represent a deal. We assume that the agent proposes first its absolute best offer, which assumes an ideally performant and ideally collaborative negotiation partner. At each round the agent can "accept" the opponent's offer, "confirm" the acceptance, "propose" a counter-offer, and "quit" the negotiation.

Upon receipt of an offer, the agent proceeds to *evaluate* it. If the offer is feasible and rational, the evaluation is the offer itself: $E = O$. If the offer is not feasible (for instance, because the agent can not reach the join location in time, or it can not match the required speed during the common path), the agent can extend the temporal components of the offer such that they become feasible for the agent. If the resulting offer is rational for the agent, it will become the evaluation. If it is not feasible, the evaluation is considered to be the empty set $E = \emptyset$. The evaluation will paired with a counter offer to form the return message. Thus the response of the agent A at negotiation round i will be the pair $(O_i^A,\ E_{i-1}^B)$.

While the offers are binding, the evaluations are not. An empty evaluation intuitively means "the proposed spatial coordinates are very wrong", while an evaluation returned with a counteroffer means "I would be able to accept the offer, but I am not willing to". The evaluation does not immediately disclose the utility function of the agent, but they allow the opponent to select its offer more efficiently. Thus, the EBOMNE protocol represents a simple variant of *argumentation*.

4.3 Negotiation strategies: general considerations

One the negotiation space is determined and the negotiation protocol agreed upon, the flow of the negotiation for a certain negotiation scenario is defined by the negotiation strategies of the agents. The agents have a considerable freedom in choosing the negotiation strategy, which is limited only by the requirement of conformance with the protocol.

However, the structure of the successful strategies is frequently dictated by the objective nature of the negotiation domain. In the following we present some considerations about the state of the convoy negotiation, which need to be implicitly or explicitly made by any successful strategy. Let us consider that agent A has just received a message $(O_{i-1}^B,\ E_{i-2}^A)$ from agent B. The agent A can evaluate the current state of the negotiation as follows.

A *Blind search* ($E_{i-2}^A = \emptyset$, $U^{(A)}(E_{i-1}^B) < 0$). In this case the agent A was told that its previous offer was not rational for B, but it also finds that the offer of the opponent is not rational for himself either. This situation frequently happens at the beginning of the negotiation. Being in this state does not necessarily means that there is no deal possible, but the agents need to explore the state space for areas where mutually rational offers can be found.

B *Accept or concede* ($E^A_{i-2} = \emptyset$, $U^{(A)}(E^B_{i-1}) > 0$): The agent's last offer (O^A_{i-2}) isn't rational for the opponent but opponent's last offer (O^B_{i-1}) is pragmatically rational after extending the time issues. In this situation, the agent can either accept the opponent's offer with the time components extended or concede in a counter-offer. A deal will be formed only if the opponent confirms the modified offer.

C *Unbalanced blind search* ($U^{(A)}(E^A_{i-2}) < 0$, $U^{(A)}(E^B_{i-1}) < 0$): The opponent returns an evaluation of the agent's last offer. However, this extended offer is not rational for the agent A. On the other hand, the opponent counter-offer is not rational for the agent, either. This situation can happen when the agent is near the forest while its opponent is not. Although the opponent accepts the joining and splitting locations, they fail to form a mutually rational agreement.

D *Opponent's offer acceptable* ($U^{(A)}(E^A_{i-2}) < 0$, $U^{(A)}(E^B_{i-1}) > 0$): The evaluation of the agent's previous offer was not rational for the agent, but the opponent's offer evaluation is rational. The agent can either accept the opponents offer, or create a counter-offer which it hopes to be rational to the opponent.

E *Agent's offer acceptable* ($U^{(A)}(E^A_{i-2}) > 0$, $U^{(A)}(E^B_{i-1}) < 0$): The evaluation of the agent's offer is pragmatically rational, while the opponent's offer is not. Intuitively, the agent has no motivation to conceed until the opponent comes up with a rational offer. The agent will insist on its own offer until the opponent either accepts it, or provides a rational counter-offer.

F *Mutual concessions* ($U^{(A)}(E^A_{i-2}) > 0$, $U^{(A)}(E^B_{i-1}) > 0$): Both offers are evaluated to be rational, thus the agents now need to reach a deal with mutual concessions. Other things being equal, the agents will try to minimize their concessions. However, at the same time, the agents need to consider the risk of the opponent quitting the negotiation, as well as weight the potential benefits they can obtain from further negotiation against the time t_r lost in every negotiation round.

G *Accepted offer can not be confirmed* ($U^{(A)}(E^A_{i-2}) < 0$, $O^B_{i-1} = \emptyset$): The opponent accepted the evaluated version of the agent's counter offer. This evaluation, however, is not rational for the agent. As the opponent is also interested in getting to the split point as soon as possible, this means that no deal is possible with the current set of spatial components (it is not the matter of the opponent conceeding more). The agent can either generate a spatially different counter-offer or quit the negotiation.

H *Accepted offer can be confirmed* ($U^{(A)}(E^A_{i-2}) > 0$, $O^B_{i-1} = \emptyset$): The opponent accepted the evaluated version of the offer, and this evaluation is rational for the agent. The agent can confirm the offer, which then becomes a deal. Alternatively, the agent can restart the negotiation with a new counter-offer if it considers that it can form the basis of a better deal.

4.4 Two strategies without collaborativeness analysis

4.4.1 Internal negotiation deadline (IND)

In the internal negotiation deadline strategy, the agent sets up a deadline n_{max} (expressed as a number of negotiation rounds) and adapts the speed of concession in function of the remaining negotiation rounds. In stage E, the IND agent will insist its last offer to force the opponent to concede. In the

G and H stages, the IND agent stops calculating the next offer and makes decision between "quit" or "confirm". In the other stages, the IND agent will calculate the next conceded offer described by the following values:

$$y_{join}^{(A),i} = \begin{cases} y_{join}^{(A),i-2} - c_m & \text{if } y_{join}^{(B),i-1} < y_{join}^{(A),i-2} \\ y_{join}^{(A),i-2} + c_m & \text{if } y_{join}^{(B),i-1} > y_{join}^{(A),i-2} \end{cases} \tag{15}$$

where the conceding amount in the meeting location is:

$$c_m = \frac{\left| y_{join}^{(B),i-1} - y_{join}^{(A),i-2} \right|}{\lceil (n_{max} - i)/2 \rceil}, \text{for } i < n_{max} - 2 \tag{16}$$

A similar expression for $y_{split}^{(A),i}$, the best time completion t_{join}^{A} and t_{split}^{A} are calculated accordingly. Note that there are three situations that the IND agent couldn't find the next offer: (a) the next concession break its own rationality constraint, (b) the current negotiation round i is greater or equal than $n_{max} - 2$ (one call left for the agent), and (c) the next conceded offer is worse than the evaluation of opponent's previous offer. In these situations, the IND agent, again makes decision: either "accept" or "quit" the negotiation according to $U^{(A)}(E_{i-1}^{B})$. If the evaluation E_{i-1}^{B} is the same with the opponent's offer O_{i-1}^{B} (no extension in time issues), the IND agent can "confirm" it directly.

4.4.2 Uniform concession (UC)

The advantage of the IND strategy is that it is easy to understand and simple to implement. It resembles the monotonic concession strategy from single-issue worth-oriented domains. There are, however, some important differences. Conceding in the join and split location does not necessarily mean an even concession in terms of utility. By exploring only specific combinations of meeting and splitting points, with the tight joining and splitting time according to its own speed, the strategy excludes a large part of the solution space.

In the uniform concession strategy, the agent generates a pool of all possible offers (all combinations of joining and splitting location with a certain resolution), as well as possible time buffers at the meeting time field. The splitting time is calculated based on the minimum common speed in the history of all previous offers and evaluations. The offer pool is then divided into a number of subpools. The first subpool contains offers which have the agent's absolute best utility $U_{ab}^{(A)}$. Each successive subpool $i = 1 \dots n$ groups offers whose utility $U_{sp}(i)$ is smaller by the value α than the previous one, where $\alpha \in [0, 1]$ is the *conceeding speed* of the agent:

$$U_{sp}(i) = (1 - (i \times \alpha)/2) \times U_{ab}^{(A)}, \text{for } (1 - (i \times \alpha)/2) \geq 0 \tag{17}$$

The insight is that from the agent's point of view all the offers in a given subpool are equivalent - however, for the opponent, the different offers in a subpool might provide different utilities. When conceding, the UC agent will simply pick the new offer from the next pool. Whenever the opponent's offer evaluates to a utility which is larger than the utility of the current subpool, the agent accepts the offer. Otherwise, it will calculate the next counter offer from the offer pool which is the most similar to E_{i-1}^{B}. The similarity between two offers is defined as the sum of squared difference for each issue:

$$\mathbf{O_i^A} = \arg\min_{\mathbf{O}}(||\mathbf{O} - \mathbf{E_{i-1}^B}||^2), \text{for } U^A(O) \geq U_{sp}(i) \text{ and } U^A(O) < U_{sp}(i-2) \tag{18}$$

If the agent reaches the last subpool without a deal, it quits the negotiation and takes the conflict deal.

4.5 Augmenting strategies with collaborativeness analysis

In our current setting, negotiation happens in physical time, each negotiation round taking time t_r. The decision to close the negotiation (by either accepting the current offer, or by quitting with the conflict deal), should depend on the agent's view of the possible benefits it can obtain if it continues the negotiation weighted against the time delays this would involve. If no deal is possible, the agent is wasting utility by negotiating.

The collaborativeness metric we introduced in Equation 11 was developed precisely for the purpose of characterizing the potential deals in a scenario. We need to emphasize that a high collaborativeness metric does not necessarily guarantee a negotiation success, because the agents need to find those mutually beneficial deals, which depends on the offer formation strategies. On the other hand, even if the collaborativeness is low, the agent might hope to "trick" the opponent in a deal which is only marginally rational for the opponent, but much better for the agent. Overall, however, the collaborativeness metric is a good predictor of negotiation success of certain scenario.

Thus, it makes sense to augment the negotiation strategies with collaborativeness analysis. These augmented strategies would alter their behavior in function of the collaborativeness of the current scenario, for instance, quitting earlier the negotiation for low collaborativeness scenarios.

The problem with the collaborativeness metric $\Xi^{(A)}$ is that it can be evaluated only by a full knowledge agent (e.g. a supervisor). The agent participating in a negotiation starts with zero knowledge, but it can gradually acquire information from the negotiation. At the second round of the negotiation, the agent assumes its absolute best time $C_{ab}^{(A)}$ as the rationality constrained best time $C_{rcb}^{(A)\{B\}}$. It will approximate the rationality constrained social best time $C_{rcsoc}^{(A),\{B\}}$ as the averaged utility between its first offer and evaluation of the opponent's first offer. Thus the agent will estimate the collaborativeness as:

$$\Xi \approx \Xi_{estimate}^{(A)}(E_2^B) = \frac{C_{conflict}^{(A)} - \frac{C^{(A)}(O_1^{(A)}) + C^{(A)}(E_2^B)}{2}}{C_{conflict}^{(A)} - C^{(A)}(O_1^{(A)})} = \frac{U_{ab}^{(A)} + U^{(A)}(E_2^B)}{2 \times U_{ab}^{(A)}} \tag{19}$$

If this value is negative, it can be viewed as non-collaborative scenario with collaborativeness of zero.

Let us now see how this value can be used by the negotiation strategies. The internal negotiation deadline augmented with collaborativeness analysis (IND+CA) will compare the estimated collaborativeness with a threshold $\Xi_{threshold}$. If the estimate is smaller, the agent will quit the negotiation either by accepting the opponents first offer (if it is rational) or by taking the conflict deal. For $\Xi_{estimate}^{(A)}(E_2^B) > \Xi_{threshold}$ the IND+CA agent will change its negotiation deadline according to the following formula:

$$n'_{max} = \begin{cases} 0 & \text{if } \Xi_{estimate} < \Xi_{threshold} \\ n_{max} \times \frac{\Xi_{estimate} - \Xi_{threshold}}{1 - \Xi_{threshold}} & \text{otherwise} \end{cases} \tag{20}$$

The uniform concession with collaborativeness analysis (UC+CA) agent, will also compare the estimated collaborativeness with a threshold $\Xi_{threshold}$. If the estimate is smaller, the agent will quit the negotiation either by accepting the opponents first offer (if it is rational) or by taking the conflict deal. For $\Xi^{(A)}_{estimate}(E^B_2) > \Xi_{threshold}$ the IND+CA agent will change its conceding pace α according to the following formula:

$$\alpha' = \begin{cases} 1 & \text{if } \Xi_{estimate} < \Xi_{threshold} \\ \alpha \times \frac{1-\Xi_{threshold}}{\Xi_{estimate}-\Xi_{threshold}} & \text{otherwise} \end{cases} \tag{21}$$

The intuition behind this update is the absolute best offer can get the agent to reach its destination, assuming it has an ideal opponent in an ideal scenario. Its utility should be similar with the rationality constrained best utility, as the latter one just adds two geometric restrictions in the solution space. On the other hand, taking the average between the utility of the two first offers seems to be fair, if the utility function is linear and both agents concede until they meet in the middle. In a low collaborative scenario, the two agent need long time to search the deal. Even they form an agreement at last, the utility of the deal may not compensate the cost of negotiation. In this case, it should be wise for the learning agent to drop the negotiation immediately or increase the conceding pace so that they can end the negotiation quickly.

The estimation above didn't consider the impact of time scale t_r. If each negotiation round takes too much time, the agent should continue to accelerate the negotiation. We let the agent remember the best evaluation in history which has the most pragmatical utility, and continuously check if such evaluation is irrational from baseline point of view. If the agent finds out the its baseline utility is less than a threshold, it will drop the negotiation by either sending "accept" message or "quit" the negotiation directly. The intuition behind this is when the negotiation time t_r is expensive, the baseline utilities of all un-explored deals are decreasing quickly. The best potential deal which has already been explored by agents somehow indicates the decreasing speed of all un-explored deals. If its baseline utility is less than a threshold, the agent should quit the negotiation immediately to avoid the further damage.

5 Experimental results

In this paper we defined an approach to measure the collaborativeness of convoy negotiation scenarios. We also described an approach through which an agent starting the negotiation with zero knowledge can estimate the collaborativeness in the first several rounds of negotiation. Finally, we have shown the way in which the negotiation strategies can be augmented to take into account the estimate of the collaborativeness values.

In this section we proceed to experimentally validate the proposed metric, its estimation and application in negotiation strategies. The questions we plan to answer are:

- What is the distribution of collaborativeness in scenarios?
- How good is the estimation of collaborativeness described in Equation 19?
- What is the relative performance of the IND, UC, IND+CA and UC+CA strategies over a wide range of collaborativeness settings and various negotiation partners?

We implemented the CRF game in the Yet Another Extensible Simulator (YAES) environment [1]. The proposed negotiation protocols and the negotiation strategies IND, UC, IND+CA and UC+CA have been implemented exactly as described in the previous sections, with the offers encoded as FIPA ACL messages.

5.1 The distribution of the collaborativeness

The distribution of the collaborativeness provides the answer to the question: if we pick a random scenario, is it going to be competitive or collaborative? Naturally, the distribution of the collaborativeness depends on the distribution of the source and destination locations of the scenarios, as well as the distribution of the speed of the agents. Let us assume that the source and destination are distributed uniformly in rectangular areas situated immediately on the left and right side of the forest. To study a variety of possible distributions we consider three settings corresponding to the source and destination areas shown in Figure 2. For each setting, we generate 1000 scenarios by choosing the source and destination according to a uniform spatial distribution from the corresponding source and destination rectangles. We calculate the value of collaborativeness according to Equation 11, and assemble the values in a 10-bucket histogram. The three resulting histograms are shown in Figure 3.

We can make the following observations:

Setting 1: has the source and destination areas a square of the same height as the forest. The histogram shows a U-shape, with higher number of scenarios falling at the higher and lower extremes of collaborativeness.

Setting 2: has the source and destination areas rectangles of the same height as the forest but a width of half as much. The corresponding histogram shows a similar U-shape like in the previous case, but it is shifted towards the higher collaborativeness. We conclude that the closer is the forest to the source and destination, the higher the probability that forming a coalition to traverse the forest will be advantageous.

Setting 3: has the source and destination areas square and half the height of the forest. We find that the distribution of the collaborativeness is weighted toward the higher values.

This result matches our intuitive expectations. For instance, citizens in tightly packed cities such as New York and San Francisco rely more on public transportation, as their source and destination locations are frequently correlated. In cities spread over large areas such as Orlando or Phoenix, the transportation interests are rarely collaborative.

5.2 Accuracy of collaborativeness estimation

Figure 4 shows the scatter plot and the average of the estimated collaborativeness $\Xi_{estimate}$ function of the real value of the collaborativeness. In this graph, every point represents the estimate at negotiation round 2 for a total of 1000 scenarios. The closer it is the point to the diagonal, the better the estimate. The first observation is that the estimate is by no means perfect. Quite a number of datapoints fall far from the diagonal. There are even cases where a fully collaborative scenario is estimated to have near-zero collaborativeness. There are, however, no cases where low collaborativeness scenarios are estimated to have high collaborativeness. The average value, on the other hand, is tracking the diagonal relatively well, although it is always below the diagonal. Agents using this metric will likely err on the side of safety, underestimating collaborativeness rather than overestimating it.

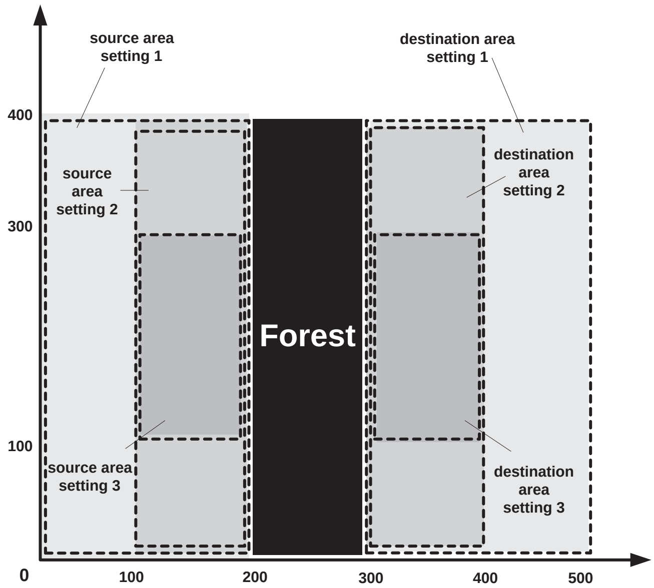

Figure 2: Three settings for the distribution of the source and destination areas for the study of the distribution of collaborativeness among scenarios.

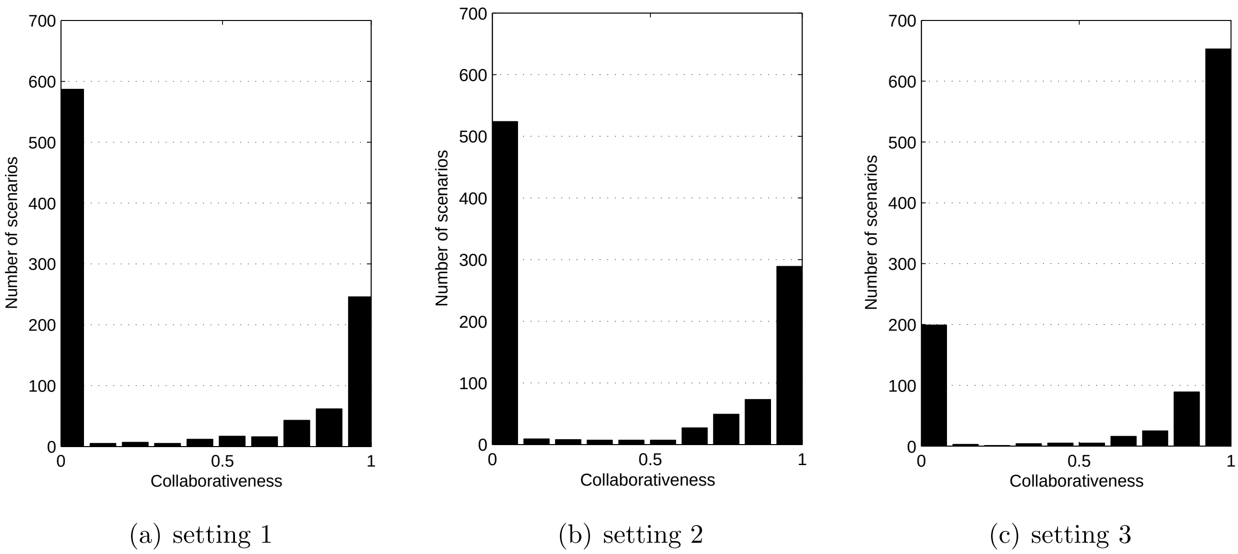

Figure 3: The comparison of collaborativeness distributions in three cases of restricted areas.

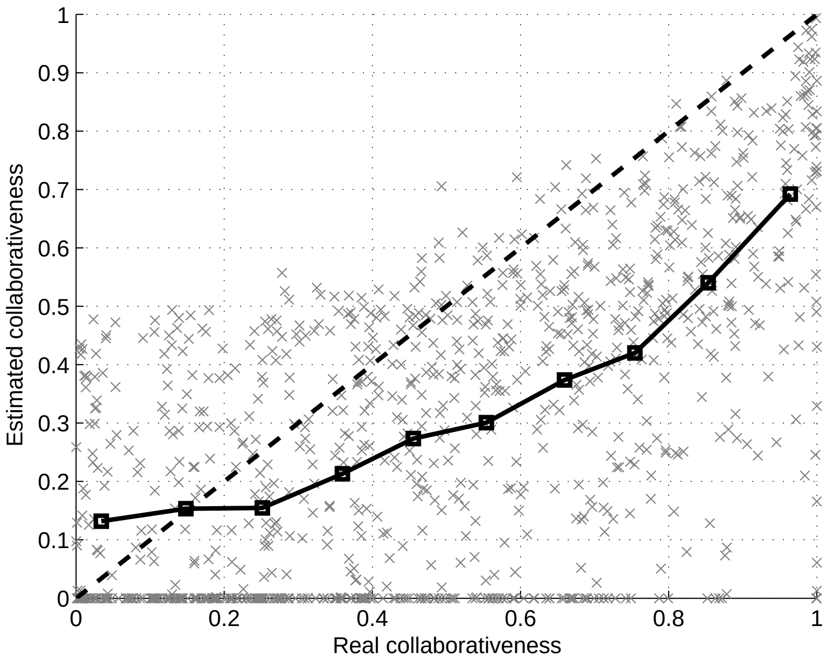

Figure 4: The estimated vs. real collaborativeness. Each point in the scatter plot corresponds to one scenario. The solid line is the average estimate for different collaborativeness values.

Overall, the estimate of collaborativeness is satisfactory, considering that we are only two negotiation rounds in a negotiation started with zero knowledge. It also opens the possibility of future work towards of a more accurate estimation based on information acquired in subsequent negotiation rounds.

5.3 Negotiation performance

In the following we investigate the performance of the negotiation strategies IND and UC and their variants augmented with collaborativeness analysis IND+CA and UC+CA. In a setting where the negotiation takes place physical time, the negotiation performance of the agents can be considered from two points of view. The pragmatic relative utility (Definition 7) measures the balance between the negotiation results of the participating agents. An agent which frequently manages to convince the opponent to concede more will have a high relative utility. This is a pragmatic measure which does not depend on the negotiation time. The baseline utility (Definition 1) on the other hand, considers the time spent during negotiation as part of the cost. Certain strategies might choose to exit the difficult negotiation scenarios early even at the cost of an unrequited concession, which damages their relative utility, but can boost their baseline utility.

For the experiments describe in this section, we consider a negotiation round to take $t_r = 0.5$. The deadline for the IND strategy is $n_{max} = 40$, the conceding speed of the UC strategy is $\alpha = 0.05$. For the strategies which use collaborativeness analysis, we let the threshold $\Xi_{threshold} = 0.3$.

In the first set of experiments we compare the IND and UC strategies in all four possible pairings

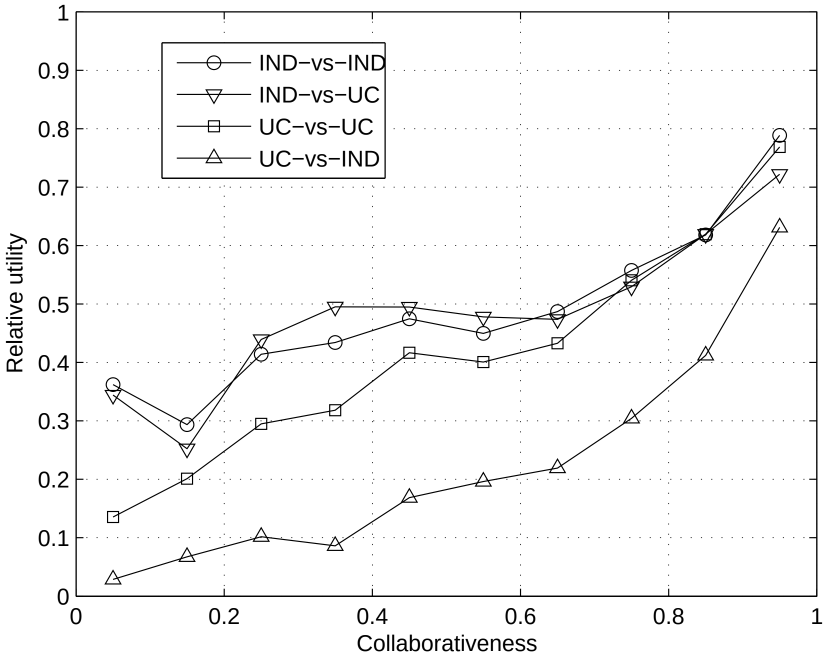

Figure 5: The relative pragmatic utility in the function of collaborativeness for the following strategy pairs: IND vs IND, IND vs UC, UC vs UC and UC vs IND.

(IND vs IND, IND vs UC, UC vs IND and UC vs UC). We plot the relative pragmatic utility in Figure 5 and the baseline utility in Figure 6. The first observation is that for all settings the utility increases monotonically with the collaborativeness, but there is a significant variation among the negotiating strategy pairs. For the relative pragmatic utility the IND strategy always outperforms UC. For the baseline utility, however, the order is different, the best performance being obtained by the UC vs UC pairing. The baseline utility graphs shows how difficult is to obtain a positive negotiation result under the settings of our problem: the UC vs UC pairings yields negative average for $\Xi < 0.45$, but IND vs UC is negative for $\Xi < 0.65$ and UC vs IND and IND vs IND is negative for $\Xi < 0.85$! These negative values are a result of long negotiation sessions trying to obtain a better concession from the opponent, while loosing more on the time spent for each negotiation round.

Figure 7 shows the baseline utility for the IND, IND+CA, UC and UC+CA strategies when negotiating with opponents using the same strategy. In addition to the averages, these graphs also show the scatter plot of the individual negotiation results. The immediate observation is the significant improvement of the IND+CA and UC+CA strategies for the low collaborativeness values. While the scatter plot shows a large number of negotiations finishing in the negative for IND and UC, there are virtually none of them for IND+CA and UC+CA.

Another noteworthy feature is the visible concentration of points on the -20 horizontal line at Figure 7-a (IND vs IND). This line corresponds to the $t_r \cdot n_{max} = 0.5 \times 40 = 20$ value of the negotiations where the IND agent was forced to take the conflict deal after reaching the internal negotiation deadline.

A similar concentration of points can be found around the line corresponding to zero utility for the IND+CA and UC+CA graphs. These points correspond to the case when the collaborativeness

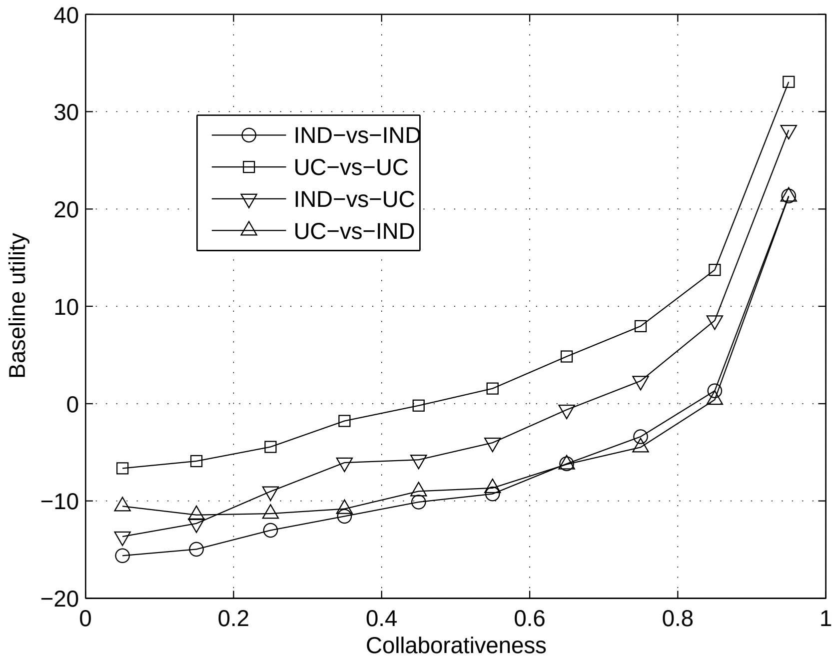

Figure 6: The baseline utility in the function of collaborativeness for the following strategy pairs: IND vs IND, IND vs UC, UC vs UC and UC vs IND.

analysis component dictated an early termination of the negotiation.

For a closer analysis of the relative performance, we ran a series of experiments where all the proposed strategies (IND, UC, IND+CA and UC+CA) negotiate against the same opponent, IND for Figure 8 and UC for Figure 9. The trend is similar for all the combinations on these graphs: the strategies augmented with collaborativeness analysis significantly outperform the other ones for low collaborativeness values, limiting their losses to the cost of the several negotiation rounds necessary to come up with an estimate. For scenarios of high collaborativeness, on the other hand, the performance is roughly equivalent. In some cases, such as the UC vs UC and UC+CA vs UC in Figure 9, the CA version might perform slightly worse for the highest collaborativeness values. This phenomena appears because of the inaccuracies of the collaborativeness estimation.

6 Conclusions

In this paper we considered the problem of negotiating convoy formation under time constraints. This is a relatively complex multi-issue negotiation with two spatial and two temporal issues. Not all the offers are feasible, the utility is a non-linear function of the issues and the offer formation is difficult, as it might require complex path calculations. We developed a collaborativeness metric which allows us to put a quantitative value on our intuition of “easy” and “hard” negotiation scenarios. The metric is not dependent on the negotiation scenario, but can be evaluated only by a full knowledge supervisor. We describe an approach through which agents starting with zero knowledge can estimate the collaborativeness of the scenario using information acquired from the first several negotiation rounds. Finally, we show how this estimate can be used to augment negotiation scenarios with

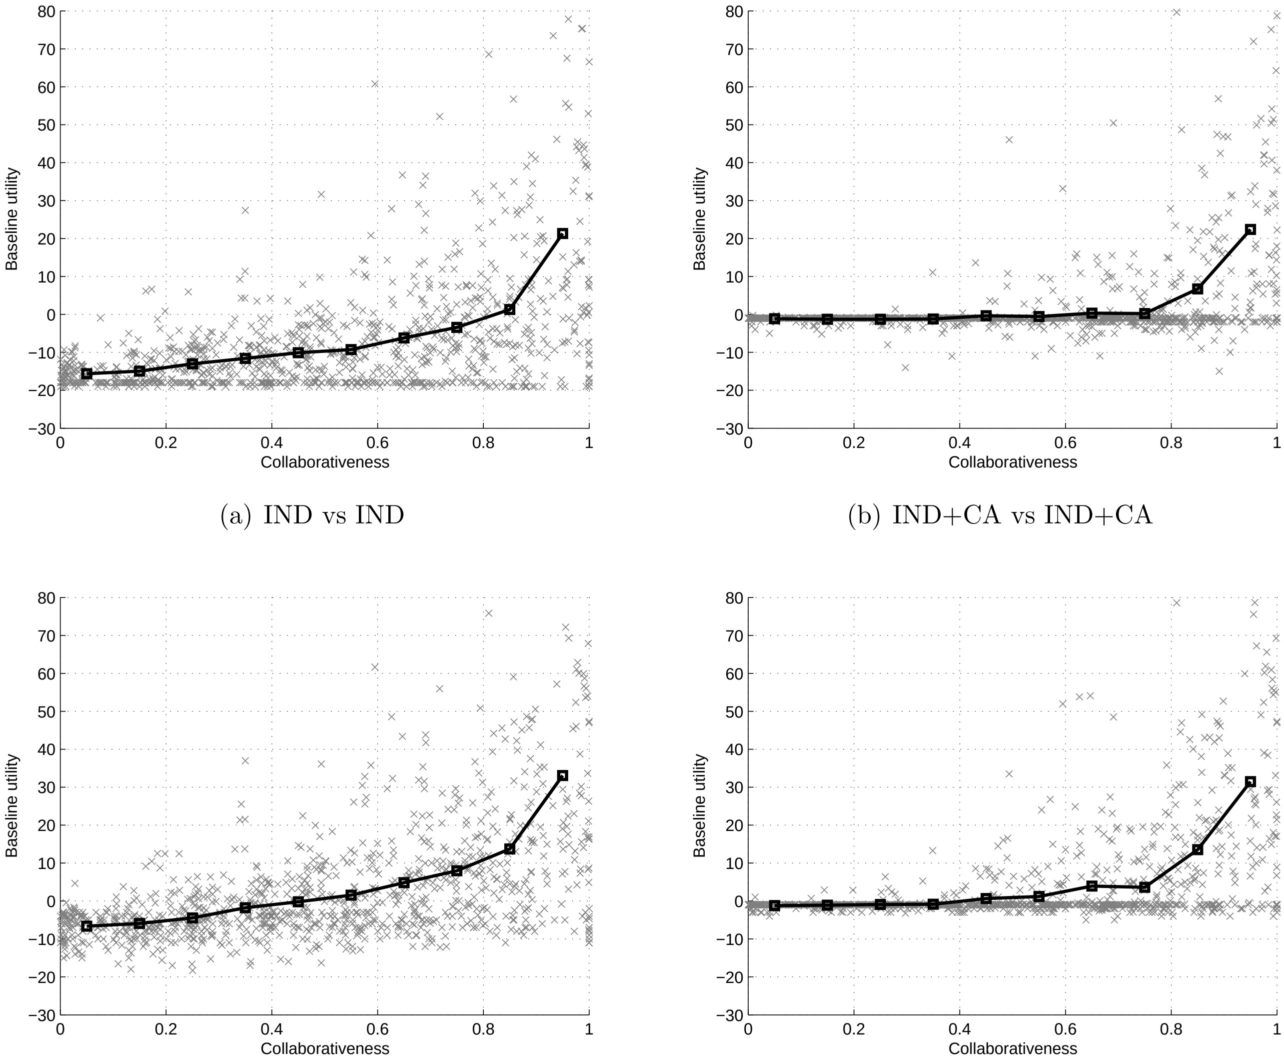

(a) IND vs IND

(b) IND+CA vs IND+CA

(c) UC vs UC

(d) UC+CA vs UC+CA

Figure 7: The baseline utility of the IND, IND+CA, UC and UC+CA agents negotiating with opponents using the same strategy.

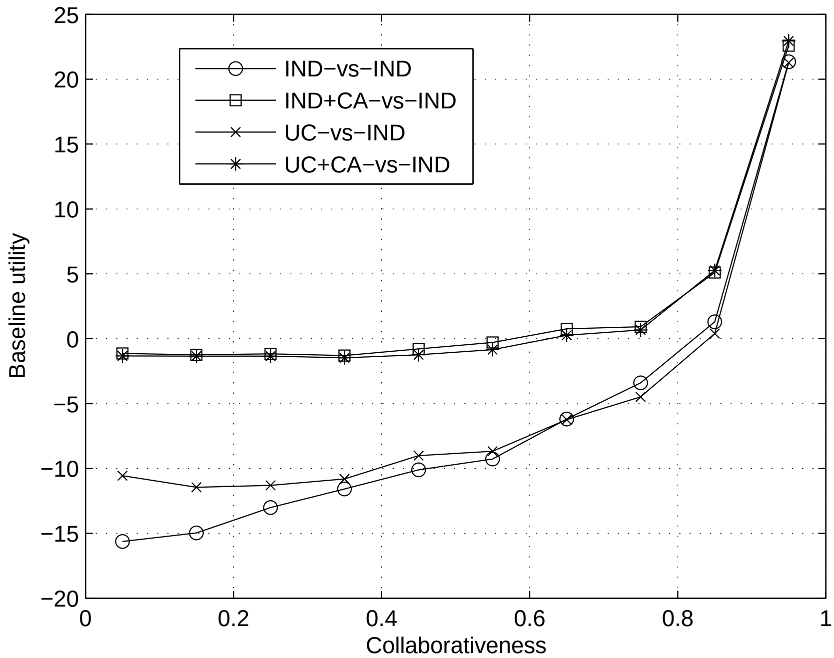

Figure 8: Baseline utility of the IND, IND+CA, UC and UC+CA strategies negotiating against an agent using the IND strategy.

collaborativeness analysis. In a series of experimental studies, we have shown that the augmented strategies significantly outperform the original strategies for low collaborativeness scenarios, and closely match them for high collaborativeness scenarios.

Future work is planned in several directions. We need to study the influence of the accuracy of the collaborativeness estimate on the performance of the agent. Agents might employ strategies to mislead the opponent into believing that a scenario is more collaborative than it is in reality.

Another direction of future work is when agents are acting while negotiating. The assumption that agents hold still during negotiation might not be the right model for real-world convoy formation. While negotiating, agents might keep moving either on the conflict deal trajectory, or on the trajectory of the predicted deal. Such agents are controlled by the pair of negotiation and action strategies. In addition, the utility values and the collaborativeness dynamically change in time. Another direction is the exploration of other negotiation protocols, including more sophisticated argumentation models. Such models would also need to explicitly model the temporal cost associated with the creation and selection of arguments.

Acknowledgements

This work was partially funded by the NSF Information and Intelligent Systems division under award 0712869.

This research was sponsored in part by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-06-2-0041. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official

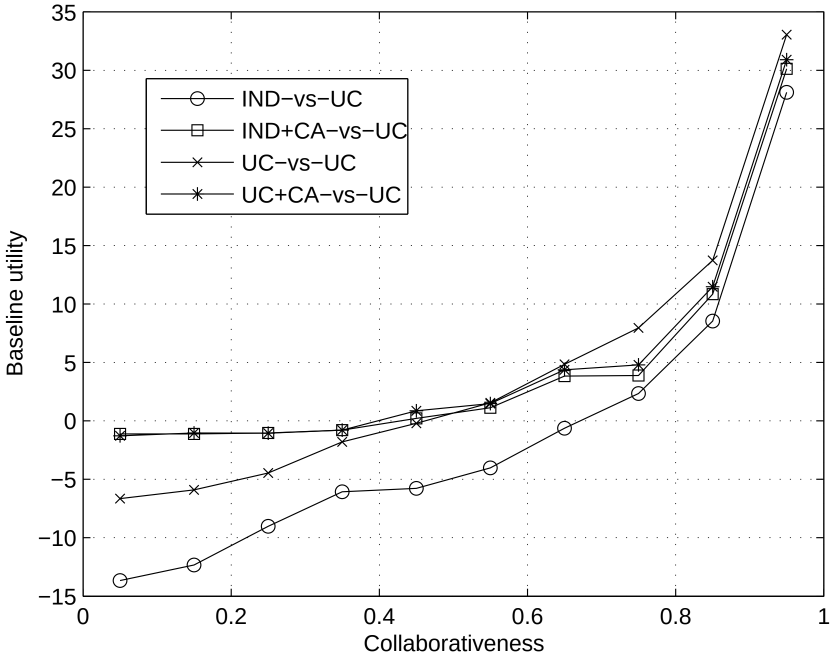

Figure 9: Baseline utility of the IND, IND+CA, UC and UC+CA strategies negotiating against an agent using the UC strategy.

policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

References

[1] L. Bölöni and D. Turgut. YAES - a modular simulator for mobile networks. In *Proceedings of the 8-th ACM/IEEE International Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems MSWIM 2005*, pages 169–173, October 2005.

[2] E. Crawford and M. Veloso. An experts approach to strategy selection in multiagent meeting scheduling. *Autonomous Agents and Multi-Agent Systems*, 15(1):5–28, 2007.

[3] S. Fatima, M. Wooldridge, and N. Jennings. Multi-issue negotiation under time constraints. In *Proceedings of the first international joint conference on Autonomous agents and multiagent systems: part 1*, pages 143–150. ACM New York, NY, USA, 2002.

[4] S. S. Fatima, M. Wooldridge, and N. R. Jennings. Multi-issue negotiation with deadlines. *J. Artif. Intell. Res. (JAIR)*, 27:381–417, 2006.

[5] M. Golfarelli and S. R. D. Maio. Multi-agent path planning based on task-swap negotiation. In *Proceedings 16th UK Planning and Scheduling SIG Workshop*, pages 69–82, 1997.

[6] M. Golfarelli and S. Rizzi. Spatio-temporal clustering of tasks for swap-based negotiation protocols in multi-agent systems. In *Proceedings 6th International Conference on Intelligent Autonomous Systems*, pages 172–179, 2000.

[7] T. Ito, M. Klein, and H. Hattori. A multi-issue negotiation protocol among agents with nonlinear utility functions. *Multiagent and Grid Systems*, 4(1):67–83, 2008.

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

[9] Y. Luo and L. Bölöni. Children in the forest: towards a canonical problem of spatio-temporal collaboration. In *The Sixth Intl. Joint Conf. on Autonomous Agents and Multi-Agent Systems (AAMAS 07)*, pages 986–993, 2007.

[10] Y. Luo and L. Bölöni. Collaborative and competitive scenarios in spatio-temporal negotiation with agents of bounded rationality. In *Proceedings of the 1st International Workshop on Agent-based Complex Automated Negotiations , in conjunction with the The Seventh Intl. Joint Conf. on Autonomous Agents and Multi-Agent Systems (AAMAS 08)*, pages 40–47, 2008.

[11] M. J. Osborne and A. Rubinstein. *A Course in Game Theory.* The MIT Press, 1994.

[12] S. Saha and S. Sen. Negotiating efficient outcomes over multiple issues. In *5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2006)*, pages 423–425, 2006.

[13] T. Sandholm and N. Vulkan. Bargaining with deadlines. In *AAAI/IAAI*, pages 44–51, 1999.