Collaborative and competitive scenarios in spatio-temporal negotiation with agents of bounded rationality

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

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

ABSTRACT

In spatio-temporal negotiation evaluating an offer for feasibility or utility often requires computationally expensive path planning, thus practical negotiation strategies can evaluate only a small subset of the possible offers during offer formation. 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 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 off with a loss of utility for the other agent).

Using the Children in the Rectangular Forest (CRF) game as a canonical model of spatio-temporal negotiation, we develop a series of quantitative metrics for the characterization of deals in relation to the possibilities of the scenario and the interest of the other agent. We also develop a metric for the collaborativeness of the scenario. Through an experimental study involving three negotiation strategies of increasing complexity, we show that the proposed metrics match our intuition about the scenarios and can serve as a tool in analyzing and developing strategies as well as in designing negotiation mechanisms promoting cooperative behavior.

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 problems appear as sub-problems in many practical applications such as transportation and disaster rescue. In previous work [6] we have identified five differentiating features of the spatio-temporal negotiation problems: (1) heterogeneous issues (which include spatial locations and time points), (2) non-monotonic valuation of issues, (3) an evolving environment, (4) offers need to be verified for feasibility (usually by both parties) and (5) there is an interaction between the negotiation time and the physical time. As this category of problems can not be conveniently represented through the "split the pie" model of negotiation, we proposed a model, Children in the Rectangular Forest (CRF), which captures the main features of these problems and can be used as an canonical problem for the study of spatio-temporal negotiation strategies.

The evaluation (and creation) of offers in spatio-temporal negotiation 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).

To show the intuition behind collaborativeness in a negotiation scenario, let us first consider the simpler scenario of the split the pie game. Here two agents are negotiating over the partitioning of a pie. As the parts allocated to one agent are lost for the other agent, the single pie game is fully competitive. This is true even for the cases when we are partitioning over multiple pies. Note that although all zero-sum games are fully competitive, not all fully competitive games are zero sum[1]. For instance, a "split multiple

ACAN-08 May 12–16, 2008, Estoril, Portugal
.

[1]Some game theory texts, such as [7] equate fully competitive with zero sum, by making the assumption that the utility function is just a convenient expression of the preference ordering. In our case, 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 language the first scenario is fully competitive and zero sum, while the second

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.

For spatio-temporal negotiations similar considerations apply; however both the definition of collaborativeness and the calculation of the optimal deal is more difficult.

2. RELATED WORK

While automated negotiation [5] 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 [9] 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 [2] 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.

Golfarelli et al. [3] 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 [4], 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 [8] 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 [1] 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.

scenario is fully competitive but *not* zero sum.

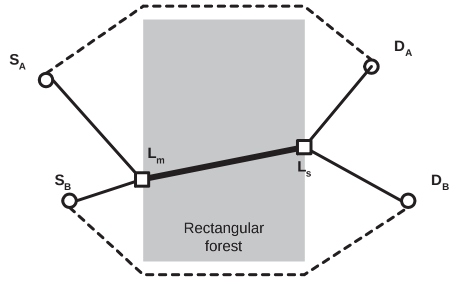

Figure 1: The Children in the Rectangular Forest problem. The trajectories associated with the conflict deal are shown with an interrupted line, while the trajectories corresponding to a possible agreement are shown with a continuous line.

3. METRICS FOR THE CHARACTERIZATION OF THE CRF PROBLEM

3.1 The CRF problem

The Children in the Rectangular Forest (CRF) problem considers a world in which two children A and B are moving from source points S_A and S_B on one end of a rectangular forest to their respective destination points D_A and D_B on the other side of the forest (see Figure 1). Traveling alone, the children can not enter the forest and they need to go around it. However, if they join together in a coalition, the children can traverse the forest, thus possibly shorten their trip. The children travel at different velocities v_A and v_B; when traveling together, they move with the velocity of the slower child. The selfish objective of each child is to reach his or her respective destination as early as possible.

The two children are using negotiation to agree on the parameters of the deal. This is a 4-issue negotiation, with two issues being spatial locations (meeting point L_m and split point L_s), and two issues time points (meeting time t_m and split time t_s). If the negotiation is unsuccessful, the agents are taking the conflict deal; that is, they go around the forest.

Several properties of rational deals in the CRF problem, are summarized below:

- The optimal trajectories of the conflict deal and the collaborative deal are sequences of straight segments. However, the trajectory of an agent *before* making a deal might be curvilinear, reflecting, for instance, the evolving estimation of the agent about the likelihood of the deal.
- The meeting point and the split point of any Pareto optimal deal is at the edge of the forest in non-degenerate cases. In degenerate cases, there is a deal with the same time to destination for both agents where the points are the edge of the forest.

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

We can make a more general observation concerning negotiation problems which involve meeting and splitting locations and times of convoys. If (a) the utility of all participants depends only on the time to reach the destination and (b) the negotiation is full knowledge, then the negotiation can be reduced to the spatial components, as the temporal components can always be calculated by considering the time it takes for the last participant to reach the given point.

However, if the negotiation is not full-knowledge, then the agent might find it impossible to form an offer where the temporal values are guaranteed to be feasible for both agents. In the simplest case, an agent might not know the current location of the negotiation partner, and thus, naturally, it cannot calculate the time at which it can reach a certain point.

We call a *fully specified offer* a quadruple $O = \{L_m, t_m, L_s, t_s\}$ which specifies both the spatial and temporal components of an offer. A *spatially specified offer* specifies only the spatial components of the offer: $O = \{L_m, ?, L_s, ?\}$. An agent A can complete a spatially specified offer by calculating the timepoints $t_m^{(A)}$ and $t_s^{(A)}$ 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:

$$BTC^{(A)}(O) = BTC^{(A)}(\{L_m, ?, L_s, ?\}) = \{L_m, t_m^{(A)}, L_s, t_s^{(A)}\} \quad (1)$$

3.2 Metrics

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 game (such as the location and destination of the agents).

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* $C^{(A)}(O)$ of agent A for a particular offer $O = \{L_m, t_m, L_s, t_s\}$ 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.

$$C^{(A)}(O) = \max\left(\frac{|S_A, L_m|}{v_A}, \frac{|S_B, L_m|}{v_B}\right) + \frac{|L_m, L_s|}{\min(v_A, v_B)} + \frac{|L_s, D_A|}{v_A} \quad (2)$$

The *time to destination of the conflict deal* $C_{conflict}^{(A)}$ is the time for the agent to reach its destination if it does not make any deal. This value of the baseline of the negotiation; a rational agent will not accept an offer which will yield a time to destination later than the conflict deal.

Definition 1. *The* **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.*

$$U^{(A)}(O) = C_{conflict}^{(A)} - C^{(A)}(O) \quad (3)$$

Definition 2. *We define the* **absolute best time to destination** *$C_{ab}^{(A)}$ 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.

$$C_{ab}^{(A)} = \frac{|S_A, D_A|}{v_A} \quad (4)$$

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** *$C_{acb}^{(A),\{B\}}$, 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.

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.*

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

$$C_{conflict}^{(A)} \geq C_{rcb}^{(A)\{B\}} \geq C_{acb}^{(A)\{B\}} \geq C_{ab}^{(A)} \quad (5)$$

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

$$C^{(A)}\left(O_{rcb}^{(A),\{B\}}\right) = C_{rcb}^{(A)\{B\}} \quad (6)$$

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 $C^{(B)}$:*

$$\forall C^{(B)}, C_1^{(A)} \geq C_2^{(A)} \Rightarrow C_{social}(C_1^{(A)}, C^{(B)}) \geq C_{social}(C_2^{(A)}, C^{(B)})$$
$$\forall C^{(A)}, C_1^{(B)} \geq C_2^{(B)} \Rightarrow C_{social}(C^{(A)}, C_1^{(B)}) \geq C_{social}(C^{(A)}, C_2^{(B)}) \quad (7)$$

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

$$O_{social} = \underset{O}{\operatorname{argmin}}\,(C_{social}(O)) \quad (8)$$

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.

$$C_{social}(O) = C^{A+B}(O) = C^{(A)}(O) + C^{(B)}(O) \quad (9)$$

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:

$$C_{rcsoc}(O) = \begin{cases} +\infty & \left(C^{(A)}(O) > C^{(A)}_{conflict}\right) \vee \left(C^{(B)}(O) > C^{(B)}_{conflict}\right) \\ C_{social}(O) & \text{otherwise} \end{cases} \quad (10)$$

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

$$O_{rcsoc} = \underset{O}{\operatorname{argmin}} \left(C_{rcsoc}(O)\right) \quad (11)$$

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:*

$$\Xi^{(A),\{B\}} = \frac{C^{(A)}_{conflict} - C^{(A),\{B\}}_{rcsoc}}{C^{(A)}_{conflict} - C^{(A),\{B\}}_{rcb}} \quad (12)$$

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:*

$$U^{(A),\{B\}}_{rel}(O) = \frac{C^{(A)}_{conflict} - C^{(A)}(O)}{C^{(A)}_{conflict} - C^{(A)}_{rcb}, \{B\}} \quad (13)$$

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.

Definition 8. *We define the* **competitive utility of an offer** *for agent A as the ratio of the utility of the offer to the utility of the rationality-constrained social best offer:*

$$U^{(A),\{B\}}_{comp}(O) = \frac{C^{(A)}_{conflict} - C^{(A)}(O)}{C^{(A)}_{conflict} - C^{(A)}_{rcsoc}, \{B\}} = \frac{U^{(A),\{B\}}_{rel}(O)}{\Xi^{(A),\{B\}}} \quad (14)$$

The competitive utility can range from 0 to $\frac{1}{\Xi^{(A),\{B\}}} > 1$. Intuitively, $U^{(A),\{B\}}_{comp}(O) = 1$ means that the agent obtained the social deal. If $U^{(A),\{B\}}_{comp}(O) < 1$ and $U^{(B),\{A\}}_{comp}(O) > 1$ it can be interpreted that agent B "outsmarted" agent A in the negotiation. If is not possible that both competitive utilities to be above 1, as the specific offer would become the new social deal. However, it is possible that both values are below 1, which means that the negotiating agents agreed on a deal which is not Pareto optimal. As finding a Pareto optimal deal in spatio-temporal problems is a non-trivial collaborative search task, this can happen quite often for real-world negotiations.

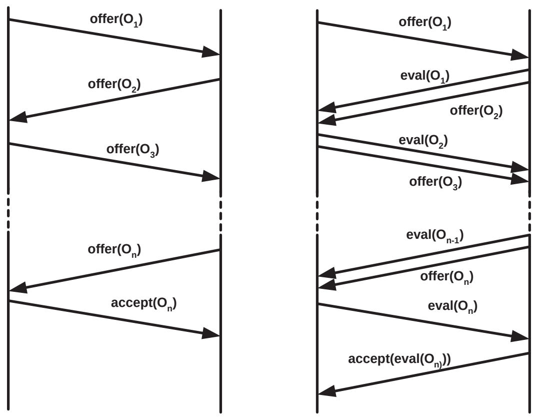

Figure 2: Example runs of negotiation protocols. (left) Exchange of binding offers (EBO). (right) Exchange of offers with mandatory, non-binding evaluations (EBOMNE).

3.3 Negotiation protocols for CRF

For most negotiation settings, it is assumed that the complexity of the negotiation lies in the strategy, while the protocol is a relatively trivial alternating exchange of offers by the two parties. Such a simple protocol would still work well for the CRF game with full knowledge. In the case of incomplete knowledge, however, the difficulty of forming a feasible offer as well as evaluating whether a given offer represents a concession or not, make simple offer-exchange protocols little better than random search. The simple protocol can be enhanced by schemes in which the agents add additional information of the negotiation flow to aid the negotiation partner in the offer formation. In the following we illustrate the design space for the CRF negotiation protocols through several examples.

Simple exchange of binding offers (EBO). In this simplest negotiation protocol, the agents are alternating in making fully specified offers in the form $O = \{y_m, t_m, y_s, t_s\}$. The offers are binding for the agents who made the offer, in the sense that once made by an agent and accepted by the other agent, the offer will be the outcome of the negotiation. An example run of this protocol is illustrated in Figure 2-left.

Exchange of binding offers with mandatory, non-binding evaluations (EBOMNE). In this protocol the agents are exchanging pairs of offers and evaluations. Agent A first chooses a spatially specified offer $O = \{y_m, ?, y_s, ?\}$, and computes the associated best time completion $BTC^{(A)}(O) = \{y_m, t^{(A)}_m, y_s, t^{(A)}_s\}$. This is the offer which A will send to agent B, which is guaranteed to be feasible for A and is binding for A. Agent B will calculate its own best time completion $BTC^{(B)}(O) = \{y_m, t^{(B)}_m, y_s, t^{(B)}_s\}$ for the same spatially specified offer. Using the two best time offers, B will form an *evaluation* of the initial offer

$$E(O) = \left\{y_m, \max(t^{(A)}_m, t^{(B)}_m), y_s, \max(t^{(A)}_s, t^{(B)}_s)\right\} \quad (15)$$

This evaluation has the form of an offer which is feasible for both agents, but it is not binding for the evaluating agent. Rather it represents a *critique* of the original offer, and such it helps the other agent in the formation of feasible offers. Note that the evaluation provides some information about the evaluating agent, but it does not immediately disclose the source, destination or maximum speed (for instance, it is providing less information then the exchange of the BTC would). Also, if the evaluation amounts to an offer which is not feasible for the evaluating agent, an empty evaluation $\emptyset$ will be returned instead.

In practice, agents would couple the evaluation of the received offer with the sending of new offer. An example run of this protocol is illustrated in Figure 2-right.

Exchange of binding offers with optional, non-binding evaluations (EBOONE). A variation of the previous protocol removes the requirement that the agents evaluate every received offer. The advantage of this protocol is that agents would not be required to disclose information in response to offers which they would not consider. An agent would normally evaluate only offers which are satisfactory from the point of view of the spatial components.

Other combinations are also possible. For instance, exchange of offers with optional but binding evaluations is the (near) equivalent of a simple exchange of offers strategy where one of the agents is choosing as its next offer the evaluation of the opponents' offer.

There is an interdependence between the negotiation protocol and the strategies of the agents. A negotiation strategy created for the EBO protocol can be trivially extended to the EBOMNE as the evaluation can be created automatically - but the strategy would not take advantage of the information contained in the evaluations. The same strategy can be also trivially extended to EBOONE, by choosing not to send any evaluation. It is more difficult to "downgrade" strategies which rely on information from evaluations such as in EBOMNE protocol, to protocols where this information might not be available, such as EBO.

The protocols described in the next section are expected to operate in the EBOMNE protocol, although the first protocol presented does not take advantage of the evaluations and it would work similarly under the EBO protocol.

4. NEGOTIATION STRATEGIES

In the following we briefly introduce three negotiation strategies for the CRF problem. The goal is to illustrate a broad range of strategies which the agents might deploy. For all these strategies the assumption is the "no initial information" setting - that is the agents start with no information about the parameters of the negotiation partner. The negotiation protocol is assumed to be the exchange of binding offers with mandatory, non-binding evaluations (EBOMNE); however, not all protocols are taking advantage of the information in the evaluations. Due to lack of space, the description of the strategies will be inevitably cursory.

4.1 Monotonic concession in space (MCS)

Although monotonic concession is one of the basic strategies in most negotiation settings, for the CRF game with incomplete information, monotonic concession is not possible. One compromise is to limit the concession to the spatial domain. This will usually, but not always represent a concession in terms of the utility function of the opponent.

The monotonic concession in space agent is parameterized by the pair (c_m, c_s) representing the concession rate in the meeting point and splitting point respectively.

The agent will start its negotiation by proposing an offer corresponding to its absolute best $O^{(A),1} = O^{(A)}_{ab} = \{y_m^{(A),1}, t_m^{(A),1}, y_s^{(A),1}, t_s^{(A),1}\}$. In response to this, the agent will receive an evaluation (which will contain the offer corresponding to the ability constrained best $O^{(A)}_{(acb)}$ and the counteroffer of agent B, $O^{(B),1} = \{y_m^{(B),1}, t_m^{(B),1}, y_s^{(B),1}, t_s^{(B),1}\}$.

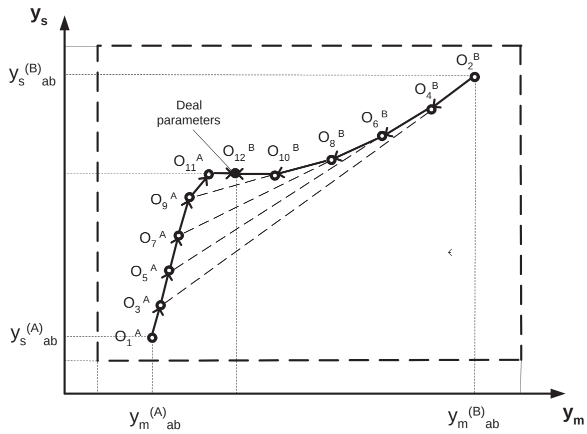

Figure 3: A negotiation trace between the a monotonic concession in space agent A and an internal negotiation deadline agent B.

The next offer of agent A is described by the following values:

$$y_m^{(A),n+1} = \begin{cases} y_m^{(A),n} - c_m & \text{if } y_m^{(B),n} < y_m^{(A),n} - c_m \\ y_m^{(B),n} & \text{if } y_m^{(B),n} < y_m^{(A),n} \leq y_m^{(B),n} + c_m \\ y_m^{(B),n} & \text{if } y_m^{(B),n} > y_m^{(A),n} \geq y_m^{(B),n} - c_m \\ y_m^{(A),n} + c_m & \text{if } y_m^{(B),n} > y_m^{(A),n} + c_m \end{cases} \quad (16)$$

with a similar expression for $y_s^{(A),n+1}$. Using the resulting spatially specified offer $\{y_m^{(A),n+1}, ?, y_s^{(A),n+1}, ?\}$, the agent will calculate the best time completion. The resulting fully specified offer will be evaluated for rationality. If the offer is not rational, the agent will break the negotiation. The agent will accept the opponent's offer if it evaluates to a utility which is higher than the next counteroffer the agent is about to make.

4.2 Internal negotiation deadline (IND)

In the internal deadline negotiation strategy the agent sets to itself a deadline (expressed as a number of negotiation rounds) and adapts the speed of concession in function of the remaining negotiation rounds and the difference between the current offers. If the number of rounds have expired without an agreement being reached the agent breaks the negotiation. The negotiation strategy is parameterized by the negotiation deadline.

Similarly to the MCS strategy, the agent starts by offering the absolute best $O^{(A)}_{ab}$. At every step the agent A will calculate the relative utility of its own previous offer based on the evaluation made by the opponent agent. Then it adapts the concession rate based on its own and the opponents' previous concession such that the deal will be reached at the negotiation deadline. Naturally, a deal can be reached sooner if the opponent agent accepts an offer. The agent will insist on an offer (by repeating it without change), if the offer is evaluated to be rational by the opponent while the opponents' offer is not rational for the agent. Figure 3 shows a negotiation trace between a MCS and IND agent. Note the adaptation of the concession speed by the IND agent.

4.3 Estimation of the opponents parameters (EOP)

The main problem with the previous two strategies is that it is difficult to create attractive offers, because the agents do not know whether a particular offer represents a concession. Furthermore, the offering agent can not even evaluate the utility of a spatially specified offer for itself, as the offer needs to be ability and rationality constrained by the opponent agent.

The EOP strategy tries to improve its offer formation by estimating the opponent's speed and current location based on the offers and evaluations made by the opponent. With this estimate, the agent can calculate the ability constrained t_m and t_s values for a spatially specified offer. The EOP strategy maintains an *offer pool* which contains a set of pre-calculated offers, which are believed to be rational for both agents. The agent will make offers from the pool in the decreasing order of its own utility. The agent accepts the opponents' offer if it is higher than the next offer in the pool, and breaks the negotiation if the offer pool is empty. Naturally, the offer pool needs to be updated whenever the estimates change.

Note that will all these negotiation strategies it is possible that two agents will not reach an agreement despite the fact that a mutually rational offer exists.

5. EXPERIMENTAL RESULTS

In the following we describe the results of a series of empirical studies which study the influence of the collaborativeness of the scenario on the results of negotiations. In particular, the three main questions we plan to answer is:

- What is the distribution of collaborativeness in scenarios?
- How is the efficiency of the negotiation process affected by the collaborativeness of the scenario?
- How does the performance of a negotiation strategy affected by the collaborativeness of the scenario?

The negotiation protocol considered was exchange of offers with mandatory, non-binding evaluations (EBOMNE). We have used three agents: (a) monotonic concession in space (MCS), with concession size $(2,2)$, (b) internal negotiation deadline (IND) with internal deadline 30 and (c) estimation of the opponents' parameters (EOP), with the offer generation resolution of 5.

5.1 The distribution of the collabo

The distribution of the collaborativeness provid the question: if we pick a random scenario, is it g petitive or collaborative? Naturally, the distributio rativeness depends on the distribution of the source locations of the scenarios, as well as the distributio the agents. Let us assume that the source and des tributed uniformly in rectangular areas situated im left and right side of the forest. To study a variety o butions we consider three settings corresponding to the source and destination areas shown in Figure 4. 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 12, and assemble the values in a 10-bucket histogram. The three resulting histograms are shown in 5.

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.

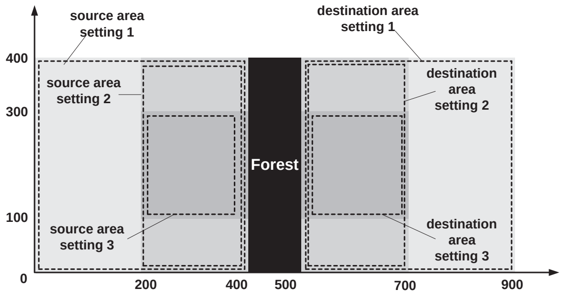

Figure 4: Three settings for the distribution of the source and destination areas for the study of the distribution of collaborativeness among scenarios.

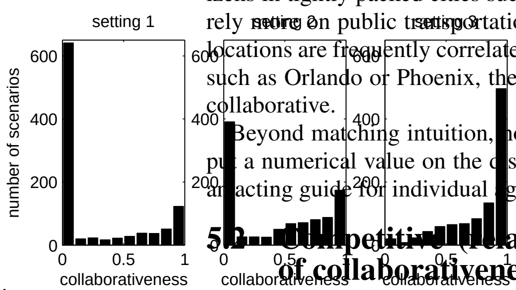

Figure 5: The comparison of collaborativeness distributions in three cases of restricted areas.

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.

Beyond matching intuition, however, our approach allows us to put a numerical value on the distribution, which is useful for both acting guide for individual agents

5.2 Competitive (relative) utility in function of collaborativeness

In our next set of experiments we study the relative utility achieved by specific strategies against specific opponents under in scenarios with various levels of collaborativeness. Naturally, if a deal falls through, the relative utility is zero.

One of the difficulties of our study is that we cannot generate directly random scenarios with predefined collaborativeness level. Thus we rely on *rejection sampling*, a technique borrowed from Monte Carlo simulation methods: we generate scenarios by picking source and distribution points according to a uniform distribution, calculate their collaborativeness, group them into 20 buckets, and then randomly reject scenarios from the buckets which are too full until they are uniformly filled with 500 scenarios each. The resulting collection of 10,000 scenarios was used in simulations. The agents are using the EBOMNE negotiation protocol. Once the negotiation terminated (either with an agreement or a conflict), we measured the relative utility of the deal.

To present both the variability in the results of negotiation, as well as the underlying trends, we choose to superimpose a scatterplot of the simulation results with a plot of the average values

Figure 6: The scatterplot and average values function of collaborativeness for two monotonic conceding in space (MCS) agents with identical parameters negotiating with each other.

Figure 8: The scatterplot and average values function of collaborativeness for two estimating the opponents' parameters (EOP) agents with identical parameters negotiating with each other.

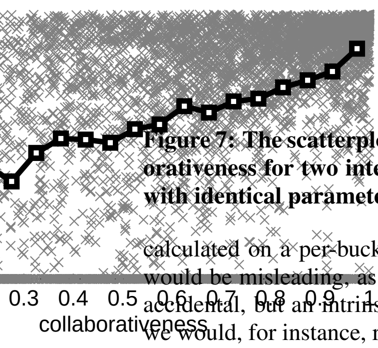

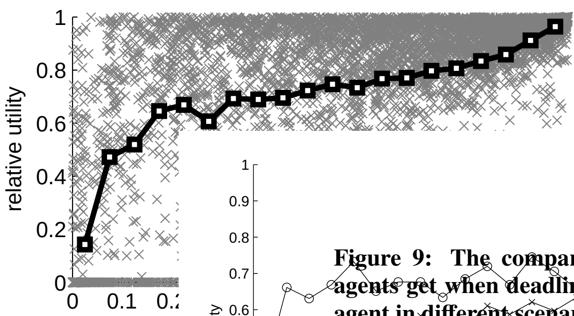

Figure 7: The scatterplot and avera… orativeness for two internal negoti… with identical parameters negotiati…

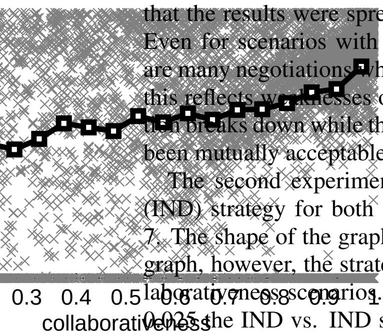

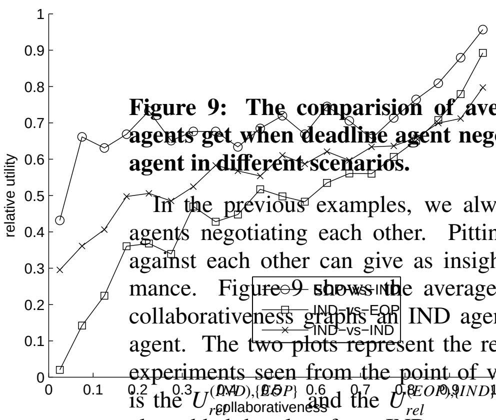

calculated on a per-bucket basis. Pl… would be misleading, as the spread o… accidental, but an intrinsic property … we would, for instance, run a larger number of experimen…

Figure 6 shows the scatterplot and average values of rel… ity for the case of two MCS agents negotiating against e… We find that, as expected, the relative utility is increasing… collaborativeness of the scenario. However, the scatterp… that the results were spread over a large range of relative… Even for scenarios with very high collaborativeness lev… are many negotiations which end without an agreement. I… this reflects weaknesses of the negotiation strategy, as the… tion breaks down while there were possible deals which would have been mutually acceptable.

The second experiment uses the internal negotiation deadline (IND) strategy for both agents. The results are shown in Figure 7. The shape of the graph is roughly similar to the MCS vs. MCS graph, however, the strategy shows better results for very low collaborativeness scenarios. For instance, for a collaborativeness of 0.025 the IND vs. IND settings obtains an average relative utility of 0.295, compared to 0.09 for the MCS vs. MCS case.

In the third experiment, we plot the same values for two agents using the estimating the opponents' parameters (EOP) strategy negotiating with each other. The results are shown in Figure 8, and the overall trend remains the same as in previous cases: the relative utility grows with the collaborativeness of the scenarios. We find that the overall relative utility of the EOP agents is much higher than the other two cases. Looking at the scatterplot, we can see that this was accomplished both through higher utility values for deals, as well as the reduction of the scenarios which ended in conflict.

Figure 9: The comparision of average relative profits both agents get when deadline agent negotiates with the estimating agent in different scenarios.

In the previous examples, we always had the same types of agents negotiating each other. Pitting agents of different types against each other can give as insight into their relative performance. Figure 9 shows the average relative utility function of collaborativeness graphs an IND agent negotiating with an EOP agent. The two plots represent the results for the same series of experiments seen from the point of view of the two agents (that is the $U_{rel}^{(IND),(EOP)}$ and the $U_{rel}^{(EOP),(IND)}$ values). For reference, we also added the plot of two IND agents negotiating each other. The overall shape of the curves is what we expected, the relative utility increases with the collaborativeness. However, the agent using the EOP strategy is able to consistently achieve higher utility values than the IND agent, which shows the (expected) superiority of the EST strategy in this setting.

An interesting observation can be made comparing the results of the IND agent versus the EOP agent and versus another IND agent. For low collaborativeness levels the IND strategy performs worse against the EOP strategy than against the EOP strategy, which we can interpret that EOP "outsmarted" its less sophisticated opponent, and pushed it into less advantageous deals. For high collaborativeness levels, however, the situation is reversed; for these scenarios the two agents have largely aligned interests, thus the efforts of EOP to find better deals for itself also improves the deals for the IND agent.

6. CONCLUSIONS

The contribution of this paper is the introduction of a set of metrics which allows us to measure the performance of a negotiation strategy in comparison to its peers, as well as to evaluate, without assuming any particular negotiation strategy whether a negotiation will be easy or difficult.

Measuring the performance of a negotiation strategy against a particular type of partner is not trivial. A weak performance in a particular scenario does not necessarily mean that the strategy is weak; it is possible that in the particular scenario "the odds were against the agent". The relative utility and competitive utility metrics allows us to evaluate the performance from a small number of negotiation runs in the cases when we are interested in the absolute value of gained utility or in the competitive advantage over the negotiation partner.

The collaborativeness metric we introduced allows us to put a quantitative value on the intuition of "easy" vs. "hard" negotiation scenarios. This measurement is independent of the strategy, but, naturally, some strategies might be better in exploiting the advantage of collaborative scenarios, or to retain acceptable utility in competitive scenarios. We had also seen some empirical evidence for the claim that negotiating against a sophisticated strategy leads to lower relative utility for scenarios with low collaborativeness levels, but it actually becomes an advantage when negotiating in scenarios with high collaborativeness levels.

One can envision strategies which perform a collaborativeness analysis on the scenarios they encounter and make a prediction of the success of the negotiation. One way to use this information might be agents which can "act while negotiating". If an agent is almost sure that a deal will be reached, it might start to move towards the predicted location of the deal while the negotiation is in progress. Alternatively, if an agent is able to choose which agent it will launch into negotiation for collaboration, the agent might pick the negotiation with the highest collaborativeness value as the negotiation most likely to succeed and reach a high relative utility.

Appendix: Generalized CRF

An important question which might be asked is whether the results in this paper can be applied to practical applications. Naturally, the CRF game itself does not map one-to-one to any practical application.

Let us now introduce a generalized version of the CRF game. As before, we assume that two agents A and B are moving from source points S_A and S_B to their respective destination points D_A and D_B. Their velocities in vector form are $\mathbf{v}_A = (v_{A,x}, v_{A,y})$ and $\mathbf{v}_B = (v_{B,x}, v_{B,y})$ respectively. They can form a coalition, in which case they have the common velocity $\mathbf{v}_{A+B}$. The velocity of the agents is subject of restrictions of the form $R_A(\mathbf{v}_A, x_A, y_A)$, $R_B(\mathbf{v}_B, x_B, y_B)$ and $R_{A+B}(\mathbf{v}_{A+B}, x_{A+B}, y_{A+B})$

The CRF game is a special version of the generalized CRF where:

$$R_A : \begin{cases} |\mathbf{v}_A| \leq c_A & \text{if } (x_A, y_A) \notin \text{forest} \\ |\mathbf{v}_A| = 0 & \text{if } (x_A, y_A) \in \text{forest} \end{cases} \tag{17}$$

$$R_{A+B} : |\mathbf{v}_{A+B}| \leq |\min(\mathbf{v}_A, \mathbf{v}_B)| \tag{18}$$

Notice, that this immediately generalizes to situations where the forest is of arbitrary size and position. It can also generalize to multiple obstacles, including those which produce only a slowdown as opposed to acting as a barrier. The approach also generalizes to cases where the agents are of different type. For instance, if A is a human, and B is a vehicle which the human can board, R_{A+B} is changes as follows:

$$R_{A+B} : |\mathbf{v}_{A+B}| \leq |\mathbf{v}_B| \tag{19}$$

We can also develop restrictions for cases when, for instance, B is a ferry, while A is a vehicle. In this case B and the coalition of A+B has mobility on water, while A has mobility on solid ground. These examples all have restrictions on the scalar component of the velocity; this, however, might not always be the case. We can model roads, where lateral movement is impossible or slower, or rivers where the speed of movement is greater in the direction of the flow.

The question is, how much of the results of this paper remain applicable to the generalized CRF? The metrics defined in Section 3.2 are applicable without modification. Similarly, the proposed negotiation protocols 3.3 remain valid and functional. The negotiation strategies, however, need to be adapted for the specific restriction models. In a generalized setting it is even more difficult to define what constitutes a concession. With a more complex map, offer generation becomes more expensive, which limits how many offers an agent can maintain in its offer pool. Nevertheless, the general outline of negotiation strategies will be still functional; it will need, however, to be adapted to cover the specifics of the restrictions.

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 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.

7. REFERENCES

[1] 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.

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

[3] 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.

[4] 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.

[5] 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.

[6] 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.

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

[8] 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.

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